

ADSP-BF51x Blackfin[®] Processor Hardware Reference

***(Includes ADSP-BF512, ADSP-BF514,
ADSP-BF516, ADSP-BF518)***

Revision 1.0, September 2010

Part Number
82-100109-01

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



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, 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	lxi
Intended Audience	lxi
What's New in This Manual	lxii
Technical or Customer Support	lxii
Product Information	lxiii
Analog Devices Web Site	lxiii
VisualDSP++ Online Documentation	lxiii
Technical Library CD	lxiv
EngineerZone	lxv
Social Networking Web Sites	lxv
Notation Conventions	lxvi

INTRODUCTION

Manual Contents	1-1
Peripherals	1-5
Memory Architecture	1-7
Internal Memory	1-9
External Memory	1-9

Contents

I/O Memory Space	1-9
One-Time-Programmable (OTP) Memory	1-10
DMA Support	1-11
External Bus Interface Unit	1-12
SDRAM Controller	1-12
Asynchronous Controller	1-13
General-Purpose I/O (GPIO)	1-13
Two-Wire Interface	1-14
Ethernet MAC	1-15
IEEE 1588 Support	1-16
RSI Interface	1-17
General-Purpose (GP) Counter	1-17
3-Phase PWM Unit	1-18
Parallel Peripheral Interface	1-19
SPORT Controllers	1-21
Serial Peripheral Interface (SPI) Ports	1-23
Timers	1-23
UART Ports	1-24
Security	1-25
Real-Time Clock	1-26
Watchdog Timer	1-27
Clock Signals	1-28

Dynamic Power Management	1-28
Full-On Mode (Maximum Performance)	1-28
Active Mode (Moderate Power Savings)	1-29
Sleep Mode (High Power Savings)	1-29
Deep Sleep Mode (Maximum Power Savings)	1-29
Hibernate State	1-30
Instruction Set Description	1-30
Development Tools	1-31

MEMORY

Memory Architecture	2-1
L1 Instruction SRAM	2-2
L1 Data SRAM	2-3
L1 Data Cache	2-4
Boot ROM	2-4
External Memory	2-5
Processor-Specific MMRs	2-5
DMEM_CONTROL Register	2-6
DTEST_COMMAND Register	2-6

ONE-TIME PROGRAMMABLE MEMORY

OTP Memory Overview	3-1
OTP Memory Map	3-2
Error Correction	3-7
Error Correction Policy	3-8

Contents

OTP Access	3-10
OTP Timing Parameters	3-11
OTP_TIMING Register	3-14
Callable ROM Functions for OTP ACCESS	3-14
Initializing OTP	3-14
bfrom_OtpCommand	3-15
Programming and Reading OTP	3-17
bfrom_OtpRead	3-17
bfrom_OtpWrite	3-18
Error Codes	3-22
Write-Protecting OTP Memory	3-24
Accessing Private OTP Memory	3-26
OTP Programming Examples	3-26
 CHIP BUS HIERARCHY	
Chip Bus Hierarchy Overview	4-1
Interface Overview	4-3
Internal Clocks	4-4
Core Bus Overview	4-4
Peripheral Access Bus (PAB)	4-6
PAB Arbitration	4-6
PAB Agents (Masters, Slaves)	4-6
PAB Performance	4-7

DMA Access Bus (DAB), DMA Core Bus (DCB), DMA External Bus (DEB)	4-8
DAB, DCB, and DEB Arbitration	4-8
DAB Bus Agents (Masters)	4-9
DAB, DCB, and DEB Performance	4-10
External Access Bus (EAB)	4-10
Arbitration of the External Bus	4-11
DEB/EAB Performance	4-11

SYSTEM INTERRUPTS

Specific Information for the ADSP-BF51x	5-1
Overview	5-1
Features	5-2
Description of Operation	5-2
Events and Sequencing	5-2
System Peripheral Interrupts	5-4
Programming Model	5-7
System Interrupt Initialization	5-8
System Interrupt Processing Summary	5-8
System Interrupt Controller Registers	5-10
System Interrupt Assignment (SIC_IAR) Register	5-11
System Interrupt Mask (SIC_IMASK) Register	5-12
System Interrupt Status (SIC_ISR) Register	5-12
System Interrupt Wakeup-Enable (SIC_IWR) Register	5-12

Contents

Programming Examples	5-13
Clearing Interrupt Requests	5-13
Unique Information for the ADSP-BF51x Processor	5-15
Interfaces	5-15
System Peripheral Interrupts	5-18

DIRECT MEMORY ACCESS

Specific Information for the ADSP-BF51x	6-1
Overview and Features	6-2
DMA Controller Overview	6-4
External Interfaces	6-4
Internal Interfaces	6-4
Peripheral DMA	6-5
Memory DMA	6-6
Handshaked Memory DMA (HMDMA) Mode	6-8
Modes of Operation	6-9
Register-Based DMA Operation	6-9
Stop Mode	6-11
Autobuffer Mode	6-11
Two-Dimensional DMA Operation	6-11
Examples of Two-Dimensional DMA	6-12
Descriptor-Based DMA Operation	6-13
Descriptor List Mode	6-14
Descriptor Array Mode	6-15

Variable Descriptor Size	6-15
Mixing Flow Modes	6-16
Functional Description	6-17
DMA Operation Flow	6-17
DMA Startup	6-17
DMA Refresh	6-22
Work Unit Transitions	6-24
DMA Transmit and MDMA Source	6-25
DMA Receive	6-26
Stopping DMA Transfers	6-28
DMA Errors (Aborts)	6-29
DMA Control Commands	6-31
Restrictions	6-34
Transmit Restart or Finish	6-34
Receive Restart or Finish	6-35
Handshaked Memory DMA Operation	6-36
Pipelining DMA Requests	6-37
HMDMA Interrupts	6-40
DMA Performance	6-40
DMA Throughput	6-42
Memory DMA Timing Details	6-44
Static Channel Prioritization	6-44
Temporary DMA Urgency	6-45

Contents

Memory DMA Priority and Scheduling	6-46
Traffic Control	6-48
Programming Model	6-50
Synchronization of Software and DMA	6-50
Single-Buffer DMA Transfers	6-53
Continuous Transfers Using Autobuffering	6-53
Descriptor Structures	6-55
Descriptor Queue Management	6-56
Descriptor Queue Using Interrupts on Every Descriptor	6-57
Descriptor Queue Using Minimal Interrupts	6-58
Software Triggered Descriptor Fetches	6-60
DMA Registers	6-62
DMA Channel Registers	6-63
DMA Peripheral Map Registers (DMAx_PERIPHERAL_MAP/ MDMA_yy_PERIPHERAL_MAP)	6-66
DMA Configuration Registers (DMAx_CONFIG/MDMA_yy_CONFIG)	6-67
DMA Interrupt Status Registers (DMAx_IRQ_STATUS/MDMA_yy_IRQ_STATUS)	6-71
DMA Start Address Registers (DMAx_START_ADDR/ MDMA_yy_START_ADDR)	6-74
DMA Current Address Registers (DMAx_CURR_ADDR/ MDMA_yy_CURR_ADDR)	6-74

DMA Inner Loop Count Registers (DMAx_X_COUNT/MDMA_yy_X_COUNT)	6-75
DMA Current Inner Loop Count Registers (DMAx_CURR_X_COUNT /MDMA_yy_CURR_X_COUNT)	6-76
DMA Inner Loop Address Increment Registers (DMAx_X_MODIFY/MDMA_yy_X_MODIFY)	6-77
DMA Outer Loop Count Registers (DMAx_Y_COUNT/MDMA_yy_Y_COUNT)	6-78
DMA Current Outer Loop Count Registers (DMAx_CURR_Y_COUNT/ MDMA_yy_CURR_Y_COUNT)	6-78
DMA Outer Loop Address Increment Registers (DMAx_Y_MODIFY/MDMA_yy_Y_MODIFY)	6-79
DMA Next Descriptor Pointer Registers (DMAx_NEXT_DESC_PTR/ MDMA_yy_NEXT_DESC_PTR)	6-80
DMA Current Descriptor Pointer Registers (DMAx_CURR_DESC_PTR/ MDMA_yy_CURR_DESC_PTR)	6-81
HMDMA Registers	6-82
Handshake MDMA Control Registers (HMDMAx_CONTROL)	6-82
Handshake MDMA Initial Block Count Registers (HMDMAx_BCINIT)	6-84
Handshake MDMA Current Block Count Registers (HMDMAx_BCOUNT)	6-85
Handshake MDMA Current Edge Count Registers (HMDMAx_ECOUNT)	6-86

Contents

Handshake MDMA Initial Edge Count Registers (HMDMAx_ECINIT)	6-87
Handshake MDMA Edge Count Urgent Registers (HMDMAx_ECURGENT)	6-87
Handshake MDMA Edge Count Overflow Interrupt Registers (HMDMAx_ECOVERFLOW)	6-88
DMA Traffic Control Registers (DMA_TC_PER and DMA_TC_CNT)	6-88
DMA_TC_PER Register	6-89
DMA_TC_CNT Register	6-89
Programming Examples	6-91
Register-Based 2-D Memory DMA	6-91
Initializing Descriptors in Memory	6-94
Software-Triggered Descriptor Fetch Example	6-97
Handshaked Memory DMA Example	6-100
Unique Information for the ADSP-BF51x Processor	6-102
DMA Control Commands	6-102
Static Channel Prioritization	6-102

EXTERNAL BUS INTERFACE UNIT

EBIU Overview	7-1
Block Diagram	7-4
Internal Memory Interfaces	7-4
Registers	7-5
Shared and Multiplexed Pins	7-6

System Clock	7-7
Error Detection	7-7
AMC Overview and Features	7-7
Features	7-8
Asynchronous Memory Interface	7-8
Asynchronous Memory Address Decode	7-9
AMC Pin Description	7-9
AMC Description of Operation	7-10
Avoiding Bus Contention	7-10
External Access Extension	7-11
AMC Functional Description	7-11
Programmable Timing Characteristics	7-11
Asynchronous Reads	7-12
Asynchronous Writes	7-13
Adding External Access Extension	7-15
Byte Enables	7-15
AMC Programming Model	7-17
AMC Registers	7-19
EBIU_AMGCTL Register	7-20
EBIU_AMBCTL0 and EBIU_AMBCTL1 Registers	7-20
AMC Programming Examples	7-23
SDC Overview and Features	7-24
Features	7-24
SDRAM Configurations Supported	7-25

Contents

SDRAM External Bank Size	7-26
SDC Address Mapping	7-26
Internal SDRAM Bank Select	7-27
Parallel Connection of SDRAMs	7-27
SDC Interface Overview	7-28
SDC Pin Description	7-28
SDRAM Performance	7-29
SDC Description of Operation	7-30
Definition of SDRAM Architecture Terms	7-31
Refresh	7-31
Row Activation	7-31
Column Read/Write	7-31
Row Precharge	7-31
Internal Bank	7-31
External Bank	7-32
Memory Size	7-32
Burst Length	7-32
Burst Type	7-32
CAS Latency	7-32
Data I/O Mask Function	7-33
SDRAM Commands	7-33
Mode Register Set (MRS) Command	7-33
Extended Mode Register Set (EMRS) Command	7-33
Bank Activate Command	7-33

Read/Write Command	7-33
Precharge/Precharge All Command	7-34
Auto-Refresh command	7-34
Enter Self-Refresh Mode	7-34
Exit Self-Refresh Mode	7-34
SDC Timing Specs	7-34
t_{MRD}	7-35
t_{RAS}	7-35
CL	7-35
t_{RCD}	7-36
t_{RRD}	7-36
t_{WR}	7-36
t_{RP}	7-37
t_{RC}	7-37
t_{RFC}	7-37
t_{XSR}	7-38
t_{REF}	7-38
t_{REFI}	7-38
SDC Functional Description	7-39
SDC Operation	7-39
SDC Address Muxing	7-41
Multibank Operation	7-42
Core and DMA Arbitration	7-43

Contents

Changing System Clock During Runtime	7-43
Changing Power Management During Runtime	7-45
Deep Sleep Mode	7-45
Hibernate State	7-45
SDC Commands	7-45
Mode Register Set Command	7-47
Extended Mode Register Set Command (Mobile SDRAM)	7-48
Bank Activation Command	7-48
Read/Write Command	7-49
Partial Write	7-49
Single Precharge Command	7-50
Precharge All Command	7-50
Auto-Refresh Command	7-51
Self-Refresh Mode	7-51
Self-Refresh Entry Command	7-51
Self-Refresh Exit Command	7-52
No Operation Command	7-53
SDC SA10 Pin	7-53
SDC Programming Model	7-53
SDC Configuration	7-54
Example SDRAM System Block Diagrams	7-55

SDC Register Definitions	7-57
EBIU_SDRRC Register	7-58
EBIU_SDBCTL Register	7-60
Using SDRAMs With Systems Smaller Than 16M Byte	7-62
EBIU_SDGCTL Register	7-64
SDRAM Clock Enable (SCTLE)	7-64
CAS Latency (CL)	7-66
Partial Array Self Refresh (PASR)	7-67
Bank Activate Command Delay (TRAS)	7-67
Bank Precharge Delay (TRP)	7-68
RAS to CAS Delay (TRCD)	7-68
Write to Precharge Delay (TWR)	7-69
Power-Up Start Delay (PUPSD)	7-69
Power-Up Sequence Mode (PSM)	7-69
Power-Up Sequence Start Enable (PSSE)	7-70
Self-Refresh Setting (SRFS)	7-71
Enter Self-Refresh Mode	7-71
Exit Self-Refresh Mode	7-72
External Buffering Enabled (EBUFE)	7-72
Fast Back-to-Back Read to Write (FBBRW)	7-73
Extended Mode Register Enabled (EMREN)	7-73
Temperature Compensated Self-Refresh (TCSR)	7-73
EBIU_SDSTAT Register	7-74
SDC Programming Examples	7-75

DYNAMIC POWER MANAGEMENT

Phase Locked Loop and Clock Control	8-1
PLL Overview	8-2
PLL Clock Multiplier Ratios	8-3
Core Clock/System Clock Ratio Control	8-5
Dynamic Power Management Controller	8-7
Operating Modes	8-7
Dynamic Power Management Controller States	8-8
Full-On Mode	8-8
Active Mode	8-8
Sleep Mode	8-9
Deep Sleep Mode	8-9
Hibernate State	8-10
Operating Mode Transitions	8-10
Programming Operating Mode Transitions	8-13
Dynamic Supply Voltage Control	8-15
Power Supply Management	8-15
Changing Voltage	8-16
Powering Down the Core (Hibernate State)	8-17
PLL and VR Registers	8-19
PLL_DIV Register	8-21
PLL_CTL Register	8-21
PLL_STAT Register	8-22

PLL_LOCKCNT Register	8-22
VR_CTL Register	8-23
System Control ROM Function	8-24
Programming Model	8-26
Accessing the System Control ROM Function in C/C++	8-26
Accessing the System Control ROM Function in Assembly	8-27
Programming Examples	8-30
Full-on Mode to Active Mode and Back	8-31
Transition to Sleep Mode or Deep Sleep Mode	8-33
Set Wakeups and Entering Hibernate State	8-35
Perform a System Reset or Soft-Reset	8-37
In Full-on Mode, Change VCO Frequency, Core Clock Frequency, and System Clock Frequency	8-38
Changing Voltage Levels	8-40

GENERAL-PURPOSE PORTS

Overview	9-1
Features	9-2
Interface Overview	9-3
External Interface	9-4
Port F Structure	9-4
Port G Structure	9-5
Port H Structure	9-6
Input Tap Considerations	9-7

Contents

PWM Unit Considerations	9-8
RSI Considerations	9-9
Internal Interfaces	9-10
SPI0 and Internal Flash Usage	9-10
GP Timer Interaction With Other Blocks	9-11
Buffered CLKIN (CLKBUF)	9-11
GP Counter	9-11
PPI	9-12
UART	9-12
SPORT	9-12
Performance/Throughput	9-13
Description of Operation	9-13
Operation	9-14
General-Purpose I/O Modules	9-14
GPIO Interrupt Processing	9-18
Programming Model	9-24
Memory-Mapped GPIO Registers	9-26
PORTx Hysteresis Control (PORTx_HYSTERESIS) Register	9-26
Non-GPIO Drive Strength Control Register	9-28
Non-GPIO Hysteresis (NONGPIO_HYSTERESIS) Register	9-29
Port Multiplexer Control Register (PORTF_MUX)	9-30
Port Multiplexer Control Register (PORTG_MUX)	9-31
Port Multiplexer Control Register (PORTH_MUX)	9-32

Function Enable Registers (PORTx_FER)	9-32
GPIO Direction Registers (PORTxIO_DIR)	9-33
GPIO Input Enable Registers (PORTxIO_INEN)	9-33
GPIO Data Registers (PORTxIO)	9-34
GPIO Set Registers (PORTxIO_SET)	9-34
GPIO Clear Registers (PORTxIO_CLEAR)	9-35
GPIO Toggle Registers (PORTxIO_TOGGLE)	9-35
GPIO Polarity Registers (PORTxIO_POLAR)	9-36
Interrupt Sensitivity Registers (PORTxIO_EDGE)	9-36
GPIO Set on Both Edges Registers (PORTxIO_BOTH)	9-37
GPIO Mask Interrupt Registers (PORTxIO_MASKA/B)	9-37
GPIO Mask Interrupt Set Registers (PORTxIO_MASKA/B_SET)	9-39
GPIO Mask Interrupt Clear Registers (PORTxIO_MASKA/B_CLEAR)	9-41
GPIO Mask Interrupt Toggle Registers (PORTxIO_MASKA/B_TOGGLE)	9-43
Programming Examples	9-44

GENERAL-PURPOSE TIMERS

Specific Information for the ADSP-BF51x	10-1
Overview	10-2
External Interface	10-3
Internal Interface	10-4

Contents

Description of Operation	10-4
Interrupt Processing	10-5
Illegal States	10-7
Modes of Operation	10-10
Pulse Width Modulation (PWM_OUT) Mode	10-10
Output Pad Disable	10-12
Single Pulse Generation	10-12
Pulse Width Modulation Waveform Generation	10-13
PULSE_HI Toggle Mode	10-15
Externally Clocked PWM_OUT	10-19
Using PWM_OUT Mode With the PPI	10-20
Stopping the Timer in PWM_OUT Mode	10-21
Pulse Width Count and Capture (WDTH_CAP) Mode	10-23
Autobaud Mode	10-31
External Event (EXT_CLK) Mode	10-31
Programming Model	10-33
Timer Registers	10-34
Timer Enable Register (TIMER_ENABLE)	10-35
Timer Disable Register (TIMER_DISABLE)	10-36
Timer Status Register (TIMER_STATUS)	10-37
Timer Configuration Register (TIMER_CONFIG)	10-38
Timer Counter Register (TIMER_COUNTER)	10-40
Timer Period (TIMER_PERIOD) and Timer Width (TIMER_WIDTH) Registers	10-41
Summary	10-44

Programming Examples 10-47

Unique Information for the ADSP-BF51x Processor 10-56

 Interface Overview 10-56

 External Interface 10-56

CORE TIMER

Specific Information for the ADSP-BF51x 11-1

Overview and Features 11-1

Timer Overview 11-2

 External Interfaces 11-2

 Internal Interfaces 11-3

Description of Operation 11-3

 Interrupt Processing 11-3

Core Timer Registers 11-4

 Core Timer Control Register (TCNTL) 11-5

 Core Timer Count Register (TCOUNT) 11-5

 Core Timer Period Register (TPERIOD) 11-6

 Core Timer Scale Register (TSCALE) 11-7

Programming Examples 11-7

Unique Information for the ADSP-BF51x Processor 11-9

WATCHDOG TIMER

Specific Information for the ADSP-BF51x 12-1

Overview and Features 12-1

Contents

Interface Overview	12-3
External Interface	12-3
Internal Interface	12-3
Description of Operation	12-4
Register Definitions	12-5
Watchdog Count (WDOG_CNT) Register	12-5
Watchdog Status (WDOG_STAT) Register	12-6
Watchdog Control (WDOG_CTL) Register	12-7
Programming Examples	12-8
Unique Information for the ADSP-BF51x Processor	12-11

GENERAL-PURPOSE COUNTER

Specific Information for the ADSP-BF51x	13-1
Overview	13-2
Features	13-2
Interface Overview	13-3
Description of Operation	13-4
Quadrature Encoder Mode	13-4
Binary Encoder Mode	13-5
Up/Down Counter Mode	13-6
Direction Counter Mode	13-6
Timed Direction Mode	13-7
Functional Description	13-7
Input Noise Filtering (Debouncing)	13-7
Zero Marker (Push Button) Operation	13-9

Boundary Comparison Modes	13-10
Control and Signaling Events	13-12
Illegal Gray/Binary Code Events	13-12
Up/Down Count Events	13-12
Zero-Count Events	13-13
Overflow Events	13-13
Boundary Match Events	13-13
Zero Marker Events	13-14
Capturing Timing Information	13-14
Capturing Time Interval Between Successive Counter Events	13-15
Capturing Counter Interval and CNT_COUNTER Read Timing	13-16
Programming Model	13-18
Registers	13-19
Counter Module Register Overview	13-19
Counter Configuration Register (CNT_CONFIG)	13-20
Counter Interrupt Mask Register (CNT_IMASK)	13-21
Counter Status Register (CNT_STATUS)	13-21
Counter Command Register (CNT_COMMAND)	13-22
Counter Debounce Register (CNT_DEBOUNCE)	13-24
Counter Count Value Register (CNT_COUNTER)	13-25
Counter Boundary Registers (CNT_MIN and CNT_MAX)	13-26

Contents

Programming Examples	13-27
Unique Information for the ADSP-BF51x Processor	13-38

PWM CONTROLLER

Specific Information for the ADSP-BF51x	14-1
Overview	14-1
General Operation	14-8
Functional Description	14-9
Three-Phase PWM Timing Unit and Dead Time Control Unit	14-10
PWM Switching Frequency (PWM_TM) Register	14-10
PWM Switching Dead Time (PWM_DT) Register	14-11
PWM Operating Mode (PWM_CTRL and PWM_STAT) Registers	14-12
PWM Duty Cycle (PWM_CHA, PWM_CHB, and PWM_CHC) Registers	14-14
Special Consideration for PWM Operation in Over-Modulation	14-19
Three-Phase PWM Timing Unit Operation	14-22
Effective PWM Accuracy	14-22
Switched Reluctance Mode	14-24
Output Control Unit	14-24
Crossover Feature	14-24
Mode Bits (POLARITY and SRMODE)	14-25

Output Enable Function	14-25
Brushless DC Motor (Electronically Commutated Motor) Control	14-26
Gate Drive Unit	14-27
High-Frequency Chopping	14-28
PWM Polarity Control	14-29
Output Control Feature Precedence	14-30
Switched Reluctance (SR) Mode	14-30
PWM Sync Operation	14-33
Internal PWM SYNC Generation	14-34
External PWM SYNC Generation	14-34
PWM Shutdown and Interrupt Control Unit	14-35
PWM Registers	14-36
PWM Control (PWM_CTRL) Register	14-37
PWM Status (PWM_STAT) Register	14-39
PWM Period (PWM_TM) Register	14-40
PWM Dead Time (PWM_DT) Register	14-41
PWM Chopping Control (PWM_GATE) Register	14-41
PWM Channel A, B, C Duty Control (PWM_CHA, PWM_CHB, PWM_CHC) Registers	14-42
PWM Crossover and Output Enable (PWM_SEG) Register	14-44
PWM Sync Pulse Width Control (PWM_SYNCWT) Register	14-46

Contents

PWM Channel AL, BL, CL Duty Control (PWM_CHAL, PWM_CHBL, PWM_CHCL) Registers	14-46
PWM Low Side Invert (PWM_LSI) Register	14-48
PWM Simulation Status (PWM_STAT2) Register	14-49
Unique Information for the ADSP-BF51x Processor	14-50
UART PORT CONTROLLERS	
Specific Information for the ADSP-BF51x	15-1
Overview	15-2
Features	15-2
Interface Overview	15-3
External Interface	15-3
Internal Interface	15-4
Description of Operation	15-4
UART Transfer Protocol	15-5
UART Transmit Operation	15-6
UART Receive Operation	15-7
IrDA Transmit Operation	15-8
IrDA Receive Operation	15-9
Interrupt Processing	15-11
Bit Rate Generation	15-12
Autobaud Detection	15-13

Programming Model	15-15
Non-DMA Mode	15-15
DMA Mode	15-17
Mixing Modes	15-18
UART Registers	15-19
UART Line Control (UART_LCR) Register	15-21
UART Modem Control (UART_MCR) Register	15-23
UART Line Status (UART_LSR) Register	15-24
UART Transmit Holding (UART_THR) Register	15-25
UART Receive Buffer (UART_RBR) Register	15-26
UART Interrupt Enable (UART_IER) Register	15-26
UART Interrupt Identification (UART_IIR) Register	15-28
UART Divisor Latch (UART_DLL and UART_DLH) Registers	15-30
UART Scratch (UART_SCR) Register	15-31
UART Global Control (UART_GCTL) Register	15-31
Programming Examples	15-32
Unique Information for the ADSP-BF51x Processor	15-42

TWO WIRE INTERFACE CONTROLLER

Specific Information for the ADSP-BF51x	16-1
Overview	16-2
Interface Overview	16-3
External Interface	16-4
Serial Clock Signal (SCL)	16-4

Contents

Serial Data Signal (SDA)	16-4
TWI Pins	16-5
Internal Interfaces	16-5
Description of Operation	16-6
TWI Transfer Protocols	16-6
Clock Generation and Synchronization	16-7
Bus Arbitration	16-8
Start and Stop Conditions	16-8
General Call Support	16-9
Fast Mode	16-10
Functional Description	16-10
General Setup	16-10
Slave Mode	16-11
Master Mode Clock Setup	16-12
Master Mode Transmit	16-12
Master Mode Receive	16-14
Repeated Start Condition	16-15
Transmit/Receive Repeated Start Sequence	16-15
Receive/Transmit Repeated Start Sequence	16-16
Clock Stretching	16-17
Clock Stretching During FIFO Underflow	16-17
Clock Stretching During FIFO Overflow	16-19
Clock Stretching During Repeated Start Condition	16-20
Programming Model	16-22

Register Descriptions	16-24
TWI CONTROL Register (TWI_CONTROL)	16-24
SCL Clock Divider Register (TWI_CLKDIV)	16-25
TWI Slave Mode Control Register (TWI_SLAVE_CTL)	16-26
TWI Slave Mode Address Register (TWI_SLAVE_ADDR)	16-28
TWI Slave Mode Status Register (TWI_SLAVE_STAT)	16-28
TWI Master Mode Control Register (TWI_MASTER_CTL)	16-29
TWI Master Mode Address Register (TWI_MASTER_ADDR)	16-32
TWI Master Mode Status Register (TWI_MASTER_STAT)	16-33
TWI FIFO Control Register (TWI_FIFO_CTL)	16-36
TWI FIFO Status Register (TWI_FIFO_STAT)	16-38
TWI FIFO Status	16-38
TWI Interrupt Mask Register (TWI_INT_MASK)	16-39
TWI Interrupt Status Register (TWI_INT_STAT)	16-40
TWI FIFO Transmit Data Single Byte Register (TWI_XMT_DATA8)	16-43
TWI FIFO Transmit Data Double Byte Register (TWI_XMT_DATA16)	16-43
TWI FIFO Receive Data Single Byte Register (TWI_RCV_DATA8)	16-44
TWI FIFO Receive Data Double Byte Register (TWI_RCV_DATA16)	16-45

Contents

Programming Examples	16-46
Master Mode Setup	16-46
Slave Mode Setup	16-51
Electrical Specifications	16-58
Unique Information for the ADSP-BF51x Processor	16-58

SPI-COMPATIBLE PORT CONTROLLER

Specific Information for the ADSP-BF51x	17-1
Overview	17-2
Features	17-2
Interface Overview	17-3
External Interface	17-4
SPI Clock Signal (SCK)	17-4
Master-Out, Slave-In (MOSI) Signal	17-5
Master-In, Slave-Out (MISO) Signal	17-5
SPI Slave Select Input Signal (SPISS)	17-6
SPI Slave Select Enable Output Signals	17-7
Slave Select Inputs	17-8
Use of FLS Bits in SPI_FLG for Multiple Slave SPI Systems	17-8
Internal Interfaces	17-10
DMA Functionality	17-10
Description of Operation	17-11
SPI Transfer Protocols	17-11
SPI General Operation	17-14

Clock Signals	17-15
Interrupt Output	17-16
Functional Description	17-16
Master Mode Operation (Non-DMA)	17-17
Transfer Initiation From Master (Transfer Modes)	17-18
Slave Mode Operation (Non-DMA)	17-19
Slave Ready for a Transfer	17-21
Programming Model	17-21
Beginning and Ending an SPI Transfer	17-21
Master Mode DMA Operation	17-23
Slave Mode DMA Operation	17-26
SPI Registers	17-33
SPI Baud Rate (SPI_BAUD) Register	17-33
SPI Control (SPI_CTL) Register	17-34
SPI Flag (SPI_FLG) Register	17-37
SPI Status (SPI_STAT) Register	17-39
Mode Fault Error (MODF)	17-40
Transmission Error (TXE)	17-41
Reception Error (RBSY)	17-41
Transmit Collision Error (TXCOL)	17-41
SPI Transmit Data Buffer (SPI_TDBR) Register	17-41
SPI Receive Data Buffer (SPI_RDBR) Register	17-42
SPI RDBR Shadow (SPI_SHADOW) Register	17-43

Contents

Programming Examples	17-43
Core-Generated Transfer	17-44
Initialization Sequence	17-44
Starting a Transfer	17-45
Post Transfer and Next Transfer	17-46
Stopping	17-47
DMA-Based Transfer	17-47
DMA Initialization Sequence	17-47
SPI Initialization Sequence	17-48
Starting a Transfer	17-50
Stopping a Transfer	17-50
Unique Information for the ADSP-BF51x Processor	17-52

SPI SERIAL FLASH

Memory Organization	18-4
Device Operation	18-4
Reset Mode	18-5
Status Register	18-7
Busy	18-8
Write Enable Latch (WEL)	18-8
Auto Address Increment (AAI)	18-9
Block-Protection (BP2, BP1, BP0)	18-9
Block-Protection Lock-Down (BPL)	18-9

Instructions	18-10
Read (20 MHz)	18-13
High-Speed-Read (25 MHz)	18-14
Byte-Program	18-15
Auto Address Increment (AAI) Word Program	18-16
End-of-Write Detection	18-17
Hardware End-of-Write Detection	18-17
Sector-Erase	18-20
32K Byte Block-Erase	18-21
64K Byte Block-Erase	18-22
Chip-Erase	18-23
Read-Status-Register (RDSR)	18-24
Write-Enable (WREN)	18-25
Write-Disable (WRDI)	18-26
Enable-Write-Status-Register (EWSR)	18-26
Write-Status-Register (WRSR)	18-27
Read-ID	18-28
JEDEC Read-ID	18-29

SPORT CONTROLLER

Specific Information for the ADSP-BF51x	19-1
Overview	19-2
Features	19-2
Interface Overview	19-4
SPORT Pin/Line Terminations	19-8

Contents

Description of Operation	19-9
SPORT Disable	19-9
Setting SPORT Modes	19-10
Stereo Serial Operation	19-10
Multichannel Operation	19-14
Multichannel Enable	19-16
Frame Syncs in Multichannel Mode	19-18
The Multichannel Frame	19-19
Multichannel Frame Delay	19-20
Window Size	19-20
Window Offset	19-20
Other Multichannel Fields in SPORT_MCMC2	19-21
Channel Selection Register	19-21
Multichannel DMA Data Packing	19-23
Support for H.100 Standard Protocol	19-23
2× Clock Recovery Control	19-24
Functional Description	19-24
Clock and Frame Sync Frequencies	19-24
Maximum Clock Rate Restrictions	19-26
Word Length	19-26
Bit Order	19-26
Data Type	19-27
Companding	19-27
Clock Signal Options	19-28

Frame Sync Options	19-29
Framed Versus Unframed	19-29
Internal Versus External Frame Syncs	19-31
Active Low Versus Active High Frame Syncs	19-32
Sampling Edge for Data and Frame Syncs	19-32
Early Versus Late Frame Syncs (Normal Versus Alternate Timing)	19-34
Data Independent Transmit Frame Sync	19-36
Moving Data Between SPORTs and Memory	19-37
SPORT RX, TX, and Error Interrupts	19-37
Peripheral Bus Errors	19-38
Timing Examples	19-38
SPORT Registers	19-44
Register Writes and Effective Latency	19-45
SPORT Transmit Configuration (SPORT_TCR1 and SPORT_TCR2) Registers	19-46
SPORT Receive Configuration (SPORT_RCR1 and SPORT_RCR2) Registers	19-51
Data Word Formats	19-55
SPORT Transmit Data (SPORT_TX) Register	19-56
SPORT Receive Data (SPORT_RX) Register	19-58
SPORT Status (SPORT_STAT) Register	19-60
SPORT Transmit and Receive Serial Clock Divider (SPORT_TCLKDIV and SPORT_RCLKDIV) Registers	19-61

Contents

SPORT Transmit and Receive Frame Sync Divider (SPORT_TFSDIV and SPORT_RFSDIV) Registers	19-62
SPORT Multichannel Configuration (SPORT_MCMC1 and SPORT_MCMC2) Registers	19-63
SPORT Current Channel (SPORT_CHNL) Register	19-64
SPORT Multichannel Receive Selection (SPORT_MRCSn) Registers	19-65
SPORT Multichannel Transmit Selection (SPORT_MTCSn) Registers	19-66
Programming Examples	19-67
SPORT Initialization Sequence	19-68
DMA Initialization Sequence	19-70
Interrupt Servicing	19-72
Starting a Transfer	19-73
Unique Information for the ADSP-BF51x Processor	19-74

PARALLEL PERIPHERAL INTERFACE

Specific Information for the ADSP-BF51x	20-1
Overview	20-2
Features	20-2
Interface Overview	20-3
Description of Operation	20-4
Functional Description	20-5
ITU-R 656 Modes	20-5
ITU-R 656 Background	20-5
ITU-R 656 Input Modes	20-9

Entire Field	20-9
Active Video Only	20-10
Vertical Blanking Interval (VBI) Only	20-10
ITU-R 656 Output Mode	20-11
Frame Synchronization in ITU-R 656 Modes	20-11
General-Purpose PPI Modes	20-12
Data Input (RX) Modes	20-14
No Frame Syncs	20-15
1, 2, or 3 External Frame Syncs	20-15
2 or 3 Internal Frame Syncs	20-16
Data Output (TX) Modes	20-17
No Frame Syncs	20-17
1 or 2 External Frame Syncs	20-18
1, 2, or 3 Internal Frame Syncs	20-18
Frame Synchronization in GP Modes	20-19
Modes With Internal Frame Syncs	20-19
Modes With External Frame Syncs	20-20
Programming Model	20-21
DMA Operation	20-22
PPI Registers	20-25
PPI Control Register (PPI_CONTROL)	20-25
PPI Status Register (PPI_STATUS)	20-29
PPI Delay Count Register (PPI_DELAY)	20-32

Contents

PPI Transfer Count Register (PPI_COUNT)	20-32
PPI Lines Per Frame Register (PPI_FRAME)	20-33
Programming Examples	20-34
Unique Information for the ADSP-BF51x Processor	20-37

REMOVABLE STORAGE INTERFACE

Overview	21-1
Interface Overview	21-2
Description of Operation	21-6
Functional Description	21-8
RSI Clock Configuration	21-9
RSI Interface Configuration	21-9
Card Detection	21-11
RSI Power Saving Configuration	21-13
RSI Commands and Responses	21-14
IDLE State	21-18
PEND State	21-19
SEND State	21-19
WAIT State	21-20
RECEIVE State	21-20
CEATA_INT_WAIT State	21-21
CEATA_INT_DIS State	21-21
RSI Command Path CRC	21-21
RSI Data	21-22
RSI Data Transmit Path	21-25

RSI Data Receive Path	21-26
RSI Data Path CRC	21-28
RSI Data FIFO	21-29
SDIO Interrupt and Read Wait Support	21-30
Programming Model	21-31
Card Identification	21-31
SD Card Identification Procedure	21-32
MMC Identification Procedure	21-34
Single Block Write Operations	21-34
Using Core	21-35
Using DMA	21-37
Single Block Read Operation	21-38
Using Core	21-39
Using DMA	21-41
Multiple Block Write Operation	21-43
Using Core	21-43
Using DMA	21-45
Multiple Block Read Operation	21-47
Using Core	21-48
Using DMA	21-50
RSI Registers	21-52
RSI Power Control Register (RSI_PWR_CONTROL)	21-53
RSI Clock Control Register (RSI_CLK_CONTROL)	21-54
RSI Argument Register (RSI_ARGUMENT)	21-56

Contents

RSI Command Register (RSI_COMMAND)	21-57
RSI Response Command Register (RSI_RESP_CMD)	21-58
RSI Response Registers (RSI_RESPONSEx)	21-59
RSI Data Timer Register (RSI_DATA_TIMER)	21-60
RSI Data Length Register (RSI_DATA_LGTH)	21-61
RSI Data Control Register (RSI_DATA_CONTROL)	21-62
RSI Data Counter Register (RSI_DATA_CNT)	21-63
RSI Status Register (RSI_STATUS)	21-64
RSI Status Clear Register (RSI_STATUSCL)	21-68
RSI Interrupt Mask Registers (RSI_MASKx)	21-70
RSI FIFO Counter Register (RSI_FIFO_CNT)	21-73
RSI CE-ATA Control Register (RSI_CEATA_CONTROL)	21-73
RSI Data FIFO Register (RSI_FIFO)	21-74
RSI Exception Status Register (RSI_ESTAT)	21-75
RSI Exception Mask Register (RSI_EMASK)	21-76
RSI Configuration Register (RSI_CONFIG)	21-77
RSI Read Wait Enable Register (RSI_RD_WAIT_EN)	21-79
RSI Peripheral ID Registers (RSI_PIDx)	21-80

ETHERNET MAC

Specific Information for the ADSP-BF51x	22-1
Overview	22-2
Features	22-2

Interface Overview	22-3
External Interface	22-4
Clocking	22-4
Pins	22-5
Internal Interface	22-7
Power Management	22-7
Description of Operation	22-7
Protocol	22-8
MII Management Interface	22-8
Operation	22-10
MII Management Interface Operation	22-10
Receive DMA Operation	22-11
Frame Reception and Filtering	22-13
RX Automatic Pad Stripping	22-17
RX DMA Data Alignment	22-17
RX DMA Buffer Structure	22-18
RX Frame Status Buffer	22-19
RX Frame Status Classification	22-19
RX IP Frame Checksum Calculation	22-21
RX DMA Direction Errors	22-22
Transmit DMA Operation	22-23
Flexible Descriptor Structure	22-26
TX DMA Data Alignment	22-27
Late Collisions	22-27

Contents

TX Frame Status Classification	22-28
TX DMA Direction Errors	22-29
Power Management	22-30
Ethernet Operation in the Sleep State	22-32
Magic Packet Detection	22-33
Remote Wake-up Filters	22-34
Ethernet Event Interrupts	22-38
RX/TX Frame Status Interrupt Operation	22-41
RX Frame Status Register Operation at Startup and Shutdown	22-41
TX Frame Status Register Operation at Startup and Shutdown	22-42
MAC Management Counters	22-42
Programming Model	22-45
Configure MAC Pins	22-45
Multiplexing Scheme	22-45
CLKBUF	22-46
Configure Interrupts	22-46
Configure MAC Registers	22-47
MAC Address	22-47
MII Station Management	22-47
Configure PHY	22-48
Receive and Transmit Data	22-49
Receiving Data	22-49
Transmitting Data	22-50

Ethernet MAC Register Definitions	22-50
Control-Status Register Group	22-59
MAC Operating Mode (EMAC_OPMODE) Register	22-60
MAC Address Low (EMAC_ADDRLO) Register	22-67
MAC Address High Register (EMAC_ADDRHI) Register	22-68
MAC Multicast Hash Table High (EMAC_HASHHI) and Low (EMAC_HASHLO) Registers	22-68
MAC Station Management Address (EMAC_STAADD) Register	22-72
MAC Station Management Data (EMAC_STADAT) Register	22-74
MAC Flow Control (EMAC_FLC) Register	22-74
MAC VLAN1 Tag (EMAC_VLAN1) and MAC VLAN2 Tag (EMAC_VLAN2)Registers	22-76
MAC Wakeup Frame Control and Status (EMAC_WKUP_CTL) Register	22-77
MAC Wakeup Frame0 Byte Mask (EMAC_WKUP_FFMSK0) MAC Wakeup Frame1 Byte Mask (EMAC_WKUP_FFMSK1) MAC Wakeup Frame2 Byte Mask (EMAC_WKUP_FFMSK2) MAC Wakeup Frame3 Byte Mask (EMAC_WKUP_FFMSK3) Registers	22-80
MAC Wakeup Frame Filter Commands (EMAC_WKUP_FFCMD) Register	22-85
Ethernet MAC Wakeup Frame Filter Offsets (EMAC_WKUP_FFOFF) Register	22-87
MAC Wakeup Frame Filter CRC0/1 (EMAC_WKUP_FFCRC0) and CRC2/3 (EMAC_WKUP_FFCRC1) Registers	22-87

Contents

System Interface Register Group	22-88
MAC System Control (EMAC_SYSCTL) Register	22-89
MAC System Status (EMAC_SYSTAT) Register	22-90
Ethernet MAC Frame Status Registers	22-92
Ethernet MAC RX Current Frame Status (EMAC_RX_STAT) Register	22-93
Ethernet MAC RX Sticky Frame Status (EMAC_RX_STKY) Register	22-99
Ethernet MAC RX Frame Status Interrupt Enable (EMAC_RX_IRQE) Register	22-104
Ethernet MAC TX Current Frame Status (EMAC_TX_STAT) Register	22-105
Ethernet MAC TX Sticky Frame Status (EMAC_TX_STKY) Register	22-109
Ethernet MAC TX Frame Status Interrupt Enable (EMAC_TX_IRQE) Register	22-112
Ethernet MAC MMC RX Interrupt Status (EMAC_MMC_RIRQS) Register	22-112
Ethernet MAC MMC RX Interrupt Enable (EMAC_MMC_RIRQE) Register	22-114
Ethernet MAC MMC TX Interrupt Status (EMAC_MMC_TIRQS) Register	22-114
Ethernet MAC MMC TX Interrupt Enable (EMAC_MMC_TIRQE) Register	22-117
MAC Management Counter Registers	22-117
MAC Management Counters Control (EMAC_MMC_CTL) Register	22-119

Programming Examples	22-120
Ethernet Structures	22-121
MAC Address Setup	22-124
PHY Control Routines	22-124
Unique Information for the ADSP-BF51x Processor	22-127

IEEE 1588 PTP ENGINE

PTP_TSYNC Overview	23-1
Features	23-1
General Operation	23-2
PTP_TSYNC Module Description of Operation	23-4
Clock Source Selection	23-4
Clock Adjustment	23-4
Event Message (Timestamping)	23-6
Transmit Packet Detection	23-7
Receive Packet Detection	23-8
Alarm	23-9
Pulse-Per-Second (PPS)	23-10
Auxiliary Snapshot	23-10
PTP_TSYNC Module Registers	23-11
Control Register (EMAC_PTP_CTL)	23-13
Interrupt Enable Register (EMAC_PTP_IE)	23-15
Interrupt Status Register (EMAC_PTP_ISTAT)	23-17
Message Filter Offset Register (EMAC_PTP_FOFF)	23-18
Message Filter Value Register 1 (EMAC_PTP_FV1)	23-20

Contents

Message Filter Value Register 2 (EMAC_PTP_FV2)	23-21
Message Filter Value Register 3 (EMAC_PTP_FV3)	23-21
Addend Register (EMAC_PTP_ADDEND)	23-23
Accumulator Register (EMAC_PTP_ACCR)	23-24
Time Offset Register (EMAC_PTP_OFFSET)	23-24
Local Clock Time Low Register (EMAC_PTP_TIMELO)	23-25
Local Clock Time High Register (EMAC_PTP_TIMEHI)	23-26
Receive Snapshot Low Register (EMAC_PTP_RXSNAPLO)	23-27
Receive Snapshot High Register (EMAC_PTP_RXSNAPHI)	23-28
Transmit Snapshot Low Register (EMAC_PTP_TXSNAPLO)	23-29
Transmit Snapshot High Register (EMAC_PTP_TXSNAPHI)	23-30
Target Alarm Time Low Register (EMAC_PTP_ALARMLO)	23-31
Target Alarm Time High Register (EMAC_PTP_ALARMHI)	23-32
Source ID Offset Register (EMAC_PTP_ID_OFF)	23-33
Source ID Snapshot Register (EMAC_PTP_ID_SNAP)	23-34
PPS Start Low Register (EMAC_PTP_PPS_STARTLO)	23-35
PPS Start High Register (EMAC_PTP_PPS_STARTHI)	23-36
PPS Period Register (EMAC_PTP_PPS_PERIOD)	23-37

PTP_TSYNC Module Programming Model	23-37
IEEE 1588-2002 Implementation Over IP/UDP	23-38
IEEE 1588-2008 Implementation Over IP/UDP	23-39
IEEE 1588-2008 Implementation Over MAC Layer	23-39
Pulse-Per-Second (PPS) Signal Generation	23-40

REAL-TIME CLOCK

Specific Information for the ADSP-BF51x	24-1
Overview	24-1
Interface Overview	24-3
Description of Operation	24-3
RTC Clock Requirements	24-3
Prescaler Enable	24-4
RTC Programming Model	24-6
Register Writes	24-8
Write Latency	24-8
Register Reads	24-9
Deep Sleep	24-9
Event Flags	24-10
Setting Time of Day	24-12
Using the Stopwatch	24-13
Interrupts	24-14
State Transitions Summary	24-16

Contents

Register Definitions	24-18
RTC Status (RTC_STAT) Register	24-20
RTC Interrupt Control (RTC_ICTL) Register	24-20
RTC Interrupt Status (RTC_ISTAT) Register	24-21
RTC Stopwatch Count (RTC_SWCNT) Register	24-21
RTC Alarm (RTC_ALARM) Register	24-22
RTC Prescaler Enable (RTC_PREN) Register	24-22
Programming Examples	24-23
Enable RTC Prescaler	24-23
RTC Stopwatch For Exiting Deep Sleep Mode	24-24
RTC Alarm to Come Out of Hibernate State	24-26
Unique Information for the ADSP-BF51x Processor	24-27

SECURITY

Overview	25-1
Features	25-4
Description of Operation	25-6
Secure State Machine	25-7
Open Mode	25-8
Secure Entry Mode	25-8
Secure Mode	25-9
Secure Mode Control	25-10

Security Features	25-12
Digital Signature Authentication	25-13
Digital Signature Authentication Performance Measurement	25-15
Protection Features	25-16
Operating in Secure Mode	25-19
Entering Secure Mode	25-20
Exiting Secure Mode	25-20
Reset Handling in Secure Mode	25-20
Hardware Reset	25-20
Clearing Private Data	25-21
Public Key Requirements	25-23
Storing Public Cipher Key in Public OTP	25-25
Cryptographic Ciphers	25-26
Keys	25-26
Debug Functionality	25-26
Programming Examples	25-31
Programming Model	25-32
Secure Entry Service Routine (SESR) API	25-32
Starting Authentication	25-33
Memory Configuration	25-34
Message Placement	25-35
Digital Signature	25-35
Message Size Constraints	25-35

Contents

Memory Usage	25-36
Memory Protection	25-36
Secure Function and Secure Entry Service Routine	
Arguments	25-36
Secure Function Arguments	25-37
Secure Entry Service Routine Arguments	25-38
usFlags	25-38
uslRQMask	25-39
ulMessageSize	25-40
ulSFEntryPoint	25-40
ulMessagePtr	25-40
Secure Message Execution	25-40
Return Codes	25-41
SECURE HASH ALGORITHM (SHA-1) API	25-43
ADI_SHA1 Data Type	25-43
bfrom_Sha1Init ROM Routine	25-44
bfrom_Sha1Hash ROM Routine	25-44
Security Registers	25-45
Secure System Switch (SECURE_SYSSWT) Register	25-45
Secure Control (SECURE_CONTROL) Register	25-53
Secure Status (SECURE_STATUS) Register	25-55

SYSTEM RESET AND BOOTING

Overview	26-1
Reset and Power-up	26-3
Hardware Reset	26-5
Software Resets	26-6
Reset Vector	26-7
Servicing Reset Interrupts	26-7
Preboot	26-9
Factory Page Settings (FPS)	26-10
Preboot Page Settings (PBS)	26-13
Alternative PBS Pages	26-14
Programming PBS Pages	26-15
Recovering From Misprogrammed PBS Pages	26-15
Customizing Power Management	26-16
Customizing Booting Options	26-17
Customizing the Asynchronous Port	26-18
Customizing the Synchronous Port	26-18
Basic Booting Process	26-19
Block Headers	26-21
Block Code	26-22
DMA Code Field	26-22
Block Flags Field	26-25
Header Checksum Field	26-26
Header Sign Field	26-27

Contents

Target Address	26-27
Byte Count	26-28
Argument	26-29
Boot Host Wait (HWAIT) Feedback Strobe	26-29
Using HWAIT as Reset Indicator	26-30
Boot Termination	26-31
Single Block Boot Streams	26-32
Direct Code Execution	26-32
Advanced Boot Techniques	26-35
Initialization Code	26-35
Quick Boot	26-39
Indirect Booting	26-40
Callback Routines	26-41
Error Handler	26-44
CRC Checksum Calculation	26-44
Load Functions	26-44
Calling the Boot Kernel at Runtime	26-46
Debugging the Boot Process	26-46
Boot Management	26-49
Booting a Different Application	26-49
Multi-DXE Boot Streams	26-50
Determining Boot Stream Start Addresses	26-55
Initialization Hook Routine	26-55

Specific Boot Modes	26-56
No Boot Mode	26-57
Flash Boot Modes	26-57
SDRAM Boot Mode	26-60
SPI Master Boot Modes	26-60
SPI Device Detection Routine	26-63
SPI Slave Boot Mode	26-65
UART Slave Mode Boot	26-68
OTP Boot Mode	26-70
Reset and Booting Registers	26-71
Software Reset (SWRST) Register	26-71
System Reset Configuration (SYSCR) Register	26-73
Boot Code Revision Control (BK_REVISION)	26-75
Boot Code Date Code (BK_DATECODE)	26-76
Zero Word (BK_ZEROS)	26-77
Ones Word (BK_ONES)	26-78
OTP Memory Pages for Booting	26-78
Lower PBS00 Half Page	26-78
Upper PBS00 Half Page	26-81
Lower PBS01 Half Page	26-82
Upper PBS01 Half Page	26-82
Lower PBS02 Half Page	26-85
Upper PBS02 Half Page	26-86
Reserved Half Pages	26-86

Contents

Data Structures	26-86
ADI_BOOT_HEADER	26-86
ADI_BOOT_BUFFER	26-87
ADI_BOOT_DATA	26-87
dFlags Word	26-91
Callable ROM Functions for Booting	26-92
BFROM_FINALINIT	26-92
BFROM_PDMA	26-93
BFROM_MDMA	26-93
BFROM_MEMBOOT	26-93
BFROM_SPIBOOT	26-95
BFROM_OTPBOOT	26-97
BFROM_BOOTKERNEL	26-98
BFROM_CRC32	26-99
BFROM_CRC32POLY	26-99
BFROM_CRC32CALLBACK	26-100
BFROM_CRC32INITCODE	26-100
Programming Examples	26-101
System Reset	26-101
Exiting Reset to User Mode	26-102
Exiting Reset to Supervisor Mode	26-103
Initcode (SDRAM Controller Setup)	26-104
Initcode (Power Management Control)	26-106
Quickboot With Restore From SDRAM	26-109

XOR Checksum	26-110
Direct Code Execution	26-111
Managing PBS Pages in OTP Memory	26-113

SYSTEM DESIGN

Pin Descriptions	27-1
Managing Clocks	27-1
Managing Core and System Clocks	27-2
Configuring and Servicing Interrupts	27-2
Semaphores	27-2
Example Code for Query Semaphore	27-3
Data Delays, Latencies and Throughput	27-4
Bus Priorities	27-4
External Memory Design Issues	27-5
Example Asynchronous Memory Interfaces	27-5
Avoiding Bus Contention	27-7
High-Frequency Design Considerations	27-8
Signal Integrity	27-8
Decoupling Capacitors and Ground Planes	27-9
5 Volt Tolerance	27-11
Test Point Access	27-12
Oscilloscope Probes	27-12
Recommended Reading	27-12
Resetting the Processor	27-13
Recommendations for Unused Pins	27-14

Contents

Programmable Outputs	27-14
Voltage Regulation Interface	27-14

SYSTEM MMR ASSIGNMENTS

Processor-Specific Memory Registers	A-3
Core Timer Registers	A-3
System Reset and Interrupt Control Registers	A-4
DMA/Memory DMA Control Registers	A-5
Handshake MDMA Control Registers	A-7
External Bus Interface Unit Registers	A-9
Ports Registers	A-9
Timer Registers	A-13
Watchdog Timer Registers	A-15
GP Counter Registers	A-15
Real-Time Clock Registers	A-16
OTP and Security Registers	A-17
Dynamic Power Management Registers	A-17
Ethernet MAC Registers	A-18
IEEE 1588 PTP Registers	A-23
PPI Registers	A-25
SPI Controller Registers	A-25
SPORT Controller Registers	A-27
UART Controller Registers	A-30
Motor Control PWM Registers	A-31

Removable Storage Interface (RSI) Registers	A-32
TWI Registers	A-34

TEST FEATURES

JTAG Standard	B-1
Boundary-Scan Architecture	B-2
Instruction Register	B-4
Public Instructions	B-6
EXTEST – Binary Code 00000	B-6
SAMPLE/PRELOAD – Binary Code 10000	B-6
BYPASS – Binary Code 11111	B-6
Boundary-Scan Register	B-7

INDEX

Contents

PREFACE

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

Purpose of This Manual

The *ADSP-BF51x Blackfin Processor Hardware Reference* provides architectural information about the ADSP-BF512, ADSP-BF514, ADSP-BF516, ADSP-BF518 processors. This hardware reference provides architectural information about these processors and the peripherals contained within the ADSP-BF51x Blackfin packages. The architectural descriptions cover functional blocks, buses, and ports, including all features and processes that they support. For programming information, see the *Blackfin Processor Programming Reference*. For timing, electrical, and package specifications, see the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

Intended Audience

The primary audience for this manual is a programmer who is familiar with Analog Devices 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 instruction set reference manuals and data sheets) that describe your target architecture.

What's New in This Manual

This is Revision 1.0 of the ADSP-BF51x Blackfin Processor Hardware Reference. Peripheral chapters have been expanded and reorganized, and modifications and corrections based on errata reports against this manual have been made.

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

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

Product Information

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.

EngineerZone

EngineerZone is a technical support forum from Analog Devices. It allows you direct access to ADI technical support engineers. You can search FAQs and technical information to get quick answers to your embedded processing and DSP design questions.

Use EngineerZone to connect with other DSP developers who face similar design challenges. You can also use this open forum to share knowledge and collaborate with the ADI support team and your peers. Visit <http://ez.analog.com> to sign up.

Social Networking Web Sites

You can now follow Analog Devices processor development on Twitter and LinkedIn. To access:

- Twitter: <http://twitter.com/blackfin>
- LinkedIn: Network with the LinkedIn group, Analog Devices Blackfin: <http://www.linkedin.com>

Notation Conventions

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

Example	Description
Close command (File menu)	Titles in reference sections indicate the location of an item within the VisualDSP++ environment's menu system (for example, the Close command appears on the File menu).
{this that}	Alternative required items in syntax descriptions appear within curly brackets and separated by vertical bars; read the example as <i>this</i> or <i>that</i> . One or the other is required.
[this that]	Optional items in syntax descriptions appear within brackets and separated by vertical bars; read the example as an optional <i>this</i> or <i>that</i> .
[this,...]	Optional item lists in syntax descriptions appear within brackets delimited by commas and terminated with an ellipse; read the example as an optional comma-separated list of <i>this</i> .
.SECTION	Commands, directives, keywords, and feature names are in text with letter gothic font.
<i>filename</i>	Non-keyword placeholders appear in text with italic style format.
	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.
	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 the devices users. In the online version of this book, the word Warning appears instead of this symbol.

1 INTRODUCTION

The ADSP-BF51x processors are members of the Blackfin processor family that offer significant high performance and low power features while retaining their ease-of-use benefits. All parts within the family are pin-compatible. However, the ADSP-BF512 does not include the RSI module, the ADSP-BF512 and ADSP-BF514 do not include an Ethernet MAC, and the ADSP-BF516 does not include an IEEE-1588 block.

Manual Contents

This manual consists of one volume:

- [Chapter 1, “Introduction”](#)
Provides a high level overview of the processor, including peripherals, power management, and development tools.
- [Chapter 2, “Memory”](#)
Describes processor-specific memory topics, including L1 memories and processor-specific memory MMRs.
- [Chapter 3, “One-Time Programmable Memory”](#)
Describes the one-time-programmable memory features.
- [Chapter 4, “Chip Bus Hierarchy”](#)
Describes on-chip buses, including how data moves through the system.

Manual Contents

- [Chapter 5, “System Interrupts”](#)
Describes the system peripheral interrupts, including setup and clearing of interrupt requests.
- [Chapter 6, “Direct Memory Access”](#)
Describes the peripheral DMA and Memory DMA controllers. Includes performance, software management of DMA, and DMA errors.
- [Chapter 7, “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 8, “Dynamic Power Management”](#)
Describes the clocking, including the PLL, and the dynamic power management controller.
- [Chapter 9, “General-Purpose Ports”](#)
Describes the general-purpose I/O ports, including the structure of each port, multiplexing, configuring the pins, and generating interrupts.
- [Chapter 10, “General-Purpose Timers”](#)
Describes the eight general-purpose timers.
- [Chapter 11, “Core Timer”](#)
Describes the core timer.
- [Chapter 12, “Watchdog Timer”](#)
Describes the watchdog timer.

- [Chapter 13, “General-Purpose Counter”](#)
Describes the Rotary (up/down) Counter. This counter provides support for manually controlled rotary controllers, such as the volume wheel on a radio device. This unit also supports industrial or motor-control type of wheels.
- [Chapter 14, “PWM Controller”](#)
Describes the programmable, three-phase PWM waveform generator which can generate switching patterns to drive a three-phase voltage source inverter for ac induction motor (ACIM) or permanent magnet synchronous motor (PMSM) control.
- [Chapter 15, “UART Port Controllers”](#)
Describes the Universal Asynchronous Receiver/Transmitter port that converts data between serial and parallel formats. The UART supports the half-duplex IrDA® SIR protocol as a mode-enabled feature.
- [Chapter 16, “Two Wire Interface Controller”](#)
Describes the Two Wire Interface (TWI) controller, which allows a device to interface to an Inter IC bus as specified by the *Philips I²C Bus Specification version 2.1* dated January 2000.
- [Chapter 17, “SPI-Compatible Port Controller”](#)
Describes the Serial Peripheral Interface (SPI) port that provides an I/O interface to a variety of SPI compatible peripheral devices.
- [Chapter 18, “SPI Serial Flash”](#)
Describes the SPI serial flash memory contained within the package of the processor and connected to the SPI0.
- [Chapter 19, “SPORT Controller”](#)
Describes the independent, synchronous Serial Port Controller which provides an I/O interface to a variety of serial peripheral devices.

Manual Contents

- [Chapter 20, “Parallel Peripheral Interface”](#)
Describes the Parallel Peripheral Interface (PPI) of the processor. The PPI is a half-duplex, bidirectional port accommodating up to 16 bits of data and is used for digital video and data converter applications.
- [Chapter 21, “Removable Storage Interface”](#)
Describes the RSI interface for multimedia cards (MMC), secure digital memory cards (SD), secure digital input/output cards (SDIO) and consumer electronic ATA devices (CE-ATA).
- [Chapter 22, “Ethernet MAC”](#)
Describes the Ethernet Media Access Controller (MAC) peripheral that provides a 10/100M bit/s Ethernet interface, compliant to IEEE Std. 802.3-2002, between an MII (Media Independent Interface) and the Blackfin peripheral subsystem.
- [Chapter 23, “IEEE 1588 PTP Engine”](#)
Describes the IEEE 1588 engine module (PTP_TSYNC for ADSP-BF518 processors) and the module’s operation.
- [Chapter 24, “Real-Time Clock”](#)
The RTC provides a set of digital watch features to the processor, including time of day, alarm, and stopwatch countdown. It is typically used to implement either a real-time watch or a life counter, which counts the elapsed time since the last system reset.
- [Chapter 25, “Security”](#)
Describes the Lockbox™ Secure Technology for Analog Devices Blackfin processors. This comprises a mix of hardware and software mechanisms designed to prevent unauthorized accesses and allow trusted code to execute on the processor.
- [Chapter 26, “System Reset and Booting”](#)
Describes the booting methods, booting process and specific boot modes for the processor.

- [Chapter 27, “System Design”](#)
Describes how to use the processor as part of an overall system. It includes information about bus timing and latency numbers, semaphores, and a discussion of the treatment of unused pins.
- [Appendix A, “System MMR Assignments”](#)
Lists the memory-mapped registers included in this manual, their addresses, and cross-references to text.
- [Appendix B, “Test Features”](#)
Describes test features for the processor, discusses the JTAG standard, boundary-scan architecture, instruction and boundary registers, and public instructions.



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

Peripherals

The processor system peripherals include:

- Two memory-to-memory DMAs with handshake DMA
- Event handler with 56 interrupt inputs
- 12 peripheral DMAs (2 mastered by the Ethernet MAC)
- Removable Storage Interface (RSI) (not available on ADSP-BF512)
- 40 General-Purpose I/Os (GPIOs)
- Eight 32-bit timer/counters with PWM support
- 32-bit core timer
- Real-Time Clock (RTC) and watchdog timer
- Rotary counter

Peripherals

- Lockbox security controller
- One-time Programmable (OTP) Memory
- On-chip PLL capable of 0.5× to 64× frequency multiplication
- Debug/JTAG interface
- IEEE 802.3-compliant 10/100 Ethernet MAC (only on the ADSP-BF516 and ADSP-BF518)
- IEEE-1588 precision clock synchronization protocol for 10/100 Ethernet MAC (only on the ADSP-BF518)
- Parallel Peripheral Interface (PPI), supporting ITU-R 656 video data formats
- Two Serial Peripheral Interface (SPI)-compatible ports
- Two-Wire Interface (TWI) controller
- Two dual-channel, full-duplex synchronous Serial Ports (SPORTs), supporting eight stereo I²S channels
- Two UARTs with IrDA® support
- 3-phase PWM generation unit

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

Most of the peripherals are supported by a flexible DMA structure. There are also two separate memory DMA channels dedicated to data transfers between the processor'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.

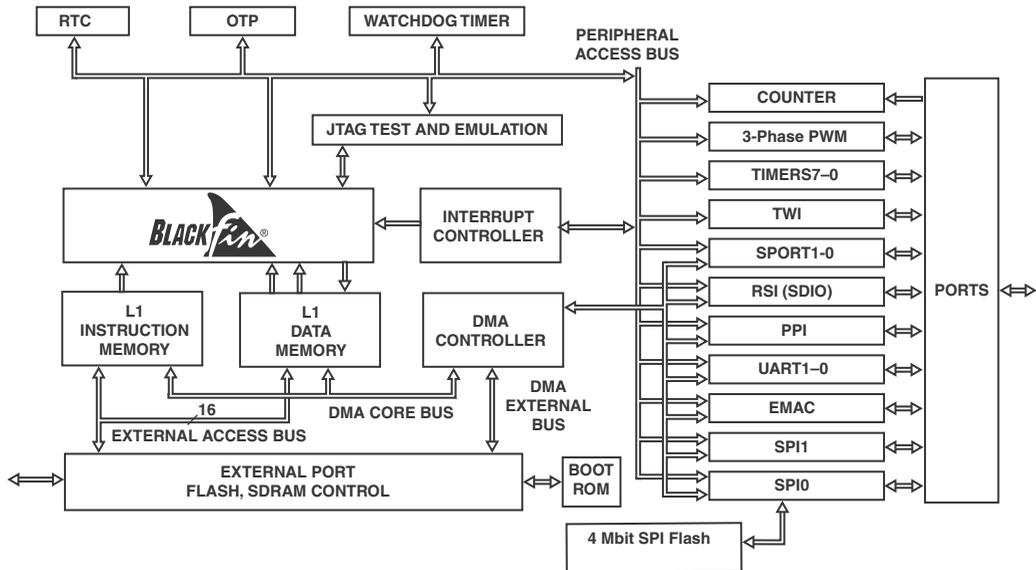


Figure 1-1. ADSP-BF51x Processor Block Diagram

Memory Architecture

The Blackfin processor architecture structures memory as a single, unified 4G byte 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. [Table 1-1](#) shows the memory for the ADSP-BF51x processors.

Memory Architecture

Table 1-1. Memory Configurations

Type of Memory	ADSP-BF51x
Instruction SRAM/cache, lockable by way or line	16K byte
Instruction SRAM	32K byte
Data SRAM/cache	32K byte
Data SRAM	32K byte
Data scratchpad SRAM	4K byte
L3 Boot ROM	32K byte
Total	148K byte

The L1 memory system is the primary highest performance memory available to the core. The off-chip memory system, accessed through the external bus interface unit (EBIU), provides expansion with SDRAM, flash memory, and SRAM, optionally accessing up to 132M bytes of physical memory.

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.

Internal Memory

The processor has three blocks of on-chip memory that provide high bandwidth access to the core:

- L1 instruction memory, consisting of SRAM and a 4-way set-associative cache. This memory is accessed at full processor speed.
- L1 data memory, consisting of SRAM and/or a 2-way set-associative cache. This memory block is accessed at full processor speed.
- L1 scratchpad RAM, which runs at the same speed as the L1 memories but is only accessible as data SRAM and cannot be configured as cache memory.

External Memory

External (off-chip) memory is accessed via the external bus interface unit (EBIU). This 16-bit interface provides a glueless connection to a bank 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 SDRAM controller can be programmed to interface to up to 128M bytes of SDRAM.

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

I/O Memory Space

Blackfin processors do not define a separate I/O space. All resources are mapped through the flat 32-bit address space. Control registers for on-chip I/O devices are mapped into memory-mapped registers (MMRs)

Memory Architecture

at addresses near the top of the 4G byte 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.

One-Time-Programmable (OTP) Memory

ADSP-BF51x processors also include an on-chip OTP memory array which provides 64K bits of non-volatile memory that can be programmed by the developer one time only. It includes the array and logic to support read access and programming. A mechanism for error correction is provided. Additionally, its pages can be write protected.

The OTP is not part of the Blackfin linear memory map. OTP memory is not accessed directly using the Blackfin memory map; rather, it is accessed via four 32-bit-wide registers that act as the OTP memory read/write buffer.

This memory is organized into 512 pages, each comprised of 128 bits and equally separated into two distinct areas with privileged access dependant upon modes of operation when security features are utilized. Approximately 400 pages are available for developer use. The remaining 100 pages are utilized for page protection bits, error correction, and Analog Devices factory-reserved areas. One area is read/write accessible at all time (Public OTP Memory). The second area maintains privileged access and can only be accessed (read/write) upon entry to Secure Mode when security features are utilized (Private OTP Memory).

All together, OTP memory provides a means to store Public Keys in Public OTP Memory or secrets such as Private Keys or Symmetric Keys in Private OTP Memory. One page of the Public OTP Memory is initialized in the Analog Devices factory with a Unique Chip ID.

This OTP memory provides a means to store public and private cipher keys as well as chip, customer, and factory identification data.

DMA Support

The processor has a DMA controller which supports automated data transfers with minimal overhead for the core. DMA transfers can occur between the 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 ports, UARTs, RSI, Ethernet, and PPI. Each individual DMA-capable peripheral has at least one dedicated DMA channel.

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

The 2-D 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

External Bus Interface Unit

In addition to the dedicated peripheral DMA channels, there are two separate pairs of memory DMA channels 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.

The ADSP-BF51x processors also include a handshake DMA capability via dual external DMA request pins when used in conjunction with the external bus interface unit (EBIU). This functionality can be used when a high speed interface is required for external FIFOs and high bandwidth communications peripherals such as USB 2.0. It allows control of the number of data transfers for MDMA. The number of transfers per edge is programmable. This feature can be programmed to allow MDMA to have an increased priority on the external bus relative to the core.

External Bus Interface Unit

The external bus interface unit (EBIU) on the processor interfaces with a wide variety of industry-standard memory devices. The controller consists of an SDRAM controller and an asynchronous memory controller.

SDRAM Controller

The SDRAM controller provides an interface to a single bank of industry-standard SDRAM devices or DIMMs. The bank can be configured to contain between 16M and 128M bytes of memory.

A set of programmable timing parameters is available to configure the SDRAM bank to support slower memory devices. The memory bank is 16 bits wide for minimum device count and lower system cost.

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 1M byte window in the processor address space, but if not fully populated, these are not made contiguous by the memory controller. The banks are 16 bits wide, for interfacing to a range of memories and I/O devices.

General-Purpose I/O (GPIO)

The ADSP-BF51x processors have 40 bi-directional, general-purpose I/O (GPIO) pins allocated across three separate GPIO modules—PORTFIO, PORTGIO, and PORTHIO, associated with port F, port G, and port H, respectively. Port J does not provide GPIO functionality. Each GPIO-capable pin shares functionality with other ADSP-BF51x processor peripherals via a multiplexing scheme; however, the GPIO functionality is the default state of the device upon powerup. Neither GPIO output or input drivers are active by default. Each general-purpose port pin can be individually controlled by manipulation of the port control, status, and interrupt registers:

- GPIO direction control register – Specifies the direction of each individual GPIO pin as input or output.
- GPIO control and status registers – The ADSP-BF51x processors employ a “write one to modify” mechanism that allows any combination of individual GPIO pins to be modified in a single instruction, without affecting the level of any other GPIO pins. Four control registers are provided. One register is written in order to set pin values, one register is written in order to clear pin values,

Two-Wire Interface

one register is written in order to toggle pin values, and one register is written in order to specify a pin value. Reading the GPIO status register allows software to interrogate the sense of the pins.

- GPIO interrupt mask registers – The two GPIO interrupt mask registers allow each individual GPIO pin to function as an interrupt to the processor. Similar to the two GPIO control registers that are used to set and clear individual pin values, one GPIO interrupt mask register sets bits to enable interrupt function, and the other GPIO interrupt mask register clears bits to disable interrupt function. GPIO pins defined as inputs can be configured to generate hardware interrupts, while output pins can be triggered by software interrupts.
- GPIO interrupt sensitivity registers – The two GPIO interrupt sensitivity registers specify whether individual 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.

Two-Wire Interface

The Two-Wire Interface (TWI) is fully compatible with the widely used I²C bus standard. It was designed with a high level of functionality and is compatible with multi-master, multi-slave bus configurations. To preserve processor bandwidth, the TWI controller can be set up and a transfer initiated with interrupts only to service FIFO buffer data reads and writes. Protocol related interrupts are optional.

The TWI externally moves 8-bit data while maintaining compliance with the I²C bus protocol. The *Philips I²C Bus Specification version 2.1* covers many variants of I²C. The TWI controller includes these features:

- Simultaneous master and slave operation on multiple device systems
- Support for multi-master data arbitration
- 7-bit addressing
- 100K bits/second and 400K bit/second data rates
- General call address support
- Master clock synchronization and support for clock low extension
- Separate multiple-byte receive and transmit FIFOs
- Low interrupt rate
- Individual override control of data and clock lines in the event of bus lock-up
- Input filter for spike suppression
- Serial camera control bus support as specified in the *OmniVision Serial Camera Control Bus (SCCB) Functional Specification version 2.1*

Ethernet MAC

The Ethernet Media Access Controller (MAC) peripheral provides a 10/100M bit/second Ethernet interface, compliant with IEEE Std. 802.3-2002, between a Media Independent Interface (MII) and the Blackfin peripheral subsystem. The MAC operates in both half-duplex and full-duplex modes. It provides programmable enhanced features designed

IEEE 1588 Support

to minimize bus utilization and pre- or post-message processing. The connection to the external physical layer device (PHY) is achieved via the MII or a Reduced Media Independent Interface (RMII). The RMII provides data buses half as wide (2 bit vs. 4 bit) as those of an MII, operating at double the frequency.

The MAC is clocked internally from the `CLKIN` pin on the processor. A buffered version of this clock can also be used to drive the external PHY via the `CLKBUF` pin. A 25 MHz source should be used with an MII PHY. A 50 MHz clock source is required to drive an RMII PHY.

IEEE 1588 Support

The IEEE 1588 standard is a precision clock synchronization protocol for networked measurement and control systems. The ADSP-BF518 processors include hardware support for IEEE 1588 with an integrated precision time protocol synchronization engine (PTP_TSYNC). This engine provides hardware assisted time stamping to improve the accuracy of clock synchronization between PTP nodes. The main features of the PTP_SYNC engine are:

- Support for both IEEE 1588-2002 and IEEE 1588-2008 protocol standards
- Hardware assisted time stamping capable of 12.5 ns resolution
- Lock adjustment
- Programmable PTM message support
- Dedicated interrupts
- Programmable alarm
- Multiple input clock sources (SCLK, MII clock, external clock up to 50 MHz)

- Programmable pulse per second (PPS) output
- Auxiliary snapshot to time stamp external events

RSI Interface

The removable storage interface (RSI) controller acts as the host interface for multi-media cards (MMC), secure digital memory cards (SD Card), secure digital input/output cards (SDIO), and CE-ATA hard disk drives. The following list describes the main features of the RSI controller:

- Support for a single MMC, SD memory, SDIO card or CE-ATA hard disk drive
- Support for 1-bit and 4-bit SD modes
- Support for 1-bit, 4-bit and 8-bit MMC modes
- Support for 4-bit and 8-bit CE-ATA hard disk drives
- A ten-signal external interface with clock, command, and up to eight data lines
- Card detection using one of the data signals
- Card interface clock generation from SCLK
- SDIO interrupt and read wait features
- CE-ATA command completion signal recognition and disable

General-Purpose (GP) Counter

A 32-bit GP counter is provided that can sense 2-bit quadrature or binary codes as typically emitted by industrial drives or manual thumb wheels. The counter can also operate in general-purpose up/down count modes.

3-Phase PWM Unit

Then, count direction is either controlled by a level-sensitive input signal or by two edge detectors. A third input can provide flexible zero marker support and can alternatively be used to input the push-button signal of thumb wheels. All three signals have a programmable debouncing circuit. An internal signal forwarded to the GP timer unit enables one timer to measure the intervals between count events. Boundary registers enable auto-zero operation or simple system warning by interrupts when programmable count values are exceeded.

3-Phase PWM Unit

The processors integrate a flexible and programmable 3-phase PWM waveform generator that can be programmed to generate the required switching patterns to drive a 3-phase voltage source inverter for ac induction (ACIM) or permanent magnet synchronous (PMSM) motor control. In addition, the PWM block contains special functions that considerably simplify the generation of the required PWM switching patterns for control of the electronically commutated motor (ECM) or brushless dc motor (BDCM). Software can enable a special mode for switched reluctance motors (SRM).

Features of the 3-phase PWM generation unit are:

- 16-bit center-based PWM generation unit
- Programmable PWM pulse width
- Single/double update modes
- Programmable dead time and switching frequency
- Twos-complement implementation which permits smooth transition to full ON and full OFF states
- Possibility to synchronize the PWM generation to an external synchronization

- Special provisions for BDCM operation (crossover and output enable functions)
- Wide variety of special switched reluctance (SR) operating modes
- Output polarity and clock gating control
- Dedicated asynchronous PWM shutdown signal

The six PWM output signals consist of three high-side drive signals (PWM_AH, PWM_BH, and PWM_CH) and three low-side drive signals (PWM_AL, PWM_BL, and PWM_CL). The polarity of the generated PWM signal be set with software, so that either active high or active low PWM patterns can be produced. The switching frequency of the generated PWM pattern is programmable. The PWM generator can operate in single update mode or double update mode. In single update mode the duty cycle values are programmable only once per PWM period, so that the resultant PWM patterns are symmetrical about the midpoint of the PWM period. In the double update mode, a second updating of the PWM registers is implemented at the midpoint of the PWM period. In this mode, it is possible to produce asymmetrical PWM patterns that produce lower harmonic distortion in 3-phase PWM inverters.

Parallel Peripheral Interface

The processor provides a Parallel Peripheral Interface (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 and three multiplexed frame sync 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.

Parallel Peripheral Interface

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 2-D 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.

SPORT Controllers

The processor incorporates two dual-channel synchronous serial ports (SPORT0 and SPORT1) for serial and multiprocessor communications. The SPORTs support these features:

- Bidirectional, I²S capable operation

Each SPORT has two sets of independent transmit and receive pins, which enable eight channels of I²S stereo audio.

- Buffered (eight-deep) transmit and receive ports

Each port has a data register for transferring data words to and from other processor 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.

SPORT Controllers

- 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) Ports

The processor has two SPI-compatible ports that enable the processor to communicate with multiple SPI-compatible devices.

Each 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 several SPI chip select output pins let the processor select other SPI devices. The SPI select pins are reconfigured general-purpose I/O pins. Using these pins, the SPI port provides a full-duplex, synchronous serial interface, which supports both master and slave modes and multimaster environments.

Each 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 datastreams. 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 nine general-purpose programmable timer units in the processor. Eight timers have 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 TMRCLK/PPI_CLK pin or to the internal SCLK.

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

UART Ports

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 eight general-purpose programmable timers, a 9th timer is also provided. 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 Ports

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

- Programmed I/O

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. Each of the two UARTs have 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 UARTs' baud rate, serial data format, error code generation and status, and interrupts can be programmed to support:

- 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 ports are further extended with support for the Infrared Data Association (IrDA[®]) Serial Infrared Physical Layer Link Specification (SIR) protocol.

Security

ADSP-BF51x processors provides security features (Blackfin Lockbox[™] Secure Technology) that enable customer applications to use secure protocols, consisting of code authentication and execution of code within a secure environment. Implementing secure protocols on Blackfin processors involves a combination of hardware and software components. Together these components protect secure memory spaces and restrict control of security features to authenticated developer code.

- Blackfin Lockbox Secure Technology incorporates a secure hardware platform for *confidentiality* and *integrity* protection of secure code and data with *authenticity* maintained by secure software.
- This secure platform provides:
- A secure execution mode
- Secure storage for on-chip keys

Real-Time Clock

- On-chip secure ROM
- Secure RAM
- Access to code and data in the secure domain is monitored by the hardware and any unauthorized access to the secure domain is prevented.
- The secure ROM code establishes the *root of trust* for the secure software in the system.
- The secure RAM provides *integrity* protection and *confidentiality* for authenticated code and data.
- User-defined cipher key(s) and ID(s) can be securely stored in the on-chip OTP memory.
- Every processor ships from the ADI factory with a unique chip ID value stored in publicly accessible OTP memory area.

Real-Time Clock

The processor's Real-Time Clock (RTC) provides a robust set of digital watch features, including current time, stopwatch, and alarm. The RTC is clocked by a 32.768 kHz crystal external to the processor. The RTC peripheral has dedicated power supply pins, so that it can remain powered up and clocked even when the rest of the processor is in a low power state. The RTC provides several programmable interrupt options, including interrupt per second, minute, hour, or day clock ticks, interrupt on programmable stopwatch countdown, or interrupt at a programmed alarm time.

The 32.768 kHz input clock frequency is divided down to a 1 Hz signal by a prescaler. The counter function of the timer consists of four counters: a 60 second counter, a 60 minute counter, a 24 hours counter, and a 32768 day counter.

When enabled, the alarm function generates an interrupt when the output of the timer matches the programmed value in the alarm control register. There are two alarms. The first alarm is for a time of day. The second alarm is for a day and time of that day.

The stopwatch function counts down from a programmed value, with one minute resolution. When the stopwatch is enabled and the counter underflows, an interrupt is generated.

Like the other peripherals, the RTC can wake up the processor from sleep mode or deep sleep mode upon generation of any RTC wakeup event. An RTC wakeup event can also wake up the on-chip internal voltage regulator from a powered down state.

Watchdog Timer

The processor 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, nonmaskable 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 that 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 CPU and the 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 control register.

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

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 ($0.5\times$ to $64\times$) multiplication factor (bounded by specified minimum and maximum `VCO` frequencies). The default multiplier is $10\times$, 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 processor 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 peripherals also reduces power consumption.

Full-On Mode (Maximum Performance)

In the full-on mode, the 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 core and all enabled peripherals run at full speed.

Active Mode (Moderate Power Savings)

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. 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 Mode (High Power Savings)

The sleep mode reduces power dissipation by disabling the clock to the processor core (CCLK). The PLL and system clock (SCLK), however, continue to operate in this mode. Typically an external event or RTC activity will wake up the processor. When in the sleep mode, assertion of any interrupt causes the processor to sense the value of the bypass bit (BYPASS) in the PLL control register (PLL_CTL). If bypass is disabled, the processor transitions to the full on mode. If bypass is enabled, the processor transitions to the active mode.

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

Deep Sleep Mode (Maximum Power Savings)

The deep sleep mode maximizes dynamic power savings by disabling the processor core and synchronous system clocks (CCLK and SCLK). Asynchronous systems, such as the RTC, may still be running, but cannot access

Instruction Set Description

internal resources or external memory. This powered-down mode can only be exited by assertion of the reset interrupt or by an asynchronous interrupt generated by the RTC. When in deep sleep mode, an RTC asynchronous interrupt causes the processor to transition to the active mode. Assertion of $\overline{\text{RESET}}$ while in deep sleep mode causes the processor to transition to the full on mode.

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.

Instruction Set Description

The Blackfin processor family assembly language instruction set employs an algebraic syntax designed for ease of coding and readability. Refer to the *Blackfin Processor Programming Reference* for detailed information. 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 core resources.

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

- Embedded 16/32-bit 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
- Seamlessly integrated DSP/CPU features 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 + two load/store + two pointer updates per cycle
- All registers, I/O, and memory mapped into a unified 4G byte memory space, providing a simplified programming model

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 processor 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 products also fully emulates the Blackfin processor family.

The VisualDSP++ project management 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++ runtime library that includes DSP and mathematical functions. A key point for these tools is

Development Tools

C/C++ code efficiency. The compiler has been developed for efficient translation of C/C++ code to Blackfin processor 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++ Integrated Development and Debugging Environment (IDDE) 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 emulators use the IEEE 1149.1 JTAG test access port of the processor 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. Nonintrusive 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 processor family. Hardware tools include the ADSP-BF51x EZ-KIT Lite[®] standalone evaluation/development cards. Third party software tools include DSP libraries, real-time operating systems, and block diagram design tools.

2 MEMORY

This chapter discusses memory population specific to the ADSP-BF51x processors. Functional memory architecture is described in the *Blackfin Processor Programming Reference*.

Memory Architecture

[Figure 2-1](#) provides an overview of the ADSP-BF51x processor system memory map. For a detailed discussion of how to use them, see the *Blackfin Processor Programming Reference*. Note 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.

The upper portion of internal memory space is allocated to the core and system MMRs. Accesses to this area are allowed only when the processor is in supervisor or emulation mode (see the Operating Modes and States chapter of the *Blackfin Processor Programming Reference*).

Within the external memory map, four banks of asynchronous memory space and one bank of SDRAM memory are available. Each of the asynchronous banks is 1M byte and the SDRAM bank is up to 128M byte.

L1 Instruction SRAM

Table 2-1. ADSP-BF51x Memory Map

Starting Address	Ending Address	Description
0xFFE0 0000	0xFFFF FFFF	Core MMR (2M bytes)
0xFFC0 0000	0xFFDF FFFF	System MMR (2M bytes)
0xFFB0 1000	0xFFBF FFFF	reserved
0xFFB0 0000	0xFFB0 0FFF	Scratchpad SRAM (4K bytes)
0xFFA1 4000	0xFFAF FFFF	reserved
0xFFA1 0000	0xFFA1 3FFF	Instruction SRAM/Cache (16K bytes)
0xFFA0 8000	0xFFA0 FFFF	reserved
0xFFA0 4000	0xFFA0 7FFF	Instruction Bank B SRAM (16K bytes)
0xFFA0 0000	0xFFA0 3FFF	Instruction Bank A SRAM (16K bytes)
0xFF90 8000	0xFF9F FFFF	reserved
0xFF90 4000	0xFF90 7FFF	Data Bank B SRAM/Cache (16K bytes)
0xFF90 0000	0xFF90 3FFF	Data Bank B SRAM (16K bytes)
0xFF80 8000	0xFF9F FFFF	reserved
0xFF80 4000	0xFF80 7FFF	Data Bank A SRAM/Cache (16K bytes)
0xFF80 0000	0xFF80 3FFF	Data Bank A SRAM (16K bytes)
0xEF00 8000	0xFF7F FFFF	reserved
0xEF00 0000	0xEF00 7FFF	BOOT ROM (32K bytes)
0x2040 0000	0xEEFF FFFF	reserved
0x2030 0000	0x203F FFFF	Async Bank 3 (1M bytes)
0x2020 0000	0x202F FFFF	Async Bank 2 (1M bytes)
0x2010 0000	0x201F FFFF	Async Bank 1 (1M bytes)
0x2000 0000	0x200F FFFF	Async Bank 0 (1M bytes)
0x0800 0000	0x1FFF FFFF	reserved
0x0000 0000	0x07FF FFFF	SDRAM (128M bytes)

L1 Instruction SRAM

The processor 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).

Table 2-2 lists the memory start locations of the L1 instruction memory subbanks.

Table 2-2. L1 Instruction Memory Subbanks

Memory Subbank	Memory Start Location for ADSP-BF51x Processors
0	0xFFA0 0000
1	0xFFA0 1000
2	0xFFA0 2000
3	0xFFA0 3000
4	0xFFA0 4000
5	0xFFA0 5000
6	0xFFA0 6000
7	0xFFA0 7000
8	0xFFA1 0000
9	0xFFA1 1000
10	0xFFA1 2000
11	0xFFA1 3000

L1 Data SRAM

Table 2-3 shows how the subbank organization is mapped into memory.

Table 2-3. L1 Data Memory SRAM Subbank Start Addresses

Memory Bank and Subbank	ADSP-BF51x Processors
Data Bank A, Subbank 0	0xFF80 0000
Data Bank A, Subbank 1	0xFF80 1000
Data Bank A, Subbank 2	0xFF80 2000
Data Bank A, Subbank 3	0xFF80 3000

L1 Data Cache

Table 2-3. L1 Data Memory SRAM Subbank Start Addresses (Continued)

Memory Bank and Subbank	ADSP-BF51x Processors
Data Bank A, Subbank 4	0xFF80 4000
Data Bank A, Subbank 5	0xFF80 5000
Data Bank A, Subbank 6	0xFF80 6000
Data Bank A, Subbank 7	0xFF80 7000
Data Bank B, Subbank 0	0xFF90 0000
Data Bank B, Subbank 1	0xFF90 1000
Data Bank B, Subbank 2	0xFF90 2000
Data Bank B, Subbank 3	0xFF90 3000
Data Bank B, Subbank 4	0xFF90 4000
Data Bank B, Subbank 5	0xFF90 5000
Data Bank B, Subbank 6	0xFF90 6000
Data Bank B, Subbank 7	0xFF90 7000

L1 Data Cache

When data cache is enabled (controlled by bits `DMC[1:0]` in the `DMEM_CONTROL` register), either 16K byte of data bank A or 16K byte of both data bank A and data bank B can be set to serve as cache.

Boot ROM

The lowest 32K byte of internal memory space is occupied by the boot ROM starting from address `0xEF00 0000`. This 16-bit boot ROM is not part of the L1 memory module. Read accesses take one `SCLK` cycle and no wait states are required. The read-only memory can be read by the core as well as by DMA. It can be cached and protected by CPLD blocks like external memory. The boot ROM not only contains boot-strap loader

code, it also provides some subfunctions that are user-callable at runtime. For more information, see [Chapter 26, “System Reset and Booting”](#).

External Memory

The external memory space is shown in [Figure 2-1 on page 2-2](#). One of the memory regions is dedicated to SDRAM support. The size of the SDRAM bank is programmable and can range in size from 16M byte to 128M byte. The start address of the bank is 0x0000 0000.

Each of the next four banks contains 1M byte and is dedicated to support asynchronous memories. The start address of the asynchronous memory bank is 0x2000 0000.

Processor-Specific MMRs

The complete set of memory-related MMRs is described in the *Blackfin Processor Programming Reference*. Several MMRs have bit definitions specific to the processors described in this manual. These registers are described in the following sections.

DMEM_CONTROL Register

The data memory control register (DMEM_CONTROL), shown in Figure 2-1, contains control bits for the L1 data memory.

Data Memory Control Register (DMEM_CONTROL)

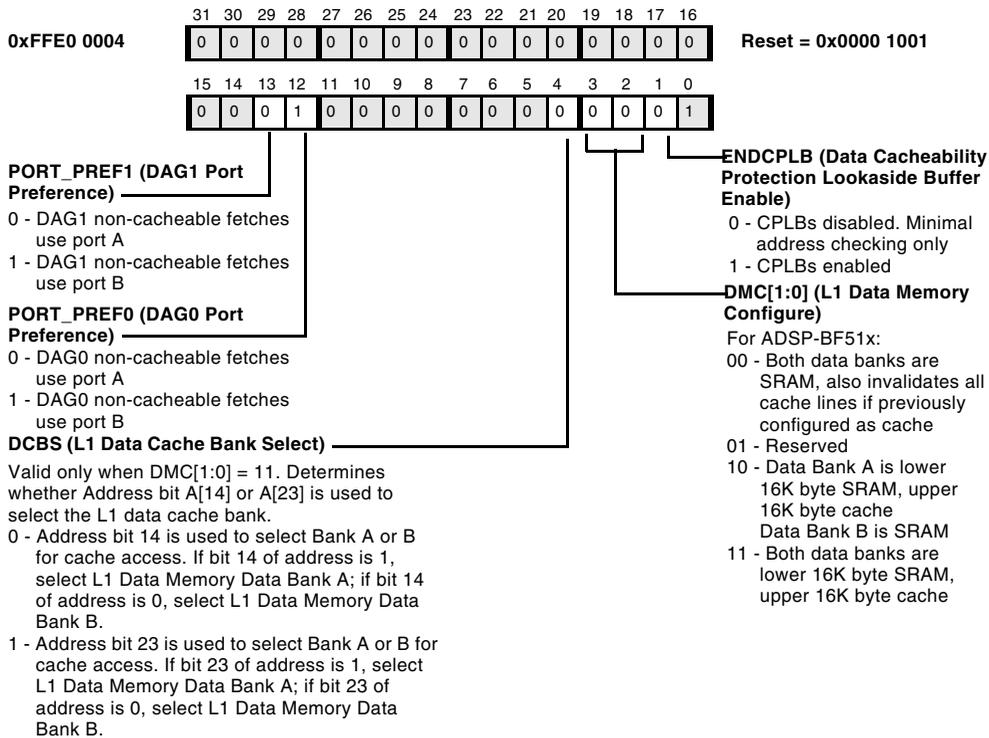


Figure 2-1. L1 Data Memory Control Register

DTEST_COMMAND Register

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]). This register is shown in Figure 2-2.

i The data/instruction access bit allows direct access via the DTEST_COMMAND MMR to L1 instruction SRAM.

Data Test Command Register (DTEST_COMMAND)

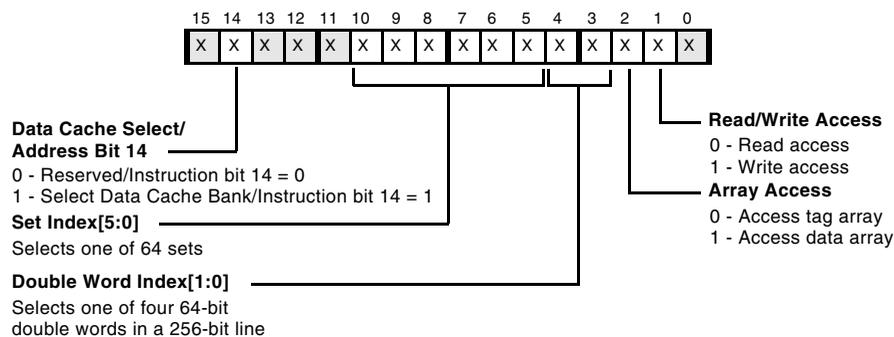
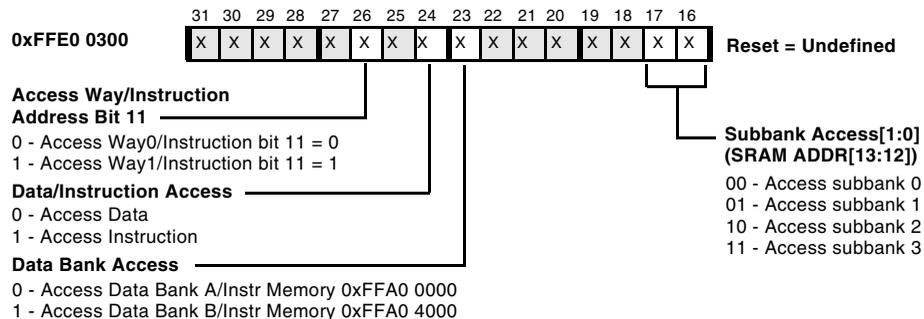


Figure 2-2. Data Test Command Register

i Bit 14 and bit 23 of DTEST_COMMAND must be cleared to correctly access the L1 Instruction Memory Bank A. Likewise, for accesses to L1 Instruction Memory Bank B, bits 14 and 23 must both be set.

Processor-Specific MMRs

3 ONE-TIME PROGRAMMABLE MEMORY

This chapter describes One-Time-Programmable (OTP) memory features of the ADSP-BF51x Blackfin processor.

The chapter includes the following sections:

- [“OTP Memory Overview”](#)
- [“OTP Memory Map” on page 3-2](#)
- [“Error Correction” on page 3-7](#)
- [“OTP Access” on page 3-10](#)
- [“Error Correction Policy” on page 3-8](#)
- [“OTP Timing Parameters” on page 3-11](#)
- [“Callable ROM Functions for OTP ACCESS” on page 3-14](#)
- [“Programming and Reading OTP” on page 3-17](#)
- [“Write-Protecting OTP Memory” on page 3-24](#)
- [“Accessing Private OTP Memory” on page 3-26](#)
- [“OTP Programming Examples” on page 3-26](#)

OTP Memory Overview

The ADSP-BF51x processors include an on-chip, one-time-programmable memory array which provides 64k-bits of non-volatile memory. This

OTP Memory Map

includes the array and logic to support read access and programming. A mechanism for error correction is also provided. Additionally, pages can be write protected.

OTP memory can be programmed through various methods including software running on the Blackfin processor. The ADSP-BF51x processors provide C and assembly callable functions in the on-chip ROM to help the developer access the OTP memory.

The one-time-programmable memory is divided into two main regions. A 32-k bit “public” unsecured region which has no access restrictions and a 32-k bit “private” secured region with access restricted to authenticated code when operating in Secure Mode (For information about these modes, see [Chapter 25, “Security”](#) in this volume of the *ADSP-BF51x Blackfin Processor Hardware Reference*.)

OTP enables developers to store both public and private data on-chip. A 64K x 1 bit array is available as shown in [Figure 3-2](#). In addition to storing public and private data, it allows developers to store completely user-definable data such as customer ID, product ID, MAC address, etc.

 The public portion of OTP memory contains many “factory set only” values. Users are urged to exercise caution when writing to OTP memory and to consult the OTP memory map for details of Customer Programmable Settings (CPS) and factory reserved areas of this memory. See also Factory Page Settings (FPS) and Preboot Page Settings (PBS) in [Chapter 26, “System Reset and Booting”](#) in this volume of the *ADSP-BF51x Blackfin Processor Hardware Reference*.

OTP Memory Map

The OTP is not part of the Blackfin linear memory map. It has a separate memory map that is shown in [Figure 3-2](#). OTP memory is not accessed directly using the Blackfin memory map, rather, it is accessed via four

One-Time Programmable Memory

32-bit wide registers (OTP_DATA3:0) which act as the OTP memory read/write buffer.

In the case of an OTP memory read, the OTP_DATA_x registers will contain the 16-byte result of the OTP memory access. In the case of an OTP memory write, the OTP_DATA_x registers will contain 16 bytes of data to be written to the OTP memory.

OTP_DATA3–0 registers are organized into a 128 bit page as shown in [Figure 3-1](#).

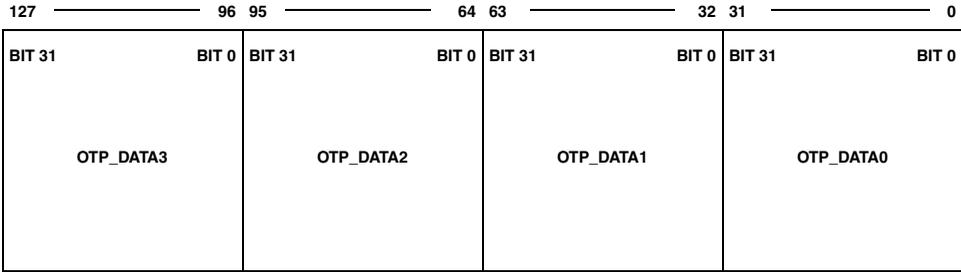
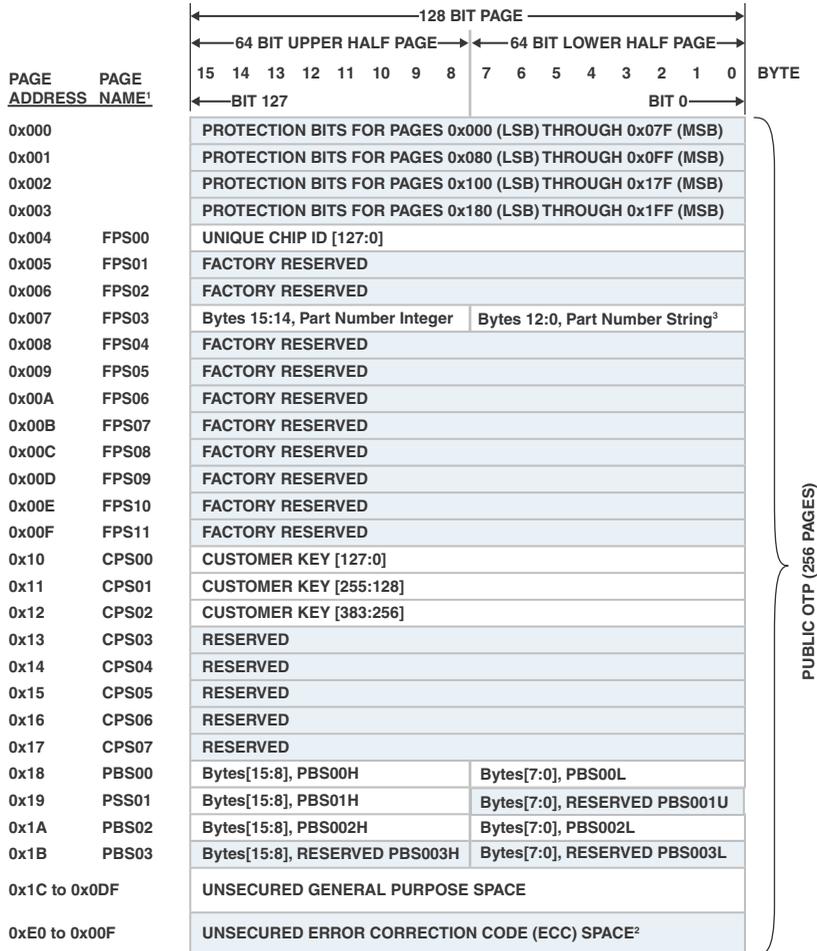


Figure 3-1. OTP_DATA_x Registers

OTP Memory Map

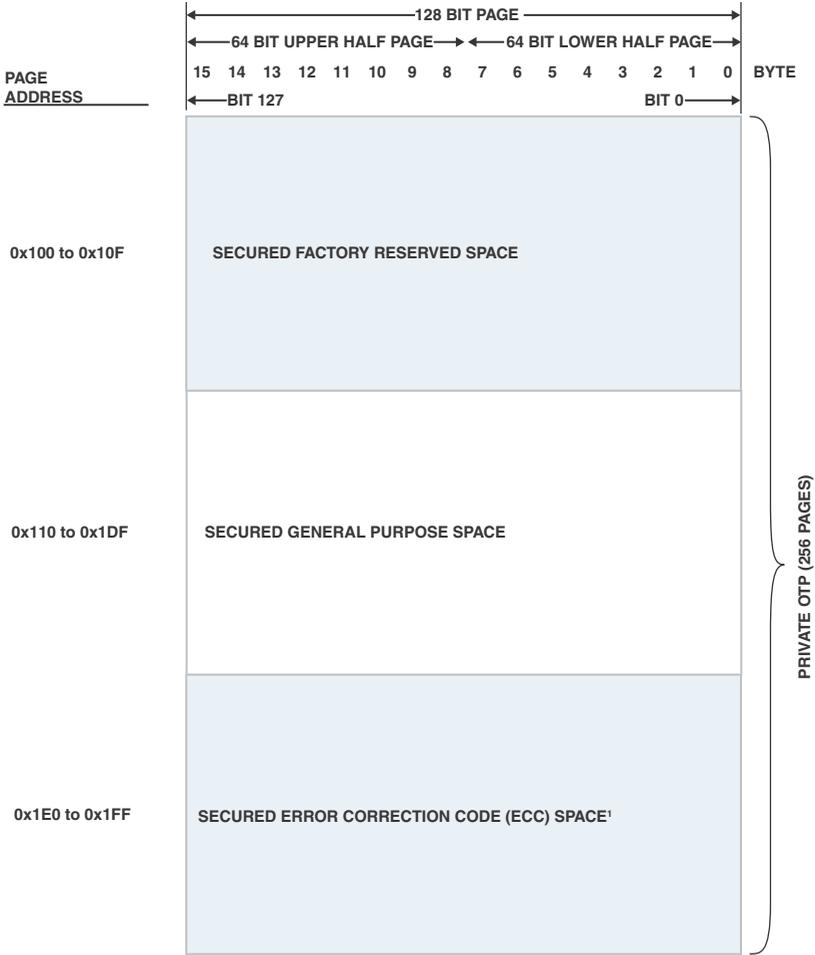


Footnotes

1. Factory Programmable Settings (FPS) are programmed at the factory. Customer Programmable Settings (CPS) are programmed by the customer.
2. This space should NOT be written by the customer. 8-bit error correction codes are automatically generated by firmware and stored in this region.
3. Part Number Field Definition. A string indicating the model number of the product is programmed into this location. Each character is represented by standard 8-bit ASCII code. A termination character of 0x00000000 terminates the string. The field supports up to 12 alphanumeric characters plus one termination character. The first string character resides in bits[7:0] and the string grows to the left with the left most character being the termination character. Integer representation of the part number is shown in Table 3-1. Byte 13 in FPS03 is reserved.

Figure 3-2. Public OTP Memory Map

One-Time Programmable Memory



Footnotes
 1. This space should NOT be written by the customer. 8-bit error correction codes are automatically generated by firmware and stored in this region.

Figure 3-3. Private OTP Memory Map

OTP Memory Map

Table 3-1. Part Number Field Definition

Part #	Code
ADSP-BF512	0x0200
ADSP-BF512F	0xF200
ADSP-BF514	0x0202
ADSP-BF514F	0xF202
ADSP-BF516	0x0204
ADSP-BF516F	0xF204
ADSP-BF518	0x0206
ADSP-BF518F	0xF206

OTP memory ranges marked as Factory Reserved, Reserved and Error Correction Code Space, in [Figure 3-2](#), must not be programmed by the user. Customer Programmable Settings are optionally programmed by the developer.

Page-Protection bits provide protection for each 128-bit page within the OTP. By default, the OTP array bits are not set and will read back as zero values if left unprogrammed. Programmed data values consist of zeroes and ones, therefore, after programming OTP memory, some bits will intentionally remain as zero values. The write-protect bits provide protection for the zero value bits to remain as zeroes and prevent future programming (inadvertent or malicious) from changing bit values from zero to one.

Pages 0x10, 0x11, and 0x12 hold the customer public key which is used for Lockbox digital signature authentication. Please refer to [Chapter 25, “Security”](#) for more information on Lockbox and how the public key is used.

OTP memory is logically arranged in a sequential set of 128-bit pages. Each OTP memory address refers to a 128-bit page. The ADSP-BF51x processor thus provides 512 pages of OTP memory.

In order to read or program the OTP memory, a set of functions are provided in the on-chip ROM. These functions include `bfrom_OtpRead()`, `bfrom_OtpWrite()` and `bfrom_OtpCommand()`.

Error Correction

To meet strict quality goals, error correction is used to ensure data integrity. `bfrom_OtpRead()` and `bfrom_OtpWrite()`, provided in the on-chip ROM, support error correction.

Error correction works by calculating an 8-bit Error Correction Code (ECC) for each 64-bit data word (half page) when it is programmed into the OTP. When this word is later read from OTP, its corresponding ECC is also read and a data integrity check is performed. If the check fails, error correction on the data word can be attempted using the ECC. Depending on the type of error, the error correction algorithm will perform as shown in [Table 3-2](#).

Table 3-2. Hamming Code Single Error Corrections, Double Error Detection

Number of Bad Bits in Data Word	Error(s) Detected?	Error(s) Corrected?
0	N/A	N/A
1	Yes	Yes
2	Yes	No
3 or more	No	No

Error Correction Policy

1. Error correction requires that OTP space is written and read in 64-bit widths. Firmware will only support writing or reading half of an OTP page.
2. Error correction is used to correct data in all pages of OTP space except the protection pages (0x0 to 0x3) and ECC pages themselves. See “[OTP Access](#)” on [page 3-10](#) for more information.
3. Firmware will generate and program the 8-bit ECC fields as mapped in [Table 3-3](#) and [Table 3-4](#).
4. The developer is responsible for locking both the data page(s) AND the ECC page(s) after all programming is complete.
5. Pages 0x04 to 0x0F are reserved for ADI factory use. Therefore, pages 0x004 to 0x00F, 0x0E0, and 0x0E1 will be locked coming out of the Analog Devices factory.

Table 3-3. Mapping for Storage of Error Correction Codes for Unsecured OTP Space

Page	Byte							
	15	14	13	12	11	10	9	8
0x0E0	0x007U	0x007L	0x006U	0x006L	0x005U	0x005L	0x004U	0x004L
0x0E1	0x00FU	0x00FL	0x00EU	0x00EL	0x00DU	0x00DL	0x00CU	0x00CL
0x0E2	0x017U	0x017L	0x016U	0x016L	0x015U	0x015L	0x014U	0x014L
....								
0x0FB	0x0DFU	0x0DFL	0x0DEU	0x0DEL	0x0DDU	0x0DDL	0x0DCU	0x0DCL

One-Time Programmable Memory

Table 3-3. Mapping for Storage of Error Correction Codes for Unsecured OTP Space (Continued)

Page	Byte							
	15	14	13	12	11	10	9	8
Page	7	6	5	4	3	2	1	0
0x0E0	Unused							
0x0E1	0x00BU	0x00BL	0x00AU	0x00AL	0x009U	0x009L	0x008U	0x008L
0x0E2	0x013U	0x013L	0x012U	0x012L	0x011U	0x011L	0x010U	0x010L
....								
0x0FB	0x0DBU	0x0DBL	0x0DAU	0x0DAL	0x0D9U	0x0D9L	0x0D8U	0x0D8L

Table 3-4. Mapping for Storage of Error Correction Codes for Secured OTP Space

Page	Byte							
	15	14	13	12	11	10	9	8
0x1E0	0x107U	0x107L	0x106U	0x106L	0x105U	0x105L	0x104U	0x104L
0x1E1	0x10FU	0x10FL	0x10EU	0x10EL	0x10DU	0x10DL	0x10CU	0x10CL
0x1E2	0x117U	0x117L	0x116U	0x116L	0x115U	0x115L	0x114U	0x114L
....								
0x1FB	0x1DFU	0x1DFL	0x1DEU	0x1DEL	0x1DDU	0x1DDL	0x1DCU	0x1DCL
Page	7	6	5	4	3	2	1	0
0x1E0	0x103U	0x103L	0x102U	0x102L	0x101U	0x101L	0x100U	0x100L
0x1E1	0x10BU	0x10BL	0x10AU	0x10AL	0x109U	0x109L	0x108U	0x108L
0x1E2	0x113U	0x113L	0x112U	0x112L	0x111U	0x111L	0x110U	0x110L
....								
0x1FB	0x1DBU	0x1DBL	0x1DAU	0x1DAL	0x1D9U	0x1D9L	0x1D8U	0x1D8L

OTP Access

The ADSP-BF51x on-chip ROM contains functions for initializing OTP timing parameters, reading and programming the OTP memory. These functions include `bfrom_OtpRead()`, `bfrom_OtpWrite()` and `bfrom_OtpCommand()`.

 These functions are callable from C or assembly application code. Use only these functions for accessing OTP memory. Directly accessing memory locations within OTP memory by other means is not supported.

The existing ECC in ROM is known as “Hamming [72,64]” - This is specifically a 64-bit Data, +8-bit ECC Field, for 1-bit correction and 2-bit error detection scheme.

 The ROM-based OTP read/write API *must* be used for all OTP data accesses (see limited exceptions below). The ROM code incorporates the ONLY ECC method supported by Analog Devices. Analog Devices does not support direct access of OTP data without using error correction.

Exceptions: The only bits that do not use ECC are page lock bits (first four pages) and the preboot invalidate bits. See the Preboot section in [Chapter 26, “System Reset and Booting”](#).

ADI does not support any ECC other than the ECC provided by ADI within the ROM API. All attempts to implement other schemes are not guaranteed or supported by Analog Devices.

OTP memory programming is done serially under software control. Since the unprogrammed OTP memory value defaults to zero, only bits whose value is intended to be “1” have to be programmed. In order to protect areas of OTP memory that have been programmed or areas which have intentionally been left unprogrammed which end users wish to remain unchanged, write-protect bits can be set for each 128-bit page within

OTP memory. Each write-protect bit, when set, will prevent further programming attempts to OTP memory on a per page basis. Refer to the OTP memory map ([Figure 3-2](#)) for details.

The ADSP-BF51x Blackfin processor can program OTP through software code executing directly on the Blackfin processor. No on-chip charge pump exists, therefore, an externally applied voltage is required to apply the voltage levels appropriate for programming OTP memory. Refer to the processor data sheet for VPPOTP specifications. OTP programming code can be loaded into the processor via JTAG emulation, DMA, and all supported boot methods.

OTP memory can only be written once (changing a bit from 0 to 1). Once a bit has been changed from a 0 to a 1, it cannot be changed back to 0. The write-protect bits prevent OTP memory from having any bits that are meant to remain as 0 value later programmed to a value of 1.

Prior to accessing OTP memory, refer to the product data sheet for specifications on VDDOTP and VPPOTP voltage levels to ensure reliable OTP programming. OTP timing parameter settings must be set prior to attempting any write accesses to OTP.

OTP Timing Parameters

In order to read and program the OTP memory reliably, set the OTP timing parameters prior to accessing OTP memory. All of the timing parameters are bitfields within the `OTP_TIMING` register (see [“OTP_TIMING Register” on page 3-14](#)). The `bfrom_OtpCommand()` function (detailed in the following sections) is provided in the on-chip ROM to program the timing parameters.

 OTP timing parameters must be set with `bfrom_OtpCommand()`. OTP read accesses can use the OTP timing default reset value (Reset: `OTP_TIMING = 0x00001485`).

Use of the OTP timing default reset value for writes will result in write errors as this timing value is not appropriate for performing write accesses.

Insufficient voltage/current provided to OTP during write access or incorrect OTP timing parameters may result in the following error returned during OTP writes:

0x11: error code returned (multiple bad bits in 64 bit data), and subsequent reads from this page return 0.

The OTP timing parameters consist of several concatenated fields and form one value, which then is passed as an argument to the `bfrom_OtpCommand()` function. There is one field for which the developer must calculate a value based upon the desired `SCLK` frequency of operation at which the OTP access will be performed. This calculated value then is combined with a constant value field whose value is provided by Analog Devices to arrive at the setting appropriate for the access.

The OTP timing parameters are comprised of two values as follows.

`OTP_TIMING[7:0] = OTP_TP1 = 1000 / sclk_period` (in nanoseconds)

`OTP_TIMING[31:8] = OTP_TP2 = 0x145487`

The `OTP_TP2` field is specified by Analog Devices and must be used to ensure reliable OTP write accesses. The user-calculated field must be combined with the `OTP_TP2` value as shown in [Listing 3-1](#) and [Listing 3-2](#).

Example calculations shown in [Listing 3-1](#) and [Listing 3-2](#) are based upon `VDDOTP` and `VPPOTP` voltage values specified in the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

The OTP timing parameter calculations are dependent upon user-defined

One-Time Programmable Memory

SCLK frequency of operation. (Refer to the processor data sheet for actual VDDOTP and VPPOTP voltage and SCLK specifications, do not rely on the specifications quoted in these examples.)

Listing 3-1. OTP Timing Calculations for SCLK = 80 MHz

For SCLK = 12.5ns (80 MHz), the following field calculations are needed to determine the OTP timing argument for the `bfrom_otpCommand()` call.

$OTP_TP1 = 1000 / sclk_period = 1000 / 12.5 = 0x50$	0x00000050
OTP_TP2 = (constant)	0x145487xx
Calculated OTP timing parameter value	0x14548750

The code for the API call (in C) is:

```
// Initialize OTP access settings
// Proper access settings for SCLK = 80 MHz
const u32 OTP_init_value = 0x14548750;
return_code = bfrom_otpCommand( OTP_INIT, OTP_init_value);
```

Listing 3-2. OTP Timing Calculations for SCLK = 50 MHz

For SCLK = 20.0ns (50 MHz), the following field calculations are needed to determine the OTP timing argument for the `bfrom_otpCommand()` call.

$OTP_TP1 = 1000 / sclk_period = 1000 / 20.0 = 0x32$	0x00000032
OTP_TP2 = (constant)	0x145487xx
Calculated OTP timing parameter value	0x14548732

OTP Access

The code for the API call (in C) is:

```
// Initialize OTP access settings
// Proper access settings for SCLK = 50 MHz
const u32 OTP_init_value = 0x14548732;
return_code = bfrom_otpCommand( OTP_INIT, OTP_init_value);
```

OTP_TIMING Register

OTP_TIMING Register

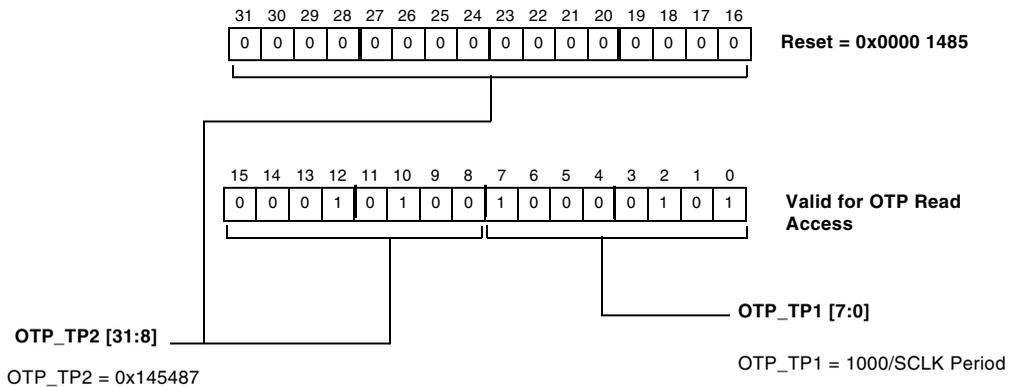


Figure 3-4. OTP_TIMING Register

Callable ROM Functions for OTP ACCESS

The following functions support OTP access.

Initializing OTP

This section describes the usage of the `bfrom_otpCommand()` function for the OTP memory controller setup provided in the ADSP-BF51x processor's on-chip ROM. The prototype and macros to help decode the function's return codes are supplied in the `bfrom.h` header file located in

the VisualDSP++ installation directory. The meaning of the error code is described in [“Error Codes” on page 3-22](#).

bfrom_OtpCommand

This function is used to implement various “commands” to setup the OTP controller. The first input parameter is a mnemonic label specifying the command. The second parameter is a generic value that is passed as argument for the requested command. The second parameter is optional and can be an integer value or (via opportune casting) a pointer or a pointer to an extension structure. There are two commands:

- `OTP_INIT`: sets the required timing value (register `OTP_TIMING`) to “value”.
- `OTP_CLOSE`: reinitializes the OTP controller. This can be called by the user before exiting Secure Mode if desired. The Value parameter may be specified as “0” or “NULL” with `OTP_CLOSE`.

Entry address: `0xEF00 0018`

Arguments:

R0: **command** (dCommand)

`OTP_INIT`

`OTP_CLOSE`

R1: **timing value to be programmed** (dValue), not used for `OTP_CLOSE`

C Prototype: `u32 bfrom_otpCommand(u32 dCommand, u32 dValue);`

Return code:

`bfrom_otpCommand()` currently always returns with “0”.

From the examples above, the OTP timing parameter was calculated to be `0x14548750` processor with `SCLK = 80 MHz`. Shown below is a sample of

OTP Access

C code that uses the `bfrom_OtpCommand()` function to program this timing parameter.

```
#include <bfrom.h>
#define OTP_TIMING_PARAM (0x14548750)

u32 Otp_Timing_Param_Init()
{
    u32 otp_timing_parameter;
    u32 = RetVal;
    otp_timing_parameter = OTP_TIMING_PARAM;
    RetVal = bfrom_OtpCommand(OTP_INIT, otp_timing_parameter);
    // (equivalently, with a variable):
    RetVal = bfrom_OtpCommand(OTP_INIT, OTP_TIMING_PARAM);

    return RetVal;
}
```

More examples:

```
//timing parameter
const u32 init_value = 0x14548750;

// call sets OTP_TIMING register
RetVal = bfrom_OtpCommand(OTP_INIT, init_value);

// call sets OTP_TIMING register
RetVal = bfrom_OtpCommand(OTP_INIT, 0x14548750);

// call clears OTP controller and data registers
RetVal = bfrom_OtpCommand(OTP_CLOSE, NULL);
```

The prototype of `bfrom_OtpCommand()` is also included in the `bfrom.h` header file installed with VisualDSP++ 5.0 and later releases. The `OTP_INIT` macro is defined in `bfrom.h` as well.

Programming and Reading OTP

This section describes the usage of `bfrom_otpRead()` and `bfrom_otpWrite()` read and write functions for OTP memory provided in the ADSP-BF51x processor's on-chip ROM. The prototypes and macros to help decode their return codes are supplied in the `bfrom.h` header file located in the VisualDSP++ installation directory. The meaning of the error code is described in [“Error Codes” on page 3-22](#).

`bfrom_otpRead`

This function is used to read 64-bit OTP half-pages using error correction.

Entry address: 0xEF00 001A

Arguments:

R0: OTP page address (`dPage`)

R1: Flags (`dFlags`)

OTP_LOWER_HALF

OTP_UPPER_HALF

OTP_NO_ECC

R2: Pointer to 64-bit memory struct (`long long`) to put read data (`*pPageContent`)

C prototype:

```
u32 bfrom_otpRead (u32 dPage, u32 dFlags, u64 *pPageContent);
```

Return code:

R0: error or warning code, see [Table 3-5](#).

This function reads a half-page and stores the content in the 64-bit variable pointed to by its last parameter. The page parameter defines the

OTP Access

address. The flags parameter defines whether the upper or the lower half page is to be read. The default reset `OTP_TIMING` value may be used for all read accesses without requiring any new setting value to be programmed prior to performing read accesses. Programming a valid value suitable for write accesses will also allow read accesses.

The use of flag parameter `OTP_NO_ECC` is not recommended for use with any OTP read access as it will bypass error correction code support. It is available only for diagnostic purposes.

bfrom_OtpWrite

This function attempts to write to (program) a half-page with the content in the 64-bit variable pointed to by its last parameter. The page parameter defines the address.

Entry address: 0xEF00 001C

Arguments:

R0: OTP page address (dFlag)

R1: Flags (dFlags)

`OTP_LOWER_HALF`

`OTP_UPPER_HALF`

`OTP_NO_ECC`

`OTP_LOCK`

`OTP_CHECK_FOR_PREV_WRITE`

R2: Pointer to 64-bit memory struct (long long) that contains the data to be written to OTP memory (*pPageContent)

C Prototype:

```
u32 bfrom_OtpWrite (u32 dPage, u32 dFlags, u64 *pPageContent);
```

Return code:

R0: error or warning code, see [Table 3-5](#).

The `dFlags` parameter defines whether the upper or the lower half page is to be written to and also whether the target half page should be checked for a previously written value before any write attempt is made. Additionally, a page can optionally be locked (permanently protected against further writes).

When performing pure lock operations (only locking a page without writing any data values to it), the half-page parameter is not required and it makes no difference which half-page is specified if this parameter is included in the function call.

In order to reduce the probability of inadvertent writes to OTP pages, the `bfrom_OtpWrite()` function checks for a valid OTP write timing setting in the `OTP_TIMING` register. More specifically, bits [31:15] must not be equal to zero. Calls to the write routine, when this field (bits [31:15]) is equal to zero cause an access violation error and the requested action is not performed. The user can use this mechanism to protect against inadvertent writes by calling the `bfrom_OtpCommand (OTP_init, ...)` function with appropriate values for reads only and for read/write accesses. Users are free to ignore this mechanism by calling `bfrom_OtpCommand (OTP_init, ...)` only once for read/write access.

When the flag `OTP_CHECK_FOR_PREV_WRITE` is NOT specified, a previously written value will be overwritten, both in the ECC and data fields for any unlocked page where a write access is performed. Of course, once a bit was set to “1” it cannot be reset to “0” by the new write operation. This means that, in all likelihood, if the new value is different from the previous one, the result will have multiple bit errors, in either or both the ECC and data fields.



Since the ECC field is written first by the ROM function, a multiple bit error will abort the operation without writing the new data value to the OTP data page.

Also note that multiple bit errors have a statistical chance of not being detected as such. So this default mode of operation is not recommended to be used, or used with appropriate caution.

The flag, `OTP_CHECK_FOR_PREV_WRITE`, should always be used by default when performing write accesses to OTP with the `bfrom_otpWrite()` function.

If the flag `OTP_CHECK_FOR_PREV_WRITE` is specified in the call, a write to a previously programmed page causes dedicated error messages and will not be undertaken. More specifically, the criterion for generating errors is as follows: the 64-bit data and the 8-bit ECC field are read and the total number of “1” is counted. If this number is equal to or greater than 2, the error flag `OTP_PREV_WR_ERR0` is returned and the write operation is not performed. If the number is 0, the page is certainly blank and the write is performed. If the number is one, a more thorough check is performed. If the “1” is in the ECC field, an error flag `OTP_SB_DEFECT_ERROR` is returned and the write is not performed. If the “1” is in the data field, it is determined whether the value to be written contains a “1” in the same position. If so, the write is performed. If not, the error flag `OTP_SB_DEFECT_ERROR` is returned and the write is not performed. This error code warns the user that it could be a single-bit defect in the page. The user can then decide whether to use this page regardless (by repeating the call without the `OTP_CHECK_FOR_PREV_WRITE` flag) or skip this page.

The `OTP_CHECK_FOR_PREV_WRITE` flag is ignored when a pure lock operation is requested (for example, a `OTP_LOCK` flag is set and `*pPageContent = NULL`). It is therefore unnecessary and harmless to specify this flag. The `OTP_CHECK_FOR_PREV_WRITE` flag is not ignored when doing a lock operation after a write (for example, `OTP_LOCK + write` in the same call and `*pPageContent = NULL`).

If the flag parameter for the write operation is augmented by the OR with `OTP_LOCK` flag, the write operation, if successful, will be immediately followed by setting the protection bit for the requested full 128-bit page.

A special case is the following (`OTP_LOCK`): if the third parameter is `NULL`, this call will lock a page without writing any data value to it (pure lock function). Note that in this case, “page” can span all pages from `0x000` to `0x1FF`. **This is the only way to lock the ECC pages themselves.**



The use of flag parameter `OTP_NO_ECC` is only supported in write operations when used to implement write-protection/ page-locking (use of `OTP_LOCK` parameter in `bfrom_Otp_Write` function is preferred method of locking pages, see Write Protecting OTP Memory section below) or to set the preboot invalidate bits (see the Preboot section in [Chapter 26, “System Reset and Booting”](#)). Bypassing error correction in OTP writes may result in loss of OTP data integrity and is not supported for any other OTP access.

The use of ECC in all OTP accesses other than the limited exceptions described previously is mandatory.

OTP Access

Error Codes

This section describes the returned error codes from the API functions. [Figure 3-5](#) and [Table 3-5](#) demonstrate and list the returned error codes from API functions.

Returned Error Codes from API Functions

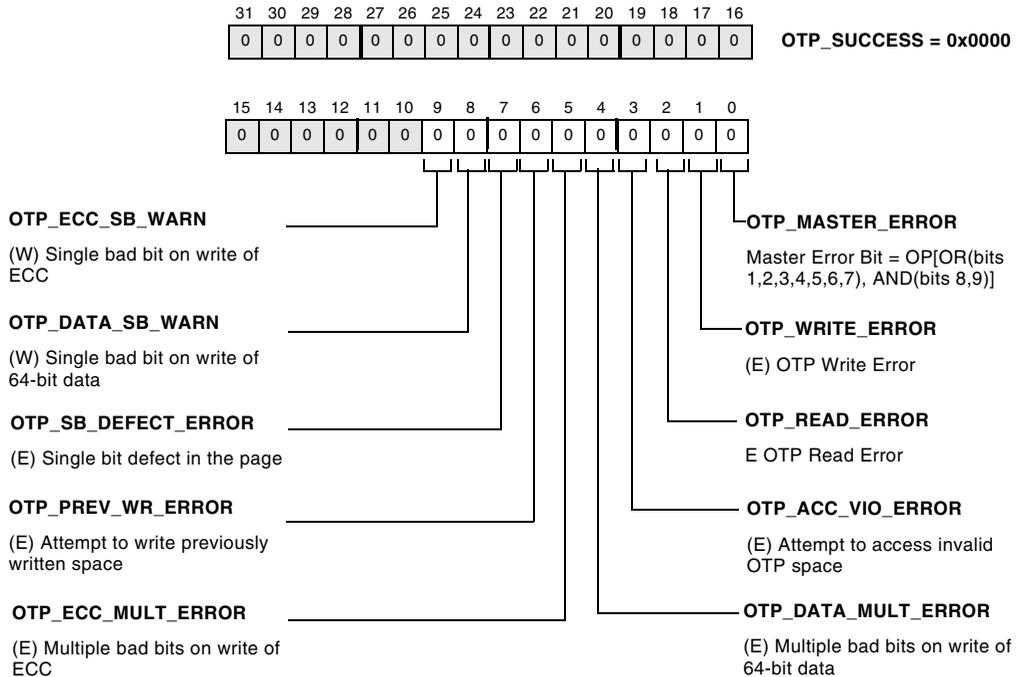


Figure 3-5. Returned Error Codes from API Functions

Table 3-5. Returned Error Codes from API Functions

Bit Position	Name	Example Return Value	Definition
N/A	OTP_SUCCESS	0x0	No Error
0	OTP_MASTER_ERROR	0x1	Master Error Bit = OR [OR (bits 1,2,3,4,5,6,7), AND (bits 8,9)]
1	OTP_WRITE_ERROR	0x3	(E) OTP Write Error
2	OTP_READ_ERROR	0x5	(E) OTP Read Error
3	OTP_ACC_VIO_ERROR	0x9	(E) Attempt to access invalid OTP space
4	OTP_DATA_MULT_ERROR	0x11	(E) Multiple bad bits on write of 64 bit data
5	OTP_ECC_MULT_ERROR	0x21	(E) Multiple bad bits on write of ECC
6	OTP_PREV_WR_ERROR	0x41	(E) Attempt to write previously written space
7	OTP_SB_DEFECT_ERROR	0x81	(E) Single-bit defect in the page
8	OTP_DATA_SB_WARN	0x100	(W) Single bad bit on write of 64 bit data
9	OTP_ECC_SB_WARN	0x200	(W) Single bad bit on write of ECC

`bfrom_OtpRead()` returns with an error when any of the bits [6:2] are set or both bits [9:8] are set. In this case, the `OTP_MASTER_ERROR` bit is also set. It returns with a warning if only one of the bits [9:8] is set.

`bfrom_OtpWrite()` returns with an error when any of the bits [7:1] are set or both bits [9:8] are set. In this case, the `OTP_MASTER_ERROR` bit is also set. It returns with a warning if only one of the bits [9:8] is set.

`bfrom_OtpCommand()` currently always returns with “0”.

Write-Protecting OTP Memory

As shown in [Figure 3-2](#), a small portion of OTP memory is reserved for write-protect bits (“write-protect” is synonymous with “page-protect” in the context of this discussion). After programming OTP memory, the programmer can use these protection bits to “lock” the page that was just programmed by setting the write-protect bit corresponding to the OTP data page. Once the write-protect bit is set and the lock is in place, further attempts to write to that page will not be allowed, resulting in an error. Page protect bits can also be set in order to prevent programming of unwritten OTP pages as well. Once an OTP page is page-protected, the write protection can not be reversed and no further write accesses can be made to the protected page(s).

There are four pages reserved for the write-protection bits. Pages 0x0 through 0x3 contain the 512 write-protect bits, one bit for each of the 512 data pages within OTP memory. The first two write-protect bit pages (pages 0x0 and 0x1) correspond to the public (non-secure) regions of the OTP map. The other two write-protect bit pages (0x2 and 0x3) correspond to the protection of private (secure) regions of the OTP map. The processor does not need to be operating in Secure Mode in order to be able to program protection pages associated with secure OTP regions. All protection bits can be written in any security state including Open Mode.

 Note that while reads and writes access a half-page at a time, setting a protection bit for a page will effectively lock an entire page for future write accesses (lower and upper half page). The programmer must ensure that all required programming is completed on a full 128-bit OTP data page prior to setting the write-protect bit for that page. In other words, the programmer must make sure that a full 128-bit OTP page is programmed, or that no future programming is required to be performed to the unprogrammed portion of the page before locking the page.

If P is the OTP page that is needed to be write-protected, the write-protect bit and its page can be calculated as follows:

Let WPP be the write-protect page where the write-protect bit resides and let WPB be the write-protect bit that needs to be set in order to lock page P .

```
The write-protect page can be calculated by:  
WPP = P >> 7;  
and the write-protect bit can be calculated by:  
WPB = P & 0x7f;
```

Manual calculation is largely unnecessary due to the fact that the `bfrom_otpWrite()` function can be used to lock pages (see [“OTP Programming Examples” on page 3-26](#) for details).

```
// lock page (note third parameter equals NULL)  
return_code = bfrom_otpWrite( 0x01C, OTP_LOCK, NULL);
```

Locking a single ECC (error correction code) page results in locking the correction codes which correspond to eight OTP data pages (16 half pages). This is due to the fact that a 64-bit half-page access must be performed when write protecting the ECC page and every 8-bits within an ECC page is a parity correction code which corresponds to a 64-bit half-page of data in OTP. Therefore, a full 128-bit ECC page holds the correction codes for eight full 128-bit pages of data in OTP, or 16 half-pages. Pages can only be locked as full 128-bit pages even though read/write accesses may occur at 64-bit half-page granularity. Locking a single ECC page will prevent further write access to the corresponding eight OTP data pages.

ECC (error correction code) space is not permitted to be written to directly.

OTP Programming Examples

For example, locking ECC page 0xFB will result in locking the error correction parity data associated with the 16 data pages in the range of 0x0D8 – 0x0DF.

```
// Only Lock ECC code page
return_code = bfrom_otpWrite(0xFB, OTP_LOCK, NULL);
```

No further write accesses to the ECC page 0xFB or corresponding data pages 0x0D8 – 0x0DF will be allowed following write protection of the ECC page in this example.

 Bits [3:0] of OTP page 0 are the write-protect bits for the first four OTP pages, which contain the write-protect bits. If these bits are set, it will prevent the other write-protect bits from being set, thus disabling the write protection mechanism. But this does not prevent the user from programming the other user-programmable OTP pages.

Accessing Private OTP Memory

In order to read or write to the private area of OTP memory, the processor must be operating in Secure Mode and the `OTPSEN` bit within the `SECURE_SYSSWT` register must be set to a value of 1 to enable secured OTP access. For information about Security, Secure Mode, and the Secure State Machine, see the Secure State Machine section of [Chapter 25, “Security”](#).

OTP Programming Examples

The recommended sequence of steps when accessing OTP memory is as follows:

1. Initialize OTP array by calling `bfrom_otpCommand()`.
2. Perform OTP read or write access by calling `bfrom_otpRead()` or `bfrom_otpWrite()`.

3. Call `bfrom_otpCommand()` with `OTP_CLOSE` parameter to re-initialize the OTP controller when OTP read/write access is complete.
4. Initialize OTP array by calling `bfrom_otpCommand()` for next OTP access.
5. Repeat steps 1–3 for subsequent OTP accesses.

In general, it is recommended to use `OTP_CLOSE` if sensitive data has been written/read in some secure mode, and the processor is subsequently returned to Open Mode operation. For information about these modes, see [Chapter 25, “Security”](#).

To enable access to private OTP memory space while operating in Secure Mode, use the code shown in [Listing 3-3](#).

Listing 3-3. Enable Access to Private OTP

```
// Enable private OTP access
*pSECURE_SYSSWT = ~EMUDABL | OTPSEN;
SSYNC();
...
```

To enable access to private OTP memory space via `OTPSSEN` while operating in Secure Mode, use the code shown in [Listing 3-4](#).

Listing 3-4. Enable Access to Private OTP and Enable JTAG Emulation in Secure Mode

```
// Enable JTAG and private OTP access
*pSECURE_SYSSWT = *pSECURE_SYSSWT & (~EMUDABL) | OTPSEN;
SSYNC(0);
...
```

To read pages 0x4 through 0xDF in public OTP memory space and print results to VisualDSP++ console, use the code shown in [Listing 3-5](#).

OTP Programming Examples

Listing 3-5. Read Pages 0x4 Through 0xDF in Public OTP Memory Space and Print Results to the VisualDSP++ Console

```
#include <blackfin.h>
#include <bfrom.h>

u32 return_code, i;
u64 value;

// Initialize OTP timing parameter
// Proper timing for OTP read access
const u32 OTP_init_value = 0x00001485;
return_code = bfrom_0tpCommand( OTP_INIT, OTP_init_value);
...
for (i= 0x004; <0x0xE0; i++)
{
return_code = bfrom_0tpRead(i, OTP_LOWER_HALF, &value);

printf("page: 0x%03xL, Content ECC: 0x%01611x, returncode:
      0x%03x \n", i, value, return_code);

return_code = bfrom_0tpRead(i, OTP_UPPER_HALF, &value);

printf("page: 0x%03xH, Content ECC: 0x%01611x, returncode:
      0x%03x \n", i, value, return_code);
}
```

To write and lock a single OTP page and return the results to the VisualDSP++ console via `printf`, use the code shown in [Listing 3-6](#).

Listing 3-6. Perform OTP Write to a Single Page via Two 64-Bit (Half-Page) Accesses

```
#include <blackfin.h>
#include <bfrom.h>

u64 value;
u32 return_code;

// Initialize OTP timing parameter
// Proper timing for SCLK = 80 MHz
const u32 OTP_init_value = 0x0A548850;

return_code = bfrom_OtpCommand( OTP_INIT, OTP_init_value);
return_code = bfrom_OtpWrite(0x01C, OTP_LOWER_HALF |
OTP_CHECK_FOR_PREV_WRITE, &testdata);

printf("WRITE page: 0x%03xL, Content ECC: 0x%01611x,
        returncode: 0x%03x \n", 0x1C, testdata, return_code);
return_code = bfrom_OtpWrite(0x01C, OTP_UPPER_HALF |
OTP_CHECK_FOR_PREV_WRITE | OTP_LOCK, &testdata);

printf("WRITE page: 0x%03xH, Content ECC: 0x%01611x,
        returncode: 0x%03x \n", 0x1C, testdata, return_code);
```

Note that locking a page will lock the full 128-bit page, whereas the previous examples perform OTP access on a 64-bit half-page granularity. This is the finest level of granularity that is allowed due to the OTP error correction implementation. The page lock should occur only after both the lower and upper portion of the page have been written. Note that the page lock operation is performed on the second and final access to the page in the code in [Listing 3-6](#).

It may be desired to lock some specific OTP pages in a separate access after writing of data values is already complete.

OTP Programming Examples

OTP pages are typically locked in order to protect them from being overwritten or to prevent inadvertent or malicious tampering. This can be performed by the following instructions in [Listing 3-7](#).

Listing 3-7. Perform Pure Page Lock Operation Without Writing any Data Values

```
#include <blackfin.h>
#include <bfrom.h>

u64 value;
u32 return_code;

// Initialize OTP timing parameter
// Proper timing for SCLK = 80 MHz
const u32 OTP_init_value = 0x14548750;
return_code = bfrom_OtpCommand( OTP_INIT, OTP_init_value);
return_code = bfrom_OtpWrite(0x01C, OTP_LOCK, NULL);
```

Listing 3-8. Read Unique Chip ID Stored in OTP Memory

```
#include <bfrom.h>
#include <stdio.h>
#include <cdefBF518.h>
#include <ccb1kfn.h> // contains intrinsics for Blackfin
                    // assembler commands

void main()
{
    u32 return_code; // 32-bit element to hold return code
    u64 idupper, idlower; // Two 64-bit elements to hold the
                        // upper & lower halves of the unique chip id
```

One-Time Programmable Memory

```
// Code to read the unique chip ID
return_code = bfrom_OtpRead(0x4, OTP_LOWER_HALF, &idlower);

printf("page: 0x%03xL, Content ECC: 0x%016llx, returncode:
      0x%03x\n", 0x4, idlower, return_code);

return_code = bfrom_OtpRead(0x4, OTP_UPPER_HALF, &idupper);

printf("page: 0x%03xH, Content ECC: 0x%016llx, returncode:
      0x%03x\n", 0x4, idupper, return_code);

    return;
}
```

OTP Programming Examples

4 CHIP BUS HIERARCHY

This chapter discusses on-chip buses, how data moves through the system, and other factors that determine the system organization. Following an overview and a list of key features is a block diagram of the chip bus hierarchy and a description of its operation. The chapter concludes with details about the system interconnects and associated system buses.

This chapter provides

- [“Chip Bus Hierarchy Overview”](#)
- [“Interface Overview”](#) on page 4-3

Chip Bus Hierarchy Overview

ADSP-BF51x Blackfin processors feature a powerful chip bus hierarchy on which all data movement between the processor core, internal memory, external memory, and its rich set of peripherals occurs. The chip bus hierarchy includes the controllers for system interrupts, test/emulation, and clock and power management. Synchronous clock domain conversion is provided to support clock domain transactions between the core and the system.

Chip Bus Hierarchy Overview

The processor system includes:

- The peripheral set including timers, GP Counter, RTC, TWI, RSI (except ADSP-BF512), 10/100 Ethernet MAC (ADSP-BF516 and ADSP-BF518), IEEE 1588-2008 (ADSP-BF518), UARTs, SPORTs, SPIs, PPI, watchdog timer, and PWM
- The External Bus Interface Unit (EBIU)
- The Direct Memory Access (DMA) controller
- The interfaces between these, the system, and the optional external (off-chip) resources

The following sections describe the on-chip interfaces between the system and the peripherals via the:

- Peripheral Access Bus (PAB)
- DMA Access Bus (DAB)
- DMA Core Bus (DCB)
- DMA External Bus (DEB)
- External Access Bus (EAB)

The External Bus Interface Unit (EBIU) is the primary chip pin bus and is discussed in [Chapter 7, “External Bus Interface Unit”](#).

Interface Overview

Figure 4-1 shows the core processor and system boundaries as well as the interfaces between them.

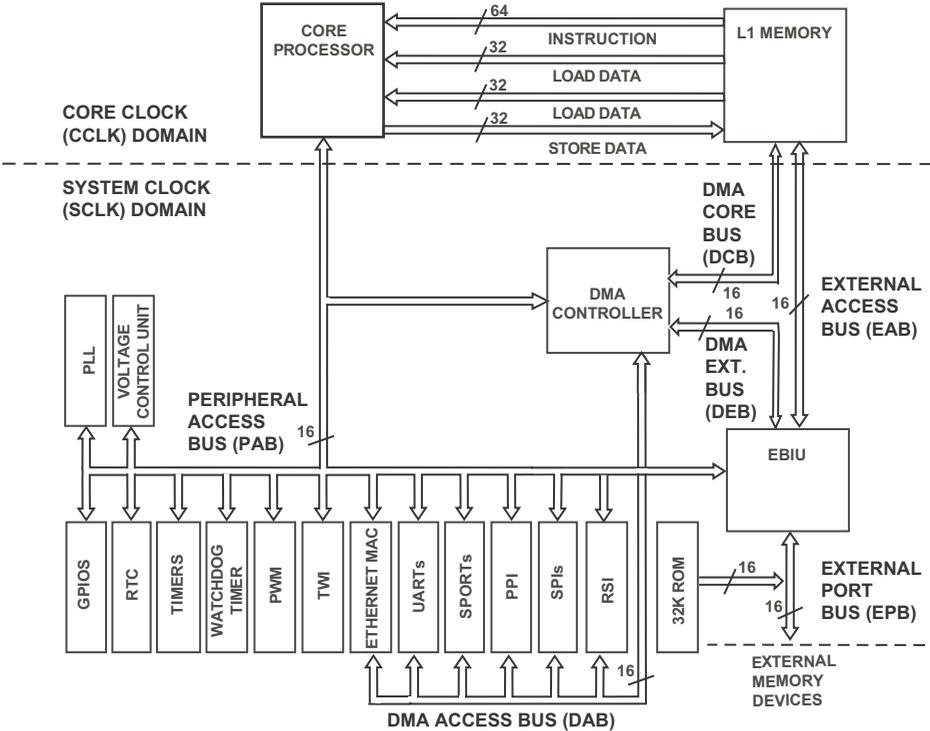


Figure 4-1. Processor Bus Hierarchy

Internal Clocks

The core processor clock (CCLK) rate is highly programmable with respect to CLKIN. The CCLK rate is divided down from the Phase Locked Loop (PLL) output rate. This divider ratio is set using the CSEL parameter of the PLL divide register.

The PAB, the DAB, the EAB, the DCB, the DEB, the EPB, and the EBIU run at system clock frequency (SCLK domain). This divider ratio is set using the SSEL parameter of the PLL divide (PLL_DIV) register and must be set so that these buses run as specified in the processor data sheet, and slower than or equal to the core clock frequency.

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 all synchronous peripherals derive their timing from the SCLK. For example, the UART clock rate is determined by further dividing this clock frequency.

Core Bus Overview

For the purposes of this discussion, level 1 memories (L1) are included in the description of the core; they have full bandwidth access from the processor core with a 64-bit instruction bus and two 32-bit data buses.

Figure 4-2 shows the core processor and its interfaces to the peripherals and external memory resources.

The core can generate up to three simultaneous off-core accesses per cycle.

The core bus structure between the processor and L1 memory runs at the full core frequency and has data paths up to 64 bits.

When the instruction request is filled, the 64-bit read can contain a single 64-bit instruction or any combination of 16-, 32-, or 64-bit (partial) instructions.

When cache is enabled, four 64-bit read requests are issued to support 32-byte line fill burst operations. These requests are pipelined so that each transfer after the first is filled in a single, consecutive cycle.

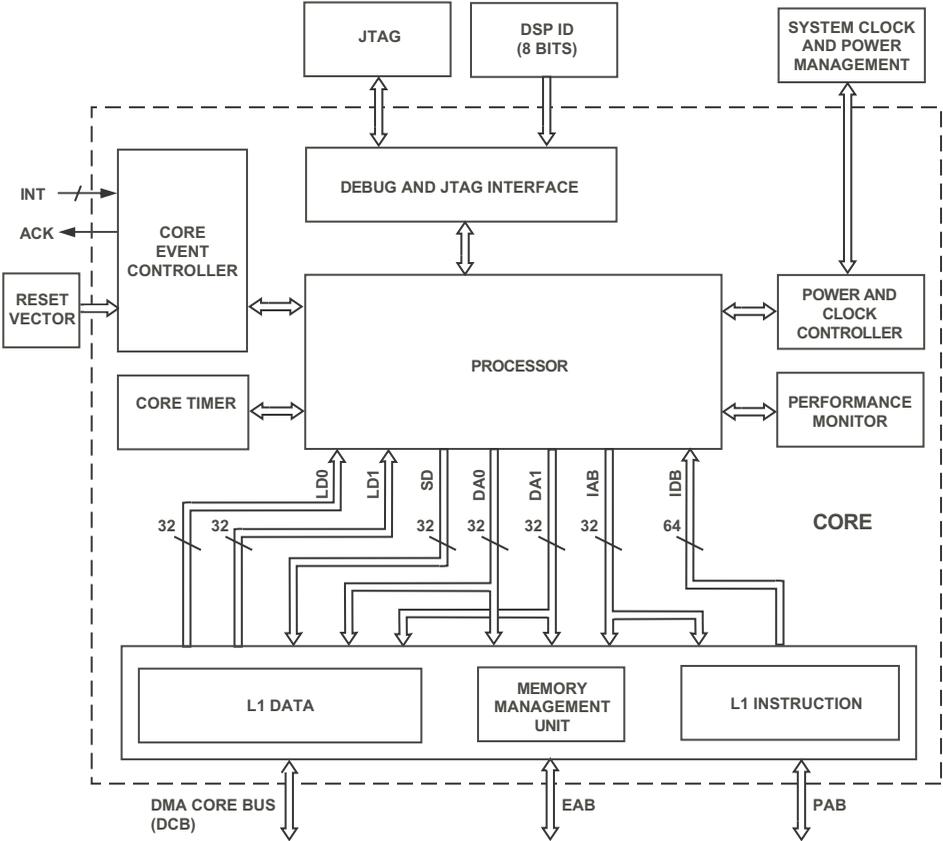


Figure 4-2. Core Block Diagram

Peripheral Access Bus (PAB)

The processor has a dedicated low latency peripheral bus that keeps core stalls to a minimum and allows for manageable interrupt latencies to time-critical peripherals. All peripheral resources accessed through the PAB are mapped into the system MMR space of the processor memory map. The core accesses system MMR space through the PAB bus.

The core processor has byte addressability, but the programming model is restricted to only 32-bit (aligned) access to the system MMRs. Byte accesses to this region are not supported.

PAB Arbitration

The core is the only master on this bus. No arbitration is necessary.

PAB Agents (Masters, Slaves)

The processor core can master bus operations on the PAB. All peripherals have a peripheral bus slave interface which allows the core to access control and status state. These registers are mapped into the system MMR space of the memory map. Appendix B lists system MMR addresses.

The slaves on the PAB bus are:

- System event controller
- Clock and power management controller
- Watchdog timer
- Real-time clock (RTC)
- Timer 0–7
- SPORT0–1
- SPI0–1

- Ports
- UART0–1
- PPI
- TWI
- Ethernet MAC IEEE 1588-2008
- PWM
- RSI
- Asynchronous memory controller (AMC)
- SDRAM controller (SDC)
- DMA controller

PAB Performance

For the PAB, the primary performance criteria is latency, not throughput. Transfer latencies for both read and write transfers on the PAB are two `SCLK` cycles.

For example, the core can transfer up to 32 bits per access to the PAB slaves. With the core clock running at 2x the frequency of the system clock, the first and subsequent system MMR read or write accesses take four core clocks (`CCLK`) of latency.

The PAB has a maximum frequency of `SCLK`.

DMA Access Bus (DAB), DMA Core Bus (DCB), DMA External Bus (DEB)

The DAB, DCB, and DEB buses provide a means for DMA-capable peripherals to gain access to on-chip and off-chip memory with little or no degradation in core bandwidth to memory.

DAB, DCB, and DEB Arbitration

Sixteen DMA channels and bus masters support the DMA-capable peripherals in the processor system. The twelve peripheral DMA channel controllers can transfer data between peripherals and internal or external memory. Both the read and write channels of the dual-stream memory DMA controller access their descriptor lists through the DAB.

The DCB has priority over the core processor on arbitration into L1 configured as SRAM. For off-chip memory, the core (by default) has priority over the DEB for accesses to the EPB. The processor has a programmable priority arbitration policy on the DAB. [Table 4-1](#) shows the default arbitration priority. In addition, by setting the `CDPRIO` bit in the `EBIU_AMGCTL` register, all DEB transactions to the EPB have priority over core accesses to external memory. Use of this bit is application-dependent. For example, if you are polling a peripheral mapped to asynchronous memory with long access times, by default the core will “win” over DMA requests. By setting the `CDPRIO` bit, the core would be held off until DMA requests were serviced.

Table 4-1. DAB, DCB, and DEB Arbitration Priority

DAB, DCB, DEB Master	Default Arbitration Priority
PPI receive/transmit	0 - highest
Ethernet receive	1
Ethernet transmit	2
SPORT0 receive	3
SPORT0 transmit RSI	4
SPORT1 receive SPI1 transmit/receive	5
SPORT1 transmit	6
SPI0 receive/transmit	7
UART0 receive	8
UART0 transmit	9
UART1 receive	10
UART1 transmit	11
MDMA stream 0 destination	12
MDMA stream 0 source	13
MDMA stream 1 destination	14
MDMA stream 1 source	15 - lowest

DAB Bus Agents (Masters)

All peripherals capable of sourcing a DMA access are masters on this bus, as shown in [Table 4-1](#). A single arbiter supports a programmable priority arbitration policy for access to the DAB.

When two or more DMA master channels are actively requesting the DAB, bus utilization is considerably higher due to the DAB's pipelined design. Bus arbitration cycles are concurrent with the previous DMA access's data cycles.

Interface Overview

DAB, DCB, and DEB Performance

The processor DAB supports data transfers of 16 bits or 32 bits. The data bus has a 16-bit width with a maximum frequency as specified in the processor data sheet.

The DAB has a dedicated port into L1 memory. No stalls occur as long as the core access and the DMA access are not to the same memory bank (4K byte size for L1). 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. Latencies caused by these stalls are in addition to any arbitration latencies.



The core processor and the DAB must arbitrate for access to external memory through the EBIU. This additional arbitration latency added to the latency required to read off-chip memory devices can significantly degrade DAB throughput, potentially causing peripheral data buffers to underflow or overflow. If you use DMA peripherals other than the memory DMA controller, and you target external memory for DMA accesses, you need to carefully analyze your specific traffic patterns. Make sure that isochronous peripherals targeting internal memory have enough allocated bandwidth and the appropriate maximum arbitration latencies.

External Access Bus (EAB)

The EAB provides a way for the processor core to directly access off-chip memory.

Arbitration of the External Bus

Arbitration for use of external port bus interface resources is required because of possible contention between the potential masters of this bus. A fixed-priority arbitration scheme is used. That is, core accesses via the EAB will be of higher priority than those from the DMA external bus (DEB).

DEB/EAB Performance

The DEB and the EAB support single word accesses of either 8-bit or 16-bit data types. The DEB and the EAB operate at the same frequency as the PAB and the DAB, up to the maximum `SCLK` frequency specified in the processor data sheet.

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 4-2](#) shows many types of 16-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) is always more efficient than executing two 16-bit accesses (of the form $R0 = W[P0++]$; where $P0$ points to an address in SDRAM). In

Interface Overview

this example, a 32-bit SDRAM read takes 10 SCLK cycles while two 16-bit reads take 9 SCLK cycles each.

Table 4-2. Performance of DMA Access to External Memory

Source	Destination	Approximate SCLKs For n Words (from start of DMA to interrupt at end)
16-bit SDRAM	L1 data memory	$n + 14$
L1 data memory	16-bit SDRAM	$n + 11$
16-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	16-bit async memory	$xn + 9$, where x is the number of wait states + setup/hold SCLK cycles (minimum $x = 2$)
16-bit SDRAM	16-bit SDRAM	$10 + (17n/7)$
16-bit async memory	16-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	$2n + 12$

5 SYSTEM INTERRUPTS

This chapter discusses the system interrupt controller (SIC). While this chapter does refer to features of the core event controller (CEC), it does not cover all aspects of it. Please refer to the *Blackfin Processor Programming Reference* for more information on the CEC.

Specific Information for the ADSP-BF51x

For details regarding the number of system interrupts for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

To determine how each of the system interrupts is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each DMA, refer to [Chapter A, “System MMR Assignments”](#).

System interrupt behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 5-15](#).

Overview

The processor system has numerous peripherals, which therefore require many supporting interrupts.

Description of Operation

Features

The Blackfin architecture provides a two-level interrupt processing scheme:

- The core event controller (CEC) runs in the `CCLK` clock domain. It interacts closely with the program sequencer and manages the event vector table (EVT). The CEC processes not only core-related interrupts such as exceptions, core errors, and emulation events; it also supports software interrupts.
- The system interrupt controller (SIC) runs in the `SCLK` clock domain. It masks, groups, and prioritizes interrupt requests signalled by on-chip or off-chip peripherals and forwards them to the CEC.

Description of Operation

The following sections describe the operation of the system interrupts.

Events and Sequencing

The processor employs a two-level event control mechanism. The processor SIC works with the 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 of the processor manages five types of activities or events:

- Emulation
- Reset
- Nonmaskable interrupts (NMI)

- Exceptions
- Interrupts

Note the word *event* describes all five types of activities. The CEC manages fifteen different events in all: emulation, reset, NMI, exception, and eleven interrupts.

An interrupt is an event that changes the 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 pre-empted by one of higher priority.

The CEC supports nine general-purpose interrupts (IVG7 – IVG15) in addition to the dedicated interrupt and exception events that are described in [Table 5-1](#). It is common for applications to reserve the lowest or the two lowest priority interrupts (IVG14 and IVG15) for software interrupts, leaving eight or seven prioritized interrupt inputs (IVG7 – IVG13) for peripheral purposes. Refer to [Table 5-1](#).

Table 5-1. 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

Description of Operation

Table 5-1. System and Core Event Mapping (Continued)

Event Source	Core Event Name
System interrupts	IVG7–IVG13
Software interrupt 1	IVG14
Software interrupt 2 (lowest priority)	IVG15

System Peripheral Interrupts

To service the rich set of peripherals, the SIC has multiple interrupt request inputs and outputs that go to the CEC. The primary function of the SIC is to mask, group, and prioritize interrupt requests and to forward them to the nine general-purpose interrupt inputs of the CEC (IVG7–IVG15). Additionally, the SIC controller can enable individual peripheral interrupts to wake up the processor from Idle or power-down state.

The nine general-purpose interrupt inputs (IVG7–IVG15) of the core event controller have fixed priority. Of this group, the IVG7 channel has the highest priority and IVG15 has the lowest priority. Therefore, the interrupt assignment in the SIC_IAR registers not only groups peripheral interrupts; it also programs their priority by assigning them to individual IVG channels. However, 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 SIC_IAR register settings shown in [Figure 5-2 on page 5-11](#) and the tables in [Appendix A, “System MMR Assignments”](#). 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.

 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.

The core timer has a dedicated input to the CEC controller. Its interrupt is not routed through the SIC controller and always has higher priority than requests from all peripherals.

The `SIC_IMASK` register allows software to mask any peripheral interrupt source at the SIC level. This functionality is independent of whether the particular interrupt is enabled at the peripheral itself. At reset, the contents of the `SIC_IMASK` register are all 0s to mask off all peripheral interrupts. Turning off a system interrupt mask and enabling the particular interrupt is performed by writing a 1 to a bit location in the `SIC_IMASK` register.

The SIC includes one or more read-only `SIC_ISR` registers with individual bits which correspond to the interrupt status of one of the peripheral interrupt sources. When the SIC detects the interrupt, the bit is asserted. When the SIC detects that the peripheral interrupt input has been deasserted, the respective bit in the system interrupt status register is cleared. Note for some peripherals, such as general-purpose I/O asynchronous input interrupts, many cycles of latency may pass from the time an interrupt service routine initiates the clearing of the interrupt (usually by writing a system MMR) to the time the SIC senses that the interrupt has been deasserted.

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 the `SIC_ISR` register 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.

Description of Operation

 When an interrupt's service routine is finished, the RTI instruction clears the appropriate bit in the `IPEND` register. However, the relevant `SIC_ISR` 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, the `SIC_ISR` register will seldom, if ever, need to be interrogated.

The `SIC_ISR` register is not affected by the state of the `SIC_IMASK` register and can be read at any time. Writes to the `SIC_ISR` register have no effect on its contents.

Peripheral DMA channels are mapped in a fixed manner to the peripheral interrupt IDs. However, the assignment between peripherals and DMA channels is freely programmable with the `DMA_PERIPHERAL_MAP` registers. [Table 5-1 on page 5-3](#) and [Table 5-2 on page 5-11](#) show the default DMA assignment. Once a peripheral has been assigned to any other DMA channel it uses the new DMA channel's interrupt ID regardless of whether DMA is enabled or not. Therefore, clean `DMA_PERIPHERAL_MAP` management is required even if the DMA is not used. The default setup should be the best choice for all non-DMA applications.

For dynamic power management, 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 `SIC_IWR` register (refer to [Table 5-1 on page 5-3](#) and [Table 5-2 on page 5-11](#)). If a peripheral interrupt source is enabled in `SIC_IWR` and the core is idled, the interrupt causes the DPMC to initiate the core wakeup sequence in order to process the interrupt. Note 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 the Dynamic Power Management chapter.

The `SIC_IWR` register has no effect unless the core is idled. By default, all interrupts generate a wakeup request to the core. However, for some

applications it may be desirable to disable this function for some peripherals, such as for a SPORT transmit interrupt. The `SIC_IWR` register can be read from or written to at any time. To prevent spurious or lost interrupt activity, this register should be written to only when all peripheral interrupts are disabled.

 The wakeup function is independent of the interrupt mask function. If an interrupt source is enabled in the `SIC_IWR` but masked off in the `SIC_IMASK` register, the core wakes up if it is idled, but it does not generate an interrupt.

The peripheral interrupt structure of the processor is flexible. Upon reset, multiple peripheral interrupts share a single, general-purpose interrupt in the core by default, as shown in [Table 5-2 on page 5-11](#).

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

Programming Model

The programming model for the system interrupts is described in the following sections.

System Interrupt Initialization

If the default peripheral-to-IVG assignments shown in [Table 5-1 on page 5-3](#) and [Table 5-2 on page 5-11](#) are acceptable, then interrupt initialization involves only:

- Initialization of the core event vector table (EVT) vector address entries
- Initialization of the IMASK register
- Unmasking the specific peripheral interrupts that the system requires in the SIC_IMASK register

System Interrupt Processing Summary

Referring to [Figure 5-1 on page 5-10](#), note when an interrupt (interrupt A) is generated by an interrupt-enabled peripheral:

1. SIC_ISR 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_IWR checks to see if it should wake up the core from an idled state based on this interrupt request.
3. SIC_IMASK 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_IAR 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_ISR` 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.

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_IWR`, `SIC_ISR`, `SIC_IMASK`, `SIC_IAR`).

If multiple interrupt sources share a single core interrupt, then the interrupt service routine (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 Interrupt Controller Registers

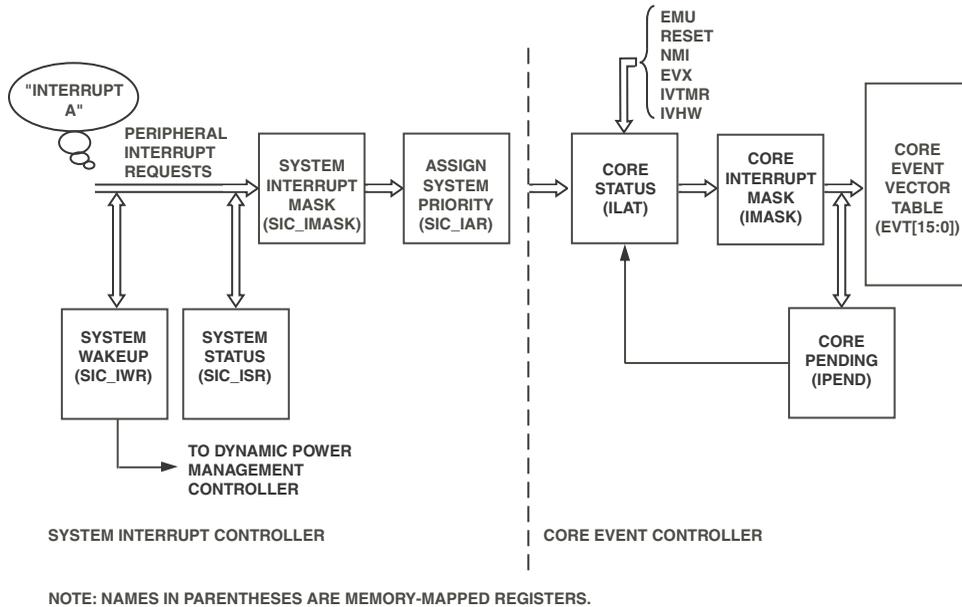


Figure 5-1. Interrupt Processing Block Diagram

System Interrupt Controller Registers

The SIC registers are described in the following sections.

These registers can be read from or written to 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 (SIC_IAR) Register

The SIC_IAR register maps each peripheral interrupt ID to a corresponding IVG priority level. This is accomplished with 4-bit groupings that translate to IVG levels as shown in Table 5-2 and Figure 5-2. In other words, Table 5-2 defines the value to write in a 4-bit field within SIC_IAR in order to configure a peripheral interrupt ID for a particular IVG priority. Refer to Table 5-1 on page 5-3 for information on SIC_IAR mappings for this specific processor.

System Interrupt Assignment Register (SIC_IAR)

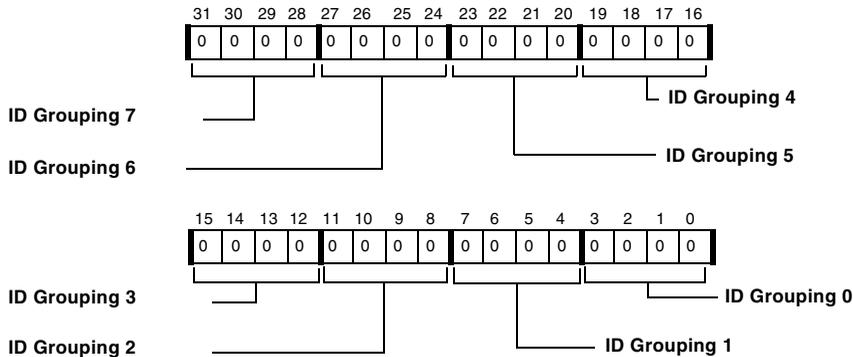


Figure 5-2. System Interrupt Assignment Register

Table 5-2. IVG Select Definition

General-Purpose Interrupt	Value in SIC_IAR
IVG7	0
IVG8	1
IVG9	2
IVG10	3
IVG11	4
IVG12	5

System Interrupt Controller Registers

Table 5-2. IVG Select Definition (Continued)

General-Purpose Interrupt	Value in SIC_IAR
IVG13	6
IVG14	7
IVG15	8

System Interrupt Mask (SIC_IMASK) Register

The `SIC_IMASK` register masks or enables peripheral interrupts at the system level. A "0" in a bit position masks off (disables) interrupts for that particular peripheral interrupt ID. A "1" enables interrupts for that interrupt ID. Refer to [Table 5-1 on page 5-3](#) and [Table 5-2](#) for information on how peripheral interrupt IDs are mapped to the `SIC_IMASK` register(s) for this particular processor.

System Interrupt Status (SIC_ISR) Register

The `SIC_ISR` register keeps track of system interrupts that are asserted but not yet serviced. A "0" in a bit position indicates that a particular interrupt is deasserted. A "1" indicates that it is asserted. Refer to [Table 5-1 on page 5-3](#) and [Table 5-2](#) for information on how peripheral interrupt IDs are mapped to the `SIC_ISR` register(s) for this particular processor.

System Interrupt Wakeup-Enable (SIC_IWR) Register

The `SIC_IWR` register allows an interrupt request to wake up the processor core from an idled state. A "0" in a bit position indicates that a particular peripheral interrupt ID is not configured to wake the core (upon assertion of the interrupt request). A "1" indicates that it is configured to do so. Refer to [Table 5-1 on page 5-3](#) and [Table 5-2](#) for information on how

peripheral interrupt IDs are mapped to the `SIC_IWR` register(s) for this particular processor.

Programming Examples

The following section provides an example for servicing interrupt requests.

Clearing Interrupt Requests

When the processor services a core event it automatically clears the requesting bit in the `ILAT` register and no further action is required by the interrupt service routine. It is important to understand that the SIC controller does not provide any interrupt acknowledgment feedback mechanism from the CEC controller back to the peripherals. Although the `ILAT` bits clear in the same way when a peripheral interrupt is serviced, the signalling peripheral does not release its level-sensitive request until it is explicitly instructed by software. If however, the peripheral keeps requesting, the respective `ILAT` bit is set again immediately and the service routine is invoked again as soon as its first run terminates by an `RTI` instruction.

Every software routine that services peripheral interrupts must clear the signalling interrupt request in the respective peripheral. The individual peripherals provide customized mechanisms for how to clear interrupt requests. Receive interrupts, for example, are cleared when received data is read from the respective buffers. Transmit requests typically clear when software (or DMA) writes new data into the transmit buffers. These implicit acknowledge mechanisms avoid the need for cycle-consuming software handshakes in streaming interfaces. Other peripherals such as timers, GPIOs, and error requests require explicit acknowledge instructions, which are typically performed by efficient `W1C` (write-1-to-clear) operations.

Programming Examples

[Listing 5-1](#) shows a representative example of how a GPIO interrupt request might be serviced.

Listing 5-1. Servicing GPIO Interrupt Request

```
#include <defBF527.h>
/*ADSP-BF527 product is used as an example*/
.section program;
_portg_a_isr:
    /* push used registers */
    [--sp] = (r7:7, p5:5);
    /* clear interrupt request on GPIO pin PG2 */
    /* no matter whether used A or B channel */
    p5.l = lo(PORTGIO_CLEAR);
    p5.h = hi(PORTGIO_CLEAR);
    r7 = PG2;
    w[p5] = r7;

    /* place user code here */

    /* sync system, pop registers and exit */
    ssync;
    (r7:7, p5:5) = [sp++];
    rti;
_portg_a_isr.end;
```

The `WIC` instruction shown in this example may require several `SCLK` cycles to complete, depending on system load and instruction history. The program sequencer does not wait until the instruction completes and continues program execution immediately. The `SSYNC` instruction ensures that the `WIC` command indeed cleared the request in the GPIO peripheral before the `RTI` instruction executes. However, the `SSYNC` instruction does not guarantee that the release of interrupt request has also been recognized by the CEC controller, which may require a few more `CCLK` cycles depending on the `CCLK-to-SCLK` ratio. In service routines consisting of a few

instructions only, two `SSYNC` instructions are recommended between the clear command and the `RTI` instruction. However, one `SSYNC` instruction is typically sufficient if the clear command performs in the very beginning of the service routine, or the `SSYNC` instruction is followed by another set of instructions before the service routine returns. Commonly, a pop-multiple instruction is used for this purpose as shown in [Listing 5-1](#).

The level-sensitive nature of peripheral interrupts enables more than one of them to share the same IVG channel and therefore the same interrupt priority. This is programmable using the assignment registers. Then a common service routine typically interrogates the `SIC_ISR` register to determine the signalling interrupt source. If multiple peripherals are requesting interrupts at the same time, it is up to the service routine to either service all requests in a single pass or to service them one by one. If only one request is serviced and the respective request is cleared by software before the `RTI` instruction executes, the same service routine is invoked another time because the second request is still pending. While the first approach may require fewer cycles to service both requests, the second approach enables higher priority requests to be serviced more quickly in a non-nested interrupt system setup.

Unique Information for the ADSP-BF51x Processor

Interfaces

[Figure 5-3](#) and [Figure 5-4](#) provide an overview of how the individual peripheral interrupt request lines connect to the SIC. These figures show how the eight `SIC_IAR` registers control the assignment to the nine available peripheral request inputs of the CEC.



The memory-mapped `ILAT`, `IMASK`, and `IPEND` registers are part of the CEC controller.

Unique Information for the ADSP-BF51x Processor

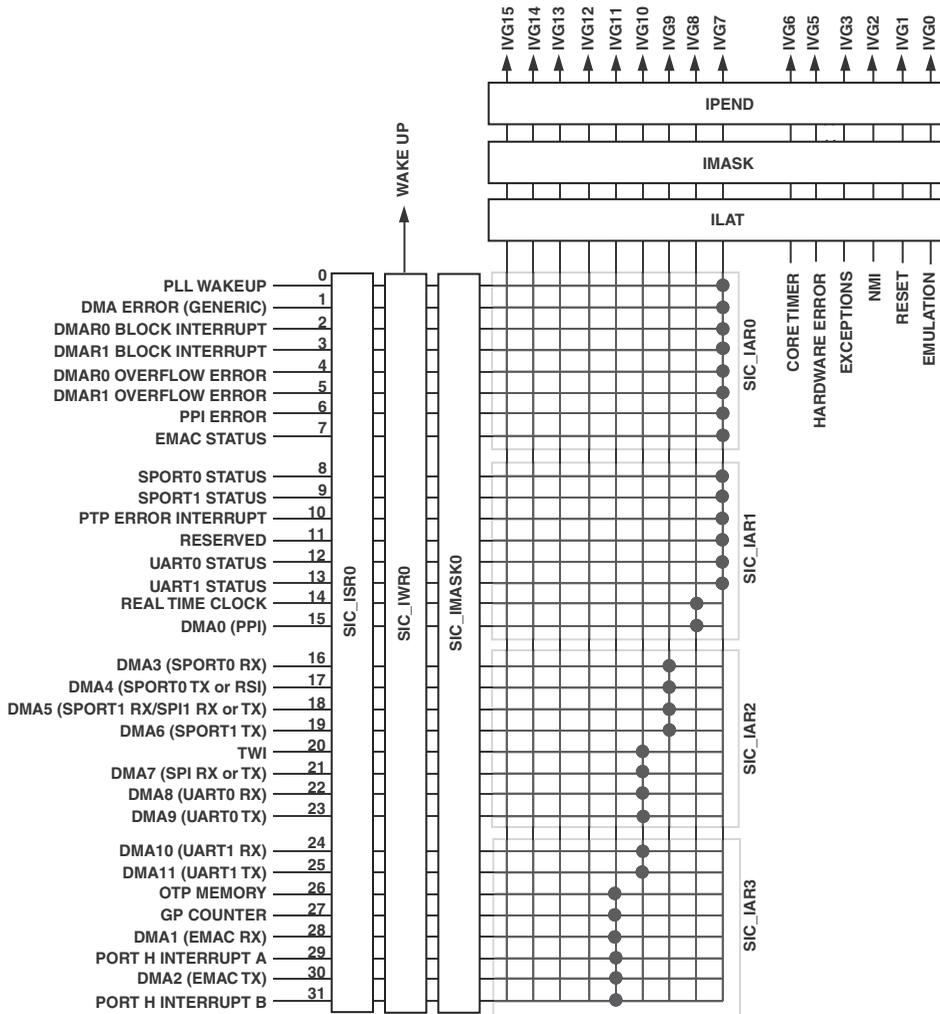


Figure 5-3. Interrupt Routing Overview (Part 1)

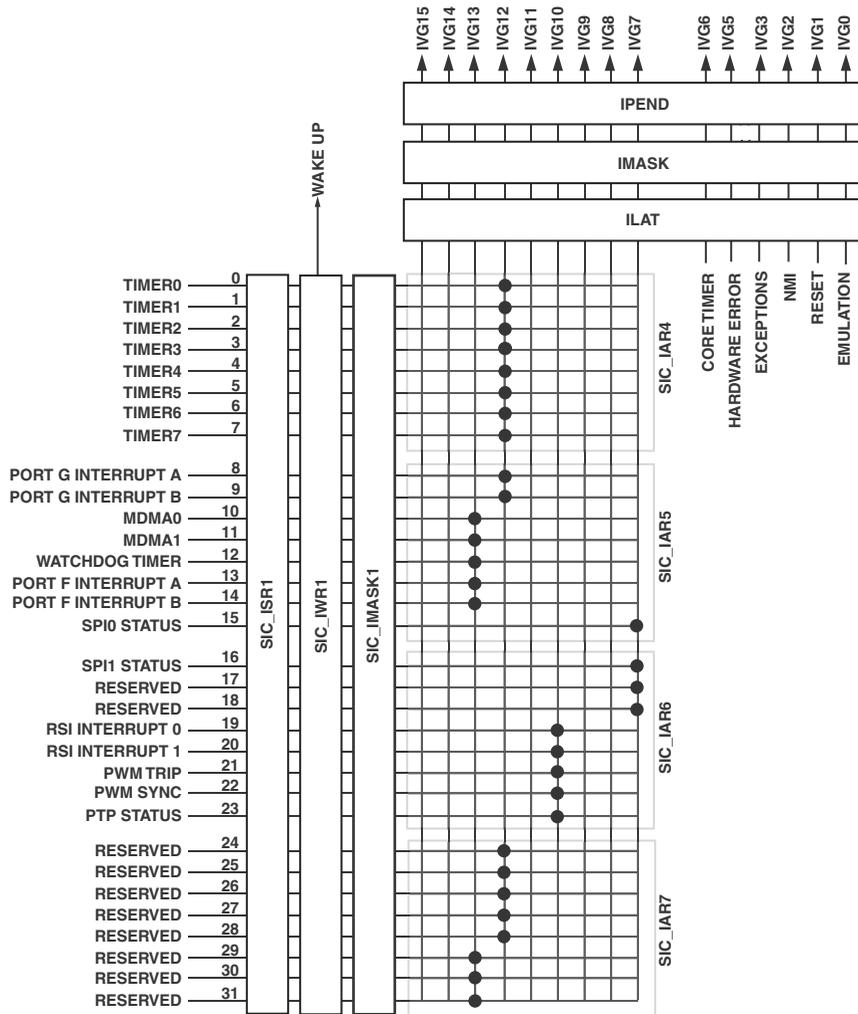


Figure 5-4. Interrupt Routing Overview (Part 2)

System Peripheral Interrupts

The MAC interrupt requests shown in [Figure 5-3 on page 5-16](#) are available only on ADSP-BF516 and ADSP-BF518 parts. However, for code compatibility, all of the default assignments for the ADSP-BF51x processors are the same.

[Table 5-3 on page 5-19](#) and [Table 5-4 on page 5-20](#) show the peripheral interrupt events, the default mapping of each event, the peripheral interrupt ID used in the system interrupt assignment registers (SIC_IAR), and the core interrupt ID.

Note that the system interrupt to core event mappings shown are the default values at reset and can be changed by software. Where there is more than one DMA interrupt source for a given interrupt ID number, the default DMA source mapping is listed first in parentheses.

The peripheral interrupt structure of the processor is flexible. Upon reset, multiple peripheral interrupts share a single, general-purpose interrupt in the core by default, as shown in [Table 5-3 on page 5-19](#) and [Table 5-4 on page 5-20](#).

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

Table 5-3. Peripheral Interrupt Events (Part 1)

Peripheral ID Number	Bit Position for SIC_ISR0, SIC_IMASK0, SIC_IWR0	SIC_IAR3-0	Interrupt Source	Default Mapping
31	Bit 31	SIC_IAR3[31:28]	Port H interrupt B	IVG11
30	Bit 30	SIC_IAR3[27:24]	DMA2 (Ethernet MAC TX)	IVG11
29	Bit 29	SIC_IAR3[23:20]	Port H interrupt A	IVG11
28	Bit 28	SIC_IAR3[19:16]	DMA1 (Ethernet MAC RX)	IVG11
27	Bit 27	SIC_IAR3[15:12]	GP Counter	IVG11
26	Bit 26	SIC_IAR3[11:8]	OTP Memory	IVG11
25	Bit 25	SIC_IAR3[7:4]	DMA11 (UART1 TX)	IVG10
24	Bit 24	SIC_IAR3[3:0]	DMA10 (UART1 RX)	IVG10
23	Bit 23	SIC_IAR2[31:28]	DMA9 (UART0 TX)	IVG10
22	Bit 22	SIC_IAR2[27:24]	DMA8 (UART0 RX)	IVG10
21	Bit 21	SIC_IAR2[23:20]	DMA7 (SPI0 RX or TX)	IVG10
20	Bit 20	SIC_IAR2[19:16]	TWI	IVG10
19	Bit 19	SIC_IAR2[15:12]	DMA6 (SPORT1 TX)	IVG9
18	Bit 18	SIC_IAR2[11:8]	DMA5 (SPORT1 RX/SPI1 RX or TX)	IVG9
17	Bit 17	SIC_IAR2[7:4]	DMA4 (SPORT0 TX/RSI)	IVG9
16	Bit 16	SIC_IAR2[3:0]	DMA3 (SPORT0 RX)	IVG9
15	Bit 15	SIC_IAR1[31:28]	DMA0 (PPI)	IVG8
14	Bit 14	SIC_IAR1[27:24]	Real-time clock	IVG8
13	Bit 13	SIC_IAR1[23:20]	UART1 status	IVG7
12	Bit 12	SIC_IAR1[19:16]	UART0 status	IVG7
11	Bit 11	SIC_IAR1[15:12]	Reserved	IVG7
10	Bit 10	SIC_IAR1[11:8]	PTP Error Interrupt	IVG7
9	Bit 9	SIC_IAR1[7:4]	SPORT1 status	IVG7
8	Bit 8	SIC_IAR1[3:0]	SPORT0 status	IVG7

Unique Information for the ADSP-BF51x Processor

Table 5-3. Peripheral Interrupt Events (Part 1) (Continued)

Peripheral ID Number	Bit Position for SIC_ISR0, SIC_IMASK0, SIC_IWR0	SIC_IAR3-0	Interrupt Source	Default Mapping
7	Bit 7	SIC_IAR0[31:28]	Ethernet MAC status	IVG7
6	Bit 6	SIC_IAR0[27:24]	PPI error	IVG7
5	Bit 5	SIC_IAR0[23:20]	DMAR1 overflow error	IVG7
4	Bit 4	SIC_IAR0[19:16]	DMAR0 overflow error	IVG7
3	Bit 3	SIC_IAR0[15:12]	DMAR1 block interrupt	IVG7
2	Bit 2	SIC_IAR0[11:8]	DMAR0 block interrupt	IVG7
1	Bit 1	SIC_IAR0[7:4]	DMA Error (generic)	IVG7
0	Bit 0	SIC_IAR0[3:0]	PLL Wakeup	IVG7

Table 5-4. Peripheral Interrupt Events (Part 2)

Peripheral ID Number	Bit Position for SIC_ISR1, SIC_IMASK1, SIC_IWR1	SIC_IAR7-4	Interrupt Source	Default Mapping
63	Bit 31	SIC_IAR7[31:28]	Reserved	IVG13
62	Bit 30	SIC_IAR7[27:24]	Reserved	IVG13
61	Bit 29	SIC_IAR7[23:20]	Reserved	IVG13
60	Bit 28	SIC_IAR7[19:16]	Reserved	IVG12
59	Bit 27	SIC_IAR7[15:12]	Reserved	IVG12
58	Bit 26	SIC_IAR7[11:8]	Reserved	IVG12
57	Bit 25	SIC_IAR7[7:4]	Reserved	IVG12
56	Bit 24	SIC_IAR7[3:0]	Reserved	IVG12
55	Bit 23	SIC_IAR6[31:28]	PTP Status	IVG10
54	Bit 22	SIC_IAR6[27:24]	PWM Sync	IVG10
53	Bit 21	SIC_IAR6[23:20]	PWM Trip	IVG10

Table 5-4. Peripheral Interrupt Events (Part 2) (Continued)

Peripheral ID Number	Bit Position for SIC_ISR1, SIC_IMASK1, SIC_IWR1	SIC_IAR7-4	Interrupt Source	Default Mapping
52	Bit 20	SIC_IAR6[19:16]	RSI Interrupt 1	IVG10
51	Bit 19	SIC_IAR6[15:12]	RSI Interrupt 0	IVG10
50	Bit 18	SIC_IAR6[11:8]	Reserved	IVG7
49	Bit 17	SIC_IAR6[7:4]	Reserved	IVG7
48	Bit 16	SIC_IAR6[3:0]	SPI1 status	IVG7
47	Bit 15	SIC_IAR5[31:28]	SPI0 status	IVG7
46	Bit 14	SIC_IAR5[27:24]	Port F interrupt B	IVG13
45	Bit 13	SIC_IAR5[23:20]	Port F interrupt A	IVG13
44	Bit 12	SIC_IAR5[19:16]	Watchdog timer	IVG13
43	Bit 11	SIC_IAR5[15:12]	MDMA1	IVG13
42	Bit 10	SIC_IAR5[11:8]	MDMA0	IVG13
41	Bit 9	SIC_IAR5[7:4]	Port G interrupt B	IVG12
40	Bit 8	SIC_IAR5[3:0]	Port G interrupt A	IVG12
39	Bit 7	SIC_IAR4[31:28]	Timer 7	IVG12
38	Bit 6	SIC_IAR4[27:24]	Timer 6	IVG12
37	Bit 5	SIC_IAR4[23:20]	Timer 5	IVG12
36	Bit 4	SIC_IAR4[19:16]	Timer 4	IVG12
35	Bit 3	SIC_IAR4[15:12]	Timer 3	IVG12
34	Bit 2	SIC_IAR4[11:8]	Timer 2	IVG12
33	Bit 1	SIC_IAR4[7:4]	Timer 1	IVG12
32	Bit 0	SIC_IAR4[3:0]	Timer 0	IVG12

Unique Information for the ADSP-BF51x Processor

6 DIRECT MEMORY ACCESS

This chapter describes the direct memory access (DMA) controller. Following an overview and list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

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 [Chapter 4, “Chip Bus Hierarchy”](#).

Specific Information for the ADSP-BF51x

For details regarding the number of DMA controllers for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For DMA interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the DMAs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each DMA, refer to [Chapter A, “System MMR Assignments”](#).

Overview and Features

DMA controller behavior for the ADSP-BF51x that differs from the general information in this chapter can be found in the section [“Unique Information for the ADSP-BF51x Processor”](#) on page 6-102.

Overview and Features

The processor uses DMA to transfer data between 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.

The DMA controller can perform several types of data transfers:

- Peripheral DMA transfers data between memory and on-chip peripherals.
- Memory DMA (MDMA) transfers data between memory and memory. The processor has two MDMA modules, each consisting of independent memory read and memory write channels.
- Handshaking memory DMA (HMDMA) transfers data between off-chip peripherals and memory. This enhancement of the MDMA channels enables external hardware to control the timing of individual data transfers or block transfers.



The HMDMA feature is not available for all products. Refer to [“Unique Information for the ADSP-BF51x Processor”](#) on page 6-102 to determine whether it applies to this product.

All DMAs can transport data to and from on-chip and off-chip memories, including L1 and SDRAM. The L1 scratchpad memory cannot be accessed by DMA.

 SDRAM and SRAM are not available on all products. Refer to [“Unique Information for the ADSP-BF51x Processor” on page 6-102](#) to determine whether it applies to this product.

DMA transfers on the processor can be descriptor-based or register-based.

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.

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.

Examples of DMA styles supported by flex descriptors include:

- A single linear buffer that stops on completion (FLOW = stop mode)
- A linear buffer with byte strides of any integer value, including negative values (DMA_X_X_MODIFY register)
- 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) (2-D DMA)
- 1-D DMA, using a set of identical ping-pong buffers defined by a linked ring of 3-word descriptors, each containing a link pointer and a 32-bit address
- 1-D DMA, using a linked list of 5-word descriptors containing a link pointer, a 32-bit address, the buffer length, and a configuration

DMA Controller Overview

- 2-D DMA, using an array of 1-word descriptors, specifying only the base DMA address within a common data page
- 2-D DMA, using a linked list of 9-word descriptors specifying everything

DMA Controller Overview

A block diagram of the DMA controller can be found in the [“Unique Information for the ADSP-BF51x Processor”](#) on page 6-102.

External Interfaces

The DMA does not connect external memories and devices directly. Rather, data is passed through the EBIU port. Any kind of device that is supported by the EBIU can also be accessed by peripheral DMA or memory DMA operation. This is typically flash memory, SRAM, SDRAM, FIFOs, or memory-mapped peripheral devices.

For products with handshaking MDMA (HMDMA), the operation is supported by two MDMA request input pins, `DMAR0` and `DMAR1`. The `DMAR0` pin controls transfer timing on the `MDMA0` destination channel. The `DMAR1` pin controls the destination channel of `MDMA1`. With these pins, external FIFO devices, ADC or DAC converters, or other streaming or block-processing devices can use the MDMA channels to exchange their data or data buffers with the Blackfin processor memory.

Internal Interfaces

[Figure 4-1 on page 4-3](#) shows the dedicated DMA buses used by the DMA controller to interconnect L1 memory, the on-chip peripherals, and the EBIU port.

The 16-bit DMA core bus (DCB) connects the DMA controller to a dedicated port of L1 memory. L1 memory has dedicated DMA ports featuring special DMA buffers to decouple DMA operation. See the *Blackfin Processor Programming Reference* for a description of the L1 memory architecture. The DCB bus operates at core clock (CCLK) frequency. It is the DMA controller's responsibility to translate DCB transfers to the system clock (SCLK) domain.

The 16-bit DMA access bus (DAB) connects the DMA controller to the on-chip peripherals. This bus operates at SCLK frequency.

The 16-bit DMA external bus (DEB) connects the DMA controller to the EBIU port. This bus is used for all peripheral and memory DMA transfers to and from external memories and devices. It operates at SCLK frequency.

Transferred data can be 8-, 16-, or 32-bits wide. The DMA controller, however, connects only to 16-bit buses.

Memory DMA can pass data every SCLK cycle between L1 memory and the EBIU. Transfers from L1 memory to L1 memory require two cycles, as the DCB bus is used for both source and destination transfers. Similarly, transfers between two off-chip devices require EBIU and DEB resources twice. Peripheral DMA transfers can be performed every other SCLK cycle.

For more details on DMA performance see [“DMA Performance” on page 6-40](#).

Peripheral DMA

The DMA controller features 12 channels that perform transfers between peripherals and on-chip or off-chip memories. The user has full control over the mapping of DMA channels and peripherals. The default DMA channel priority and mapping, shown in [Table 6-7 on page 6-104](#), can be changed by altering the 4-bit PMAP field in the DMAx_PERIPHERAL_MAP registers for the peripheral DMA channels.

DMA Controller Overview

The default configuration should suffice in most cases, but there are some cases where remapping the assignment can be helpful because of the DMA channel priorities. When competing for any of the system buses, DMA0 has higher priority than DMA1, and so on. DMA11 has the lowest priority of the peripheral DMA channels.

 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.

All peripheral DMA channels work completely independently from each other. The transfer timing is controlled by the mapped peripheral.

Every DMA channel features its own 4-deep FIFO that decouples DAB activity from DCB and DEB availability. DMA interrupt and descriptor fetch timing is aligned with the memory side (DCB/DEB side) of the FIFO. The user does, however, have an option to align interrupts with the peripheral side (DAB side) of the FIFO for transmit operations. Refer to the `SYNC` bit in the `DMAx_CONFIG` register for details.

Memory DMA

This section describes the two pairs of MDMA channels, which provide memory-to-memory DMA transfers among the various memory spaces. These include L1 memory and external synchronous/asynchronous memories.

Each MDMA channel contains a DMA FIFO, an 8-word by 16-bit FIFO block used to transfer data to and from either L1 or the DCB and DEB

buses. Typically, it is used to transfer data between external memory and internal memory. It will also support DMA from the boot ROM on the DEB bus. The FIFO can be used to hold DMA data transferred between two L1 memory locations or between two external memory locations.

Each page of MDMA channels consists of:

- A source channel (for reading from memory)
- A destination channel (for writing to memory)

A memory-to-memory transfer always requires both the source and the destination channel to be enabled. Each source/destination channel forms a “stream,” and these two streams are hardwired for DMA priorities 12 through 15.

- Priority 12: MDMA0 destination
- Priority 13: MDMA0 source
- Priority 14: MDMA1 destination
- Priority 15: MDMA1 source

MDMA0 takes precedence over MDMA1, unless round-robin scheduling is used or priorities become urgent, as programmed by the DRQ bit field in the HMDMA_CONTROL register.



It is illegal to program a source channel for memory write or a destination channel for memory read.

The channels support 8-, 16-, and 32-bit memory DMA transfers, but both ends of the MDMA connect to 16-bit buses. Source and destination channels must be programmed to the same word size. In other words, the MDMA transfer does not perform packing or unpacking of data; each read results in one write. Both ends of the MDMA FIFO for a given stream are granted priority at the same time. Each pair shares an 8-word deep 16-bit FIFO. The source DMA engine fills the FIFO, while the

DMA Controller Overview

destination DMA engine empties it. The FIFO depth allows the burst transfers of the external access bus (EAB) and DMA access bus (DAB) 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 MDMA pair, 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 transfer count is the same.

To start a MDMA transfer operation, the MMRs for the source and destination channels are written, each in a manner similar to peripheral DMA.



The `DMAx_CONFIG` register for the source channel must be written before the `DMAx_CONFIG` register for the destination channel.

Handshaked Memory DMA (HMDMA) Mode

This feature is not available for all products. Refer to the [“Unique Information for the ADSP-BF51x Processor” on page 6-102](#) to determine whether it applies to this product.

Handshaked operation applies only to memory DMA channels.

Normally, memory DMA transfers are performed at maximum speed. Once started, data is transferred in a continuous manner until either the data count expires or the MDMA is stopped. In handshake mode, the MDMA does not transfer data automatically when enabled; it waits for an external trigger on the MDMA request input signals. The `DMAR0` input is associated with MDMA0 and the `DMAR1` input with MDMA1. Once a trigger event is detected, a programmable portion of data is transferred and then the MDMA stalls again and waits for the next trigger.

Handshake operation is not only useful for controlling the timing of memory-to-memory transfers, it also enables the MDMA to operate with

asynchronous FIFO-style devices connected to the EBIU port. The Blackfin processor acknowledges a DMA request by a proper number of read or write operations. It is up to the device connected to any of the `AMSx` strobes to deassert or pulse the request signal and to decrement the number of pending requests accordingly.

Depending on HMDMA operating mode, an external DMA request may trigger individual data word transfers or block transfers. A block can consist of up to 65535 data words. For best throughput, DMA requests can be pipelined. The HMDMA controllers feature a request counter to decouple request timing from the data transfers.

See [“Handshaked Memory DMA Operation” on page 6-36](#) for a functional description.

Modes of Operation

The following sections describe the DMA operation.

Register-Based DMA Operation

Register-based DMA is the traditional kind of DMA operation. Software configures the source or destination address and the length of the data to be transferred to memory-mapped registers and then starts DMA operation.

For basic operation, the software performs these steps:

- Write the source or destination address to the 32-bit `DMAx_START_ADDR` register.
- Write the number of data words to be transferred to the 16-bit `DMAx_X_COUNT` register.

Modes of Operation

- Write the address modifier to the 16-bit `DMAx_X_MODIFY` register. This is the two's-complement value added to the address pointer after every transfer. This value must always be initialized as there is no default value. Typically, this register is set to `0x0004` for 32-bit DMA transfers, to `0x0002` for 16-bit transfers, and to `0x0001` for byte transfers.
- Write the operation mode to the `DMAx_CONFIG` register. These bits in particular need to be changed as needed:
 - The `DMAEN` bit enables the DMA channel.
 - The `WNR` bit controls the DMA direction. DMAs that read from memory (peripheral transmit DMAs and source channel MDMAs) keep this bit cleared. Peripheral receive DMAs and destination channel MDMAs set this bit because they write to memory.
 - The `WDSIZE` bit controls the data word width for the transfer. It can be 8-, 16-, or 32-bits wide.
 - The `DI_EN` bit enables an interrupt when the DMA operation has finished.
 - Set the `FLOW` field to `0x0` for stop mode or `0x1` for autobuffer mode.

Once the `DMAEN` bit is set, the DMA channel starts its operation. While running, the `DMAx_CURR_ADDR` and the `DMAx_CURR_X_COUNT` registers can be monitored to determine the current progress of the DMA operation. However they should not be used to synchronize software and hardware.

The `DMAx_IRQ_STATUS` register signals whether the DMA has finished (`DMA_DONE` bit), whether a DMA error has occurred (`DMA_ERR` bit), and whether the DMA is currently running (`DMA_RUN` bit). The `DMA_DONE` and the `DMA_ERR` bits also function as interrupt latch bits. They must be cleared by write-one-to-clear (W1C) operations by the interrupt service routine.

Stop Mode

In stop mode, the DMA operation is executed only once. When started, the DMA channel transfers the desired number of data words and stops itself when the transfer is complete. If the DMA channel is no longer used, software should clear the `DMAEN` enable bit to disable the otherwise paused channel. Stop mode is entered if the `FLOW` bit field in the DMA channel's `DMAX_CONFIG` register is 0. The `NDSIZE` field must always be 0 in this mode.

For receive (memory write) operation, the `DMA_RUN` bit functions almost the same as the inverted `DMA_DONE` bit. For transmit (memory read) operation, however, the two bits have different timing. Refer to the description of the `SYNC` bit in the `DMAX_CONFIG` register for details.

Autobuffer Mode

In autobuffer mode, the DMA operates repeatedly in a circular manner. If all data words have been transferred, the address pointer `DMAX_CURR_ADDR` is reloaded automatically by the `DMAX_START_ADDR` value. An interrupt may also be generated.

Autobuffer mode is entered if the `FLOW` field in the `DMAX_CONFIG` register is 1. The `NDSIZE` bit must be 0 in autobuffer mode.

Two-Dimensional DMA Operation

Register-based and descriptor-based DMA can operate in one-dimensional mode or two-dimensional mode.

In two-dimensional (2-D) mode, the `DMAX_X_COUNT` register is accompanied by the `DMAX_Y_COUNT` register, supporting arbitrary row and column sizes up to 64K × 64K elements, as well as arbitrary `DMAX_X_MODIFY` and `DMAX_Y_MODIFY` values up to ±32K bytes. Furthermore, `DMAX_Y_MODIFY` can be negative, allowing implementation of interleaved datastreams. The `DMAX_X_COUNT` and `DMAX_Y_COUNT` values specify the row and column sizes, where `DMAX_X_COUNT` must be 2 or greater.

Modes of Operation

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 `DMAx_CONFIG`). Misalignment causes a DMA error.

The `DMAx_X_MODIFY` value is the byte-address increment that is applied after each transfer that decrements the `DMAx_CURR_X_COUNT` register. The `DMAx_X_MODIFY` value is not applied when the inner loop count is ended by decrementing `DMAx_CURR_X_COUNT` from 1 to 0, except that it is applied on the final transfer when `DMAx_CURR_Y_COUNT` is 1 and `DMAx_CURR_X_COUNT` decrements from 1 to 0.

The `DMAx_Y_MODIFY` value is the byte-address increment that is applied after each decrement of the `DMAx_CURR_Y_COUNT` register. However, the `DMAx_Y_MODIFY` value is not applied to the last item in the array on which the outer loop count (`DMAx_CURR_Y_COUNT`) also expires by decrementing from 1 to 0.

After the last transfer completes, `DMAx_CURR_Y_COUNT = 1`, `DMAx_CURR_X_COUNT = 0`, and `DMAx_CURR_ADDR` is equal to the last item's address plus `DMAx_X_MODIFY`.

 If the DMA channel is programmed to refresh automatically (auto-buffer mode), then these registers will be loaded from `DMAx_X_COUNT`, `DMAx_Y_COUNT`, and `DMAx_START_ADDR` upon the first data transfer.

The `DI_SEL` configuration bit enables DMA interrupt requests every time the inner loop rolls over. If `DI_SEL` is cleared, but `DI_EN` is still set, only one interrupt is generated after the outer loop completes.

Examples of Two-Dimensional DMA

Example 1: Retrieve a 16×8 block of bytes from a video frame buffer of size ($N \times M$) pixels:

```
DMAx_X_MODIFY = 1
DMAx_X_COUNT = 16
```

$\text{DMAx_Y_MODIFY} = N-15$ (offset from the end of one row to the start of another)

$\text{DMAx_Y_COUNT} = 8$

This produces the following address offsets from the start address:

0, 1, 2, ... 15,
 $N, N + 1, \dots N + 15,$
 $2N, 2N + 1, \dots 2N + 15, \dots$
 $7N, 7N + 1, \dots 7N + 15,$

Example 2: Receive a video datastream of bytes, (R,G,B pixels) \times (N \times M image size):

$\text{DMAx_X_MODIFY} = (N * M)$

$\text{DMAx_X_COUNT} = 3$

$\text{DMAx_Y_MODIFY} = 1 - 2(N * M)$ (negative)

$\text{DMAx_Y_COUNT} = (N * M)$

This produces the following address offsets from the start address:

0, $(N * M), 2(N * M),$
 $1, (N * M) + 1, 2(N * M) + 1,$
 $2, (N * M) + 2, 2(N * M) + 2,$
 ...
 $(N * M) - 1, 2(N * M) - 1, 3(N * M) - 1,$

Descriptor-Based DMA Operation

In descriptor-based DMA operation, software does not set up DMA sequences by writing directly into DMA controller registers. Rather, software keeps DMA configurations, called descriptors, in memory. On demand, the DMA controller loads the descriptor from memory and overwrites the affected DMA registers by its own control. Descriptors can be fetched from L1 memory using the DCB bus or from external memory using the DEB bus.

Modes of Operation

A descriptor describes what kind of operation should be performed next by the DMA channel. This includes the DMA configuration word as well as data source/destination address, transfer count, and address modify values. A DMA sequence controlled by one descriptor is called a work unit.

Optionally, an interrupt can be requested at the end of any work unit by setting the `DI_EN` bit in the configuration word of the respective descriptor.

A DMA channel is started in descriptor-based mode by first writing the 32-bit address of the first descriptor into the `DMAx_NEXT_DESC_PTR` register (or the `DMAx_CURR_DESC_PTR` in case of descriptor array mode) and then performing a write to the `DMAx_CONFIG` register that sets the `FLOW` field to either `0x4`, `0x6`, or `0x7` and enables the `DMAEN` bit. This causes the DMA controller to immediately fetch the descriptor from the address pointed to by the `DMAx_NEXT_DESC_PTR` register. The fetch overwrites the `DMAx_CONFIG` register again. If the `DMAEN` bit is still set, the channel starts DMA processing.

The `DFETCH` bit in the `DMAx_IRQ_STATUS` register tells whether a descriptor fetch is ongoing on the respective DMA channel. The `DMAx_CURR_DESC_PTR` points to the descriptor value that is to be fetched next.

Descriptor List Mode

Descriptor list mode is selected by setting the `FLOW` bit field in the DMA channel's `DMAx_CONFIG` register to either `0x6` (small descriptor mode) or `0x7` (large descriptor mode). In either of these modes multiple descriptors form a chained list. Every descriptor contains a pointer to the next descriptor. When the descriptor is fetched, this pointer value is loaded into the `DMAx_NEXT_DESC_PTR` register of the DMA channel. In large descriptor mode this pointer is 32 bits wide. Therefore, the next descriptor may reside in any address space accessible through the DCB and DEB buses. In small descriptor mode this pointer is just 16 bits wide. For this reason, the next descriptor must reside in the same 64K byte address space

as the first one because the upper 16 bits of the `DMAx_NEXT_DESC_PTR` register are not updated.

Descriptor list modes are started by writing first to the `DMAx_NEXT_DESC_PTR` register and then to the `DMAx_CONFIG` register.

Descriptor Array Mode

Descriptor array mode is selected by setting the `FLOW` bit field in the DMA channel's `DMAx_CONFIG` register to `0x4`. In this mode, the descriptors do not contain further descriptor pointers. The initial `DMAx_CURR_DESC_PTR` value is written by software. It points to an array of descriptors. The individual descriptors are assumed to reside next to each other and, therefore, their addresses are known.

Variable Descriptor Size

In any descriptor-based mode the `NDSIZE` field in the configuration word specifies how many 16-bit words of the next descriptor need to be loaded on the next fetch. In descriptor-based operation, `NDSIZE` must be non-zero. The descriptor size can be any value from one entry (the lower 16 bits of `DMAx_START_ADDR` only) to nine entries (all the DMA parameters). [Table 6-1](#) illustrates how a descriptor must be structured in memory. The values have the same order as the corresponding MMR addresses.

If, for example, a descriptor is fetched in array mode with `NDSIZE = 0x5`, the DMA controller fetches the 32-bit start address, the DMA configuration word, and the `XCNT` and `XMOD` values. However, it does not load `YCNT` and `YMOD`. This might be the case if the DMA operates in one-dimensional mode or if the DMA is in two-dimensional mode, but the `YCNT` and `YMOD` values do not need to change.

All the other registers not loaded from the descriptor retain their prior values, although the `DMAx_CURR_ADDR`, `DMAx_CURR_X_COUNT`, and

Modes of Operation

DMAX_CURR_Y_COUNT registers are reloaded between the descriptor fetch and the start of DMA operation.

Table 6-1 shows the offsets for descriptor elements in the three modes described above. Note the names in the table describe the descriptor elements in memory, not the actual MMRs into which they are eventually loaded. For more information regarding descriptor element acronyms, see Table 6-4 on page 6-63.

Table 6-1. 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

Note that every descriptor fetch consumes bandwidth from either the DCB bus or the DEB bus and the external memory interface, so it is best to keep the size of descriptors as small as possible.

Mixing Flow Modes

The FLOW mode of a DMA is not a global setting. If the DMA configuration word is reloaded with a descriptor fetch, the FLOW and NDSIZE bit fields can also be altered. A small descriptor might be used to loop back to the first descriptor if a descriptor array is used in an endless manner. If the descriptor chain is not endless and the DMA is required to stop after a

certain descriptor has been processed, the last descriptor is typically processed in stop mode. That is, its `FLOW` and `NDSIZE` fields are 0, but its `DMAEN` bit is still set.

Functional Description

The following sections provide a functional description of DMA.

DMA Operation Flow

[Figure 6-1](#) and [Figure 6-2](#) describe the DMA flow.

DMA Startup

This section discusses starting DMA “from scratch.” This is similar to starting it after it has been paused by the `FLOW = 0` mode.

Before initiating DMA for the first time on a given channel, all parameter registers must be initialized. Be sure to initialize the upper 16 bits of the `DMAx_NEXT_DESC_PTR` (or `DMAx_CURR_DESC_PTR` register in `FLOW = 4` mode) and `DMAx_START_ADDR` registers, because they might not otherwise be accessed, depending upon the flow mode. Also note that the `DMAx_X_MODIFY` and `DMAx_Y_MODIFY` registers are not preset to a default value at reset.

The user may wish to write other DMA registers that might be static during DMA activity (for example, `DMAx_X_MODIFY`, `DMAx_Y_MODIFY`). The contents of `NDSIZE` and `FLOW` in `DMAx_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 `DMAx_CONFIG` with `DMAEN = 1`.

Functional Description

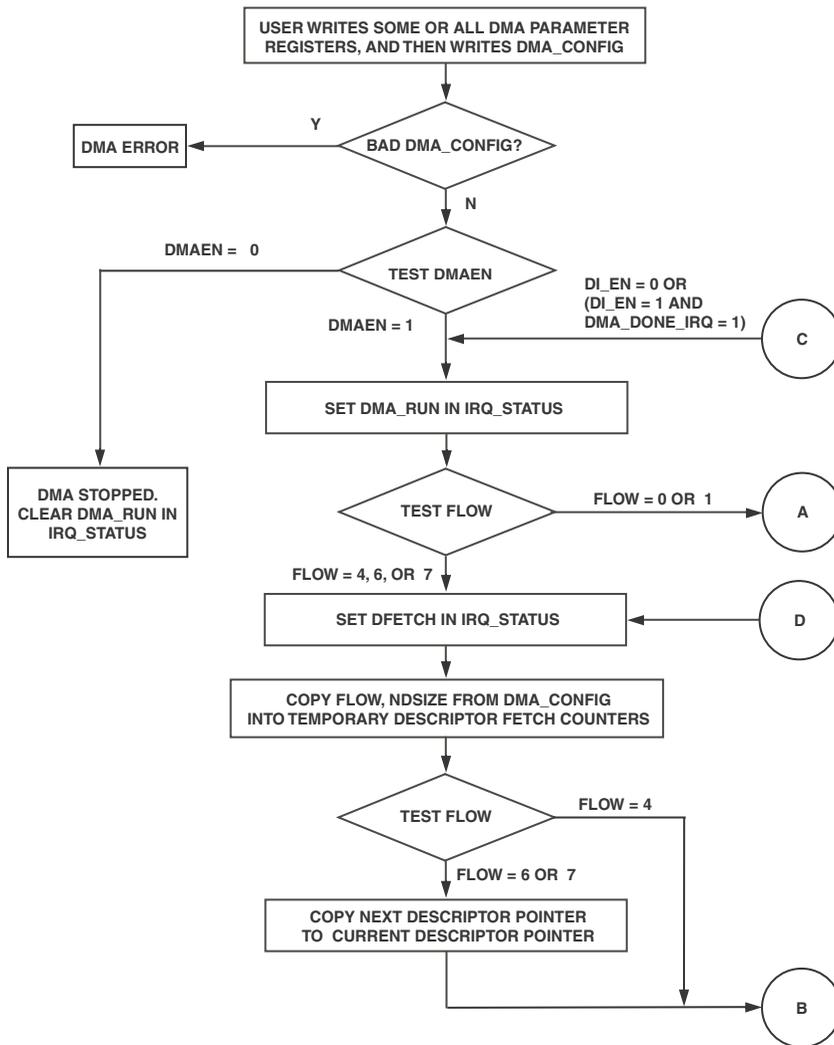


Figure 6-1. DMA Flow, From DMA Controller's Point of View (1 of 2)

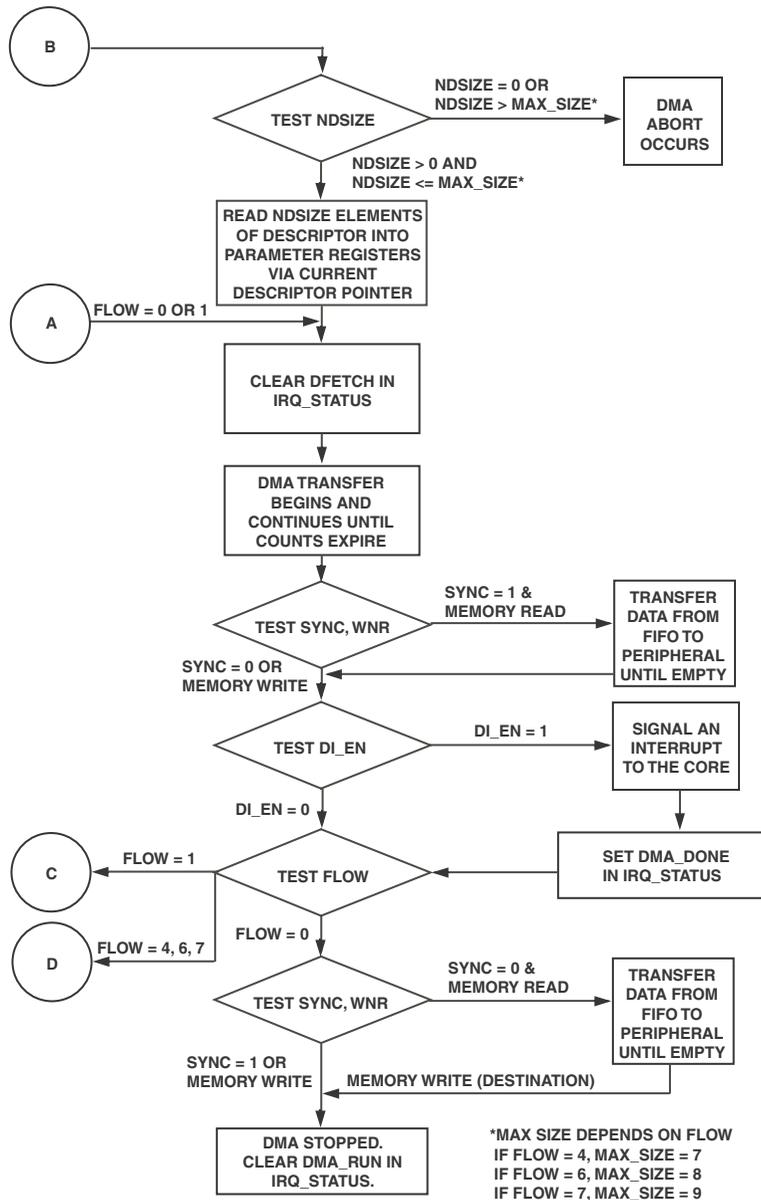


Figure 6-2. DMA Flow, From DMA Controller's Point of View (2 of 2)

Functional Description

When `DMAX_CONFIG` is written directly by software, 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 `DMAX_CONFIG` assumes control. Before this point, the direct write to `DMAX_CONFIG` had control. In other words, the `WDSIZE`, `DI_EN`, `DI_SEL`, `SYNC`, and `DMA2D` fields will be taken from the `DMACFG` value in the descriptor read from memory, while these field values initially written to the `DMAX_CONFIG` register are ignored.

As [Figure 6-1](#) and [Figure 6-2](#) show, at startup the `FLOW` and `NDSIZE` bits in `DMAX_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 `DMAX_NEXT_DESC_PTR` is copied into `DMAX_CURR_DESC_PTR`. Then, fetches of new descriptor elements from memory are performed, indexed by `DMAX_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 `DMAX_NEXT_DESC_PTR`, but the fetch of the current descriptor continues using `DMAX_CURR_DESC_PTR`. After completion of the descriptor fetch, `DMAX_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 `DMAX_CURR_DESC_PTR` does not occur. Instead, descriptor fetch indexing begins with the value in `DMAX_CURR_DESC_PTR`.

If `DMACFG` is not part of the descriptor, the previous `DMAX_CONFIG` settings (as written by MMR access at startup) control the work unit operation. If

DMACFG is part of the descriptor, then the `DMAx_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 `DMAx_IRQ_STATUS` register indicate the state of the DMA channel. After a write to `DMAx_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, DMA data transfer operation begins. The DMA channel immediately attempts to fill its FIFO, subject to channel priority—a memory write (RX) DMA channel begins accepting data from its peripheral, and a memory read (TX) DMA channel begins memory reads, provided the channel wins the grant for bus access.

When the DMA channel performs its first data memory access, its address and count computations take their input operands from the start registers (`DMAx_START_ADDR`, `DMAx_X_COUNT`, `DMAx_Y_COUNT`), and write results back to the current registers (`DMAx_CURR_ADDR`, `DMAx_CURR_X_COUNT`, `DMAx_CURR_Y_COUNT`). Note also that the current registers are not valid until the first memory access is performed, which may be some time after the channel is started by the write to the `DMA_CONFIG` register. The current registers are loaded automatically from the appropriate descriptor elements, overwriting their previous contents, as follows.

- `DMAx_START_ADDR` is copied to `DMAx_CURR_ADDR`
- `DMAx_X_COUNT` is copied to `DMAx_CURR_X_COUNT`
- `DMAx_Y_COUNT` is copied to `DMAx_CURR_Y_COUNT`

Then DMA data transfer operation begins, as shown in [Figure 6-2](#).

Functional Description

DMA Refresh

On completion of a work unit:

- The DMA controller completes the transfer of all data between memory and the DMA unit.
- If `SYNC = 1` and `WNR = 0` (memory read), the DMA controller selects a synchronized transition and transfers all data to the peripheral before continuing.
- If enabled by `DI_EN`, the DMA controller signals an interrupt to the core and sets the `DMA_DONE` bit in the channel's `DMAx_IRQ_STATUS` register.
- If `FLOW = 0` the DMA controller stops operation by clearing the `DMA_RUN` bit in `DMAx_IRQ_STATUS` register after all 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 `DMAx_IRQ_STATUS` register to 1. At this point, the DMA operation depends on whether `FLOW = 4, 6, or 7`, as follows:

If `FLOW = 4` (descriptor array) the DMA controller loads a new descriptor from memory into the DMA registers using the contents of `DMAx_CURR_DESC_PTR`, and increments `DMAx_CURR_DESC_PTR`. The descriptor size comes from the `NDSIZE` field of the `DMAx_CONFIG` register prior to the beginning of the fetch.

If `FLOW = 6` (small descriptor list) the DMA controller copies the 32-bit `DMAx_NEXT_DESC_PTR` into `DMAx_CURR_DESC_PTR`. Next, the DMA controller fetches a descriptor from memory into the DMA registers using the new contents of `DMAx_CURR_DESC_PTR`, and increments `DMAx_CURR_DESC_PTR`. The first descriptor element that is loaded is a new 16-bit value for the lower 16 bits of `DMAx_NEXT_DESC_PTR`, followed by the rest of the descriptor

elements. The high 16 bits of `DMAx_NEXT_DESC_PTR` will retain their former value. This supports a shorter, more efficient descriptor than the large descriptor list model, which is suitable whenever the application can place the channel's descriptors in the same 64K byte range of memory.

If `FLOW = 7` (large descriptor list) the DMA controller copies the 32-bit `DMAx_NEXT_DESC_PTR` into `DMAx_CURR_DESC_PTR`. Next, the DMA controller fetches a descriptor from memory into the DMA registers using the new contents of `DMAx_CURR_DESC_PTR`, and increments `DMAx_CURR_DESC_PTR`. The first descriptor element that is loaded is a new 32-bit value for the full `DMAx_NEXT_DESC_PTR`, followed by the rest of the descriptor elements. The high 16 bits of `DMAx_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.

- 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 the descriptor containing the link to the new 64K byte range needs to use `FLOW = 7`. All descriptors that reference the same 64K byte area may use `FLOW = 6`.
- If `FLOW = 4, 6, or 7` (descriptor array, small descriptor list, or large descriptor list, respectively), the DMA controller clears the `DFETCH` bit in the `DMAx_IRQ_STATUS` register.

Functional Description

- If `FLOW = any value but 0 (Stop)`, the DMA controller begins the next work unit for that channel, which must contend with other channels for priority on the memory buses. On the first memory transfer of the new work unit, the DMA controller updates the current registers from the start registers:

`DMAx_CURR_ADDR` loaded from `DMAx_START_ADDR`
`DMAx_CURR_X_COUNT` loaded from `DMAx_X_COUNT`
`DMAx_CURR_Y_COUNT` loaded from `DMAx_Y_COUNT`

The `DFETCH` bit in the `DMAx_IRQ_STATUS` register is then cleared, after which the DMA transfer begins again, as shown in [Figure 6-2](#).

Work Unit Transitions

Transitions from one work unit to the next are controlled by the `SYNC` bit in the `DMAx_CONFIG` register of the work units. In general, continuous transitions have lower latency at the cost of restrictions on changes of data format or addressed memory space in the two work units. These latency gains and data restrictions arise from the way the DMA FIFO pipeline is handled while the next descriptor is fetched. In continuous transitions (`SYNC = 0`), the DMA FIFO pipeline continues to transfer data to and from the peripheral or destination memory during the descriptor fetch and/or when the DMA channel is paused between descriptor chains.

Synchronized transitions (`SYNC = 1`), on the other hand, provide better real-time synchronization of interrupts with peripheral state and greater flexibility in the data formats and memory spaces of the two work units, at the cost of higher latency in the transition. In synchronized transitions, the DMA FIFO pipeline is drained to the destination or flushed (RX data discarded) between work units.

 Work unit transitions for MDMA streams are controlled by the SYNC bit of the MDMA source channel's DMA_x_CONFIG register. The SYNC bit of the MDMA destination channel is reserved and must be 0. In transmit (memory read) channels, the SYNC bit of the last descriptor prior to the transition controls the transition behavior. In contrast, in receive channels, the SYNC bit of the first descriptor of the next descriptor chain controls the transition.

DMA Transmit and MDMA Source

In DMA transmit (memory read) and MDMA source channels, the SYNC bit controls the interrupt timing at the end of the work unit and the handling of the DMA FIFO between the current and next work units.

If SYNC = 0, a continuous transition is selected. In a continuous transition, just after the last data item is read from memory, the following operations start in parallel:

- The interrupt (if any) is signalled.
- The DMA_DONE bit in the DMA_x_IRQ_STATUS register is set.
- The next descriptor begins to be fetched.
- The final data items are delivered from the DMA FIFO to the destination memory or peripheral.

This allows the DMA to provide data from the FIFO to the peripheral “continuously” during the descriptor fetch latency period.

When SYNC = 0, the final interrupt (if enabled) occurs when the last data is read from memory. This interrupt is at the earliest time that the output memory buffer may safely be modified without affecting the previous data transmission. Up to four data items may still be in the DMA FIFO, however, and not yet at the peripheral, so the DMA interrupt should not be used as the sole means of synchronizing the shutdown or reconfiguration of the peripheral following a transmission.

Functional Description

 If `SYNC = 0` (continuous transition) on a transmit (memory read) descriptor, the next descriptor must have the same data word size, read/write direction, and source memory (internal vs. external) as the current descriptor.

`SYNC = 0` selects continuous transition on a work unit in `FLOW = 0` mode with interrupt enabled. The interrupt service routine may begin execution while the final data is still draining from the FIFO to the peripheral. This is indicated by the `DMA_RUN` bit in the `DMAx_IRQ_STATUS` register; if it is 1, the FIFO is not empty yet. Do not start a new work unit with different word size or direction while `DMA_RUN = 1`. Further, if the channel is disabled (by writing `DMAEN = 0`), the data in the FIFO is lost.

`SYNC = 1` selects a synchronized transition in which the DMA FIFO is first drained to the destination memory or peripheral before any interrupt is signalled and before any subsequent descriptor or data is fetched. This incurs greater latency, but provides direct synchronization between the DMA interrupt and the state of the data at the peripheral.

For example, if `SYNC = 1` and `DI_EN = 1` on the last descriptor in a work unit, the interrupt occurs when the final data has been transferred to the peripheral, allowing the service routine to properly switch to non-DMA transmit operation. When the interrupt service routine is invoked, the `DMA_DONE` bit is set and the `DMA_RUN` bit is cleared.

A synchronized transition also allows greater flexibility in the format of the DMA descriptor chain. If `SYNC = 1`, the next descriptor may have any word size or read/write direction supported by the peripheral and may come from either memory space (internal or external). This can be useful in managing MDMA work unit queues, since it is no longer necessary to interrupt the queue between dissimilar work units.

DMA Receive

In DMA receive (memory write) channels, the `SYNC` bit controls the handling of the DMA FIFO between descriptor chains (not individual

descriptors), when the DMA channel is paused. The DMA channel pauses after descriptors with `FLOW = 0` mode, and may be restarted (for example, after an interrupt) by writing the channel's `DMAx_CONFIG` register with `DMAEN = 1`.

If the `SYNC` bit is 0 in the new work unit's `DMAx_CONFIG` value, a continuous transition is selected. In this mode, any data items received into the DMA FIFO while the channel was paused are retained, and they are the first items written to memory in the new work unit. This mode of operation provides lower latency at work unit transitions and ensures that no data items are dropped during a DMA pause, at the cost of certain restrictions on the DMA descriptors.

 If the `SYNC` bit is 0 on the first descriptor of a descriptor chain after a DMA pause, the DMA word size of the new chain must not change from the word size of the previous descriptor chain active before the pause, unless the DMA channel is reset between chains by writing the `DMAEN` bit to 0 and then to 1 again.

If the `SYNC` bit is 1 in the new work unit's `DMAx_CONFIG` value, a synchronized transition is selected. In this mode, only the data received from the peripheral by the DMA channel after the write to the `DMAx_CONFIG` register are delivered to memory. Any prior data items transferred from the peripheral to the DMA FIFO before this register write are discarded. This provides direct synchronization between the data stream received from the peripheral and the timing of the channel restart (when the `DMAx_CONFIG` register is written).

For receive DMAs, the `SYNC` bit has no effect in transitions between work units in the same descriptor chain (that is, when the previous descriptor's `FLOW` mode was not 0, so that DMA channel did not pause.)

If a descriptor chain begins with a `SYNC` bit of 1, there is no restriction on DMA word size of the new chain in comparison to the previous chain.

Functional Description

 The DMA word size must not change between one descriptor and the next in any DMA receive (memory write) channel within a single descriptor chain, regardless of the SYNC bit setting. In other words, if a descriptor has $WNR = 1$ and $FLOW = 4, 6, \text{ or } 7$, then the next descriptor must have the same word size. For any DMA receive (memory write) channel, there is no restriction on changes of memory space (internal vs. external) between descriptors or descriptor chains. DMA transmit (memory read) channels may have such restrictions (see [“DMA Transmit and MDMA Source” on page 6-25](#)).

Stopping 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_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).

 If a channel has been stopped abruptly by writing `DMAx_CONFIG` to 0 (or any value with $DMAEN = 0$), the user must ensure that any memory read or write accesses in the pipelines have completed before enabling the channel again. If the channel is enabled again before an “orphan” access from a previous work unit completes, the state of the DMA interrupt and FIFO is unspecified. This can generally be handled by ensuring that the core allocates several consecutive idle cycles in its usage of the relevant memory space to allow up to three pending DMA accesses to issue, plus allowing enough memory access time for the accesses themselves to complete.

DMA Errors (Aborts)

The DMA controller flags conditions that cause the DMA process to end abnormally (abort). This functionality is provided as a tool for system development and debug 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:

- Read each channel's `DMAX_IRQ_STATUS` register to look for a channel with the `DMA_ERR` bit set (bit 1).
- Clear the problem with that channel (for example, fix register values).
- Clear the `DMA_ERR` bit (write `DMAX_IRQ_STATUS` with bit 1 set).

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 6-2](#).

Functional Description

- A disallowed register write occurred while the channel was running. Only the `DMAx_CONFIG` and `DMAx_IRQ_STATUS` registers can be written when `DMA_RUN = 1`.
- An address alignment error occurred during any memory access. For example, when `DMAx_CONFIG` register `WDSIZE = 1` (16-bit) but the least significant bit (LSB) of the address is not equal to b#0, or when `WDSIZE = 2` (32-bit) but the two LSBs of the address are not equal to b#00.
- A memory space transition was attempted (internal-to-external or vice versa). For example, the value in the `DMAx_CURR_ADDR` register or `DMAx_CURR_DESC_PTR` register crossed a memory boundary.
- 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:

- `DMAx_CONFIG` direction bit (`WNR`) does not agree with the direction of the mapped peripheral.
- `DMAx_CONFIG` direction bit does not agree with the direction of the MDMA channel.
- `DMAx_CONFIG` word size (`WDSIZE`) is not supported by the mapped peripheral. See [Table 6-2](#).
- `DMAx_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 2-D DMA, $X_COUNT = 1$

Table 6-2. Legal NDSIZE Values

FLOW	NDSIZE	Note
0	0	
1	0	
4	$0 < NDSIZE \leq 7$	Descriptor array, no descriptor pointer fetched
6	$0 < NDSIZE \leq 8$	Descriptor list, small descriptor pointer fetched
7	$0 < NDSIZE \leq 9$	Descriptor list, large descriptor pointer fetched

DMA Control Commands

Advanced peripherals, such as an Ethernet MAC module, are capable of managing some of their own DMA operations, thus dramatically improving real-time performance and relieving control and interrupt demands on the Blackfin processor core. These peripherals may communicate to the DMA controller using DMA control commands, which are transmitted from the peripheral to the associated DMA channel over internal DMA request buses. Refer to [“Unique Information for the ADSP-BF51x Processor” on page 6-102](#) to determine if DMA control commands are applicable to a particular product.

The request buses consist of three wires per DMA-management-capable peripheral. The DMA control commands extend the set of operations available to the peripheral beyond the simple “request data” command used by peripherals in general.

While these DMA control commands are not visible to or controllable by the user, their use by a peripheral has implications for the structure of the DMA transfers which that peripheral can support. It is important that application software be written to comply with certain restrictions

Functional Description

regarding work units and descriptor chains (described later in this section) so that the peripheral operates properly whenever it issues DMA control commands.

MDMA channels do not service peripherals and therefore do not support DMA control commands. The DMA control commands are shown in [Table 6-3](#).

Table 6-3. DMA Control Commands

Code	Name	Description
000	NOP	No operation
001	Restart	Restarts the current work unit from the beginning
010	Finish	Finishes the current work unit and starts the next
011	-	Reserved
100	Req Data	Typical DMA data request
101	Req Data Urgent	Urgent DMA data request
110	-	Reserved
111	-	Reserved

Additional information for the control commands includes:

- **Restart**

The `Restart` command causes the current work unit to interrupt processing and start over, using the addresses and counts from `DMAx_START_ADDR`, `DMAx_X_COUNT`, and `DMAx_Y_COUNT`. No interrupt is signalled.

If a channel programmed for transmit (memory read) receives a `Restart` command, the channel momentarily pauses while any pending memory reads initiated prior to the `Restart` command are completed.

During this period of time, the channel does not grant DMA requests. Once all pending reads have been flushed from the channel's pipelines, the channel resets its counters and FIFO and starts prefetch reads from memory. DMA data requests from the peripheral are granted as soon as new prefetched data is available in the DMA FIFO. The peripheral can thus use the `Restart` command to re-attempt a failed transmission of a work unit.

If a channel programmed for receive (memory write) receives a `Restart` command, the channel stops writing to memory, discards any data held in its DMA FIFO, and resets its counters and FIFO. As soon as this initialization is complete, the channel again grants DMA write requests from the peripheral. The peripheral can thus use the `Restart` command to abort transfer of received data into a work unit and re-use the memory buffer for a later data transfer.

- **Finish**

The `Finish` command causes the current work unit to terminate and move on to the next work unit. An interrupt is signalled as usual, if selected by the `DI_EN` bit. The peripheral can thus use the `Finish` command to partition the DMA stream into work units on its own, perhaps as a result of parsing the data currently passing through its supported communication channel, without direct real-time control by the processor.

If a channel programmed for transmit (memory read) receives a `Finish` command, the channel momentarily pauses while any pending memory reads initiated prior to the `Finish` command are completed. During this period of time, the channel does not grant DMA requests. Once all pending reads have been flushed from the channel's pipelines, the channel signals an interrupt (if enabled), and begins fetching the next descriptor (if any). DMA data requests from the peripheral are granted as soon as new prefetched data is available in the DMA FIFO.

Functional Description

If a channel programmed for receive (memory write) receives a `Finish` command, the channel stops granting new DMA requests while it drains its FIFO. Any DMA data received by the DMA controller prior to the `Finish` command is written to memory. When the FIFO reaches an empty state, the channel signals an interrupt (if enabled) and begins fetching the next descriptor (if any). Once the next descriptor has been fetched, the channel initializes its FIFO and then resumes granting DMA requests from the peripheral.

- **Request Data**

The `Request Data` command is identical to the DMA request operation of peripherals that are not DMA-management-capable.

- **Request Data Urgent**

The `Request Data Urgent` command behaves identically to the DMA `Request` command, except that the DMA channel performs its memory accesses with urgent priority while it is asserted. This includes both data and descriptor-fetch memory accesses. A DMA-management-capable peripheral might use this command if an internal FIFO is approaching a critical condition.

Restrictions

The proper operation of the 4-location DMA channel FIFO leads to certain restrictions in the sequence of DMA control commands.

Transmit Restart or Finish

No `Restart` or `Finish` command may be issued by a peripheral to a channel configured for memory read unless the peripheral has already performed at least one DMA transfer in the current work unit and the current work unit has more than four items remaining in `DMAx_CURR_X_COUNT`/`DMAx_CURR_Y_COUNT` (thus not yet read from

memory). Otherwise, the current work unit may already have completed memory operations and can no longer be restarted or finished properly.

If the `DMAX_CURR_X_COUNT/ DMAX_CURR_Y_COUNT` value of the current work unit is sufficiently large that it is always at least five more than the maximum data count prior to any `Restart` or `Finish` command, the above restriction is satisfied. This implies that any work unit which might be managed by `Restart` or `Finish` commands must have `DMAX_CURR_X_COUNT/ DMAX_CURR_Y_COUNT` values representing at least five data items.

Particularly if the `DMAX_CURR_X_COUNT/ DMAX_CURR_Y_COUNT` registers are programmed to 0 (representing 65,536 transfers, the maximum value) the channel will operate properly for 1-D work units up to 65,531 data items or 2-D work units up to 4,294,967,291 data items.

Receive Restart or Finish

No `Restart` or `Finish` command may be issued by a peripheral to a channel configured for memory write unless either the peripheral has already performed at least five DMA transfers in the current work unit or the previous work unit was terminated by a `Finish` command and the peripheral has performed at least one DMA transfer in the current work unit. If five data transfers have been performed, then at least one data item has been written to memory in the current work unit, which implies that the current work unit's descriptor fetch completed before the data grant of the fifth item. Otherwise, if less than five data items have been transferred, it is possible that all of them are still in the DMA FIFO and the previous work unit is still in the process of completion and transition between work units.

Similarly, if a `Finish` command ended the previous work unit and at least one subsequent DMA data transfer has occurred, then the fact that the DMA channel issued the grant guarantees that the previous work unit has already completed the process of draining its data to memory and transitioning to the new work unit.

Functional Description

If a peripheral terminates all work units with the `Finish` opcode (effectively assuming responsibility for all work unit boundaries for the DMA channel), then the peripheral need only ensure that it performs a single transfer in each work unit before any restart or finish. This requires, however, that the user programs the descriptors for all work units managed by the channel with `DMAx_CURR_X_COUNT/ DMAx_CURR_Y_COUNT` values representing more data items than the maximum work unit size that the peripheral will encounter. For example, `DMAx_CURR_X_COUNT/ DMAx_CURR_Y_COUNT` values of 0 allow the channel to operate properly on 1-D work units up to 65,535 data items and 2-D work units up to 4,294,967,295 data items.

Handshaked Memory DMA Operation

Handshaked memory DMA operation is not available for all products. Refer to the [“Unique Information for the ADSP-BF51x Processor” on page 6-102](#) to determine whether this feature applies to this product.

Each `DMARx` input has its own set of control and status registers. Handshake operation for MDMA0 is enabled by the `HMDMAEN` bit in the `HMDMA0_CONTROL` register. Similarly, the `HMDMAEN` bit in the `HMDMA1_CONTROL` register enables handshake mode for MDMA1.

It is important to understand that the handshake hardware works completely independently from the descriptor and autobuffer capabilities of the MDMA, allowing most flexible combinations of logical data organization vs. data portioning as required by FIFO depths, for example. If, however, the connected device requires certain behavior of the address lines, these must be controlled by traditional DMA setup.



The HMDMA unit controls only the destination (memory write) channel of the memory DMA. The source channel (memory-read side) fills the 8-deep DMA buffers immediately after the receive side is enabled and issues eight read commands.

The `HMDMAX_BCINIT` registers control how many data transfers are performed upon every DMA request. If set to one, the peripheral can time every individual data transfer. If greater than one, the peripheral must have sufficient buffer size to provide or consume the number of words programmed. Once the transfer has been requested, no further handshake can hold off the DMA from transferring the entire block, except by stalling the EBIU accesses by the `ARDY` signal. Nevertheless, the peripheral may request a block transfer before the entire buffer is available by simply taking the minimum transfer time based on wait-state settings into consideration.

 The block count defines how many data transfers are performed by the MDMA engine. A single DMA transfer can cause two read or write operations on the EBIU port if the transfer word size is set to 32-bit in the `MDMA_yy_CONFIG` register (`WDSIZE = b#10`).

Since the block count registers are 16 bits wide, blocks can group up to 65,535 transfers.

Once a block transfer has been started, the `HMDMAX_BCOUNT` registers return the remaining number of transfers to complete the current block. When the complete block has been processed, the `HMDMAX_BCOUNT` register returns zero. Software can force a reload of the `HMDMAX_BCOUNT` from the `HMDMAX_BCINIT` register even during normal operation by setting the `RBC` bit in the `HMDMAX_CONTROL` register. Set `RBC` when the HMDMA module is already active, but only when the MDMA is not enabled.

Pipelining DMA Requests

The device mastering the DMA request lines is allowed to request additional transfers even before the former transfer has completed. As long as the device can provide or consume sufficient data it is permitted to pulse the `DMARx` inputs multiple times.

The `HMDMAX_ECOUNT` registers are incremented every time a significant edge is detected on the respective `DMARx` input, and they are decremented when

Functional Description

the MDMA completes the block transfer. These read-only registers use a 16-bit two's-complement data representation: if they return zero, all requested block transfers have been performed. A positive value signals up to 32767 requests that haven't been served yet and indicates that the MDMA is currently processing. Negative values indicate the number of DMA requests that will be ignored by the engine. This feature restrains initial pulses on the $\overline{\text{DMARx}}$ inputs at startup.

The HMDMAx_ECINIT registers reload the HMDMAx_ECOUNT registers every time the handshake mode is enabled (when the HMDMAEN bit changes from 0 to 1). If the initial edge count value is 0, the handshake operation starts with a settled request budget. If positive, the engine starts immediately transferring the programmed number (up to 32767) of blocks once enabled, even without detecting any activity on the DMARx pins. If negative, the engine will disregard the programmed number (up to 32768) significant edges on the DMARx inputs before starting normal operation.

Figure 6-3 illustrates how an asynchronous FIFO could be connected. In such a scenario the REP bit should be cleared to let the DMARx request pin listen to falling edges.

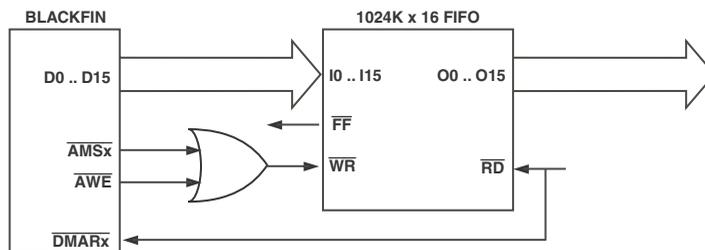


Figure 6-3. Transmit DMA Example Connection

The Blackfin processor does not evaluate the full flag such FIFOs usually provide because asynchronous polling of that signal would reduce the system throughput drastically. Moreover, the processor first fills the FIFO by initializing the HMDMAx_ECINIT register to 1024, which equals the depth of the FIFO. Once enabled, the MDMA automatically transmits 1024 data

words. Afterward it continues to transmit only if the FIFO is emptied by its read strobe again. Most likely, the `HMDMAx_BCINIT` register is programmed to 1 in this case.

In the receive example shown in Figure 6-4, the Blackfin processor again does not use the FIFO's internal control mechanism. Rather than testing the empty flag, the processor counts the number of data words available in the FIFO in its own `HMDMAx_ECOUNT` register. Theoretically, the MDMA could immediately process data as soon as it is written into the FIFO by the write strobe, but the fast MDMA engine would read out the FIFO quickly and stall soon if the FIFO was not promptly filled with new data. Streaming applications can balance the FIFO so that the producer is never held off by a full FIFO and the consumer is never held by an empty FIFO. This is accomplished by filling the FIFO halfway and then letting both consumer and producer run at the same speed. In this case the `HMDMAx_ECINIT` register can be written with a negative value, which corresponds to half the FIFO depth. Then, the MDMA does not start consuming data as long as the FIFO is not half-filled.

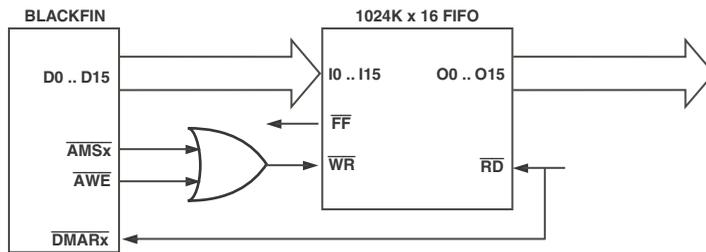


Figure 6-4. Receive DMA Example Connection

On internal system buses, memory DMA channels have lower priority than other DMAs. In busy systems, the memory DMAs may tend to starve. As this is not acceptable when transferring data through high-speed FIFOs, the handshake mode provides a high-water functionality to increase the MDMA's priority. With the `UTE` bit in the `HMDMAx_CONTROL` register set, the MDMA gets higher priority as soon as a (positive) value in

Functional Description

the `HMDMAX_ECOUNT` register becomes higher than the threshold held by the `HMDMAX_ECURGENT` register.

HMDMA Interrupts

In addition to the normal MDMA interrupt channels, the handshake hardware provides two new interrupt sources for each `DMARx` input. The `HMDMAX_CONTROL` registers provide interrupt enable and status bits. The interrupt status bits require a write-1-to-clear operation to cancel the interrupt request.

The `block done` interrupt signals that a complete MDMA block, as defined by the `HMDMAX_BCINIT` register, has been transferred (when the `HMDMAX_BCOUNT` register decrements to zero). While the `BDIE` bit enables this interrupt, the `MBDI` bit can gate it until the edge count also becomes zero, meaning that all requested MDMA transfers have been completed.

The `overflow` interrupt is generated when the `HMDMA_ECOUNT` register overflows. Since it can count up to 32767, which is much more than most peripheral devices can support, the Blackfin processor has another threshold register called `HMDMA_ECOVERFLOW`. It resets to 0xFFFF and should be written with any positive value by the user before enabling the function by the `OIE` bit. Then, the `overflow` interrupt is issued when the value of the `HMDMA_ECOUNT` register exceeds the threshold in the `HMDMA_ECOVERFLOW` register.

DMA Performance

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.

The Blackfin architecture features several mechanisms to customize system behavior for best performance. This includes DMA channel prioritization, traffic control, and priority treatment of bursted transfers. Nevertheless, the resulting performance of a DMA transfer often depends on

application-level circumstances. For best performance consider the following system software architecture questions.

- What is the required DMA bandwidth?
- Which DMA transfers have real-time requirements and which do not?
- How heavily is the DMA controller competing with the core for on-chip and off-chip resources?
- How often do competing DMA channels require the bus systems to alter direction?
- How often do competing DMA or core accesses cause the SDRAM to open different pages?
- Is there a way to distribute DMA requests nicely over time?

A key feature of the DMA architecture is the separation of the activity on the DMA access bus (DAB) used by the peripherals from the activity on the buses between the DMA and memory. For DMA to/from on-chip memory the DMA core bus (DCB) is used, and the DMA external bus (DEB) is used for DMA transfers with off-chip memory. The “Chip Bus Hierarchy” chapter explains the bus architecture.

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.

Functional Description

DMA Throughput

Each peripheral DMA channel has a maximum transfer rate of one 16-bit word per two system clocks in either direction. Like the DAB and DEB buses, the DMA controller resides in the `SCLK` domain. The controller synchronizes accesses to and from the DCB bus, which runs at the `CCLK` rate.

Each memory DMA channel has a maximum transfer rate of one 16-bit word per system clock (`SCLK`) cycle.

When the traffic on all DMA channels is taken in the aggregate:

- Transfers between the peripherals and the DMA unit have a maximum rate of one 16-bit transfer per system clock.
- Transfers between the DMA unit and internal memory (L1) have a maximum rate of one 16-bit transfer per system clock.
- Transfers between the DMA unit and external memory have a maximum rate of one 16-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 when both the core and the DMA access the same L1 bank, when SDRAM pages need to be opened/closed, or when cache lines are filled.
- Direction changes from `RX` to `TX` on the DAB bus impose 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_CONFIG`, `DMAx_IRQ_STATUS`, or `DMAx_PERIPHERAL_MAP` stall all DMA activity for one cycle per 16-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 16-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 `SYNC` bit is set in the `DMAx_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.

The MDMA channels are clocked by `SCLK`. If the 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. For example DMA typically runs at 2/3 of the system clock rate when the

Functional Description

core-to-system clock ratio is 2:1. At higher clock ratios, full bandwidth is maintained.

If the 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 three plus the memory latency at the DMA in SCLKs (which is typically seven for internal transfers and six for external transfers).

Memory DMA Timing Details

When the destination `DMAX_CONFIG` register is written, MDMA operation starts after a latency of three SCLK cycles.

If either MDMA channel has been selected to use descriptors, the descriptors are fetched from memory. The destination channel descriptors are fetched first. Then the source MDMA channel begins fetching data from the source buffer, after a latency of four SCLK cycles after the last descriptor word is returned from memory. Due to memory pipelining, this is typically eight SCLK cycles after the fetch of the last descriptor word. The resulting data is deposited in the MDMA channel's 8-location FIFO. After a latency of two SCLK cycles, the destination MDMA channel begins writing data to the destination memory buffer.

Static Channel Prioritization

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 `PMAP` field in the `DMAX_PERIPHERAL_MAP` registers. The memory DMA streams are always lower static 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.

Temporary DMA Urgency

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 become urgent when handshaked operation is enabled and the `DMARx` edge count exceeds the value stored in

Functional Description

the `HMDMAX_ECURGENT` register. If handshaked operation is disabled, software can control urgency of requests directly by altering the `DRQ` bit field in the `HMDMAX_CONTROL` register.

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). All prior incomplete memory transfers ahead of it in that memory system are also marked for expedited processing. 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.

Memory DMA 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.

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, the `MDMA_ROUND_ROBIN_PERIOD` may be programmed to select each stream in turn for a fixed number of transfers.

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 `DMA_TC_PER` register (see “[Static Channel Prioritization](#)” on page 6-44).

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 completed. 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 $1 \leq P \leq 31$, 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 `DMA_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 one, 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

Functional Description

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 the stream selection unlocks and becomes free again on the next cycle.

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 the bandwidth by a factor of $1/(P+1)$. However if both MDMA streams are used, 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). By selection of various round-robin period values P , which limit how often the MDMA streams alternate, maximal transfer efficiency can be maintained.

Traffic Control

In the Blackfin DMA architecture, 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 “Temporary DMA Urgency” on page 6-45.](#)

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. The DMA block provides a traffic control mechanism controlled by the `DMA_TC_PER` and `DMA_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 [“Memory DMA Priority and Scheduling”](#) on page 6-46.

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 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 for 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

Programming Model

maximum attainable bandwidth in congested systems, often to above 90%.

To disable preferential DMA prioritization, program the `DMA_TC_PER` register to `0x0000`.

Programming Model

Several synchronization and control methods are available for use in development of software tasks which manage peripheral DMA and memory DMA (see also “[Memory DMA](#)” on page 6-6). 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 peripheral DMA and memory DMA 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 `DMAx_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 `DMAx_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.

Due to DMA FIFOs and DMA/memory pipelining, polling of the `DMAX_CURR_ADDR`, `DMAX_CURR_DESC_PTR`, or `DMAX_CURR_X_COUNT/DMAX_CURR_Y_COUNT` registers is not recommended for precisely synchronizing DMA with data processing. 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 takes 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 of channel A. Software monitoring of channel B, based on examination of the `DMAX_CURR_ADDR` register contents, would not safely conclude whether the memory location pointed to by channel B's `DMAX_CURR_ADDR` register has or has not been written.

If allowances are made for the lengths of the DMA/memory pipeline, polling of the current address, pointer, and count registers can permit loose synchronization of DMA with software. The depth of the DMA FIFO is four locations (either four 8- or 16-bit data elements, or two 32-bit data elements) for a peripheral DMA channel, and eight locations (four 32-bit data elements) for an MDMA FIFO. 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).

Programming Model

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. This value is a maximum because 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 `DMAx_CURR_X_COUNT` register reads a value of 60 remaining elements, so that processing of the first 40 elements has at least been started. Since the total pipeline length is no greater than the sum of four (for the peripheral DMA FIFO) plus six (for the DMA/memory pipeline) or ten data elements, it is safe to conclude that the DMA transfer of the first 30 (40-10) data elements is complete.

For precise synchronization, software should either wait for an interrupt or consult the channel's `DMAx_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 a `DMAx_IRQ_STATUS` bit, it guarantees that the last memory operation of the work unit has been completed and will definitely be visible to processor 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 acknowledgement 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 `DMAx_CONFIG` and the `DMAx_NEXT_DESC_PTR` registers. Alternatively, the user may choose to write all the MMR registers directly from software, ending with the write to the `DMAx_CONFIG` register.

The simplest way to signal completion of DMA is by an interrupt. This is selected by the `DI_EN` bit in the `DMAx_CONFIG` register, and by the necessary setup of the system interrupt controller. If no interrupt is desired, the software can poll for completion by reading the `DMAx_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:

- 1-D 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.

Programming Model

- 2-D interrupt-driven (double buffering)—the DMA buffer is partitioned into two or more sub-buffers, and interrupts are selected (set `DI_SEL = 1` in `DMAX_CONFIG`) to be signaled at the completion of each DMA inner loop. In this way, a traditional double buffer or “ping-pong” scheme can be implemented.

For example, two 512-word sub-buffers inside a 1K-word buffer could be used to receive 16-bit peripheral data with these settings:

`DMAX_START_ADDR` = buffer base address

`DMAX_CONFIG` = `0x10D7` (`FLOW = 1`, `DI_EN = 1`, `DI_SEL = 1`,
`DMA2D = 1`, `WDSIZE = b#01`, `WNR = 1`, `DMAEN = 1`)

`DMAX_X_COUNT` = 512

`DMAX_X_MODIFY` = 2 for 16-bit data

`DMAX_Y_COUNT` = 2 for two sub-buffers

`DMAX_Y_MODIFY` = 2 same as `DMAX_X_MODIFY` for contiguous sub-buffers

- 2-D polled—if interrupt overhead is unacceptable but the loose synchronization of address/count register polling is acceptable, a 2-D 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 2-D DMA buffer can be allocated where each of the four sub-buffers can hold one packet with these settings:

`DMAX_START_ADDR` = buffer base address

`DMAX_CONFIG` = `0x101B` (`FLOW = 1`, `DI_EN = 0`, `DMA2D = 1`,
`WDSIZE = b#10`, `WNR = 1`, `DMAEN = 1`)

`DMAX_X_COUNT` = 16

`DMAX_X_MODIFY` = 4 for 32-bit data

`DMAx_Y_COUNT` = 4 for four sub-buffers

`DMAx_Y_MODIFY` = 4 same as `DMAx_X_MODIFY` for contiguous sub-buffers

The synchronization core might read `DMAx_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 `DMAx_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.

- 1-D 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 1-D 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 1-D or 2-D 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 by selecting the appropriate `FLOW` setting in `DMAx_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.

Programming Model

It is important to remember the meaning of the various fields in the `DMAx_CONFIG` descriptor elements when building a list or array of DMA descriptors. In particular:

- The lower byte of `DMAx_CONFIG` specifies the DMA transfer to be performed by the *current* descriptor (for example 2-D interrupt-enable mode)
- The upper byte of `DMAx_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 `DMAx_CONFIG` value written to the DMA channel's `DMAx_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`, `SYNC`, and `DMA2D` fields will be taken from the `DMAx_CONFIG` value in the descriptor read from memory. The field values initially written to the register are ignored. See [“Initializing Descriptors in Memory” on page 6-94](#) in the [“Programming Examples”](#) section for information on how descriptors can be set up.

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 back to the first.

The code that writes into this descriptor list could use the processor's circular addressing modes (`Ix`, `Lx`, `Mx`, and `Bx` 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 `DMAX_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 are equal), the software increments its count and then starts the DMA unit by writing the new descriptor's `DMAX_CONFIG` value to the DMA channel's `DMAX_CONFIG` register.

Programming Model

If the counts are unequal, the software instead modifies the next-to-last descriptor's `DMAx_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 `DMAx_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 `DMAx_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 (the modification of the next-to-last descriptor's `DMAx_CONFIG` element occurred after that element was read into the DMA unit). In this case, the interrupt handler should write the `DMAx_CONFIG` value appropriate for the last descriptor to the DMA channel's `DMAx_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 `DMAx_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. 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 with a single `DMAx_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 `DMAx_CONFIG` value to use to start the first waiting descriptor in the waiting queue (or 0 to indicate no descriptors are waiting).

Once processing by the DMA unit has started, it is critical that the software not directly modify the contents of the active descriptor queue 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 `DMAx_CONFIG` register. This single register write restarts the queue, effectively transforming the waiting queue to an

Programming Model

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 `DMAx_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 (the mailbox from the interrupt software is zero), the non-interrupt software is responsible for starting the queue (writing the first descriptor's `DMAx_CONFIG` value to the channel's `DMAx_CONFIG` register). If the queue is not stopped, the non-interrupt software must not write to the `DMAx_CONFIG` register (which would cause a DMA error). Instead the descriptor should queue to the waiting queue, and update its mailbox directed to the interrupt handler.

Software Triggered Descriptor Fetches

If a DMA has been stopped in `FLOW = 0` mode, the `DMA_RUN` bit in the `DMAx_IRQ_STATUS` register remains set until the content of the internal DMA FIFOs has been completely processed. Once the `DMA_RUN` bit clears, it is safe to restart the DMA by simply writing again to the `DMAx_CONFIG` register. The DMA sequence is repeated with the previous settings.

Similarly, a descriptor-based DMA sequence that has been stopped temporarily with a `FLOW = 0` descriptor can be continued with a new write to the configuration register. When the DMA controller detects the `FLOW = 0` condition by loading the `DMACFG` field from memory, it has already updated the next descriptor pointer, regardless of whether operating in descriptor array mode or descriptor list mode.

The next descriptor pointer remains valid if the DMA halts and is restarted. As soon as the `DMA_RUN` bit clears, software can restart the DMA

and force the DMA controller to fetch the next descriptor. To accomplish this, the software writes a value with the `DMAEN` bit set and with proper values in the `FLOW` and `NDSIZE` fields into the configuration register. The next descriptor is fetched if `FLOW` equals `0x4`, `0x6`, or `0x7`. In this mode of operation, the `NDSIZE` field should at least span up to the `DMACFG` field to overwrite the configuration register immediately.

One possible procedure is:

1. Write to `DMAX_NEXT_DESC_PTR`

2. Write to `DMAX_CONFIG` with

`FLOW = 0x8`

`NDSIZE ≥ 0xA`

`DI_EN = 0`

`DMAEN = 1`

3. Automatically fetched `DMACFG` has

`FLOW = 0x0`

`NDSIZE = 0x0`

`SYNC = 1` (for transmitting DMAs only)

`DI_EN = 1`

`DMAEN = 1`

4. In the interrupt routine, repeat step 2. The `DMAX_NEXT_DESC_PTR` is updated by the descriptor fetch.



To avoid polling of the `DMA_RUN` bit, set the `SYNC` bit in case of memory read DMAs (DMA transmit or MDMA source).

DMA Registers

If all `DMACFG` fields in a descriptor chain have the `FLOW` and `NDSIZE` fields set to zero, the individual DMA sequences do not start until triggered by software. This is useful when the DMAs need to be synchronized with other events in the system, and it is typically performed by interrupt service routines. A single MMR write is required to trigger the next DMA sequence.

Especially when applied to MDMA channels, such scenarios play an important role. Usually, the timing of MDMAs cannot be controlled (See [“Handshaked Memory DMA Operation” on page 6-36](#)). By halting descriptor chains or rings this way, the whole DMA transaction can be broken into pieces that are individually triggered by software.

 Source and destination channels of a MDMA may differ in descriptor structure. However, the total work count must match when the DMA stops. Whenever a MDMA is stopped, destination and source channels should both provide the same `FLOW = 0` mode after exactly the same number of words. Accordingly, both channels need to be started afterward.

Software-triggered descriptor fetches are illustrated in [Listing 6-7 on page 6-97](#). MDMA channels can be paused by software at any time by writing a 0 to the `DRQ` bit field in the `HMDMAX_CONTROL` register. This simply disables the self-generated DMA requests, whether or not the HMDMA is enabled.

DMA Registers

DMA registers fall into three categories:

- DMA channel registers
- Handshaked MDMA registers
- Global DMA traffic control registers

DMA Channel Registers

A processor features up to twelve peripheral DMA channels and two channel pairs for memory DMA. All channels have an identical set of registers as summarized in [Table 6-4](#).

[Table 6-4](#) 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 where applicable, the corresponding name for the data element in a DMA descriptor.

Table 6-4. Generic Names of the DMA Memory-Mapped Registers

MMR Offset	Generic MMR Name	MMR Description	Register Category	Name of Corresponding Descriptor Element in Memory
0x00	NEXT_DESC_PTR	Link pointer to next descriptor	Parameter	NDPH (upper 16 bits), NDPL (lower 16 bits)
0x04	START_ADDR	Start address of current buffer	Parameter	SAH (upper 16 bits), SAL (lower 16 bits)
0x08	CONFIG	DMA Configuration register, including enable bit	Parameter	DMACFG
0x0C	Reserved	Reserved		
0x10	X_COUNT	Inner loop count	Parameter	XCNT
0x14	X_MODIFY	Inner loop address increment, in bytes	Parameter	XMOD
0x18	Y_COUNT	Outer loop count (2-D only)	Parameter	YCNT
0x1C	Y_MODIFY	Outer loop address increment, in bytes	Parameter	YMOD
0x20	CURR_DESC_PTR	Current descriptor pointer	Current	N/A
0x24	CURR_ADDR	Current DMA address	Current	N/A

DMA Registers

Table 6-4. Generic Names of the DMA Memory-Mapped Registers (Continued)

MMR Offset	Generic MMR Name	MMR Description	Register Category	Name of Corresponding Descriptor Element in Memory
0x28	IRQ_STATUS	Interrupt status register contains completion and DMA error interrupt status and channel state (run/fetch/paused)	Control/Status	N/A
0x2C	PERIPHERAL_MAP	Peripheral to DMA channel mapping contains a 4-bit value specifying the peripheral associated with this DMA channel (read-only for MDMA channels)	Control/Status	N/A
0x30	CURR_X_COUNT	Current count (1-D) or intra-row X count (2-D); counts down from X_COUNT	Current	N/A
0x34	Reserved	Reserved		
0x38	CURR_Y_COUNT	Current row count (2-D only); counts down from Y_COUNT	Current	N/A
0x3C	Reserved	Reserved		

Channel-specific register names are composed of a prefix and the generic MMR name shown in [Table 6-4](#). For peripheral DMA channels the prefix “DMA_x” is used, where “x” stands for a channel number between 0 and 11. For memory DMA channels, the prefix is “MDMA_{yy}”, where “yy” stands for either “D0”, “S0”, “D1”, or “S1” to indicate destination and source channel registers of MDMA0 and MDMA1. For example the peripheral DMA channel 6 configuration register is called DMA6_CONFIG. The register for the MDMA1 source channel is called MDMA_S1_CONFIG.

 The generic MMR names shown in [Table 6-4](#) are not actually mapped to resources in the processor.

For convenience, discussions in this chapter use generic (non-peripheral specific) DMA and memory DMA register names.

DMA channel registers fall into three categories.

- Parameter registers such as `DMAx_CONFIG` and `DMAx_X_COUNT` that can be loaded directly from descriptor elements as shown in [Table 6-4](#)
- Current registers such as `DMAx_CURR_ADDR` and `DMAx_CURR_X_COUNT`
- Control/status registers such as `DMAx_IRQ_STATUS` and `DMAx_PERIPHERAL_MAP`

All DMA registers can be accessed as 16-bit entities. However, the following registers may also be accessed as 32-bit registers.

- `DMAx_NEXT_DESC_PTR`
- `DMAx_START_ADDR`
- `DMAx_CURR_DESC_PTR`
- `DMAx_CURR_ADDR`

 When these four registers are accessed as 16-bit entities, only the lower 16 bits can be accessed.

Because confusion might arise between descriptor element names and generic DMA register names, this chapter uses different naming conventions for physical registers and their corresponding elements in descriptors that reside in memory. [Table 6-4](#) shows the relation.

DMA Registers

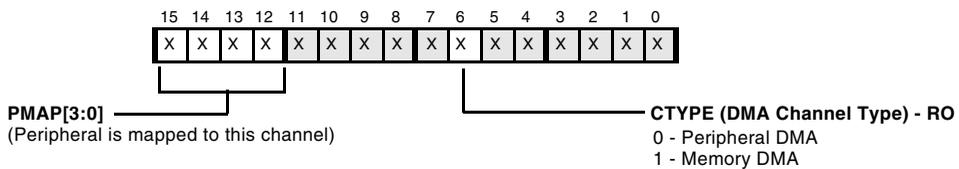
DMA Peripheral Map Registers (DMAx_PERIPHERAL_MAP/ MDMA_yy_PERIPHERAL_MAP)

Each DMA channel's DMAx_PERIPHERAL_MAP register contains bits that:

- Map the channel to a specific peripheral
- Identify whether the channel is a peripheral DMA channel or a memory DMA channel

DMA Peripheral Map Registers (DMAx_PERIPHERAL_MAP/MDMA_yy_PERIPHERAL_MAP)

R/W prior to enabling channel; RO after enabling channel



Default peripheral mappings are provided in [Table 6-7 on page 6-104](#).

Figure 6-5. DMA Peripheral Map Registers

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.

DMA Configuration Registers (DMAx_CONFIG/MDMA_yy_CONFIG)

The DMAx_CONFIG register, shown in Figure 6-6, is used to set up DMA parameters and operating modes. Writing the DMAx_CONFIG register while DMA is already running will cause a DMA error unless writing with the DMAEN bit set to 0.

DMA Configuration Registers (DMAx_CONFIG/MDMA_yy_CONFIG)

R/W prior to enabling channel; RO after enabling channel

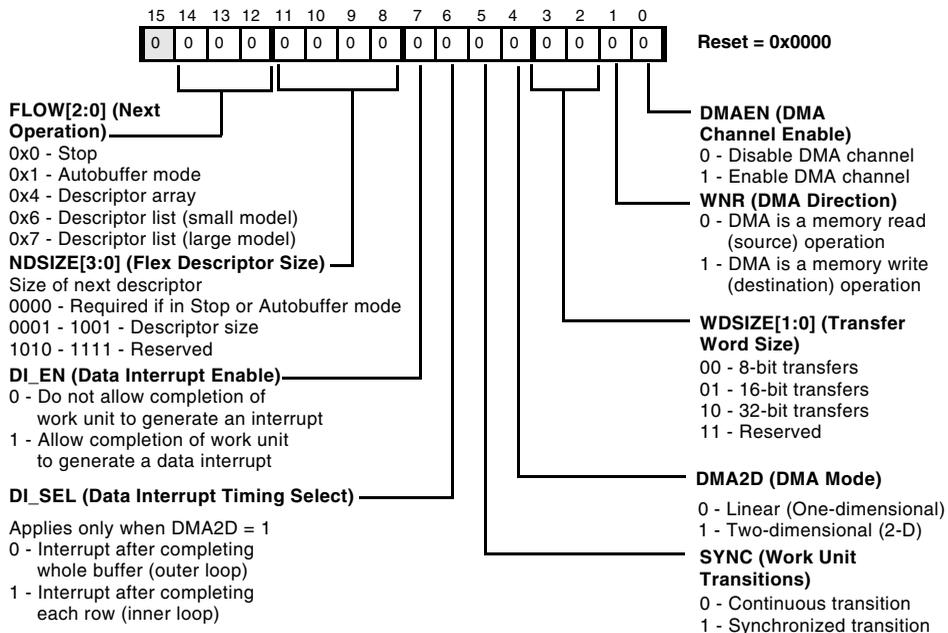


Figure 6-6. DMA Configuration Registers

DMA Registers

The fields of the `DMAx_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_IRQ_STATUS` register changes from 1 to 0, while the `DMAEN` bit in the `DMAx_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_CONFIG` register specifying the next work unit, in which 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 DMA 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_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 `DMAx_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 `DMAx_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 2-D DMA operation.
- `SYNC` (work unit transitions). This bit specifies whether the DMA channel performs a continuous transition (`SYNC = 0`) or a synchronized transition (`SYNC = 1`) between work units. For more information, see [“Work Unit Transitions” on page 6-24](#).

In DMA transmit (memory read) and MDMA source channels, the `SYNC` bit controls the interrupt timing at the end of the work unit and the handling of the DMA FIFO between the current and next work unit.



Work unit transitions for MDMA streams are controlled by the `SYNC` bit of the MDMA source channel's `DMAx_CONFIG` register. The `SYNC` bit of the MDMA destination channel is reserved and must be 0.

DMA Registers

- **DMA2D (DMA mode).** This bit specifies whether DMA mode involves only `DMAx_X_COUNT` and `DMAx_X_MODIFY` (one-dimensional DMA) or also involves `DMAx_Y_COUNT` and `DMAx_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 (although two cycles are required to transfer 32-bit data through a 16-bit memory port or through the 16-bit DMA access bus). The increment sizes (strides) of the DMA address pointer registers must be a multiple of the transfer unit size—one for 8-bit, two for 16-bit, four for 32-bit.

Only SPORT DMA and Memory DMA can operate with a transfer size of 32 bits. All other peripherals have a maximum DMA transfer size of 16 bits.

- **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.

DMA Interrupt Status Registers (DMAx_IRQ_STATUS/MDMA_yy_IRQ_STATUS)

The DMAx_IRQ_STATUS register, shown in [Figure 6-7](#), 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 MDMA transfers where an interrupt is not desired to notify when the DMA operation has ended, software should poll the DMA_DONE bit, rather than the DMA_RUN bit to determine when the transaction has completed.

DMA Registers

DMA Interrupt Status Registers (DMAx_IRQ_STATUS/MDMA_yy_IRQ_STATUS)

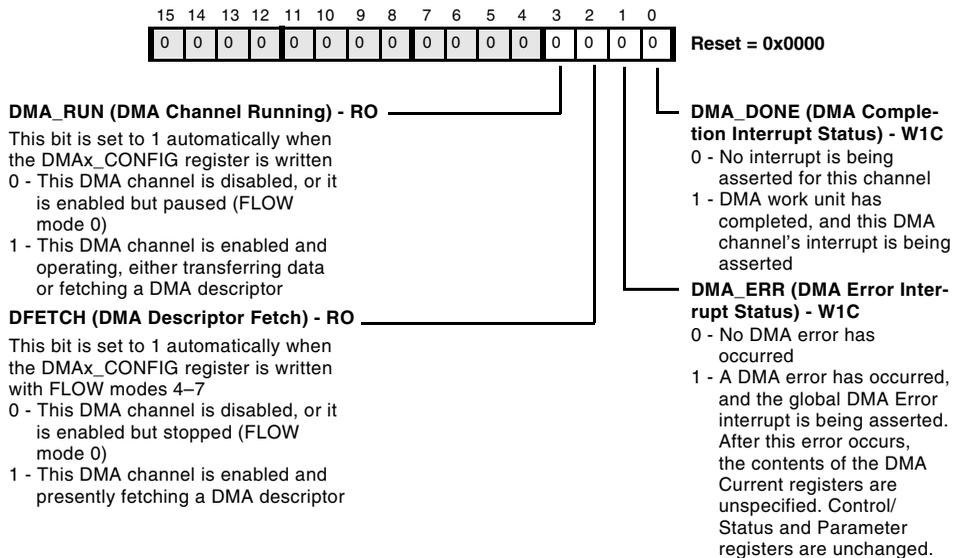


Figure 6-7. DMA Interrupt Status Registers

The processor supports a flexible interrupt control structure with three interrupt sources:

- Data driven interrupts (see [Table 6-5](#))
- Peripheral error interrupts
- DMA error interrupts (for example, bad descriptor or bus error)

Separate interrupt request (IRQ) levels are allocated for data, peripheral error, and DMA error interrupts.

Table 6-5. 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).

The DMA error conditions for all DMA channels are OR'd 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 register or `SIC_IMASK` register) so that no unintended interrupt is generated on the shared DMA/interrupt request line.

DMA Registers

DMA Start Address Registers (DMAx_START_ADDR/MDMA_yy_START_ADDR)

The DMAx_START_ADDR register, shown in [Figure 6-8](#), contains the start address of the data buffer currently targeted for DMA.

DMA Start Address Registers (DMAx_START_ADDR/MDMA_yy_START_ADDR)

R/W prior to enabling channel; RO after enabling channel

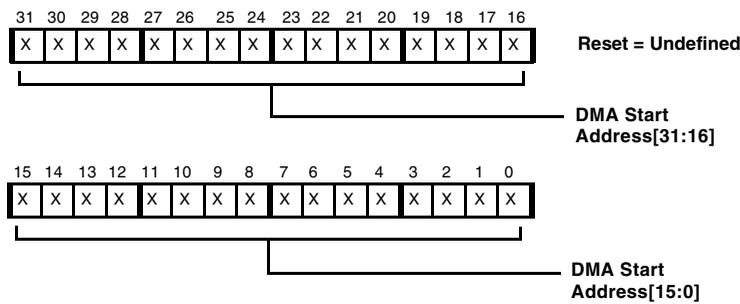


Figure 6-8. DMA Start Address Registers

DMA Current Address Registers (DMAx_CURR_ADDR/MDMA_yy_CURR_ADDR)

The 32-bit DMAx_CURR_ADDR register shown in [Figure 6-9](#), contains the present DMA transfer address for a given DMA session. On the first memory transfer of a DMA work unit, the DMAx_CURR_ADDR register is loaded from the DMAx_START_ADDR register, and it is incremented as each transfer occurs.

DMA Current Address Registers (DMAx_CURR_ADDR/MDMA_yy_CURR_ADDR)

R/W prior to enabling channel; RO after enabling channel

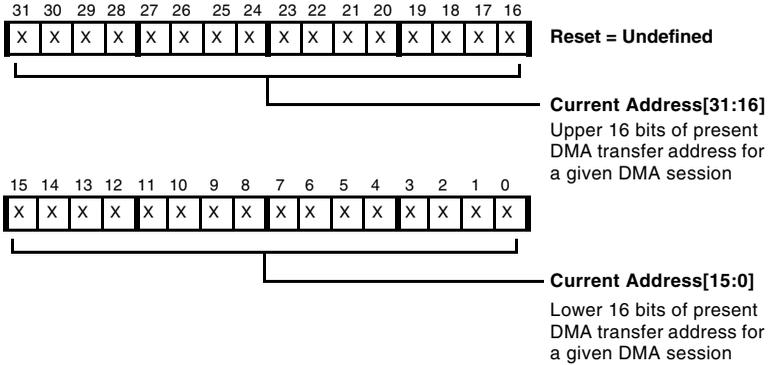


Figure 6-9. DMA Current Address Registers

DMA Inner Loop Count Registers (DMAx_X_COUNT/MDMA_yy_X_COUNT)

For 2-D DMA, the DMAx_X_COUNT register, shown in Figure 6-10, contains the inner loop count. For 1-D DMA, it specifies the number of elements to transfer. For details, see [“Two-Dimensional DMA Operation” on page 6-11](#). A value of 0 in DMAx_X_COUNT corresponds to 65,536 elements.

DMA Inner Loop Count Registers (DMAx_X_COUNT/MDMA_yy_X_COUNT)

R/W prior to enabling channel; RO after enabling channel

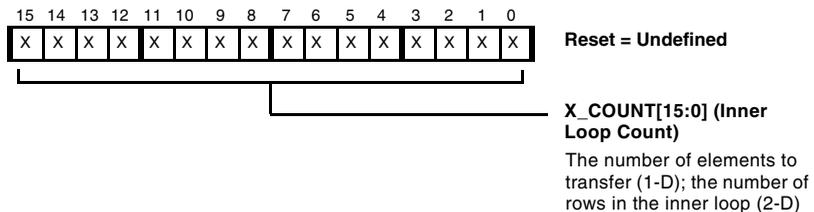


Figure 6-10. DMA Inner Loop Count Registers

DMA Registers

DMA Current Inner Loop Count Registers (DMAx_CURR_X_COUNT /MDMA_yy_CURR_X_COUNT)

The DMAx_CURR_X_COUNT register, shown in [Figure 6-11](#), holds the number of transfers remaining in the current DMA row (inner loop). On the first memory transfer of each DMA work unit, it is loaded with the value in the DMAx_X_COUNT register and then decremented. For 2-D DMA, on the last memory transfer in each row except the last row, it is reloaded with the value in the DMAx_X_COUNT register; this occurs at the same time that the value in the DMAx_CURR_Y_COUNT register is decremented. Otherwise it is decremented each time an element is transferred. Expiration of the count in this register signifies that DMA is complete. In 2-D DMA, the DMAx_CURR_X_COUNT register value is 0 only when the entire transfer is complete. Between rows it is equal to the value of the DMAx_X_COUNT register.

DMA Current Inner Loop Count Registers (DMAx_CURR_X_COUNT/ MDMA_yy_CURR_X_COUNT)

R/W prior to enabling channel; RO after enabling channel

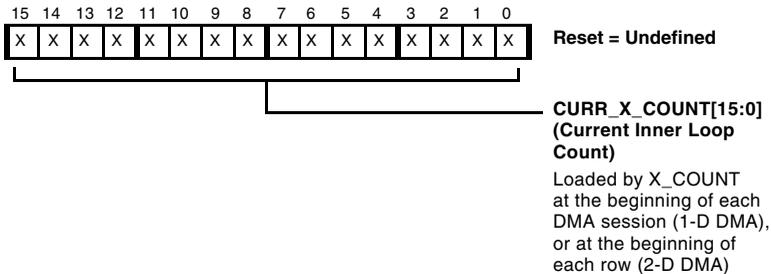


Figure 6-11. DMA Current Inner Loop Count Registers

DMA Inner Loop Address Increment Registers (DMAx_X_MODIFY/MDMA_yy_X_MODIFY)

The DMAx_X_MODIFY register, shown in [Figure 6-12](#), contains a signed, two's-complement byte-address increment. In 1-D DMA, this increment is the stride that is applied after transferring each element.



DMAx_X_MODIFY is specified in bytes, regardless of the DMA transfer size.

In 2-D 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 DMAx_Y_MODIFY register is applied instead, except on the very last transfer of each work unit. The DMAx_X_MODIFY register is always applied to the last transfer of a work unit.

The DMAx_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.

DMA Inner Loop Address Increment Registers (DMAx_X_MODIFY/MDMA_yy_X_MODIFY)

R/W prior to enabling channel; RO after enabling channel

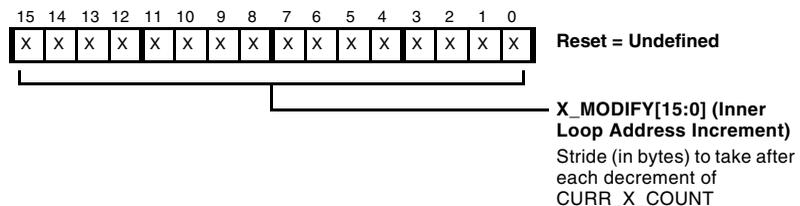


Figure 6-12. DMA Inner Loop Address Increment Registers

DMA Registers

DMA Outer Loop Count Registers (DMAx_Y_COUNT/MDMA_yy_Y_COUNT)

For 2-D DMA, the DMAx_Y_COUNT register, shown in Figure 6-13, contains the outer loop count. It is not used in 1-D DMA mode. This register contains the number of rows in the outer loop of a 2-D DMA sequence. For details, see “Two-Dimensional DMA Operation” on page 6-11.

DMA Outer Loop Count Registers (DMAx_Y_COUNT/MDMA_yy_Y_COUNT)

R/W prior to enabling channel; RO after enabling channel

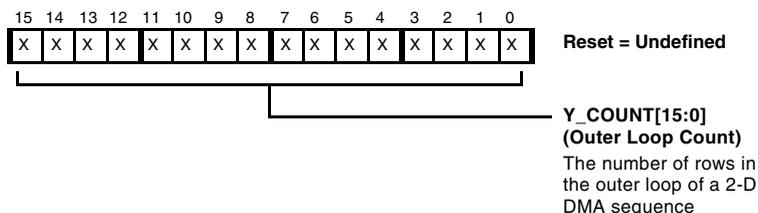


Figure 6-13. DMA Outer Loop Count Registers

DMA Current Outer Loop Count Registers (DMAx_CURR_Y_COUNT/ MDMA_yy_CURR_Y_COUNT)

The DMAx_CURR_Y_COUNT register, used only in 2-D mode, holds the number of full or partial rows (outer loops) remaining in the current work unit. See Figure 6-14. On the first memory transfer of each DMA work unit, it is loaded with the value of the DMAx_Y_COUNT register. The register is decremented each time the DMAx_CURR_X_COUNT register expires during 2-D DMA operation (1 to DMAx_X_COUNT or 1 to 0 transition), signifying completion of an entire row transfer. After a 2-D DMA session is complete, DMAx_CURR_Y_COUNT = 1 and DMAx_CURR_X_COUNT = 0.

DMA Current Outer Loop Count Registers (DMA_x_CURR_Y_COUNT/MDMA_yy_CURR_Y_COUNT)

R/W prior to enabling channel; RO after enabling channel

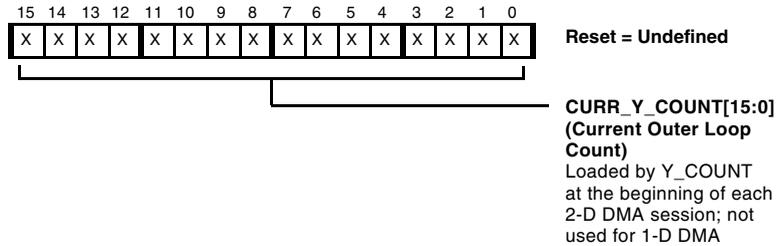


Figure 6-14. DMA Current Outer Loop Count Registers

DMA Outer Loop Address Increment Registers (DMA_x_Y_MODIFY/MDMA_yy_Y_MODIFY)

The DMA_x_Y_MODIFY register contains a signed, two’s-complement value. See [Figure 6-15](#). This byte-address increment is applied after each decrement of the DMA_x_CURR_Y_COUNT register except for the last item in the 2-D array where the DMA_x_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 Operation” on page 6-11](#).



DMA_x_Y_MODIFY is specified in bytes, regardless of the DMA transfer size.

DMA Registers

DMA Outer Loop Address Increment Registers (DMAx_Y_MODIFY/ MDMA_yy_Y_MODIFY) R/W prior to enabling channel; RO after enabling channel

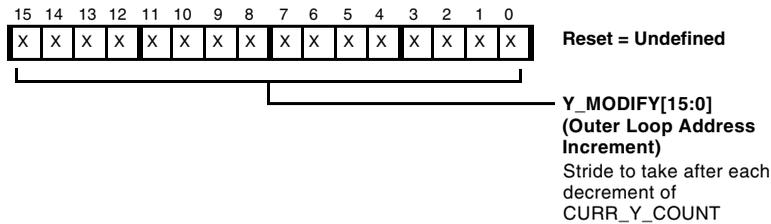


Figure 6-15. DMA Outer Loop Address Increment Registers

DMA Next Descriptor Pointer Registers (DMAx_NEXT_DESC_PTR/ MDMA_yy_NEXT_DESC_PTR)

The 32-bit DMAx_NEXT_DESC_PTR register, shown in [Figure 6-16](#), specifies where to look for the start of the next descriptor block when the DMA activity specified by the current descriptor block finishes. This register is used in small and large descriptor list modes. At the start of a descriptor fetch in either of these modes, this register is copied into the DMAx_CURR_DESC_PTR register. Then, during the descriptor fetch, the DMAx_CURR_DESC_PTR register increments after each element of the descriptor is read in.

i In small and large descriptor list modes, the DMAx_NEXT_DESC_PTR register, and not the DMAx_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 DMAx_CURR_DESC_PTR register.

DMA Next Descriptor Pointer Registers

(DMAx_NEXT_DESC_PTR/MDMA_yy_NEXT_DESC_PTR)

R/W prior to enabling channel; RO after enabling channel

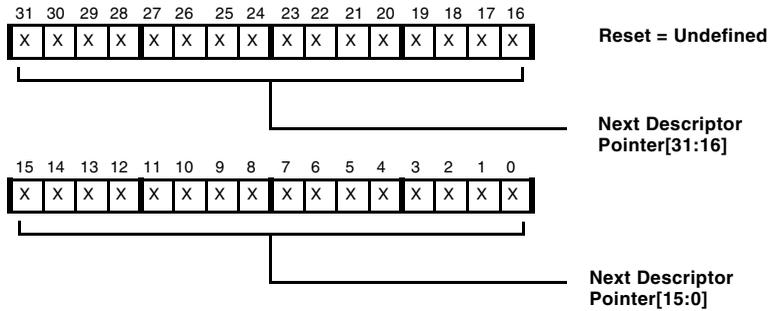


Figure 6-16. DMA Next Descriptor Pointer Registers

DMA Current Descriptor Pointer Registers

(DMAx_CURR_DESC_PTR/
MDMA_yy_CURR_DESC_PTR)

The 32-bit DMAx_CURR_DESC_PTR register, shown in [Figure 6-17](#), 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 DMAx_NEXT_DESC_PTR register before loading each descriptor. Then, the address in the DMAx_CURR_DESC_PTR register increments as each descriptor element is read in.

When the entire descriptor has been read, the DMAx_CURR_DESC_PTR register contains this value:

Descriptor Start Address + (2 × Descriptor Size) (# of elements)



For descriptor array mode (FLOW = 4), this register, and not the DMAx_NEXT_DESC_PTR register, must be programmed by MMR access before starting DMA operation.

DMA Registers

DMA Next Descriptor Pointer Registers (DMA_x_NEXT_DESC_PTR/MDMA_{yy}_NEXT_DESC_PTR)

R/W prior to enabling channel; RO after enabling channel

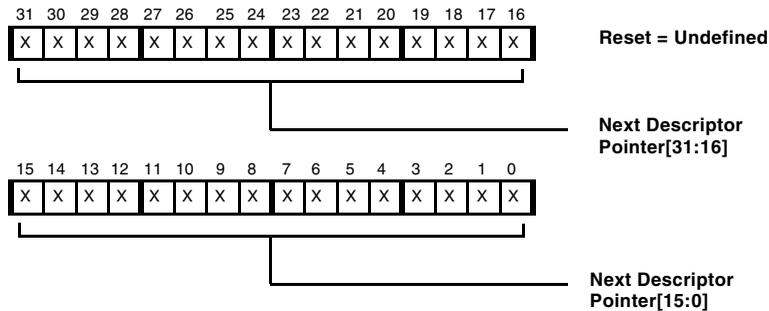


Figure 6-17. DMA Current Descriptor Pointer Registers

HMDMA Registers

Some processors have two HMDMA blocks, while others have none. See the “[Unique Information for the ADSP-BF51x Processor](#)” on page 6-102 to determine whether this feature is applicable to your product.

HMDMA0 is associated with MDMA0, and HMDMA1 is associated with MDMA1.

Handshake MDMA Control Registers (HMDMA_x_CONTROL)

The HMDMA_x_CONTROL register, shown in [Figure 6-18](#), is used to set up HMDMA parameters and operating modes.

The DRQ[1:0] field is used to control the priority of the MDMA channel when the HMDMA is disabled, that is, when handshake control is not being used (see [Table 6-6](#)).

Handshake MDMA Control Registers (HMDMAx_CONTROL)

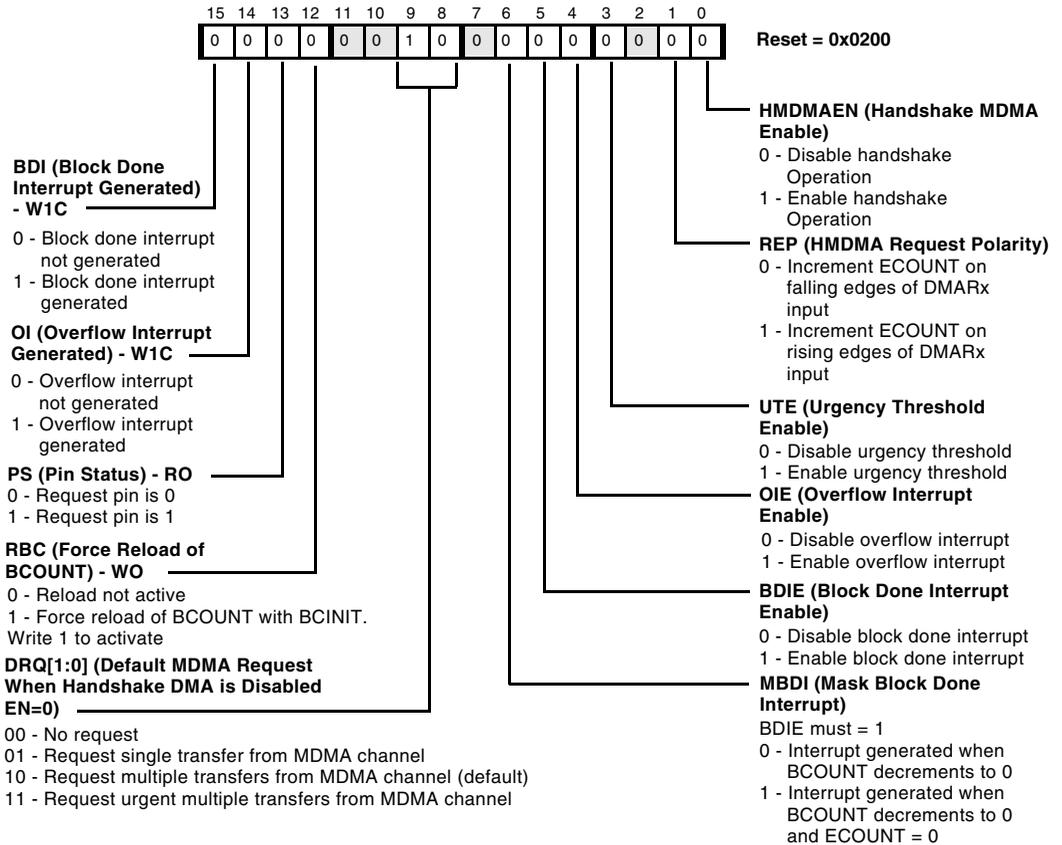


Figure 6-18. Handshake MDMA Control Registers

DMA Registers

Table 6-6. DRQ[1:0] Values

DRQ[1:0]	Priority	Description
00	Disabled	The MDMA request is disabled.
01	Enabled/S	Normal MDMA channel priority. The channel in this mode is limited to single memory transfers separated by one idle system clock. Request single transfer from MDMA channel.
10	Enabled/M	Normal MDMA channel functionality and priority. Request multiple transfers from MDMA channel (default).
11	Urgent	The MDMA channel priority is elevated to urgent. In this state, it has higher priority for memory access than non-urgent channels. If two channels are both urgent, the lower-numbered channel has priority.

The **RBC** bit forces the **BCOUNT** register to be reloaded with the **BCINIT** value while the module is already active. Do not set this bit in the same write that sets the **HMDMAEN** bit to active.

Handshake MDMA Initial Block Count Registers (HMDMAx_BCINIT)

The **HMDMAx_BCINIT** register, shown in [Figure 6-19](#), holds the number of transfers to do per edge of the **DMARx** control signal.

Handshake MDMA Initial Block Count Registers (HMDMAx_BCINIT)

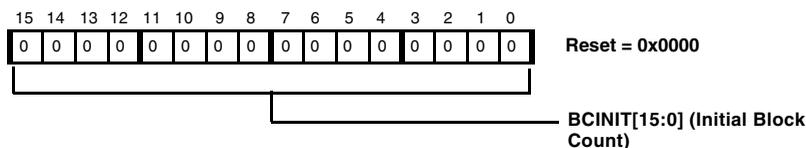


Figure 6-19. Handshake MDMA Initial Block Count Registers

Handshake MDMA Current Block Count Registers (HMDMAx_BCOUNT)

The HMDMAx_BCOUNT register, shown in [Figure 6-20](#), holds the number of transfers remaining for the current edge. MDMA requests are generated if this count is greater than 0.

Examples:

- 0000 = 0 transfers remaining
- FFFF = 65535 transfers remaining

The BCOUNT field is loaded with BCINIT when ECOUNT is greater than 0 and BCOUNT is expired (0). Also, if the RBC bit in the HMDMAx_CONTROL register is written to 1, BCOUNT is loaded with BCINIT. The BCOUNT field is decremented with each MDMA grant. It is cleared when HMDMA is disabled.

A block done interrupt is generated when BCOUNT decrements to 0. If the MBDI bit in the HMDMAx_CONTROL register is set, the interrupt is suppressed until ECOUNT is 0. If BCINIT is 0, no block done interrupt is generated, since no DMA requests were generated or grants received.

Handshake MDMA Current Block Count Register (HMDMAx_BCOUNT)

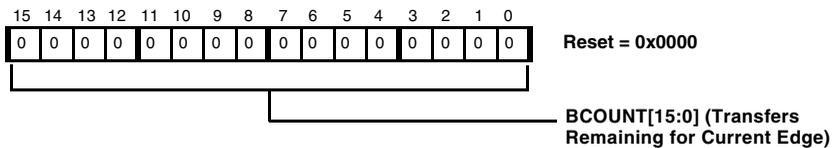


Figure 6-20. Handshake MDMA Current Block Count Registers

DMA Registers

Handshake MDMA Current Edge Count Registers (HMDMAx_ECOUNT)

The HMDMAx_ECOUNT register, shown in Figure 6-21, holds a signed number of edges remaining to be serviced. This number is in a signed two's-complement representation. When an edge is detected on the respective DMARx input, requests occur if this count is greater than or equal to 0 and BCOUNT is greater than 0.

When the handshake mode is enabled, ECOUNT is loaded and the resulting number of requests is:

Number of edges + N,

where N is the number loaded from ECINIT. The number N can be positive or negative. Examples:

- 0x7FFF = 32,767 edges remaining
- 0x0000 = 0 edges remaining
- 0x8000 = -32,768: ignore the next 32,768 edges

Each time that BCOUNT expires, ECOUNT is decremented and BCOUNT is reloaded from BCINIT. When a handshake request edge is detected, ECOUNT is incremented. The ECOUNT field is cleared when HMDMA is disabled.

Handshake MDMA Current Edge Count Register (HMDMAx_ECOUNT)

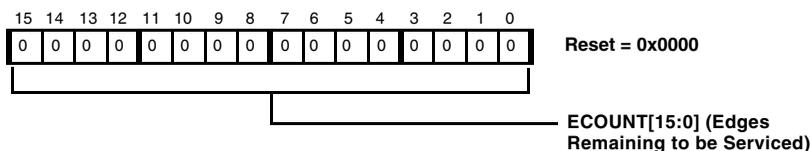


Figure 6-21. Handshake MDMA Current Edge Count Registers

Handshake MDMA Initial Edge Count Registers (HMDMA_x_ECINIT)

The HMDMA_x_ECINIT register, shown in [Figure 6-22](#), holds a signed number that is loaded into HMDMA_x_ECOUNT when handshake DMA is enabled. This number is in a signed two's complement representation.

Handshake MDMA Initial Edge Count Registers (HMDMA_x_ECINIT)

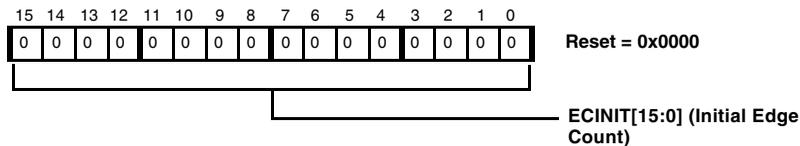


Figure 6-22. Handshake MDMA Initial Edge Count Registers

Handshake MDMA Edge Count Urgent Registers (HMDMA_x_ECURGENT)

The HMDMA_x_ECURGENT register, shown in [Figure 6-23](#), holds the urgent threshold. If the ECOUNT field in the HMDMA_x_ECOUNT register is greater than this threshold, the MDMA request is urgent and might get higher priority.

Handshake MDMA Edge Count Urgent Registers (HMDMA_x_ECURGENT)

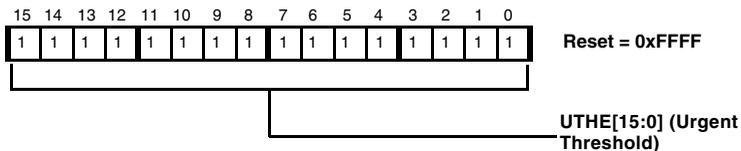


Figure 6-23. Handshake MDMA Edge Count Urgent Registers

DMA Registers

Handshake MDMA Edge Count Overflow Interrupt Registers (HMDMAx_ECOVERFLOW)

The HMDMAx_ECOVERFLOW register, shown in [Figure 6-24](#), holds the interrupt threshold. If the ECOUNT field in the HMDMAx_ECOUNTER register is greater than this threshold, an overflow interrupt is generated.

Handshake MDMA Edge Count Overflow Interrupt Registers (HMDMAx_ECOVERFLOW)

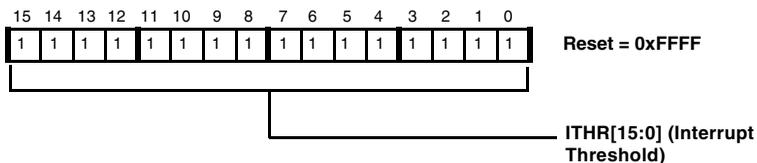


Figure 6-24. Handshake MDMA Edge Count Overflow Interrupt Registers

DMA Traffic Control Registers (DMA_TC_PER and DMA_TC_CNT)

The DMA_TC_PER register (see [Figure 6-25](#)) and the DMA_TC_CNT register (see [Figure 6-26](#)) work with other DMA registers to define traffic control.

DMA_TC_PER Register

DMA Traffic Control Counter Period Register (DMA_TC_PER)

1

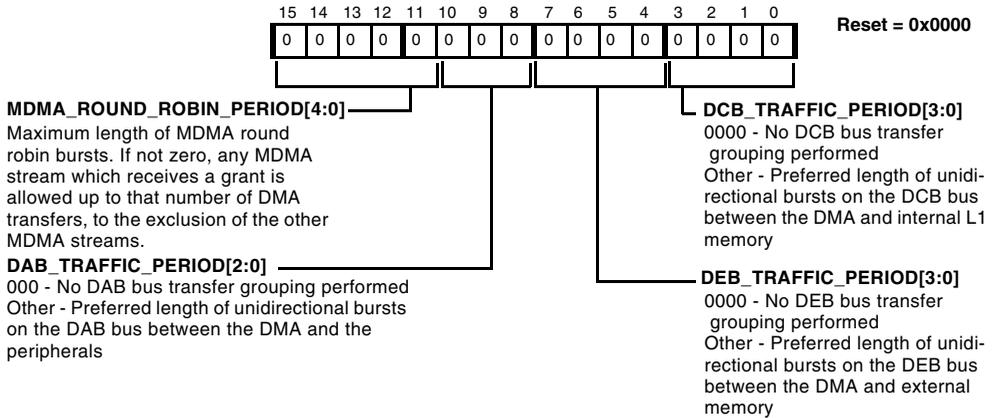


Figure 6-25. DMA Traffic Control Counter Period Register

DMA_TC_CNT Register

DMA Traffic Control Counter Register (DMA_TC_CNT)

RO

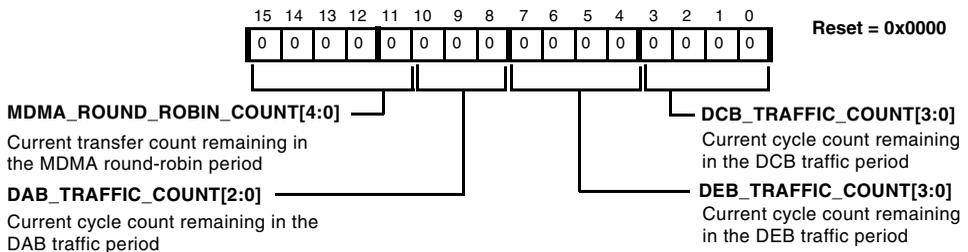


Figure 6-26. DMA Traffic Control Counter Register

The MDMA_ROUND_ROBIN_COUNT field shows the current transfer count remaining in the MDMA round-robin period. It initializes to

DMA Registers

`MDMA_ROUND_ROBIN_PERIOD` whenever `DMA_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 `DMA_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 `DMA_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 `DMA_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.

Programming Examples

The following examples illustrate memory DMA and handshaked memory DMA basics. Examples for peripheral DMAs can be found in the respective peripheral chapters.

Register-Based 2-D Memory DMA

[Listing 6-1](#) shows a register-based, two-dimensional MDMA. While the source channel processes linearly, the destination channel re-sorts elements by transposing the two-dimensional data array. See [Figure 6-27](#).

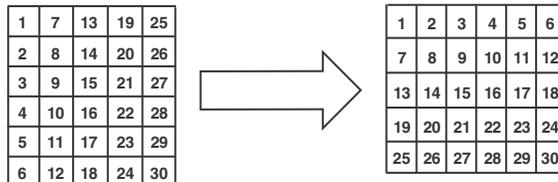


Figure 6-27. DMA Example, 2-D Array

The two arrays reside in two different L1 data memory blocks. However, the arrays could reside in any internal or external memory, including L1 instruction memory and SDRAM. For the case where the destination array resided in SDRAM, it is a good idea to let the source channel re-sort elements and to let the destination buffer store linearly.

Programming Examples

Listing 6-1. Register-Based 2-D Memory DMA

```
#include <defBF527.h> /*For ADSP-BF527 product, as an example.*/
#define X 5
#define Y 6

.section L1_data_a;
.byte2 aSource[X*Y] =
    1,  7, 13, 19, 25,
    2,  8, 14, 20, 26,
    3,  9, 15, 21, 27,
    4, 10, 16, 22, 28,
    5, 11, 17, 23, 29,
    6, 12, 18, 24, 30;

.section L1_data_b;
.byte2 aDestination[X*Y];

.section L1_code;
.global _main;
_main:
    p0.l = lo(MDMA_S0_CONFIG);
    p0.h = hi(MDMA_S0_CONFIG);
    call memdma_setup;
    call memdma_wait;
_main.forever:
    jump _main.forever;
_main.end:
```

The setup routine shown in [Listing 6-2](#) initializes either MDMA0 or MDMA1, depending on whether the MMR address of MDMA_S0_CONFIG or MDMA_S1_CONFIG is passed in the P0 register. Note that the source channel is enabled before the destination channel. Also, it is common to synchronize interrupts with the destination channel because only those interrupts indicate completion of both DMA read and write operations.

Listing 6-2. Two-Dimensional Memory DMA Setup Example

```

memdma_setup:
    [--sp] = r7;
/* setup 1D source DMA for 16-bit transfers */
    r7.l = lo(aSource);
    r7.h = hi(aSource);
    [p0 + MDMA_SO_START_ADDR - MDMA_SO_CONFIG] = r7;
    r7.l = 2;
    w[p0 + MDMA_SO_X_MODIFY - MDMA_SO_CONFIG] = r7;
    r7.l = X * Y;
    w[p0 + MDMA_SO_X_COUNT - MDMA_SO_CONFIG] = r7;
    r7.l = WDSIZE_16 | DMAEN;
    w[p0] = r7;
/* setup 2D destination DMA for 16-bit transfers */
    r7.l = lo(aDestination);
    r7.h = hi(aDestination);
    [p0 + MDMA_DO_START_ADDR - MDMA_SO_CONFIG] = r7;
    r7.l = 2*Y;
    w[p0 + MDMA_DO_X_MODIFY - MDMA_SO_CONFIG] = r7;
    r7.l = Y;
    w[p0 + MDMA_DO_Y_COUNT - MDMA_SO_CONFIG] = r7;
    r7.l = X;
    w[p0 + MDMA_DO_X_COUNT - MDMA_SO_CONFIG] = r7;
    r7.l = -2 * (Y * (X-1) - 1);
    w[p0 + MDMA_DO_Y_MODIFY - MDMA_SO_CONFIG] = r7;
    r7.l = DMA2D | DI_EN | WDSIZE_16 | WNR | DMAEN;
    w[p0 + MDMA_DO_CONFIG - MDMA_SO_CONFIG] = r7;
    r7 = [sp++];
    rts;
memdma_setup.end:

```

For simplicity the example shown in [Listing 6-3](#) polls the DMA status rather than using interrupts, which is the normal case in a real application.

Programming Examples

Listing 6-3. Polling DMA Status

```
memdma_wait:
    [--sp] = r7;
memdma_wait.test:
    r7 = w[p0 + MDMA_DO_IRQ_STATUS - MDMA_S0_CONFIG] (z);
    CC = bittst (r7, bitpos(DMA_DONE));
    if !CC jump memdma_wait.test;
    r7 = DMA_DONE (z);
    w[p0 + MDMA_DO_IRQ_STATUS - MDMA_S0_CONFIG] = r7;
    r7 = [sp++];
    rts;
memdma_wait.end:
```

Initializing Descriptors in Memory

Descriptor-based DMAs expect the descriptor data to be available in memory by the time the DMA is enabled. Often, the descriptors are programmed by software at run-time. Many times, however, the descriptors—or at least large portions of them—can be static and therefore initialized at boot time. How to set up descriptors in global memory depends heavily on the programming language and the tool set used. The following examples show how this is best performed in the VisualDSP++ tools' assembly language.

[Listing 6-4](#) uses multiple variables of either 16-bit or 32-bit size to describe DMA descriptors. This example has two descriptors in small list flow mode that point to each other. At the end of the second work unit, an interrupt is generated without discontinuing the DMA processing. The trailing `.end` label is required to let the linker know that a descriptor forms a logical unit. It prevents the linker from removing variables when optimizing.

Listing 6-4. Two Descriptors in Small List Flow Mode

```

.section sdram;
.byte2 arrBlock1[0x400];
.byte2 arrBlock2[0x800];

.section L1_data_a;
.byte2 descBlock1 = 1o(descBlock2);
.var descBlock1.addr = arrBlock1;
.byte2 descBlock1.cfg = FLOW_SMALL|NDSIZE_5|WDSIZE_16|DMAEN;
.byte2 descBlock1.len = length(arrBlock1);
descBlock1.end:

.byte2 descBlock2 = 1o(descBlock1);
.var descBlock2.addr = arrBlock2;
.byte2 descBlock2.cfg =
FLOW_SMALL|NDSIZE_5|DI_EN|WDSIZE_16|DMAEN;
.byte2 descBlock2.len = length(arrBlock2);
descBlock2.end:

```

Another method featured by the VisualDSP++ tools takes advantage of C-style structures in global header files. The header file `descriptors.h` could look like [Listing 6-5](#).

Listing 6-5. Header File to Define Descriptor Structures

```

#ifndef __INCLUDE_DESCRIPTOR__
#define __INCLUDE_DESCRIPTOR__
#ifdef _LANGUAGE_C
typedef struct {
    void *pStart;
    short dConfig;
    short dxCount;
    short dxModify;
    short dyCount;

```

Programming Examples

```
        short dYModify;
    } dma_desc_arr;

typedef struct {
    void *pNext;
    void *pStart;
    short dConfig;
    short dxCount;
    short dxModify;
    short dYCount;
    short dYModify;
} dma_desc_list;

#endif // _LANGUAGE_C
#endif // __INCLUDE_DESCRIPTOR__
```

Note that near pointers are not natively supported by the C language and, thus, pointers are always 32 bits wide. Therefore, the scheme above cannot be used directly for small list mode without giving up pointer syntax. The variable definition file is required to import the C-style header file and can finally take advantage of the structures. See [Listing 6-6](#).

Listing 6-6. Using Descriptor Structures

```
#include "descriptors.h"
.import "descriptors.h";

.section L1_data_a;
.align 4;
.var arrBlock3[N];
.var arrBlock4[N];

.struct dma_desc_list descBlock3 = {
    descBlock4, arrBlock3,
    FLOW_LARGE | NDSIZE_7 | WDSIZE_32 | DMAEN,
```

```

    length(arrBlock3), 4,
    0, 0          /* unused values */
};

.struct dma_desc_list descBlock4 = {
    descBlock3, arrBlock4,
    FLOW_LARGE | NDSIZE_7 | DI_EN | WDSIZE_32 | DMAEN,
    length(arrBlock4), 4,
    0, 0          /* unused values */
};

```

Software-Triggered Descriptor Fetch Example

[Listing 6-7](#) demonstrates a large list of descriptors that provide `FLOW = 0` (stop mode) configuration. Consequently, the DMA stops by itself as soon as the work unit has finished. Software triggers the next work unit by simply writing the proper value into the DMA configuration registers. Since these values instruct the DMA controller to fetch descriptors in large list mode, the DMA immediately fetches the descriptor, thus overwriting the configuration value again with the new settings when it is started.

Note the requirement that source and destination channels stop after the same number of transfers. Between stops, the two channels can have completely individual structures.

Listing 6-7. Software-Triggered Descriptor Fetch

```

.import "descriptors.h";

#define N 4
.section L1_data_a;
.byte2 arrSource1[N] = { 0x1001, 0x1002, 0x1003, 0x1004 };
.byte2 arrSource2[N] = { 0x2001, 0x2002, 0x2003, 0x2004 };
.byte2 arrSource3[N] = { 0x3001, 0x3002, 0x3003, 0x3004 };

```

Programming Examples

```
.byte2 arrDest1[N];
.byte2 arrDest2[2*N];

.struct dma_desc_list descSource1 = {
    descSource2, arrSource1,
    WDSIZE_16 | DMAEN,
    length(arrSource1), 2,
    0, 0          /* unused values */
};
.struct dma_desc_list descSource2 = {
    descSource3, arrSource2,
    FLOW_LARGE | NDSIZE_7 | WDSIZE_16 | DMAEN,
    length(arrSource2), 2,
    0, 0          /* unused values */
};
.struct dma_desc_list descSource3 = {
    descSource1, arrSource3,
    WDSIZE_16 | DMAEN,
    length(arrSource3), 2,
    0, 0          /* unused values */
};
.struct dma_desc_list descDest1 = {
    descDest2, arrDest1,
    DI_EN | WDSIZE_16 | WNR | DMAEN,
    length(arrDest1), 2,
    0, 0          /* unused values */
};
.struct dma_desc_list descDest2 = {
    descDest1, arrDest2,
    DI_EN | WDSIZE_16 | WNR | DMAEN,
    length(arrDest2), 2,
    0, 0          /* unused values */
};
```

```

.section L1_code;
_main:
/* write descriptor address to next descriptor pointer */
    p0.h = hi(MDMA_SO_CONFIG);
    p0.l = lo(MDMA_SO_CONFIG);
    r0.h = hi(descDest1);
    r0.l = lo(descDest1);
    [p0 + MDMA_DO_NEXT_DESC_PTR - MDMA_SO_CONFIG] = r0;
    r0.h = hi(descSource1);
    r0.l = lo(descSource1);
    [p0 + MDMA_SO_NEXT_DESC_PTR - MDMA_SO_CONFIG] = r0;

/* start first work unit */
    r6.l = FLOW_LARGE|NDSIZE_7|WDSIZE_16|DMAEN;
    w[p0 + MDMA_SO_CONFIG - MDMA_SO_CONFIG] = r6;
    r7.l = FLOW_LARGE|NDSIZE_7|WDSIZE_16|WNR|DMAEN;
    w[p0 + MDMA_DO_CONFIG - MDMA_SO_CONFIG] = r7;

/* wait until destination channel has finished and W1C latch */
_main.wait:
    r0 = w[p0 + MDMA_DO_IRQ_STATUS - MDMA_SO_CONFIG] (z);
    CC = bittst (r0, bitpos(DMA_DONE));
    if !CC jump _main.wait;
    r0.l = DMA_DONE;
    w[p0 + MDMA_DO_IRQ_STATUS - MDMA_SO_CONFIG] = r0;

/* wait for any software or hardware event here */

/* start next work unit */
    w[p0 + MDMA_SO_CONFIG - MDMA_SO_CONFIG] = r6;
    w[p0 + MDMA_DO_CONFIG - MDMA_SO_CONFIG] = r7;
    jump _main.wait;
_main.end:

```

Handshaked Memory DMA Example

The functional block for the handshaked MDMA operation can be considered completely separately from the MDMA channels themselves. Therefore the following HMDMA setup routine can be combined with any of the MDMA examples discussed above. Be sure that the HMDMA module is enabled before the MDMA channels.

[Listing 6-8](#) enables the HMDMA1 block, which is controlled by the DMAR1 pin and is associated with the MDMA1 channel pair.

Listing 6-8. HMDMA1 Block Enable

```
/* optionally, enable all four bank select strobes */
    p1.l = lo(EBIU_AMGCTL);
    p1.h = hi(EBIU_AMGCTL);
    r0.l = 0x0009;
    w[p1] = r0;

/* function enable for DMAR1 */
    p1.l = lo(PORTG_FER);
    r0.l = PG12;
    w[p1] = r0;
    p1.l = lo(PORTG_MUX);
    r0.l = 0x0000;
    w[p1] = r0;

/* every single transfer requires one DMAR1 event */
    p1.l = lo(HMDMA1_BCINIT);
    r0.l = 1;
    w[p1] = r0;

/* start with balanced request counter */
    p1.l = lo(HMDMA1_ECINIT);
    r0.l = 0;
```

```

w[p1] = r0;

/* enable for rising edges */
p1.l = 1o(HMDMA1_CONTROL);
r2.l = REP | HMDMAEN;
w[p1] = r2;

```

If the HMDMA is intended to copy from internal memory to external devices, the above setup is sufficient. If, however, the data flow is from outside the processor to internal memory, then this small issue must be considered—the HMDMA only controls the destination channel of the memory DMA. It does not gate requests to the source channel at all. Thus, as soon as the source channel is enabled, it starts filling the DMA FIFO immediately. In 16-bit DMA mode, this results in eight read strobes on the EBIU even before the first DMAR1 event has been detected. In other words, the transferred data and the DMAR1 strobes are eight positions off. The example in [Listing 6-9](#) delays processing until eight DMAR1 requests have been received. By doing so, the transmitter is required to add eight trailing dummy writes after all data words have been sent. This is because the transmit channel still has to drain the DMA FIFO.

Listing 6-9. HMDMA With Delayed Processing

```

/* wait for eight requests */
p1.l = 1o(HMDMA1_ECOUNTER);
r0 = 7 (z);
initial_requests:
r1 = w[p1] (z);
CC = r1 < r0;
if CC jump initial_requests;

/* disable and reenable to clear edge count */
p1.l = 1o(HMDMA1_CONTROL);
r0.l = 0;

```

Unique Information for the ADSP-BF51x Processor

```
w[p1] = r0;  
w[p1] = r2;
```

If the polling operation shown in [Listing 6-9](#) is too expensive, an interrupt version of it can be implemented by using the HMDMA overflow feature. Temporarily set the `HMDMAx_OVERFLOW` register to eight.

Unique Information for the ADSP-BF51x Processor

[Figure 6-28](#) provides a block diagram of the ADSP-BF51x DMA controller.

DMA Control Commands

The ADSP-BF516 and ADSP-BF518 processors both offer an Ethernet MAC module that supports DMA control commands. Refer to [Chapter 22, “Ethernet MAC”](#) for a description of how these commands are used.

Static Channel Prioritization

The default DMA channel priority and mapping shown in [Table 6-7](#) can be changed by altering the 4-bit PMAP field in the `DMAx_PERIPHERAL_MAP` registers for the peripheral DMA channels.

 Note that the ADSP-BF512 and ADSP-BF514 processors do not feature the Ethernet MAC module. Therefore, the DMA1 and DMA2 channels cannot be used with the default channel mapping on these products. Also, the ADSP-BF512 does not have an RSI module, but DMA3 can still be used for SPORT0 RX transfers in this case.

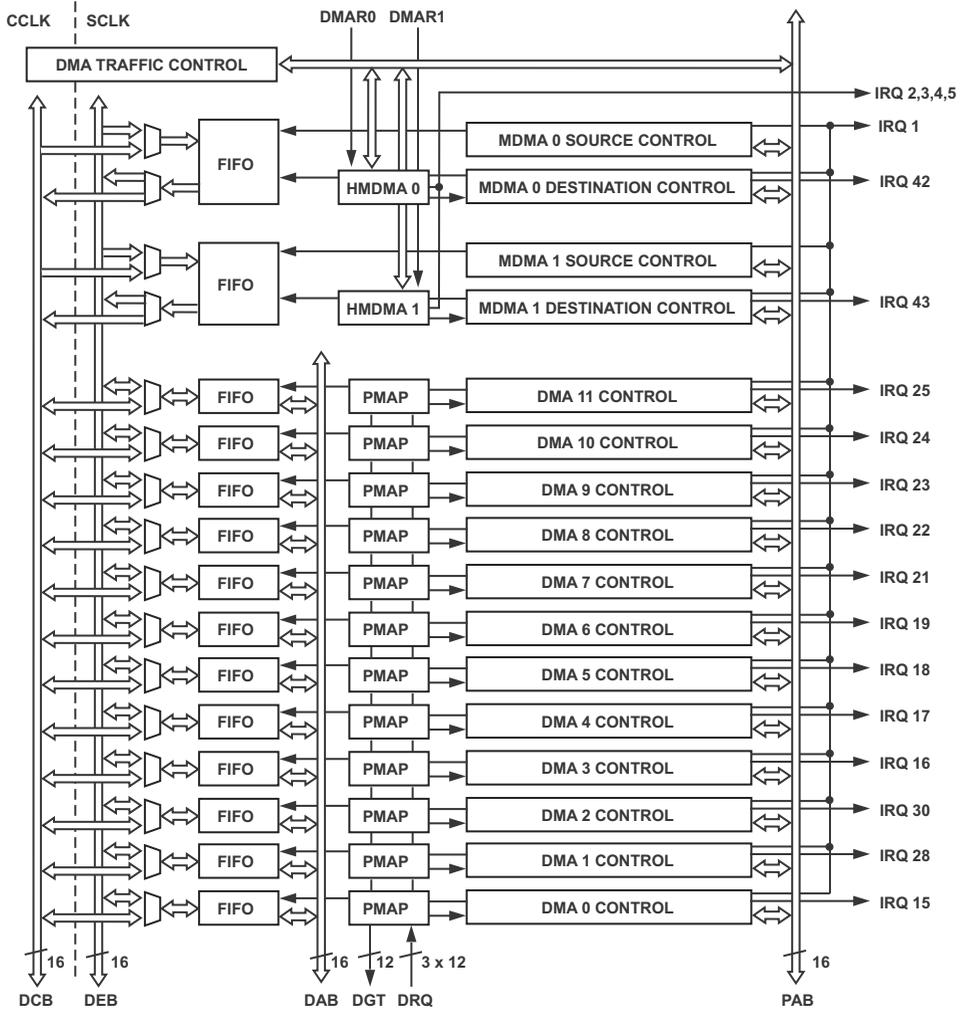


Figure 6-28. ADSP-BF51x DMA Controller Block Diagram

Unique Information for the ADSP-BF51x Processor

Table 6-7. Priority and Default Mapping of Peripheral to DMA

Priority	DMA Channel	PMP Default Value	Peripheral Mapped by Default
Highest	DMA 0	0x0	PPI receive or transmit
	DMA 1	0x1	Ethernet MAC receive
	DMA 2	0x2	Ethernet MAC transmit
	DMA 3	0x3	SPORT0 receive
	DMA 4	0x4	SPORT0 transmit or RSI
	DMA 5	0x5	SPORT1 receive or SPI1 transmit/receive
	DMA 6	0x6	SPORT1 transmit
	DMA 7	0x7	SPI0 transmit/receive
	DMA 8	0x8	UART0 receive
	DMA 9	0x9	UART0 transmit
	DMA 10	0xA	UART1 receive
	DMA 11	0xB	UART1 transmit
	MDMA D0	N/A	N/A
	MDMA S0	N/A	N/A
	MDMA D1	N/A	N/A
Lowest	MDMA S1	N/A	N/A

7 EXTERNAL BUS INTERFACE UNIT

The external bus interface unit (EBIU) provides glueless interfaces to external memories. The processor supports Synchronous DRAM (SDRAM) including mobile SDRAM. The EBIU also supports asynchronous interfaces such as SRAM, ROM, FIFOs, flash memory, and ASIC/FPGA designs.

EBIU 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 DMA controller provides high-bandwidth data movement capability. The Memory DMA (MDMA) channels can perform block transfers of code or data between the internal memory and the external memory spaces. The MDMA channels also feature a Handshake Operation mode (HMDMA) via dual external DMA request pins. When used in conjunction with the EBIU, this functionality can be used to interface high-speed external devices, such as FIFOs and USB 2.0 controllers, in an automatic manner. For more information on HMDMA and the external DMA request pins, please refer to [Chapter 6, “Direct Memory Access”](#).

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 clock frequency (CCLK) and SCLK frequency is programmable

EBIU Overview

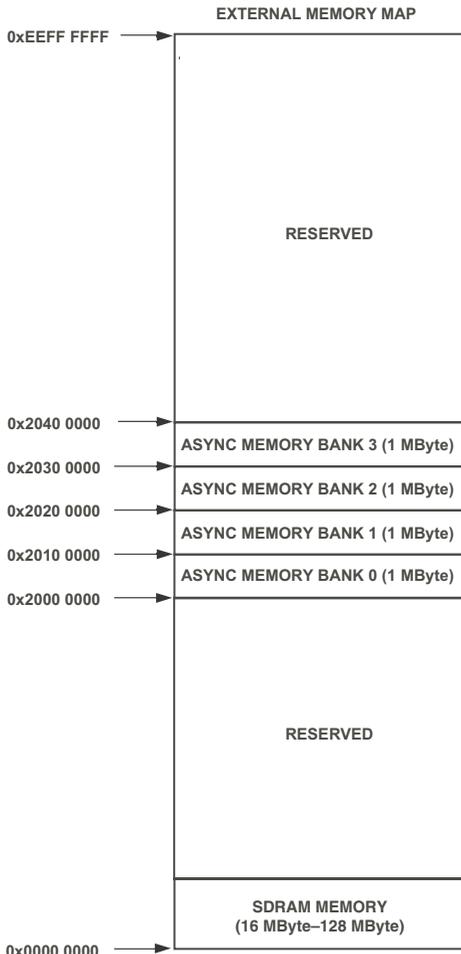
using a Phase Locked Loop (PLL) system Memory-Mapped Register (MMR). For more information, see “Core Clock/System Clock Ratio Control” on page 8-5.

The external memory space is shown in [Figure 7-1](#). 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 16M byte to 128M byte.

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 are dedicated to supporting asynchronous memories. Each asynchronous memory region 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.



NOTE: RESERVED OFF-CHIP MEMORY AREAS ARE LABELED IN THE DIAGRAM ABOVE. ALL OTHER OFF-CHIP SYSTEM RESOURCES ARE ADDRESSABLE BY BOTH THE CORE AND THE SYSTEM.

Figure 7-1. External Memory Map

Block Diagram

Figure 7-2 is a conceptual block diagram of the EBIU and its interfaces. Signal names shown with an overbar are active low signals.

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.

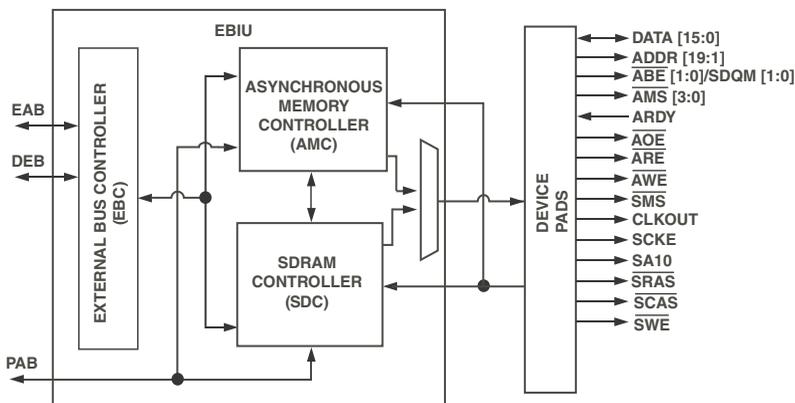


Figure 7-2. External Bus Interface Unit (EBIU)

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 is the EBIU. The EAB provides access to both asynchronous external memory and synchronous DRAM external memory. The external access is controlled by either the AMC or the 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, nor 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 on the bus selects. The AMC and SDC compete for access to the shared resources. 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.

Registers

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`)

EBIU Overview

- 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.

Shared and Multiplexed Pins

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:

- ADDR[19:1], address bus
- DATA[15:0], data bus
- $\overline{\text{ABE}}[1:0]/\text{SDQM}[1:0]$, AMC byte enables/SDC data masks
- CLKOUT, system clock for SDC and AMC

No other signals are multiplexed between the two controllers.

The following AMC signals are multiplexed. Refer to [Chapter 9, “General-Purpose Ports”](#) for the locations of these signals and information on configuring them.

- $\overline{\text{AMS}}[3:2]$ – Asynchronous memory bank selects
- $\overline{\text{AOE}}$ – Asynchronous memory output enable
- ARDY – Asynchronous memory ready response

System Clock

The `CLKOUT` pin is shared by both the SDC and AMC. Two different registers are used to control this:

- `EBIU_SDGCTL` register, `SCTLE` bit for SDC clock
- `EBIU_AMGCTL` register, `AMCKEN` bit for AMC clock

If enabling or disabling the system clock, software control for both registers is required.

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 a disabled external memory bank
- Any access to reserved SDRAM memory space
- Any access to unpopulated SDRAM space

If the core requested the faulting bus operation, the bus error response from the EBIU is gated into the hardware error interrupt (`IVHW`) 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.

AMC Overview and Features

The following sections describe the features of the AMC.

Features

The EBIU AMC features include:

- 16-bit I/O width
- 1.8, 2.5 or 3.3 V I/O supply
- Supports up to 4M bytes of SRAM in four external banks
- AMC supports 8-bit data masking writes
- AMC has control of the EBIU while auto-refresh is performed to SDRAM
- AMC supports asynchronous access extension (ARDY pin)
- Supports instruction fetch
- Allows booting from bank 0 ($\overline{\text{AMS0}}$)

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 pin select associated with it, shown in [Table 7-1](#).

Table 7-1. Asynchronous Memory Bank Address Range

Memory Bank Select	Address Start	Address End
$\overline{\text{AMS}}[3]$	0x2030 0000	0x203F FFFF
$\overline{\text{AMS}}[2]$	0x2020 0000	0x202F FFFF
$\overline{\text{AMS}}[1]$	0x2010 0000	0x201F FFFF
$\overline{\text{AMS}}[0]$	0x2000 0000	0x200F FFFF

Asynchronous Memory Address Decode

The address range allocated to each asynchronous memory bank is fixed at 1M bytes; however, not all of an enabled memory bank need be populated.

 Accesses to unpopulated memory or 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 7-2](#). The timing of these pins is programmable to allow a flexible interface to devices of different speeds. For example interfaces, see [Chapter 27, “System Design”](#).

AMC Pin Description

The following table describes the signals associated with each interface.

Table 7-2. Asynchronous Memory Interface Signals

Pad	Pin Type ¹	Description
DATA[15:0]	I/O	External data bus
CLKOUT	O	Switches at system clock frequency. Connect to the peripheral if required.
ADDR[19:1]	O	External address bus
$\overline{\text{AMS}}[3:0]$	O	Asynchronous memory bank selects
$\overline{\text{AWE}}$	O	Asynchronous memory write enable
$\overline{\text{ARE}}$	O	Asynchronous memory read enable
$\overline{\text{AOE}}$	O	Asynchronous memory output enable In most cases, the $\overline{\text{AOE}}$ pin should be connected to the $\overline{\text{OE}}$ pin of an external memory-mapped asynchronous device. Please refer to the product data sheet for specific timing information between the $\overline{\text{AOE}}$ and $\overline{\text{ARE}}$ signals to determine which interface signal should be used in your system.

AMC Description of Operation

Table 7-2. Asynchronous Memory Interface Signals (Continued)

Pad	Pin Type ¹	Description
ARDY	I	Asynchronous memory ready response
$\overline{\text{ABE}}[1:0]/\text{SDQM}[1:0]$	O	Byte enables

¹ Pin Types: I = Input, O = Output

AMC Description of Operation

The following sections describe the operation of the AMC.

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.

External Access Extension

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 the access time as required.

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 External Access Extension” on page 7-15.](#)

AMC Functional Description

The following sections provide a functional description of the AMC.

Programmable Timing Characteristics

This section describes the programmable timing characteristics for the EBIU. Timing relationships depend on the programming of the AMC, no matter whether the transaction initiation is from the core or from memory DMA, or what the sequence of transactions is (read followed by read, read followed by write, and so on).

Asynchronous Reads

Figure 7-3 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.

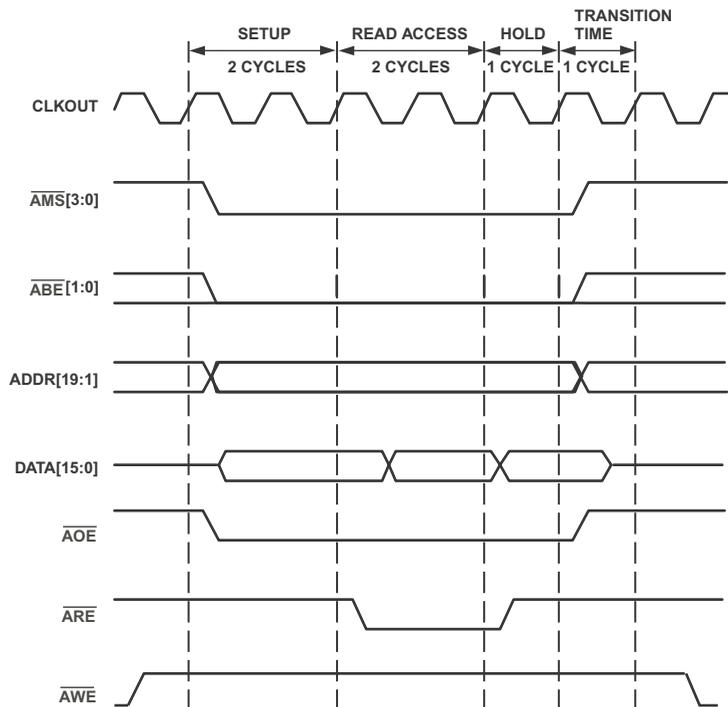


Figure 7-3. Asynchronous Read Bus Cycles

Asynchronous read bus cycles proceed as follows.

1. At the start of the setup period, $\overline{AMS[x]}$ and \overline{AOE} assert. The address bus becomes valid. The $\overline{ABE[1:0]}$ signals are low during the read.
2. At the beginning of the read access period and after the 2 setup cycles, \overline{ARE} asserts.

3. At the beginning of the hold period, read data is sampled on the rising edge of the EBIU clock. The \overline{ARE} pin deasserts after this rising edge.
4. At the end of the hold period, \overline{AOE} deasserts unless this bus cycle is followed by another asynchronous read to the same memory space. Also, $\overline{AMS[x]}$ 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.

Asynchronous Writes

Figure 7-4 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, $\overline{AMS[x]}$, the address bus, data buses, and $\overline{ABE[1:0]}$ become valid. See “Byte Enables” on page 7-15 for more information.
2. At the beginning of the write access period, \overline{AWE} asserts.
3. At the beginning of the hold period, \overline{AWE} deasserts.

Asynchronous read bus cycles proceed as follows.

1. At the start of the setup period, $\overline{AMS[x]}$ and \overline{AOE} assert. The address bus becomes valid. The $\overline{ABE[1:0]}$ signals are low during the read.
2. At the beginning of the read access period, \overline{ARE} asserts.

AMC Functional Description

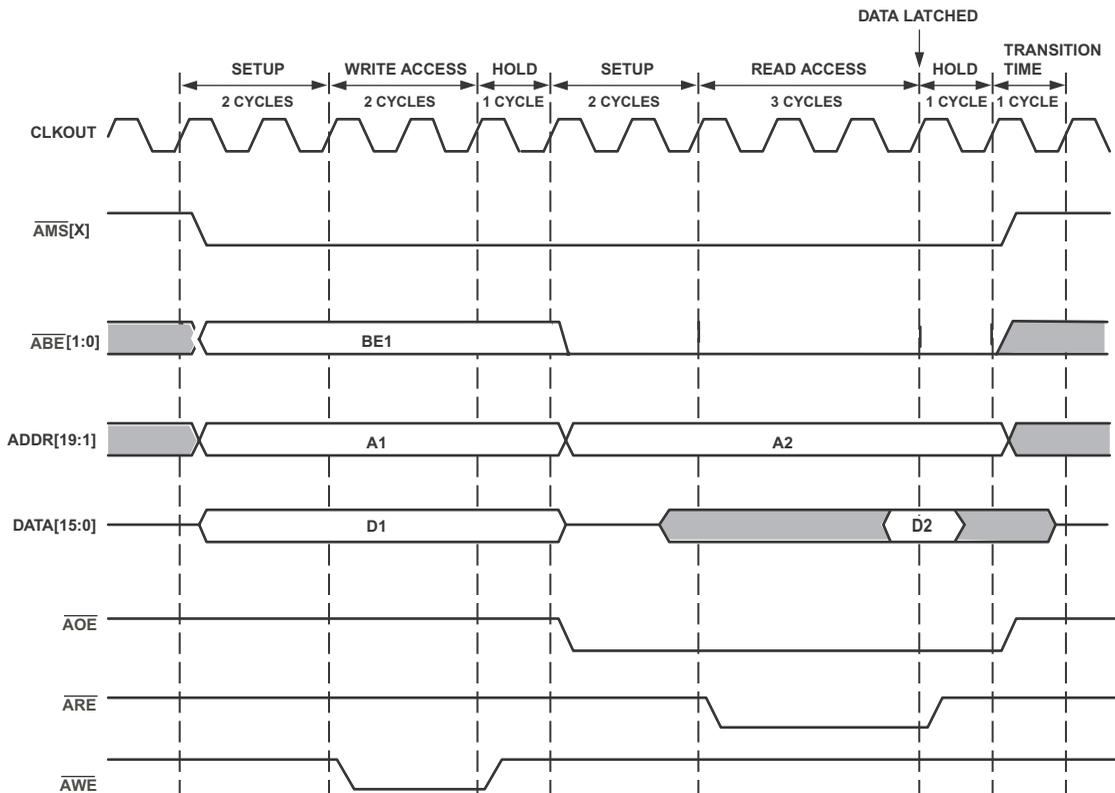


Figure 7-4. Asynchronous Write and Read Bus Cycles

3. At the beginning of the hold period, read data is sampled on the rising edge of the EBIU clock. The ARE signal deasserts after this rising edge.
4. At the end of the hold period, AOE deasserts unless this bus cycle is followed by another asynchronous read to the same memory space. Also, AMS[x] 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 External Access Extension

The ARDY pin is used to insert extra wait states. The EBIU starts sampling ARDY on the clock cycle before the end of the programmed strobe period. If ARDY is sampled as deasserted, the access period is extended. The ARDY pin is then sampled on each subsequent clock edge. Read data is latched on the clock edge after ARDY is sampled as asserted. The read- or write-enable remains asserted for one clock cycle after ARDY is sampled as asserted. An example of this behavior is shown in [Figure 7-5](#), where setup = 2 cycles, read access = 4 cycles, and hold = 1 cycle.



The read access period must be programmed to a minimum of two cycles to make use of the ARDY input.

Byte Enables

The AMC provides byte enable pins to allow the processor to perform efficient byte-wide arithmetic and byte-wide processing in external memory.

In general, there are two different ways to modify a single byte within the 16-bit interface. First, it can be done by a read/modify/write sequence. However, this is not very efficient because multiple accesses are required (that is, it takes many cycles for reads and writes to external memory). Another option is available where just a specific byte can be modified for a 16-bit devices using the ABE[1:0] pins. See [Table 7-3](#).

The ABE[1:0] pins are both low during all asynchronous reads and 16-bit asynchronous writes. When an asynchronous write is made to the upper byte of a 16-bit memory, ABE1 = 0 and ABE0 = 1. When an asynchronous write is made to the lower byte of a 16-bit memory, ABE1 = 1 and ABE0 = 0.

AMC Functional Description

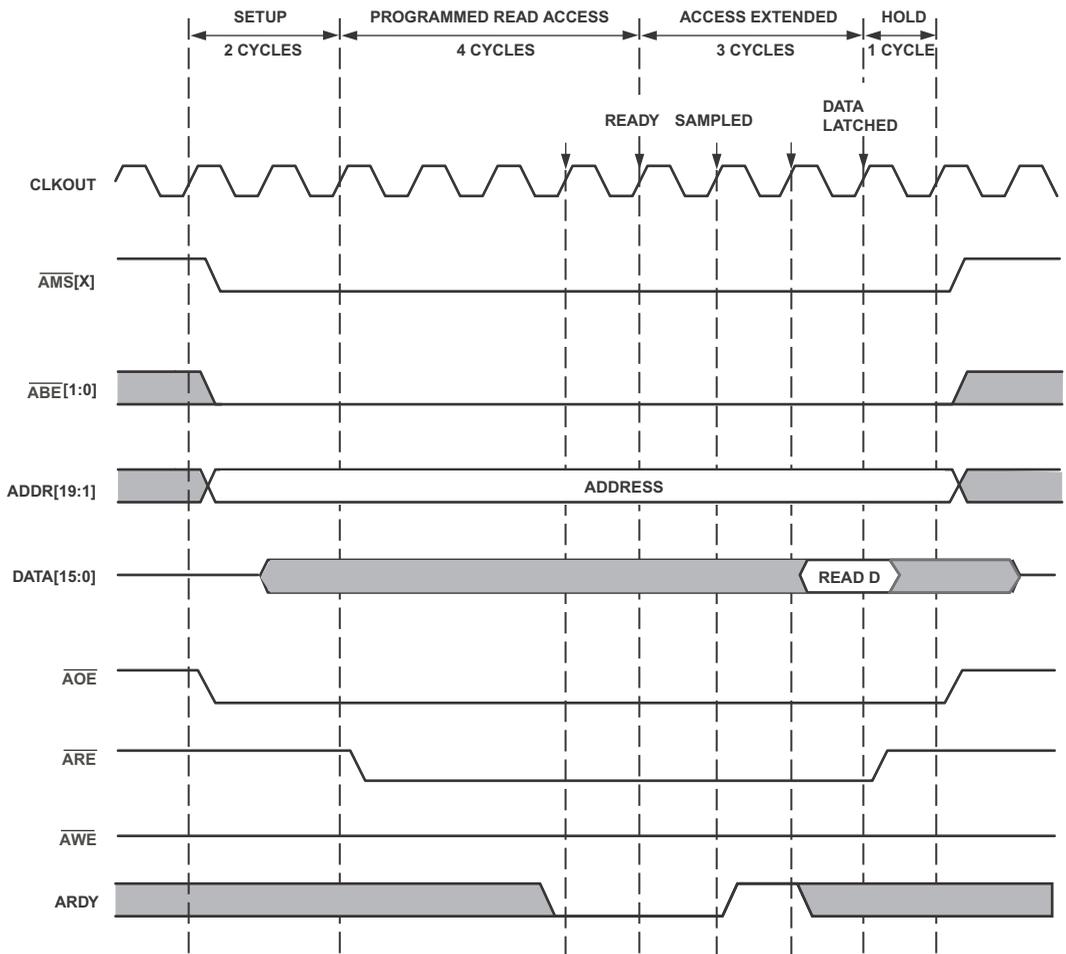


Figure 7-5. Inserting Wait States Using ARDY

Table 7-3. Byte Enables 8-Bit Write Accesses

Internal Address IA[0]	Internal Transfer Size	
	1 byte	2 bytes
0	ABE[1] = 1 ABE[0] = 0	ABE[1] = 0 ABE[0] = 0
1	ABE[1] = 0 ABE[0] = 1	ABE[1] = 0 ABE[0] = 0 This combination is invalid.

AMC Programming Model

The asynchronous memory global control register (EBIU_AMGCTL) 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.

Additional information for the EBIU_AMGCTL register bits includes:

- **Asynchronous memory clock enable** (AMCKEN)

For external devices that need a clock, CLKOUT can be enabled by setting the AMCKEN bit in the EBIU_AMGCTL register. In systems that do not use CLKOUT, set the AMCKEN bit to 0.

- **Asynchronous memory bank enable** (AMBEN).

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

AMC Programming Model

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.

- **Core/DMA priority** (CDPRIO).

This bit configures the AMC to control the priority over requests that occur simultaneously to the EBIU from either processor core or the DMA controller. When this bit is set to 0, a request from the core has 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 controller, including the memory DMAs, have priority over core accesses. For the purposes of this discussion, core accesses include both data fetches and instruction fetches.



The CDPRIO bit also applies to the SDC.

The EBIU asynchronous memory controller has two asynchronous memory bank control registers (EBIU_AMBCTL0 and EBIU_AMBCTL1). They contain bits for counters for setup, access, 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 (AMS[x] low) and the read-enable assertion (ARE low) or write-enable assertion (AWE low).
- **Read access:** the time between read-enable assertion (ARE low) and deassertion (ARE high).

- Write access: the time between write-enable assertion (AWE low) and deassertion (AWE high).
- Hold: the time between read-enable deassertion (ARE high) or write-enable deassertion (AWE high) and the end of the memory cycle ($AMS[x]$ 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

AMC Registers

The following sections describe the AMC registers.

AMC Registers

EBIU_AMGCTL Register

Figure 7-6 shows the asynchronous memory global control register (EBIU_AMGCTL).

Asynchronous Memory Global Control Register (EBIU_AMGCTL)

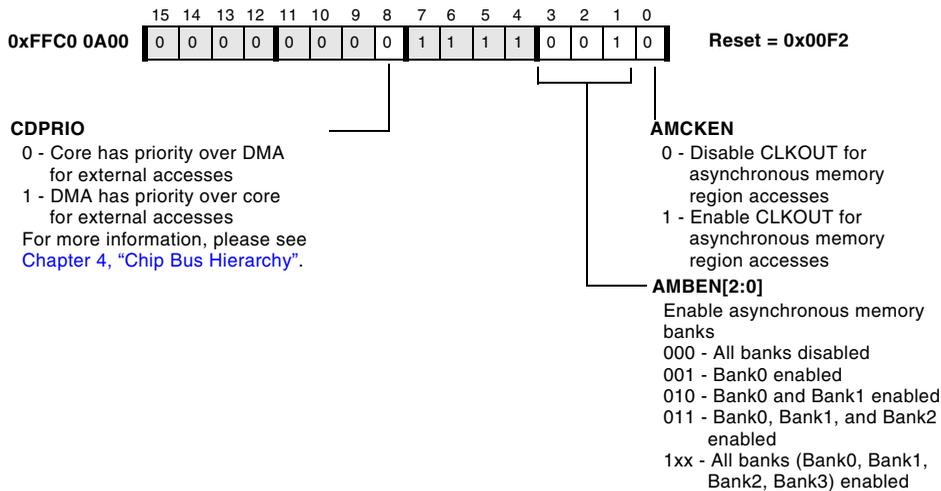


Figure 7-6. Asynchronous Memory Global Control Register

EBIU_AMBCTL0 and EBIU_AMBCTL1 Registers

Figure 7-7 and Figure 7-8 show the asynchronous memory bank control registers (EBIU_AMBCTL0 and EBIU_AMBCTL1).

Asynchronous Memory Bank Control 0 Register (EBIU_AMBCTL0)

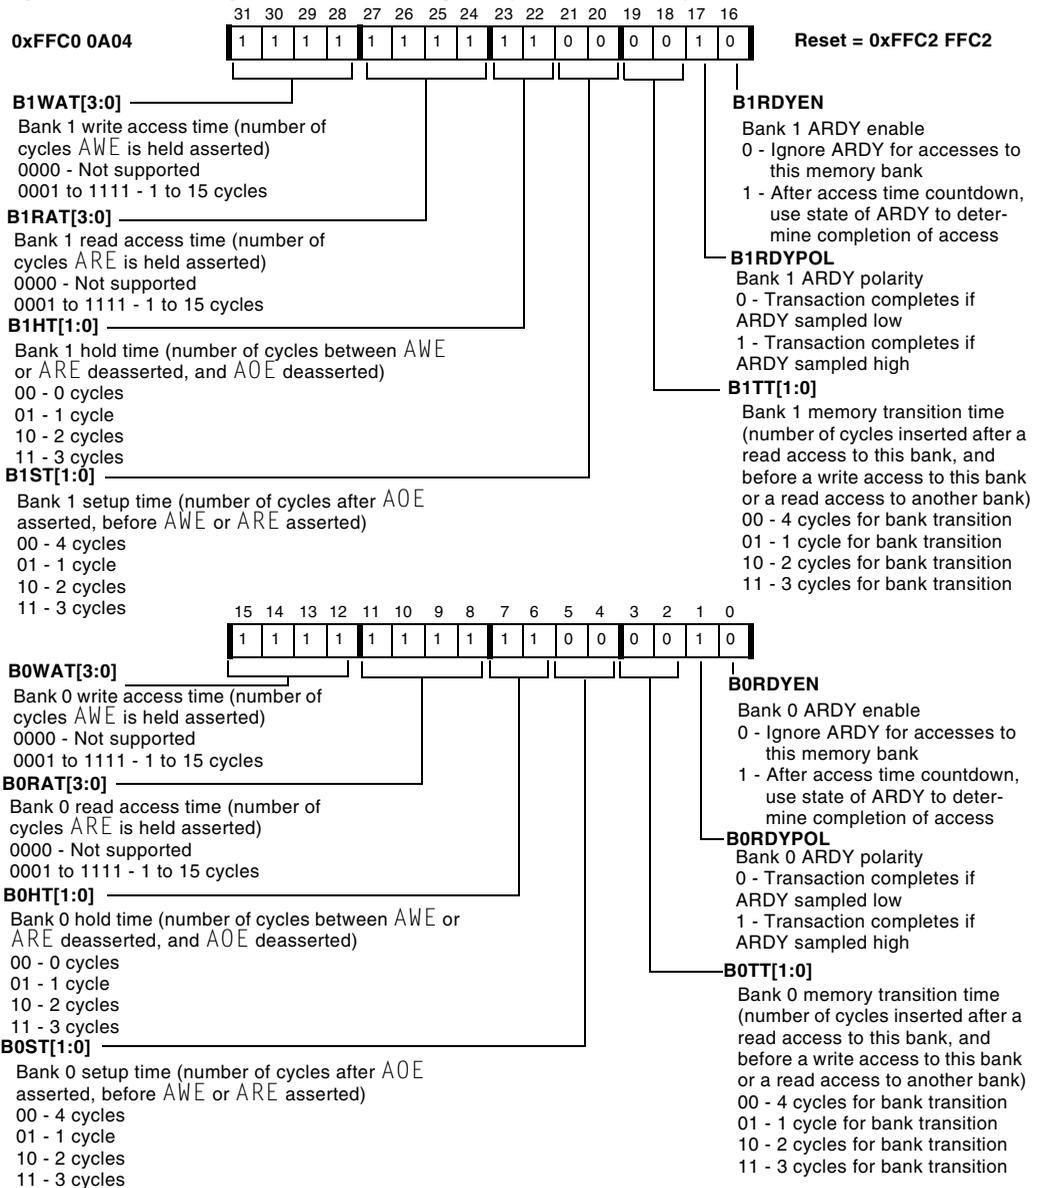


Figure 7-7. Asynchronous Memory Bank Control 0 Register

AMC Registers

Asynchronous Memory Bank Control 1 Register (EBIU_AMBCTL1)

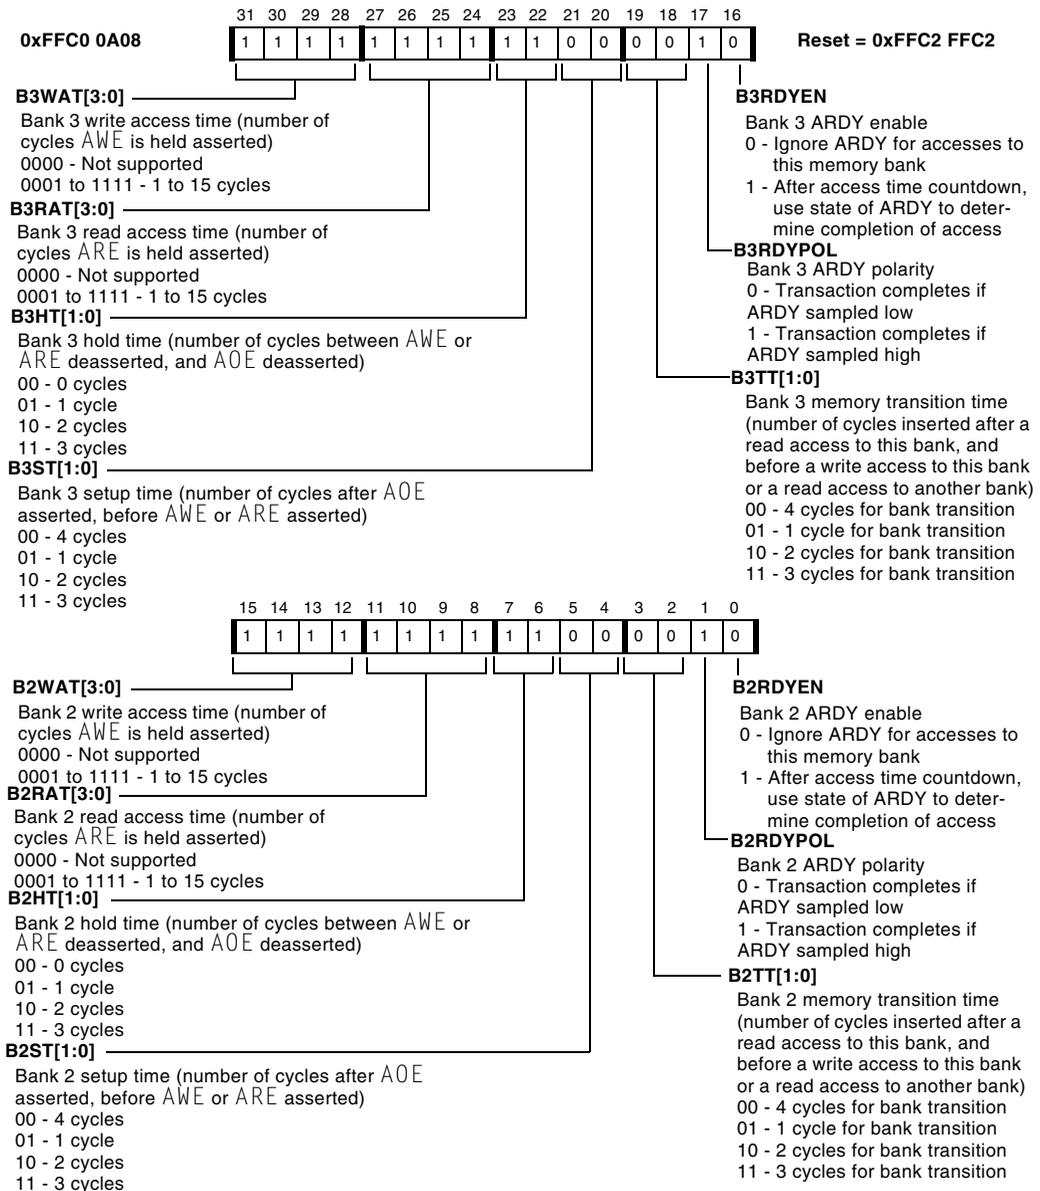


Figure 7-8. Asynchronous Memory Bank Control 1 Register

AMC Programming Examples

[Listing 7-1](#) and [Listing 7-2](#) provide examples for working with the AMC.

Listing 7-1. 16-Bit Core Transfers to SRAM

```
.section L1_data_b;
.byte2 source[N] = 0x1122, 0x3344, 0x5566, 0x7788;
.section SRAM_bank_0;
.byte2 dest[N];
.section L1_code;
I0.L = lo(source);
I0.H = hi(source);
I1.L = lo(dest);
I1.H = hi(dest);
    R0.L = w[I0++];
    P5=N-1;
    lsetup(lp, lp) LC0=P5;
lp:  R0.L = w[I0++] || w[I1++] = R0.L;
                                     w[I1++] = R0.L;
```

Listing 7-2. 8-Bit Core Transfers to SRAM Using Byte Mask ABE[1:0] Pins

```
.section L1_data_b;
.byte source[N] = 0x11, 0x22, 0x33, 0x44, 0x55, 0x66, 0x77, 0x88;
.section SRAM_bank_0;
.byte dest[N];
p0.L = lo(source);
p0.H = hi(source);
p1.L = lo(dest);
p1.H = hi(dest);
    p5=N;
    lsetup(start, end) LC0=P5;
```

SDC Overview and Features

```
start: R0 = b[p0++](z);
end:   b[p1++] = R0; /* byte data masking */
```

SDC Overview and Features

The SDRAM Controller (SDC) enables the processor to transfer data to and from Synchronous DRAM (SDRAM) with a maximum frequency specified in the product data sheet. The processor supports a glueless interface with one external bank of standard SDRAMs of 64M bit to 512M bit, with configurations x4, x8, and x16, up to a maximum total capacity of 128M bytes of SDRAM.

Features

The EBIU SDC provides a glueless interface with standard SDRAMs. Features include:

- I/O width 16-bit, I/O supply 1.8, 2.5, or 3.3 V
- Supports up to 128M byte of SDRAM in external bank
- Types of 64, 128, 256, and 512M bit with I/O of x4, x8, and x16
- Supports SDRAM page sizes of 512 byte, 1K, 2K, and 4K byte
- Supports multibank operation within the SDRAM
- Supports mobile SDRAMs
- SDC uses no-burst mode ($BL = 1$) with sequential burst type
- SDC supports 8-bit data masking writes
- SDC uses open page policy—any open page is closed only if a new access in another page of the same bank occurs
- 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
- Allows independent auto-refresh while the asynchronous memory controller has control of the EBIU port
- Supports self-refresh mode for power savings
- During hibernate state, self-refresh mode is supported
- Supports instruction fetch

SDRAM Configurations Supported

Table 7-4 shows all possible bank sizes, and SDRAM discrete component configurations that can be gluelessly interfaced to the SDC. The bank width for all cases is 16 bits.

Table 7-4. SDRAM Discrete Component Configurations Supported

System Size (M byte)	System Size (M bit)	SDRAM Configuration	Number of Chips
16	8M x 16	8M x 8	2
16	8M x 16	8M x 16	1
32	16M x 16	16M x 4	4
32	16M x 16	16M x 8	2
32	16M x 16	16M x 16	1
64	32M x 16	32M x 4	4
64	32M x 16	32M x 8	2
64	32M x 16	32M x 16	1
128	64M x 16	64M x 4	4
128	64M x 16	64M x 8	2
128	64M x 16	64M x 16	1

SDRAM External Bank Size

The total amount of external SDRAM memory addressed by the processor is controlled by the `EBSZ` bits of the `EBIU_SDBCTL` register (see [Table 7-5](#)). Accesses above the range shown for a specialized `EBSZ` value results in an internal bus error and the access does not occur. [For more information, see “Error Detection” on page 7-7.](#)

SDC Address Mapping

The address mapping scheme describes how the SDC maps the address into SDRAM. To access SDRAM, the SDC uses the bank interleaving map scheme, which fills each internal SDRAM bank before switching to the next internal bank. Since the SDRAMs have four internal banks, the entire SDRAM address space is therefore divided into four sub-address regions containing the addresses of each internal bank. (See [Figure 7-10 on page 7-40](#).) It starts with address `0x0` for internal bank A and ends with the last valid address (specified with `EBSZ` and `EBCAW` parameters) containing the internal bank D.

The internal 29-bit non-multiplexed address (See [Figure 7-9](#)) is multiplexed into:

- Byte data mask (`IA[0]`)
- SDRAM column address
- SDRAM row address
- Internal SDRAM bank address

 A good understanding of the SDC address map scheme in conjunction with the multibank operation is required to obtain optimized system performance.

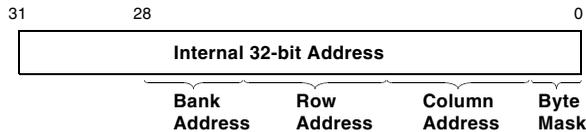


Figure 7-9. Multiplexed SDRAM Addressing Scheme

Table 7-5. External Bank Size Encodings

EBSZ	Bank Size (M byte)	Valid SDRAM Addresses
b#00	16	0x0000 0000 – 0x00FF FFFF
b#01	32	0x0000 0000 – 0x01FF FFFF
b#10	64	0x0000 0000 – 0x03FF FFFF
b#11	128	0x0000 0000 – 0x07FF FFFF

Internal SDRAM Bank Select

The internal SDRAM banks are driven by the ADSP-BF51x ADDR[19:18] which are part of the row and column address and connected to the SDRAM BA[1:0].

 Do not flip up both internal bank select connections, if using the mobile SDRAM PASR feature. If this is done, the system will not work properly because the selected internal banks are not refreshed during partial array self-refresh.

Parallel Connection of SDRAMs

To specify an SDRAM system, multiple possibilities are given based on the different architectures. (See [Table 7-13 on page 7-62.](#)) For the

SDC Interface Overview

ADSP-BF51x processors, I/O capabilities of 1 x 16-bit, 2 x 8-bit or 4 x 4-bit are given. The reason to use a system of 4 x 4-bit vs. 2 x 8-bit or 1 x 16-bit is determined by the SDRAM page size. All 3 systems have the same external bank size, but different page sizes. On one hand, the higher the page size, the higher the performance. On the other hand, the higher the page size, the higher the hardware layout requirements.

 Even if connecting SDRAMs in parallel, the SDC always considers the entire system as one external SDRAM bank (SMS pin) because all address and control lines feed the parallel parts.

However, access to a single cluster part is achieved using the mask feature (SDQM[1:0] pins). This allows masked 8-bit I/O writes to dedicated chips whereby the other 8-bit I/O is masked at its input buffer of the other chips. See [Listing 7-4 on page 7-76](#).

SDC Interface Overview

The following sections describe the SDC interface.

SDC Pin Description

The SDRAM interface signals are shown in [Table 7-6](#).

Table 7-6. SDRAM Interface Signals

Pad	Pin Type ¹	Description
DATA[15:0]	I/O	External data bus
ADDR[19:18], ADDR[16:12], ADDR[10:1]	O	External address bus Connect to SDRAM address pins. Bank address is output on ADDR[19:18] and should be connected to SDRAM BA[1:0] pins.
SRAS	O	SDRAM row address strobe pin Connect to SDRAM's RAS pin.

Table 7-6. SDRAM Interface Signals (Continued)

Pad	Pin Type ¹	Description
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.
ABE[1:0]/ SDQM[1:0]	O	SDRAM data mask pins Connect to SDRAM's DQM pins.
SMS	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 A[10] pin.
SCKE	O	SDRAM clock enable pin Connect to SDRAM's CKE pin.
CLKOUT	O	SDRAM clock output pin Switches at system clock frequency. Connect to the SDRAM's CLK pin.

¹ Pin Types: I = Input, O = Output

SDRAM Performance

On-page sequential or non-sequential accesses are from internal data memory to SDRAM. [Table 7-7](#) summarizes SDRAM performance for these on-page accesses.

Table 7-7. SDRAM Performance Between Internal Data Memory and SDRAM¹

Type of access	Performance
DAG access, write	1 SCLK cycle per 16-bit word
DAG access, read	8 SCLK cycles per 16-bit word

SDC Description of Operation

Table 7-7. SDRAM Performance Between Internal Data Memory and SDRAM¹ (Continued)

Type of access	Performance
MemDMA access, write	1 SCLK cycle per 16-bit word
MemDMA access, read	^a 1.1 SCLK cycles per 16-bit word

1 Valid for core/system clock > 2:1

On-page sequential instruction fetches from SDRAM are summarized in [Table 7-8](#).

Table 7-8. SDRAM Performance For On-Page Instruction Fetches

Type of access	Performance
Ifetch from SDRAM	^a 1.1 SCLK cycles per 16-bit word
I/Dcache line fill from SDRAM	^a 1.1 SCLK cycles per 16-bit word

Off-page accesses are summarized in [Table 7-9](#).

Table 7-9. SDRAM Stall Cycles For Off-Page Accesses

Type of access	Stall Cycles
Write	$t_{WR} + t_{RP} + t_{RCD}$
Read	$t_{RP} + t_{RCD} + CL$

SDC Description of Operation

The following sections describe the operation of the SDC.

Definition of SDRAM Architecture Terms

The following are definitions of SDRAM architecture terms used in the remainder of this chapter.

Refresh

Since the information is stored in a low-capacitance cell that suffers from leakage effects, the SDRAM must be refreshed periodically.

Row Activation

SDRAM accesses are multiplexed, which means any first access will open a row/page before the column access is performed. It stores the row in a “row cache” called row activation.

Column Read/Write

The row’s columns represent a page, which can be accessed with successive read or write commands without needing to activate another row. This is called column access and performs transfers from the “row cache.”

Row Precharge

If the next access is in a different row, the current row is closed before another is opened. The current “row cache” is written back to the row. This is called row precharge.

Internal Bank

There are up to 4 internal memory banks on a given SDRAM. Each of these banks can be accessed with the bank select lines $BA[1:0]$. The bank address can be thought of as part of the row address.

SDC Description of Operation

External Bank

This is the address region where the SDC address the SDRAM.

 Do not confuse the internal banks, which are internal to the SDRAM and are selected with the BA[1:0] pins with the external bank that is enabled by the CS pin.

Memory Size

Since the 2-D memory is based on rows and columns, the size is:
mem size = (# rows) x (# columns) x (# internal banks) x I/O (Mbit)

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 followed by a NOP (no operation) command, respectively (Number of NOPs = burst length - 1). Burst lengths of full page, 8, 4, 2, and 1 (no burst) are available. The burst length is selected by writing the BL bits in the SDRAM mode register during the SDRAM powerup sequence.

Burst Type

The burst type determines the address order in which the SDRAM delivers burst data. The burst type is selected by writing the BT bits in the SDRAM mode register during the SDRAM powerup sequence.

CAS Latency

The CAS latency, or read latency, specifies the time between latching a read address and driving the data off chip. This spec is normalized to the system clock and varies from 2 to 3 cycles based on the speed. The CAS latency is selected by writing the CL bits in the SDRAM mode register during the SDRAM powerup sequence.

Data I/O Mask Function

SDRAMs allow a data byte-masking capability on writes. The `DQM[1:0]` mask pins are used to block the data input buffer of the SDRAM during write operations.

SDRAM Commands

SDRAM commands are not based on typical read or write strobes. The pulsed `CS`, `RAS`, `CAS`, and `WE` lines determine the command on the rising clock edge by a truth table.

Mode Register Set (MRS) Command

SDRAM devices contain an internal extended configuration register which allows specification of the mobile SDRAM device's functionality.

Extended Mode Register Set (EMRS) Command

Mobile SDRAM devices contain an internal extended configuration register which allows specification of the mobile SDRAM device's functionality.

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, it 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.

Read/Write Command

For the read command, the SDRAM latches the column address. The start address is set according to the column address. For the write command,

SDC Description of Operation

SDRAM latches the column address. Data is also asserted in the same cycle. The start address is set according to the column address.

Precharge/Precharge All Command

The precharge command closes a specific active page in an internal bank and the precharge all command closes all 4 active pages in all 4 banks.

Auto-Refresh command

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 SDRAM to generate an internal auto-refresh cycle. When the internal refresh completes, all four internal SDRAM banks are precharged.

Enter Self-Refresh Mode

When the SDRAM enters self-refresh mode, the SDRAM's internal timer initiates refresh cycles periodically, without external control input.

Exit Self-Refresh Mode

When the SDRAM exits self-refresh mode, the SDRAM's internal timer stops refresh cycles and relinquishes control to external SDC.

SDC Timing Specs

The following SDRAM timing specs are used by the SDC and SDRAM. To program the SDRAM interface, see the SDRAM specific datasheet information



Any absolute timing parameter must be normalized to the system clock, which allows the SDC to adapt to the timing parameter of the device.

t_{MRD}

This is the required delay between issuing a mode register set and an activate command during powerup.

Dependency: system clock frequency

SDC setting: 3 system clock cycles

SDC usage: MRS command

t_{RAS}

This is the required delay between issuing a bank A activate command and issuing a bank A precharge command.

Dependency: system clock frequency

SDC setting: 1–15 normalized system clock cycles

SDC usage: single column read/write, auto-refresh, self-refresh command

SDC dependencies: t_{RC}, t_{RFC}, and t_{XSR}

CL

The CAS latency, or read latency, is the delay between when the SDRAM detects the read command and when it provides the data off-chip. This spec does not apply to writes.

Dependency: system clock frequency and speed grade

SDC setting: 2–3 normalized system clock cycles

SDC usage: first read command

SDC Description of Operation

t_{RCD}

This is the required delay between a bank A activate command and the first bank A read or write command.

Dependency: system clock frequency

SDC setting: 1–7 normalized system clock cycles

SDC usage: first read/write command

t_{RRD}

This is the required delay between a bank A activate command and a bank B activate command. This spec is used for multibank operation.

Dependency: system clock frequency

SDC setting: $t_{RCD} + 1$ normalized system clock cycles

SDC usage: multiple bank activation

t_{WR}

This is the required delay between a bank A write command and a bank A precharge command. This spec does not apply to reads.

Dependency: system clock frequency

SDC setting: 1–3 normalized system clock cycles

SDC usage: during off-page write command

t_{RP}

This is the required delay between a bank A precharge command and a bank A activation command.

Dependency: system clock frequency

SDC setting: 1–7 normalized system clock cycles

SDC usage: off-page read/write, auto-refresh, self-refresh command

SDC dependencies: t_{RC} , t_{RFC} , and t_{XSR}

 t_{RC}

This is the required delay between issuing successive bank activate commands.

Dependency: system clock frequency

SDC setting: User must ensure that $t_{RAS} + t_{RP} \geq t_{RC}$ (normalized system clock cycles)

SDC usage: single column read/write command

 t_{RFC}

This is the required delay between issuing successive auto-refresh commands (all banks).

Dependency: system clock frequency

SDC setting: User must ensure that $t_{RAS} + t_{RP} \geq t_{RFC}$ (normalized system clock cycles)

SDC usage: auto-refresh, exit self-refresh command

SDC Description of Operation

t_{XSR}

This is the required delay between exiting self-refresh mode and the auto-refresh command.

Dependency: system clock frequency

SDC setting: User must ensure that $t_{RAS} + t_{RP} \geq t_{XSR}$ (normalized system clock cycles)

SDC usage: exit self-refresh command

t_{REF}

This is the row refresh period, and typically takes 64 ms.

Dependency: system clock frequency

SDC setting: none

SDC usage: auto-refresh command

t_{REFI}

This is the row refresh interval and typically takes 15.6 μ s for < 8k rows and 7.8 μ s for \geq 8k rows. This spec is available by dividing t_{REF} by the number of rows. This number is used by the SDC refresh counter.

Dependency: system clock frequency

SDC setting: t_{REFI} normalized system clock cycles (RDIV register)

SDC usage: auto-refresh command

 In typical applications making sequential (not random) accesses to the SDRAM memory, the t_{RAS} timing parameter is less critical than t_{RP} . System designers should be aware that whenever the $(t_{RP} + t_{RAS})$ in their design is violating one of the other timing specifications, then they should increase the t_{RAS} parameter.

SDC Functional Description

The functional description of the SDC is provided in the following sections.

SDC Operation

The AMC normally generates an external memory address, which then asserts the corresponding CS select, along with RD and WR strobes. However these control signals are not used by the SDC. The internal strobes are used to generate pulsed commands (SMS, SCKE, SRAS, SCAS, SWE) within a truth table (see [Table 7-11 on page 7-46](#)). The memory access to SDRAM is based by mapping ADDR[28:0] causing an internal memory select to SDRAM space (see [Figure 7-10](#)).

The configuration is programmed in the SDBCTL register. The SDRAM controller can hold off the processor core or DMA controller with an internally connected acknowledge signal, as controlled by refresh, or page miss latency overhead.

A programmable refresh counter is provided which generates 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.

The SDC uses no burst mode ($BL = 1$) for read and write operations. This requires the SDC to post every read or write address on the bus as for non-sequential reads or writes, but does not cause any performance degradation. For 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 any single read transfer. Subsequent reads do not have latency.

Whenever a page miss to the same bank 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.

SDC Address Muxing

Table 7-10 shows the connection of the address pins with the SDRAM device pins.

Table 7-10. SDRAM Address Connections for 16-Bit Banks

External Address Pin	SDRAM Address Pin
ADDR[19]	BA[1]
ADDR[18]	BA[0]
ADDR[16]	A[15]
ADDR[15]	A[14]
ADDR[14]	A[13]
ADDR[13]	A[12]
ADDR[12]	A[11]
ADDR[11]	Not used
SA[10]	A[10]
ADDR[10]	A[9]
ADDR[9]	A[8]

SDC Functional Description

Table 7-10. SDRAM Address Connections for 16-Bit Banks (Continued)

External Address Pin	SDRAM Address Pin
ADDR[8]	A[7]
ADDR[7]	A[6]
ADDR[6]	A[5]
ADDR[5]	A[4]
ADDR[4]	A[3]
ADDR[3]	A[2]
ADDR[2]	A[1]
ADDR[1]	A[0]

Multibank Operation

Since an SDRAM contains 4 independent internal banks (A-D), the SDC is capable of supporting multibank operation thus taking advantage of the architecture.

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 other banks (B, C, D) the SDC leaves bank (A) open and activates any of the other banks (B, C, D). Bank (A) to bank (B) active time is controlled by $t_{RRD} = t_{RCD} + 1$. This scenario is repeated until all 4 banks (A-D) are opened and results in an effective page size up to 4 pages because no latency causes switching between these open pages (compared to 1 page in only one bank at the time). Any access to any closed page in any opened bank (A-D) forces a precharge command only to that bank. If, for example, 2 MemDMA channels are pointing to the same internal SDRAM bank, this always forces precharge and activation cycles to switch between the different pages. However, if the 2 MemDMA channels are pointing to different internal SDRAM banks, it does not cause additional overhead. See [Figure 7-11](#).

- i** The benefit of multibank operation reduces precharge and activation cycles by mapping opcode/data among different internal SDRAM banks driven by the A[19:18] pins.

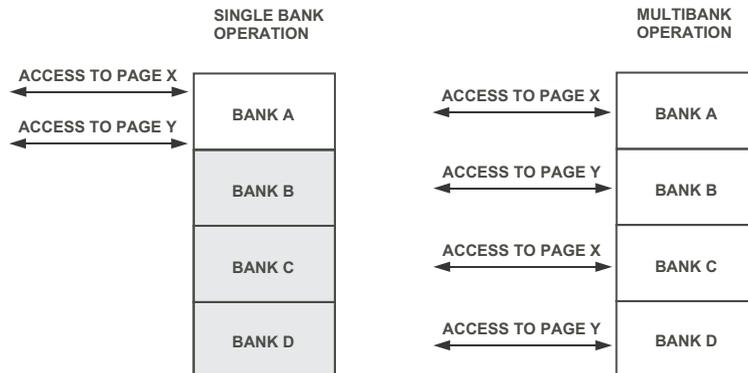


Figure 7-11. SDRAM Bank Operation Types

Core and DMA Arbitration

The `CDPRI0` bit configures the SDC to control the priority over requests that occur simultaneously to the EBIU from either the processor core or the DMA controller. When this bit is set to 0, a request from the core has priority over a request from the DMA controller to the SDC, unless the DMA is urgent. When it is set to 1, all requests from the DMA controller, including the memory DMAs, have priority over core accesses. For the purposes of this discussion, core accesses include both data fetches and instruction fetches.

Changing System Clock During Runtime

All timing specs are normalized to the system clock. Since most of them are minimum specs, except t_{REF} , which is a maximum spec, a variation of system clock will on one hand violate a specific spec and on the other hand cause a performance degradation for the other specs.

SDC Functional Description

The reduction of system clock will violate the minimum specs, while increasing system clock will violate the maximum t_{REF} spec. Therefore, careful software control is required to adapt these changes.

-  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.

The recommended procedure for changing the PLL VCO frequency is:

1. Issue an SSYNC instruction to ensure all pending memory operations have completed.
2. Set the SDRAM to self-refresh mode by writing a 1 to the SRFS bit of EBIU_SDGCTL.
3. Execute the desired PLL programming sequence. (For details, refer to [Chapter 8, “Dynamic Power Management”](#).)
4. After the wakeup occurs that signifies the PLL has settled to the new VCO frequency, reprogram the SDRAM registers (EBIU_SDRRC, EBIU_SDGCTL) with values appropriate to the new SCLK frequency, and assure that the PSSE bit is set.
5. Bring the SDRAM out of self-refresh mode by clearing the SRFS bit of EBIU_SDGCTL.

Changing the SCLK frequency using the SSEL bits in PLL_DIV, as opposed to actually changing the VCO frequency, should be done using these steps:

1. Issue an SSYNC instruction to ensure all pending memory operations have completed.
2. Set the SDRAM to self-refresh mode by writing a 1 to the SRFS bit of EBIU_SDGCTL.
3. Execute the desired write to the SSEL bits.

4. Reprogram the SDRAM registers with values appropriate to the new `SCLK` frequency, and assure that the `PSSE` bit is set.
5. Bring the SDRAM out of self-refresh mode by clearing the `SRFS` bit of `EBIU_SDGCTL`.

Changing Power Management During Runtime

Deep Sleep Mode

During deep sleep mode, the core and system clock will halt. Therefore, careful software control is required to place the SDRAM in self-refresh before the device enters deep sleep mode.

Hibernate State

In the hibernate state the core voltage is 0 (core reset), but the I/O voltage can still be applied. In order to save the SDRAM volatile data, the ADSP-BF51x processor supports driving the `SCKE` signal low during core reset. Setting the `SCKELOW` bit of `VR_CTL` keeps the `SCKE` signal low. This ensures that the self-refresh mode is not exited during the reset sequence initiated by a hibernate wake-up event. Normally, the `SCKE` pin is toggled high during reset to comply with PC-133 specifications. For details about the `SCKELOW` bit, refer to [Chapter 8, “Dynamic Power Management”](#).

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.

- MODE REGISTER SET
- EXTENDED MODE REGISTER SET

SDC Functional Description

- BANK ACTIVATION
- READ and WRITE
- SINGLE PRECHARGE
- PRECHARGE ALL
- AUTO-REFRESH
- SELF-REFRESH ENTRY and SELF-REFRESH EXIT
- NOP

Table 7-11 shows the SDRAM pin state during SDC commands.

Table 7-11. Pin State During SDC Commands

Command	SCKE (n - 1)	SCKE (n)	SMS	SRAS	SCAS	SWE	SA10	Addresses
(E)/Mode register set	High	High	Low	Low	Low	Low	Op-code	Op-code
Activate	High	High	Low	Low	High	High	Valid address bit	Valid
Read	High	High	Low	High	Low	High	Low (CMD)	Valid
Single precharge	High	High	Low	Low	High	Low	Low	Valid
Precharge all	High	High	Low	Low	High	Low	High	Don't care
Write	High	High	Low	High	Low	Low	Low (CMD)	Valid
Auto-refresh	High	High	Low	Low	Low	High	Don't care	Don't care
Self-refresh entry	High	Low	Low	Low	Low	High	Don't care	Don't care
Self-refresh	Low	Low	Don't care	Don't care				
Self-refresh exit	Low	High	High	Don't care	Don't care	Don't care	Don't care	Don't care

Table 7-11. Pin State During SDC Commands (Continued)

Command	SCKE (n - 1)	SCKE (n)	SMS	SRAS	SCAS	SWE	SA10	Addresses
NOP	High	High	Low	High	High	High	Don't care	Don't care
Inhibit	High	High	High	Don't care				

Mode Register Set Command

The `MODE REGISTER SET` (MRS) command initializes SDRAM operation parameters. This command is a part of the SDRAM power-up sequence. The MRS command uses the address bus of the SDRAM as data input. The power-up sequence is initiated by setting 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 power-up sequence. The exact order of the power-up sequence is determined by the `PSM` bit of the `EBIU_SDGCTL` register.

The MRS command initializes these parameters:

- Burst length = 1, bits `A[2-0]`, always 0
- Burst type = sequential, bit `A[3]`, always 0
- CAS latency, bits `A[6-4]`, programmable in the `EBIU_SDGCTL` register
- Bits `A[12-7]`, always 0

After power-up and before executing a read or write to the SDRAM memory space, the application must trigger the SDC to write the SDRAM mode register. The write of the SDRAM mode register is triggered by setting 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 power-up sequence to be run, which programs the SDRAM mode register with burst length, burst type, and CAS latency from the `EBIU_SDGCTL` register and

SDC Functional Description

optionally the content to the extended mode register. This initial read or write to SDRAM takes many cycles to complete.

While executing an MRS command, the unused address pins are cleared. During the two clock cycles following the MRS command (t_{MRD}), the SDC issues only NOP commands.

Extended Mode Register Set Command (Mobile SDRAM)

The extended mode register is a subset of the mode register. The EBIU enables programming of the extended mode register during power-up via the EMREN bit in the EBIU_SDGCTL register.

The extended mode register is initialized with these parameters:

- Partial array self-refresh, bits A[2-0], bit A[2] always 0, bits A[1-0] programmable in EBIU_SDGCTL
- Temperature compensated self-refresh, bits A[4-3], bit A[3] always 1, bit A[4] programmable in EBIU_SDGCTL
- Drive strength control, bits A[6-5], always 0
- Bits A[12-7], always 0, and bit A[13] always 1

 Not programming the extended mode register upon initialization results in default settings for the low-power features. The extended mode defaults with the temperature sensor enabled, full drive strength, and full array refresh.

Bank Activation Command

The BANK_ACTIVATION command is required for first access to any internal bank in SDRAM. Any subsequent access to the same internal bank but different row will be preceded by a precharge and activation command to that bank.

However, if an access to another bank occurs, the SDC leaves the current page open and issues a `BANK ACTIVATION` command before executing the read or write command to that bank. With this method, called multibank operation, one page per bank can be open at a time, which results in a maximum of four pages.

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.

Partial Write

In general, there are two different ways to modify a single byte within the 16-bit interface. First, it can be done by a read/modify/write sequence. However, this is not very efficient because multiple accesses are required.

During partial writes to SDRAM, the `SDQM[1:0]` pins are used to mask writes to bytes that are not accessed. [Table 7-12](#) shows the `SDQM[1:0]` encodings based on the internal transfer address bit `IA[0]` and the transfer size.

However, during read transfers to SDRAM banks, reads are always done of all bytes in the bank regardless of the transfer size. This means for 16-bit SDRAM banks, `SDQM[1:0]` are all zeros (0s).

SDC Functional Description

-  The SDC provides byte enable pins SDQM[1:0] to allow the processor to perform efficient byte-wide arithmetic and byte-wide processing in external memory.

Table 7-12. SDQM[1:0] Encodings During Writes

Internal Address IA[0]	Internal Transfer Size	
	1 byte	2 bytes
0	SDQM[1] = 1 SDQM[0] = 0	SDQM[1] = 0 SDQM[0] = 0
1	SDQM[1] = 0 SDQM[0] = 1	SDQM[1] = 0 SDQM[0] = 0

-  For 16-bit SDRAMs, connect SDQM[0] to DQML, and connect SDQM[1] to DQMH.

Single Precharge Command

For a page miss during reads or writes in a specific internal SDRAM bank, the SDC uses the SINGLE PRECHARGE command to that bank.

-  The SDC does not use the auto-precharge read or write command of SDRAMs, which is enabled by asserting SA10 high during a read or write command.

Precharge All Command

The PRECHARGE ALL command is used to precharge all internal banks at the same time before executing an auto-refresh. All open banks will be automatically closed. This is possible since the SDC uses a separate SA10 pin which is asserted high during this command. This command precedes the AUTO-REFRESH command.

Auto-Refresh Command

The SDRAM internally increments the refresh address counter and causes an auto-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. This command is issued to the external bank whether or not it is enabled (`EBE` in the SDRAM memory global control register). 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 after exiting self-refresh mode.

Self-Refresh Mode

The self-refresh mode is controlled by the `SELF-REFRESH ENTRY` and `SELF-REFRESH EXIT` commands. The SDC must issue a series of commands, including the `SELF-REFRESH ENTRY` command, to put the SDRAM into this low power operation, and it must issue another series of commands, including the `SELF-REFRESH EXIT` command, to re-access the SDRAM.

Self-Refresh Entry Command

The `SELF-REFRESH ENTRY` 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 commands while the SDRAM is in self-refresh mode. Before executing the `SELF-REFRESH ENTRY` command, all internal banks are precharged. The `SELF-REFRESH ENTRY` command is started by setting the `SRFS` bit of the SDRAM memory global control register (`EBIU_SDGCTL`). The SDC now drives `SCKE` low.

SDC Functional Description

-  Only the `SCKE` pin keeps control during self-refresh, all other SDRAM pins are allowed to be disabled. However the SDC still drives the `SCLK` during self-refresh mode. Software may disable the clock by clearing the `SCTLE` bit in `EBIU_SDGCTL`.

Self-Refresh Exit Command

Leaving self-refresh mode is performed with the `SELF-REFRESH EXIT` command, whereby the SDC asserts `SCKE`. Any internal core/DMA access causes the SDC to perform an `SELF-REFRESH EXIT` command. 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. 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:

Time to exit self-refresh: $2 \times (t_{RAS} + t_{RP})$

-  The minimum time between a subsequent `SELF-REFRESH ENTRY` and the `SELF-REFRESH EXIT` command is at least t_{RAS} cycles. If a self-refresh entry command is issued during any MDMA transfer, the SDC satisfies this core request with the minimum self-refresh period (t_{RAS}).

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 7-7.](#)

No Operation Command

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 to insert additional wait states, the NOP command is given. When the SDC is not accessing the SDRAM, the command inhibit command is given ($SMS = 1$).

SDC SA10 Pin

The SDRAM's $A[10]$ pin follows the truth table below:

- During the precharge command, it is used to indicate a precharge all
- During a bank activate command, it outputs the row address bit
- During read and write commands, it is used to disable auto-precharge

Therefore, the SDC uses a separate SA10 pin with these rules.

 Connect the SA10 pin with the SDRAM $A[10]$ pin. Because the ADSP-BF51x processor uses byte addressing, it starts with $A[1]$. The $A[11]$ pin is left unconnected for SDRAM accesses and is replaced by the SA10 pin.

SDC Programming Model

The following sections provide programming model information for the SDC.

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, these steps are required:

1. Ensure the clock to the SDRAM is stable after the power has stabilized for the proper amount of time (typically 100 ms).
2. Write to the SDRAM refresh rate control register (EBIU_SDRRC).
3. Write to the SDRAM memory bank control register (EBIU_SDBCTL).
4. Write to the SDRAM memory global control register (EBIU_SDGCTL) and issue an SSYNC instruction.
5. Perform SDRAM access.

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 /configuration of SDRAM memory (EBSZ and EBCAW) and to enable the external bank (EBE). Prior to the start of the SDRAM powerup sequence, 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 7-7.](#)

The powerup latency can be estimated as:

$$t_{RP} + (8 \times t_{RFC}) + t_{MRD} + t_{RCD}$$

If the external bank remains disabled after the SDRAM powerup sequence has completed, any transfers to it will result in a hardware error interrupt and the SDRAM transfer will 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)
- To select and enable the start of the SDRAM powerup sequence (PSM, PSSE)

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 7-7.](#)

Once the `PSSE` bit in the `EBIU_SDGCTL` register is set, 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.

Example SDRAM System Block Diagrams

[Figure 7-12](#) shows a block diagram of an SDRAM interface. In this example, the SDC is connected to $2 \times (8\text{M} \times 8) = 8\text{M} \times 16$ to form one external 128M bit / 16M byte bank of memory. The system’s page size is 1024 bytes. The same address and control bus feeds both SDRAM devices.

SDC Programming Model

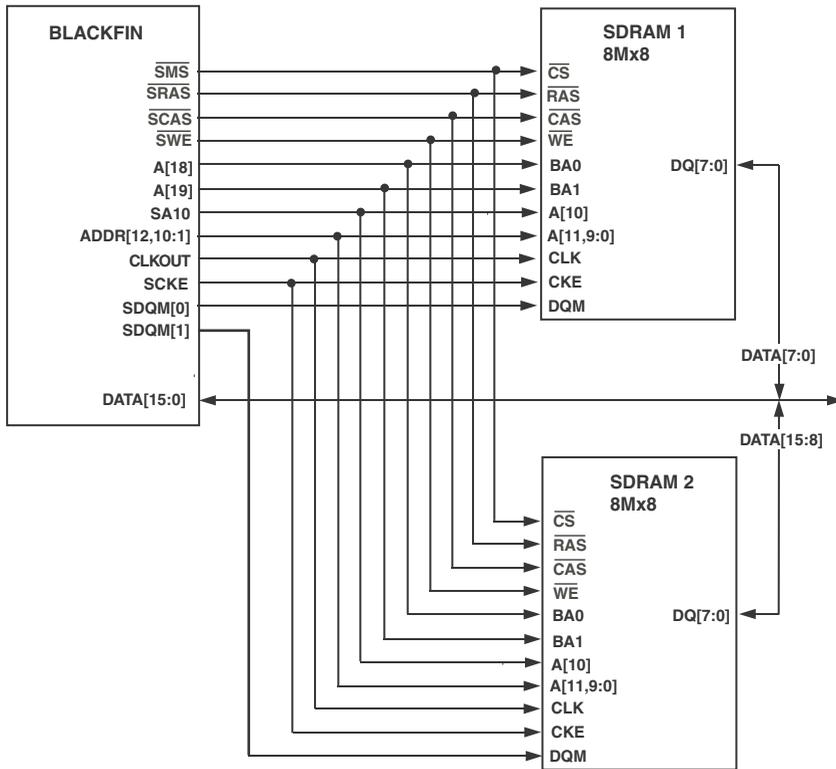


Figure 7-12. SDRAM System Block Diagram, Example 1

Figure 7-13 shows a block diagram of an SDRAM interface. In this example, the SDC is connected to $4 \times (16\text{M} \times 4) = 16\text{M} \times 16$ to form one external 256M bit / 32M byte bank of memory. The system's page size is 2048 bytes. The same address and control bus pass a registered buffer before they feed all 4 SDRAM devices.

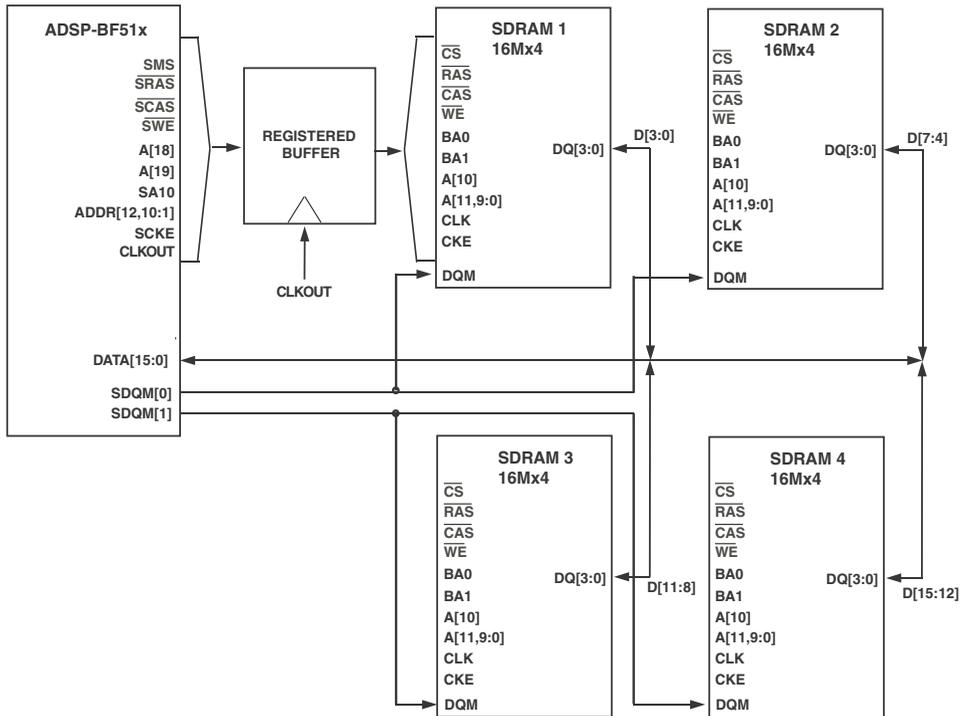


Figure 7-13. SDRAM System Block Diagram, Example 2

Furthermore, the `EBUFE` bit should be 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. Using this setting adds a cycle of data buffering to read and write accesses.

SDC Register Definitions

The following sections describe the SDC registers.

EBIU_SDRRC Register

The SDRAM refresh rate control register (EBIU_SDRRC, shown in Figure 7-14) 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)

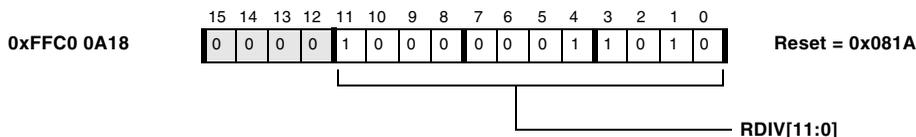


Figure 7-14. SDRAM Refresh Rate Control Register

To calculate the value that should be written to the EBIU_SDRRC register, use the following equation:

$$\begin{aligned}
 \text{RDIV} &= ((f_{\text{SCLK}} \times t_{\text{REF}}) / \text{NRA}) - (t_{\text{RAS}} + t_{\text{RP}}) \\
 &= (f_{\text{SCLK}} \times t_{\text{REF}}) - (t_{\text{RAS}} + t_{\text{RP}})
 \end{aligned}$$

Where:

- f_{SCLK} = SDRAM clock frequency (system clock frequency)
- t_{REF} = SDRAM row refresh period

- t_{REFI} = SDRAM row refresh interval
- NRA = Number of row addresses in SDRAM (refresh cycles to refresh whole SDRAM)
- t_{RAS} = Active to precharge time (t_{RAS} in the SDRAM memory global control register) in number of clock cycles
- t_{RP} = RAS to precharge time (t_{RP} in the SDRAM memory global control register) in number of clock cycles



Please see DRAM data sheet if NRA differs from the number of required refresh cycles. In this case use the refresh cycle number instead of NRA.

This equation (8192 row addresses per refresh cycle) 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 R_{DIV} for a typical SDRAM in a system with a 80 MHz clock:

$$f_{SCLK} = 80 \text{ MHz}$$

$$t_{REF} = 64 \text{ ms}$$

$$NRA = 8192 \text{ row addresses}$$

$$t_{RAS} = 6$$

$$t_{RP} = 3$$

The equation for R_{DIV} yields:

$$R_{DIV} = ((80 \times 10^6 \times 64 \times 10^{-3}) / 8192) - (6 + 3) = 616 \text{ clock cycles}$$

SDC Register Definitions

This means `RDIV` is 0x268 and the `EBIU_SDRRC` register should be written with 0x0268.



`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.

EBIU_SDBCTL Register

The SDRAM memory bank control register (`EBIU_SDBCTL`), shown in [Figure 7-15](#), includes 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.

- **External bank enable** (`EBE`)

The `EBE` bit is 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 7-7.](#)

- **External bank size** (`EBSZ`)

The `EBSZ` encoding stores the configuration information for the SDRAM bank interface. The EBIU supports 64M bit, 128M bit, 256M bit, and 512M bit SDRAM devices with x4, x8, and x16 configurations. [Table 7-13](#) maps SDRAM density and I/O width. See [“SDRAM External Bank Size” on page 7-26](#) for more information regarding the decoding of bank start addresses.

- **External bank column address width (EBCAW)**

The SDC determines the internal SDRAM page size from the EBCAW parameters. Page sizes of 512 B, 1K byte, 2K byte, and 4K byte are supported. Table 7-13 shows the page size and breakdown of the internal address (IA[31:0]), as seen from the core or DMA) into the row, bank, column, and byte address. The bank width in all cases is 16 bits. The column address and the byte address together make up the address inside the page.

SDRAM Memory Bank Control Register (EBIU_SDBCTL)

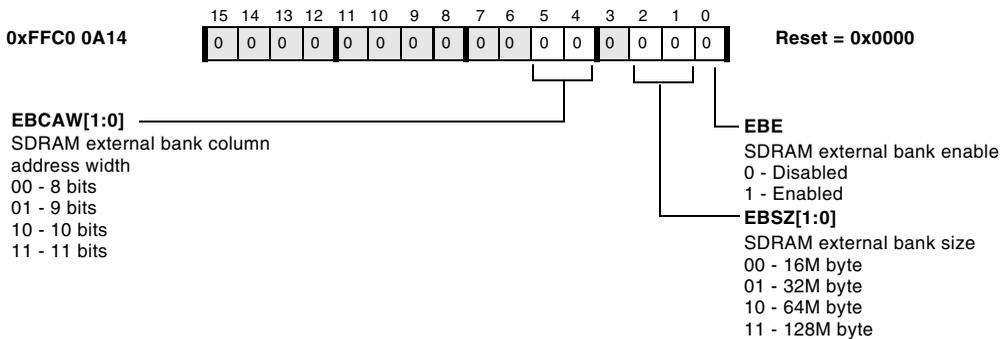


Figure 7-15. SDRAM Memory Bank Control Register

The page size can be calculated for 16-bit SDRAM banks with this formula:

$$\text{page size} = 2^{(\text{CAW} + 1)}$$

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).

SDC Register Definitions

Table 7-13. Internal Address Mapping

Bank Size (M byte) EBSZ bits	Col. Address. Width (CAW) EBCAW bits	Page Size (K Byte)	Bank Address	Row Address	Page	
					Column Address	Byte Address
128	11	4	IA[26:25]	IA[24:12]	IA[11:1]	IA[0]
128	10	2	IA[26:25]	IA[24:11]	IA[10:1]	IA[0]
128	9	1	IA[26:25]	IA[24:10]	IA[9:1]	IA[0]
128	8	.5	IA[26:25]	IA[24:9]	IA[8:1]	IA[0]
64	11	4	IA[25:24]	IA[23:12]	IA[11:1]	IA[0]
64	10	2	IA[25:24]	IA[23:11]	IA[10:1]	IA[0]
64	9	1	IA[25:24]	IA[23:10]	IA[9:1]	IA[0]
64	8	.5	IA[25:24]	IA[23:9]	IA[8:1]	IA[0]
32	11	4	IA[24:23]	IA[22:12]	IA[11:1]	IA[0]
32	10	2	IA[24:23]	IA[22:11]	IA[10:1]	IA[0]
32	9	1	IA[24:23]	IA[22:10]	IA[9:1]	IA[0]
32	8	.5	IA[24:23]	IA[22:9]	IA[8:1]	IA[0]
16	11	4	IA[23:22]	IA[21:12]	IA[11:1]	IA[0]
16	10	2	IA[23:22]	IA[21:11]	IA[10:1]	IA[0]
16	9	1	IA[23:22]	IA[21:10]	IA[9:1]	IA[0]
16	8	.5	IA[23:22]	IA[21:9]	IA[8:1]	IA[0]

Using SDRAMs With Systems Smaller Than 16M Byte

It is possible to use SDRAMs smaller than 16M byte on the ADSP-BF51x, as long as it is understood how the resulting memory map is altered.

Figure 7-16 shows an example where a 2M byte SDRAM (512K x 16 bits x 2 banks) is mapped to the external memory interface. In this example, there are 11 row addresses and eight column addresses per bank. Referring to Table 7-4 on page 7-25, the lowest available bank size (16M byte) for a device with eight column addresses has two bank address lines (IA[23:22])

and 13 row address lines (IA[21:9]). Therefore, one processor bank address line and two row address lines are unused when hooking up to the SDRAM in the example. This causes aliasing in the processor's external memory map, which results in the SDRAM being mapped into non-contiguous regions of the processor's memory space.

Referring to the table in [Figure 7-16](#), note that each line in the table corresponds to 2^{19} bytes, or 512K byte. Thus, the mapping of the 2M byte SDRAM is non-contiguous in Blackfin memory, as shown by the memory mapping in the left side of the figure.

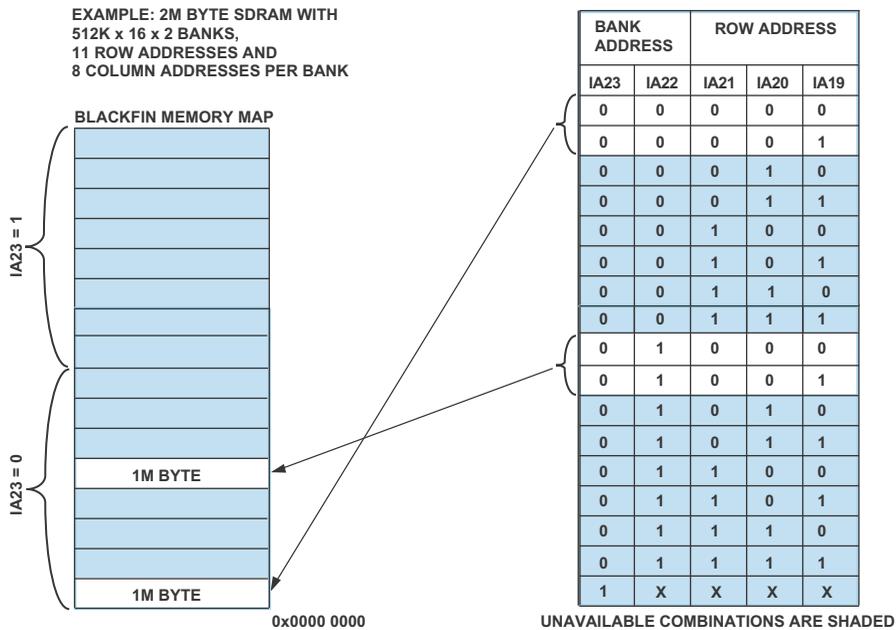


Figure 7-16. Using Small SDRAMs

EBIU_SDGCTL Register

The SDRAM memory global control register (EBIU_SDGCTL) includes all programmable parameters associated with the SDRAM access timing and configuration. [Figure 7-17](#) shows the EBIU_SDGCTL register bit definitions.

 Writes to this register should be followed by an SSYNC instruction to prevent a subsequent external access from occurring before the timing changes are properly effected.

SDRAM Clock Enable (SCTLE)

The SCTLE bit is used to enable or disable the SDC. If SCTLE is cleared, 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 7-7](#). When SCTLE is cleared, all SDC control pins are in their inactive states and the SDRAM clock is not running. The SCTLE bit must be set for SDC operation and is set by default at reset. 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.

The SCTLE bit allows software to disable all SDRAM control pins. These pins are SDQM[3:0], SCAS, SRAS, SWE, SCKE, and CLKOUT.

- SCTLE = 0
Disable all SDRAM control pins (control pins negated, CLKOUT low).
- SCTLE = 1
Enable all SDRAM control pins (CLKOUT toggles).

Note that the CLKOUT function is also shared with the AMC. Even if SCTLE is disabled, CLKOUT can be enabled independently by the CLKOUT enable in the AMC (AMCKEN in the EBIU_AMGCTL register).

SDRAM Memory Global Control Register (EBIU_SDGCTL)

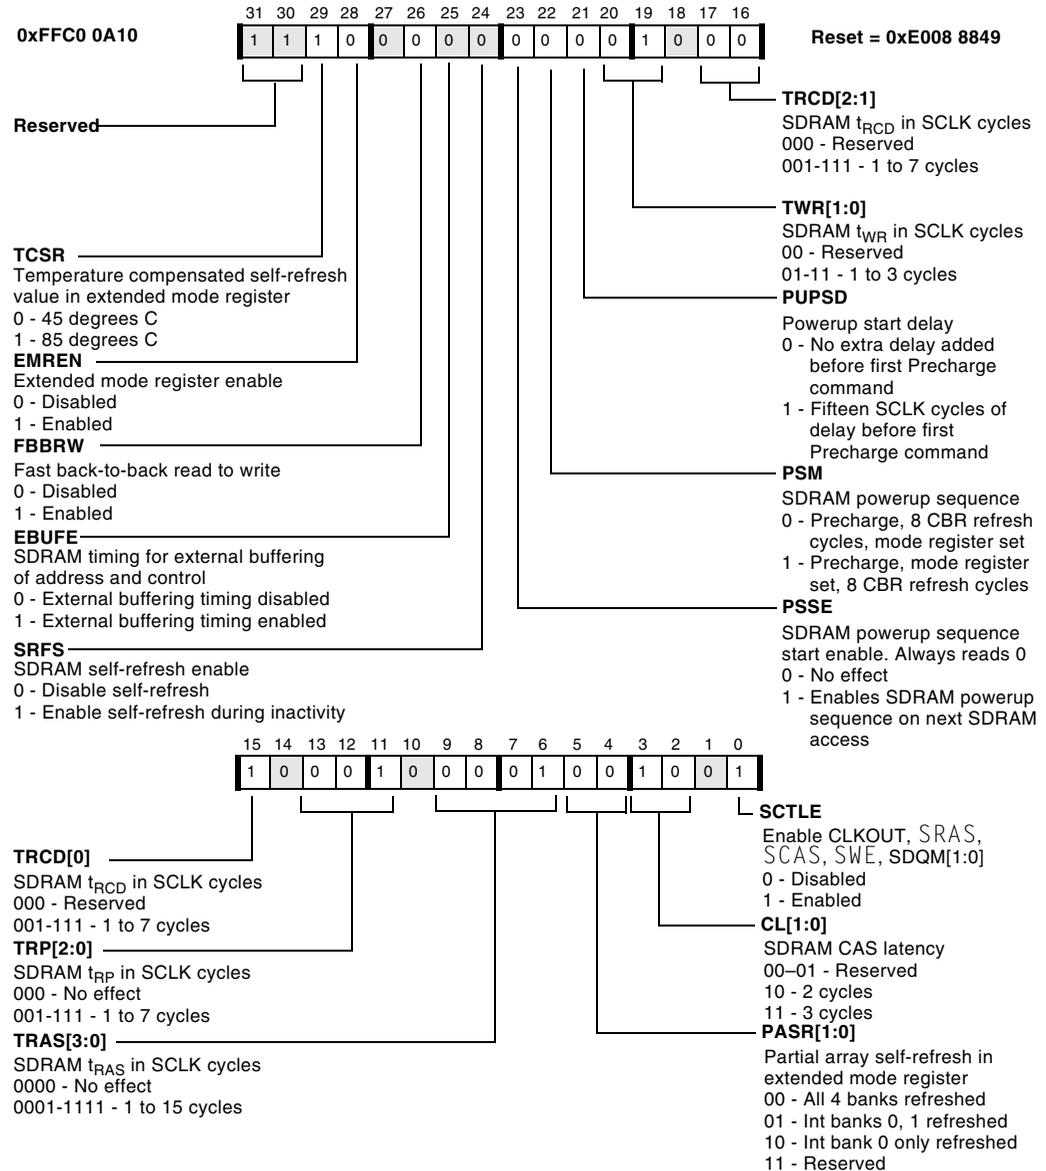


Figure 7-17. SDRAM Memory Global Control Register

SDC Register Definitions

If the system does not use SDRAM, `SCTLE` should be set to 0.

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 7-7.](#)

 With careful software control, the `SCTLE` bit can be used in conjunction with the `SRFS` bit to further lower power consumption by freezing the `CLKOUT` pin. However, `SCTLE` must remain enabled at all times when the SDC is needed to generate auto-refresh commands to SDRAM.

CAS Latency (CL)

The `CL` bits in the SDRAM memory global control register (`EBIU_SDGCTL`) select the CAS latency value:

- `CL = b#00`
Reserved
- `CL = b#01`
Reserved
- `CL = b#10`
2 clock cycles
- `CL = b#11`
3 clock cycles

Partial Array Self Refresh (PASR)

The PASR bits determine how many internal SDRAM banks are refreshed during self-refresh.

- PASR = b#00
All 4 banks
- PASR = b#01
Internal banks 0 and 1 refreshed
- PASR = b#10
Only internal bank 0 refreshed
- PASR = b#11
Reserved

Internal banks are decoded with the A[19:18] pins.



The PASR feature requires careful software control with regard to the internal bank used.

Bank Activate Command Delay (TRAS)

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 = b#0000
No effect
- TRAS = b#0001
1 clock cycle
- TRAS = b#0010
2 clock cycles
- TRAS = b#1111
15 clock cycles

SDC Register Definitions

Bank Precharge Delay (TRP)

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 = b#000
No effect
- TRP = b#001
1 clock cycle
- TRP = b#010
2 clock cycles
- TRP = b#111
7 clock cycles

RAS to CAS Delay (TRCD)

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 = b#000
Reserved, no effect
- TRCD = b#001
1 clock cycle
- TRCD = b#010
2 clock cycles
- TRCD = b#111
7 clock cycles

Write to Precharge Delay (TWR)

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 = b#00
Reserved
- TWR = b#01
1 clock cycle
- TWR = b#10
2 clock cycles
- TWR = b#11
3 clock cycles

Power-Up Start Delay (PUPSD)

The power-up start delay bit (PUPSD) optionally delays the power-up 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 power-up 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 power-up sequence in order to meet the exit self-refresh time (t_{XSR}) of the SDRAM.

Power-Up Sequence Mode (PSM)

If the PSM bit is set to 1, the SDC command sequence is:

1. Precharge all
2. Mode register set
3. 8 auto-refresh cycles

SDC Register Definitions

If the `PSM` bit is cleared, the SDC command sequence is:

1. Precharge all
2. 8 auto-refresh cycles
3. Mode register set

Power-Up Sequence Start Enable (PSSE)

The `PSM` and `PSSE` bits work together to specify and trigger an SDRAM power-up (initialization) sequence. Two events must occur before the SDC does the SDRAM power-up sequence:

- The `PSSE` bit must be set to enable the SDRAM power-up 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 power-up sequence may occur.

The SDRAM power-up sequence occurs and is followed immediately by the read or write transfer to SDRAM that was used to trigger the SDRAM power-up sequence. Note that there is a latency for this first access to SDRAM because the SDRAM power-up sequence takes many cycles to complete.

-  Before executing the SDC power-up sequence, ensure that the SDRAM receives stable power and is clocked for the proper amount of time, as described in the SDRAM specifications.

Self-Refresh Setting (SRFS)

The SRFS and SCTL bits work together in EBIU_SDGCTL for self-refresh control.

- SRFS = b#0
Disable self-refresh mode
- SRFS = b#1
Enter self-refresh mode

When SRFS is set, self-refresh mode is triggered. Once the SDC completes any active transfers, the SDC executes a sequence of commands to put the SDRAM into self-refresh mode.

When the device comes out of reset, the SCKE pin is driven high. If it is necessary to enter self-refresh mode after reset, program SRFS = b#1.

Enter Self-Refresh Mode

When SRFS is set, once the SDC enters an idle state it issues a precharge all command and then issues a self-refresh entry command. If an internal access is pending, the SDC delays issuing the self-refresh entry command until it completes the pending SDRAM access and any subsequent pending access requests.

Once the SDRAM device enters into self-refresh mode, the SDRAM controller asserts the SDSRA bit in the SDRAM control status register (EBIU_SDSTAT).

 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.

Before disabling the CLKOUT pin with the SCTL bit, be sure to place the SDC in self-refresh mode (SRFS bit). If this is not done, the SDRAM is unlocked and will not work properly.

SDC Register Definitions

Exit Self-Refresh Mode

The SDRAM device exits self-refresh mode only when the SDC receives core or DMA requests. In conjunction with the `SRFS` bit, two possibilities are given to exit self-refresh mode.

- If the `SRFS` bit remains set before the core/DMA request, the SDC exits self-refresh mode temporarily for a single request and returns back to self-refresh mode until a new request is latched.
- If the `SRFS` bit is cleared before the core/DMA request, the SDC exits self-refresh mode and returns to auto-refresh mode.

Before exiting self-refresh mode with the `SRFS` bit, be sure to enable the `CLKOUT` pin (`SCTLE` bit). If this is not done, the SDRAM is unlocked and will not work properly.

External Buffering Enabled (EBUFE)

With the total I/O width of 16 bits, a maximum of 4x4 bits can be connected in parallel in order to increase the system's overall page size.

To meet overall system timing requirements, systems that employ several SDRAM devices connected in parallel may require buffering between the processor and the 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.

 Connection of 4 x 4 bits rather than 1 x 16 bits increases the page size by a factor of four, thus resulting in fewer off-page penalties.

Fast Back-to-Back Read to Write (FBBRW)

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 cleared, a clock cycle is inserted between read accesses followed immediately by write accesses.

Extended Mode Register Enabled (EMREN)

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 PASR[1:0] bits control the value written to the extended mode register.

Temperature Compensated Self-Refresh (TCSR)

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.

 All reserved bits in this register must always be written with 0s.

EBIU_SDSTAT Register

The SDRAM control status register (EBIU_SDSTAT), shown in [Figure 7-18](#), 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.

SDRAM Control Status Register (EBIU_SDSTAT)

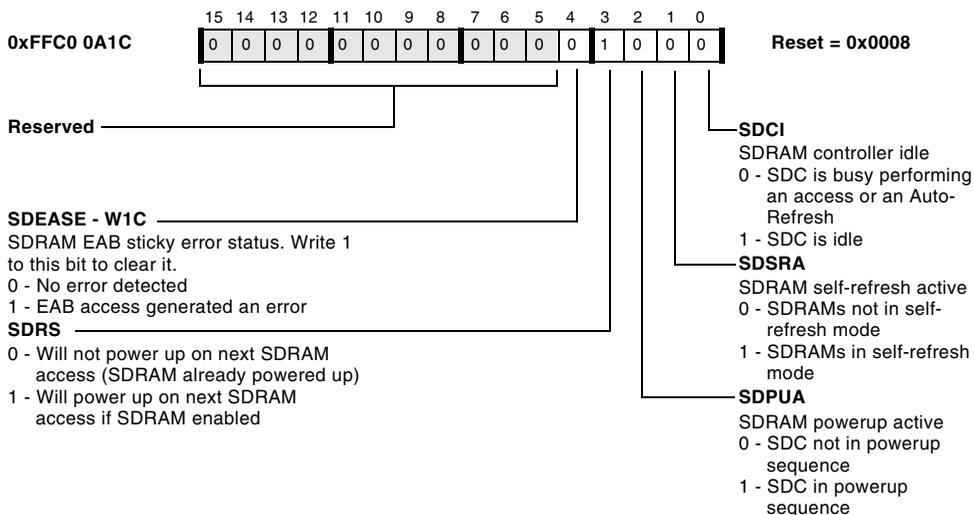


Figure 7-18. SDRAM Control Status Register

- **SDC idle (SDCI)**

If the `SDCI` bit is cleared, the SDC is performing a user access or auto-refresh. If the `SDCI` bit is set, no commands are issued and the SDC is in idle state.

- **SDC self-refresh active (SDSRA)**

If the `SDSRA` bit is cleared, the SDC is performing auto-refresh (`SCKE` pin = 0). If the `SDSRA` bit is set, the SDC performs self-refresh mode (`SCKE` pin = 1).

- **SDC powerup active** (SDPUA)

If the SDPUA bit is cleared, the SDC is not in powerup sequence. If the SDPUA bit is set, the SDC performs the powerup sequence.

- **SDC powerup delay** (SDRS)

If the SDRS bit is cleared, the SDC has already powered up. If the SDRS bit is set, the SDC will still perform the powerup sequence.

- **SDC EAB sticky error status** (SDEASE)

If the SDEASE bit is cleared, there were no errors detected on the EAB core bus. If the SDEASE bit is set, there were errors detected on the EAB core bus. The SDEASE bit 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.

SDC Programming Examples

[Listing 7-3](#) through [Listing 7-6](#) provide examples for working with the SDC.

Listing 7-3. 16-Bit Core Transfers to SDRAM

```
.section L1_data_b;
.byte2 source[N] = 0x1122, 0x3344, 0x5566, 0x7788;
.section SDRAM;
.byte2 dest[N];
.section L1_code;
    I0.L = lo(source);
    I0.H = hi(source);
    I1.L = lo(dest);
    I1.H = hi(dest);
```

SDC Programming Examples

```
R0.L = w[I0++];
p5=N-1;
lsetup(lp, lp) lc0=p5;
lp:R0.L = w[I0++] || w[I1++] = R0.L;
           w[I1++] = R0.L;
```

Listing 7-4. 8-Bit Core Transfers to SDRAM Using Byte Mask SDQM[1:0] Pins

```
.section L1_data_b;
.byte source[N] = 0x11, 0x22, 0x33, 0x44, 0x55, 0x66, 0x77, 0x88;

.section SDRAM;
.byte dest[N];

p0.L = lo(source);
p0.H = hi(source);
p1.L = lo(dest);
p1.H = hi(dest);
p5=N;
lsetup(start, end) lc0=p5;
start: R0 = b[p0++](z);
end:   b[p1++] = R0; /* byte data masking */
```

Listing 7-5. Self-Refresh Mode Power Savings With Disabled CLKOUT

```
r0.l = w[I1++]; /* SDRAM access */
ssync; /* force last SDRAM access to finish */
P0.L = lo(EBIU_SDGCTL);
P0.H = hi(EBIU_SDGCTL);
R1 = [P0];
bitset(R1, bitpos(SRFS)); /* enter self-refresh */
[P0] = R1;
```

```

ssync;

P0.L = 1o(EBIU_SDSTAT);
P0.H = hi(EBIU_SDSTAT);
R0 = [P0];
ssync;

self_refresh_status:
cc = bittst(R0, bitpos(SDSRA)); /* poll self-refresh status */
if !cc jump self_refresh_status;

P0.L = 1o(EBIU_SDGCTL);
P0.H = hi(EBIU_SDGCTL);
R1 = [P0];
bitclr(R1, bitpos(SCTLE));
/* disable CLKOUT after approx 20 cycles */
[P0] = R1;
ssync;

P5 = 30000;
LSETUP(1p,1p) LC0 = P5;
1p: nop; /* dummy loop */

R1 = [P0];
bitset(R1, bitpos(SCTLE));
/* enable CLKOUT after approx 20 cycles */
[P0] = R1;
ssync;

R1 = [P0];
bitclr(R1, bitpos(SRFS)); /* exit self-refresh */
[P0] = R1;

```

SDC Programming Examples

```
ssync;

w[I1++] = r0.l; /* SDRAM access */
```

Listing 7-6. Init

```
/******  
/* SDRAM part# Micron MT48LC32M8A2-75 (32Mx8/256Mbit) */  
/* 8k rows, 1k columns      -> EBCAW = 10 */  
/* 2xSDRAM: 32Mx16 = 64Mbytes -> EBSZ = 010 */  
  
/* populated SDRAM addresses -> 0x00000000 - 0x01FFFFFF */  
/* internal SDRAM bank A 0x00000000 - 0x007FFFFFF */  
/* internal SDRAM bank B 0x00800000 - 0x00FFFFFF */  
/* internal SDRAM bank C 0x01000000 - 0x017FFFFFF */  
  
/* internal SDRAM bank D 0x01800000 - 0x01FFFFFF */  
  
  
/* powerup: PRE-REF-MRS      -> PSM = 0 */  
/* SCLK = 80 MHz */  
/* tCK = 7.5ns min@CL=3     -> CL = 3 */  
/* tRAS = 44ns min          -> TRAS = 6 */  
/* tRP = 20ns min           -> TRP = 3 */  
/* tRCD = 20ns min          -> TRCD = 3 */  
/* tWR = 15ns min           -> TWR = 2 */  
/* tREF = 64ms max  
      ->RDIV = (80MHz*64ms)/8192-(6+3)=0x268 cycles */  
/******  
#ifdef INIT_SDRAM  
  
/* Check if already enabled */  
p0.l = lo(EBIU_SDSTAT);  
p0.h = hi(EBIU_SDSTAT);  
r0 = [p0];
```

```
cc = bittst(r0, bitpos(SDRS));
if !cc jump skip init_sdram;

/* SDRAM Refresh Rate Control Register */
P0.L = 1o(EBIU_SDRRC);
P0.H = hi(EBIU_SDRRC);
R0.L = 0x0268;
W[P0] = R0.L;

/* SDRAM Memory Bank Control Register */
P0.L = 1o(EBIU_SDBCTL);
P0.H = hi(EBIU_SDBCTL);
R0.L = 0x0025;
W[P0] = R0.L;

/* SDRAM Memory Global Control Register */
P0.L = 1o(EBIU_SDGCTL);
P0.H = hi(EBIU_SDGCTL);
R0.L = 0x998d;
R0.H = 0x8491;
[P0] = R0;
ssync; /* wait until executed */
```

SDC Programming Examples

8 DYNAMIC POWER MANAGEMENT

This chapter describes the dynamic power management functionality of the Blackfin processor and includes the following sections:

- [“Phase Locked Loop and Clock Control”](#)
- [“Dynamic Power Management Controller” on page 8-7](#)
 - [“Operating Modes” on page 8-7](#)
 - [“Dynamic Supply Voltage Control” on page 8-15](#)
- [“System Control ROM Function” on page 8-24](#)
- [“PLL and VR Registers” on page 8-19](#)
- [“Programming Examples” on page 8-30](#)

Phase Locked Loop and Clock Control

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 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 clocks the Peripheral Access Bus (PAB),

Phase Locked Loop and Clock Control

DMA Access Bus (DAB), External Access Bus (EAB), and the 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 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.

PLL Overview

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 8-7](#).

Subject to the maximum VCO frequency specified in the processor data sheet, 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.

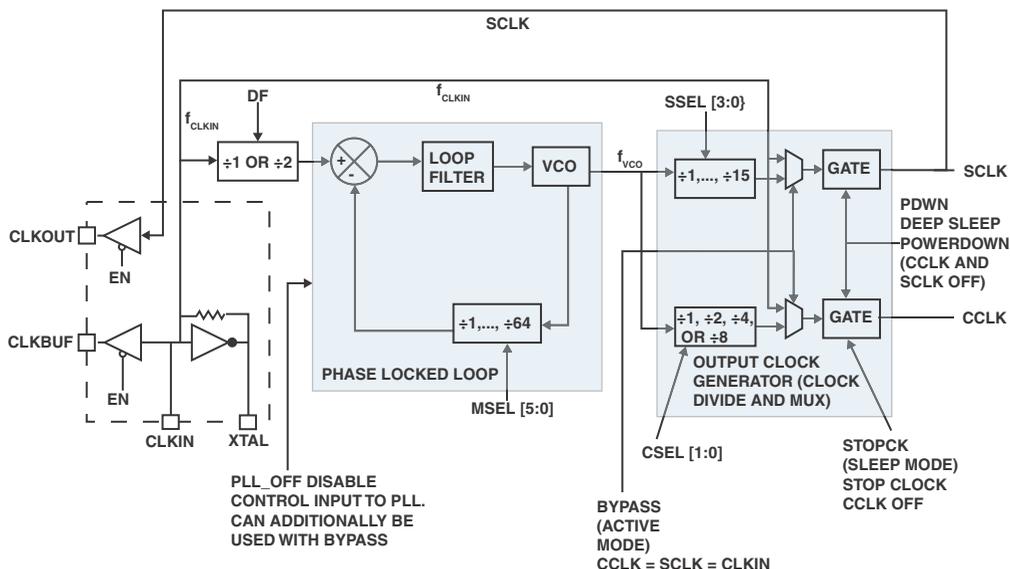


Figure 8-1. PLL Block Diagram

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_CTL Register” on page 8-21.

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

Phase Locked Loop and Clock Control

The reset value of `MSEL` is `0x5`. 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 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

The PLL control (`PLL_CTL`) register controls operation of the PLL (see [Figure 8-4 on page 8-21](#)). Note that changes to the `PLL_CTL` register do not take effect immediately. In general, the `PLL_CTL` register is first programmed with a new value, and then a specific PLL programming sequence must be executed to implement the changes. This is handled automatically by the system control ROM function (`bfrom_SysControl()`) as described in [“System Control ROM Function” on page 8-24](#).

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 the divider ratio must be chosen to limit the SCLK to a frequency specified in the processor data sheet. The SCLK drives all synchronous, system-level logic.

The divider ratio control bits, CSEL and SSEL, are in the PLL divide (PLL_DIV) register. For information about this register, see “PLL_DIV Register” on page 8-21.

The reset value of CSEL[1:0] is 0x0, and the reset value of SSEL[3:0] is 0x4. These values can be reprogrammed at startup by the boot code.

By updating PLL_DIV with an appropriate value, you can change the CSEL and SSEL value dynamically. Note 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.

Unlike writing the PLL_CTL register, the PLL_DIV register can be programmed at any time to change the CCLK and SCLK divide values without entering the PLL programming sequence.

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	300	150
10	4	400	100
11	8	400	50

Phase Locked Loop and Clock Control

As long as the MSEL and DF control bits in the PLL control (PLL_CTL) register remain constant, the PLL is locked.

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	50	50
0010	2:1	150	75
0011	3:1	150	50
0100	4:1	200	50
0101	5:1	300	60
0110	6:1	360	60
N = 7–15	N:1	400	400/N



If changing the clock ratio via writing a new SSEL value into PLL_DIV, take care that the enabled peripherals do not suffer data loss due to SCLK frequency changes.

When changing clock frequencies in the PLL, the PLL requires time to stabilize and lock to the new frequency. The PLL lock count (PLL_LOCKCNT) register 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 processor data sheet for more information about PLL stabilization time and programmed values for this register. For more information about operating modes, see [“Operating Modes” on page 8-7](#).

Dynamic Power Management Controller

The Dynamic Power Management Controller (DPMC) works in conjunction with the PLL, allowing the user to control the processor's performance characteristics and power dissipation dynamically. The DPMC provides these features that allow the user to control performance and power:

- Multiple operating modes – The processor works in four operating modes, each with different performance characteristics and power dissipation profiles. See [“Operating Modes”](#).
- Peripheral clocks – Clocks to each peripheral are disabled automatically when the peripheral is disabled.
- Voltage control – The V_{DDINT} domain must be powered by an external voltage regulator. For more information see [“Voltage Regulation Interface”](#) on page 27-14.

Operating Modes

The processor works 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 Status	PLL Bypassed	CCLK	SCLK	Allowed DMA Access
Full On	None	Enabled	No	Enabled	Enabled	L1
Active	Medium	Enabled ¹	Yes	Enabled	Enabled	L1
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 active and full-on states of the DPMC/PLL can be determined by reading the PLL status register (see “[PLL_STAT Register](#)” on page 8-22). In these modes, the core can either execute instructions or be in the `IDLE` core state. If the core is in the `IDLE` state, it can be awakened by several sources. (See [Chapter 5, “System Interrupts”](#) for details.)

The following sections describe the DPMC/PLL states in more detail, as they relate to the power management controller functions.

Full-On Mode

Full-on mode is the maximum performance mode. In this mode, the PLL is enabled and not bypassed. Full-on mode is the normal execution state of the processor, with the processor and all enabled peripherals running at full speed. The system clock (`SCLK`) frequency is determined by the `SSEL` specified ratio to `VCO`. DMA access is available to L1 and external memories. From full-on mode, the processor can transition directly to active, sleep, or deep sleep modes, as shown in [Figure 8-2 on page 8-12](#).

Active Mode

In 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 external memories.

In 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 full-on or sleep modes.

From active mode, the processor can transition directly to full-on, sleep, or deep sleep modes.



In this mode or in the transition phase to other modes, changes to MSEL are not latched by the PLL.

Sleep Mode

Sleep mode significantly reduces power dissipation by idling the processor core. The CCLK is disabled in this mode; however, SCLK continues to run at the speed configured by MSEL and SSEL bit settings. Since CCLK is disabled, DMA access is available only to external memory in sleep mode. From sleep mode, a wakeup event causes the processor to transition to one of these modes:

- Active mode if the BYPASS bit in the PLL_CTL register is set
- Full-on mode if the BYPASS bit is cleared

When sleep mode is exited, the processor resumes execution from the program counter value present immediately prior to entering sleep mode.

Deep Sleep Mode

Deep sleep mode maximizes power savings by disabling the PLL, CCLK, and SCLK. In this mode, the processor core and all peripherals except the real-time clock (RTC) are disabled. DMA is not supported in this mode.

Deep sleep mode can be exited only by a hardware reset event or an RTC interrupt. A hardware reset begins the hardware reset sequence. For more information about hardware reset, see [Chapter 5, “System Interrupts”](#). An RTC interrupt causes the processor to transition to active mode, and execution resumes from where the program counter was when deep sleep mode was entered. If an interrupt is also enabled in SIC_IMASK, the vector is taken immediately after exiting deep sleep and the ISR is executed.

Note an RTC interrupt in deep sleep mode automatically resets some fields of the PLL control (PLL_CTL) register. See [Table 8-5](#).

Dynamic Power Management Controller



When in deep sleep mode, clocking to the SDRAM is turned off. Before entering deep sleep mode, software should ensure that important information in SDRAM is saved to a non-volatile memory and/or the SDRAM is placed into self-refresh mode.

Table 8-5. PLL_CTL Values after RTC Wakeup Interrupt

Field	Value
PLL_OFF	0
STOPCK	0
PDWN	0
BYPASS	1

Hibernate State

For lowest possible power dissipation, this state allows the internal supply (V_{DDINT}) to be powered down by the external regulator, while keeping the I/O supply (V_{DDEXT} and V_{DDMEM}) 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 8-2 on page 8-12](#). This feature is discussed in detail in “[Powering Down the Core \(Hibernate State\)](#)” on [page 8-17](#).

Operating Mode Transitions

[Figure 8-2 on page 8-12](#) graphically illustrates the operating modes and transitions. In the diagram, ellipses represent operating modes and rectangles represent processor states. Arrows show the allowed transitions into and out of each mode or state.

For mode transitions, the text next to each transition arrow shows the fields in the PLL control (PLL_CTL) register that must be changed for the transition to occur. For example, the transition from 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.

For transitions to processor states, the text next to each transition arrow shows either a processor event (RTC wake up or hardware reset) or the fields in the voltage regulator control register (`VR_CTL`) that must be changed for the transition to occur.

For information about how to effect mode transitions, see [“Programming Operating Mode Transitions” on page 8-13](#).

In addition to the mode transitions shown in [Figure 8-2 on page 8-12](#), the PLL can be modified while in active operating mode. 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 [“Programming Operating Mode Transitions” on page 8-13](#)).

- **PLL disabled:** In addition to being bypassed in the active mode, the PLL can be disabled.

When the PLL is disabled, additional power savings are achieved although they are relatively small. To disable 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 disabled, it can be re-enabled later when additional performance is required.

The PLL must be re-enabled before transitioning to the full-on or sleep operating modes. To re-enable the PLL, clear the `PLL_OFF` bit in the `PLL_CTL` register, and then execute the PLL programming sequence.

- New multiplier ratio: The multiplier ratio can also be changed while in full-on mode.

The PLL state automatically transitions to active mode while the PLL is locking. After locking, the PLL returns to 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 8-13).

Table 8-6 summarizes the allowed operating mode transitions.



Attempting to cause mode transitions other than those shown in Table 8-6 causes unpredictable behavior.

Table 8-6. Allowed Operating Mode Transitions

New Mode	Current Mode			
	Full-On	Active	Sleep	Deep Sleep
Full On	–	Allowed	Allowed	Allowed
Active	Allowed	–	Allowed	Allowed
Sleep	Allowed	Allowed	–	–
Deep Sleep	Allowed	Allowed	–	–

Programming Operating Mode Transitions

The operating mode is defined by the state of the PLL_OFF, BYPASS, STOPCK, and PDWN bits of the PLL control (PLL_CTL) register. Merely modifying the bits of the PLL_CTL register does not change the operating mode or behavior of the PLL. Changes to the PLL_CTL register are realized only after a specific code sequence is executed. This sequence is managed by a user-callable routine in the on-chip ROM called `bfrom_SysControl()`. When calling this function, no further precautions have to be taken. See “System Control ROM Function” on page 8-24 for more information.

Dynamic Power Management Controller

If the `PLL_CTL` register changes include a new `CLKIN` to VCO multiplier or power is reapplied to the PLL, the PLL needs to relock. To relock, the PLL lock counter is cleared first, then starts incrementing once per `SCLK` cycle. After the PLL lock counter reaches the value programmed in the PLL lock count (`PLL_LOCKCNT`) register, the PLL sets the `PLL_LOCKED` bit in the PLL status (`PLL_STAT`) register, and the PLL asserts the PLL wake-up interrupt.

When the `bfrom_SysControl()` routine reprograms the `PLL_CTL` register with a new value, the `bfrom_SysControl()` routine executes a subsequent `IDLE` instruction and prevents all other system interrupt sources, other than the DPMC, from waking up the core from the `IDLE` state. If the lock counter expires, the PLL issues an interrupt, and the code execution continues the instruction after the `IDLE` instruction. Therefore, the system is in the new state by the time the `bfrom_SysControl()` routine returns.



- If the new value written to the `PLL_CTL` or `VR_CTL` register is the same as the previous value, the PLL wake-up occurs immediately (PLL is already locked), but the core and system clock are bypassed for the `PLL_LOCKCNT` duration. For this interval, code executes at the `CLKIN` rate instead of the expected `CCLK` rate. Software guards against this condition by comparing the current value to the new value before writing the new value.
- When the wake-up signal is asserted, the code execution continues the instruction after the `IDLE` instruction, causing a 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 sleep operating mode, the processor transitions immediately to sleep mode and waits for a wake-up signal before continuing code execution. If the

PLL_CTL register is programmed to enter the deep sleep operating mode, the processor immediately transitions to deep sleep mode and waits for an RTC interrupt or hardware reset signal:

- An RTC interrupt causes the processor to enter active operating mode and to return from the `bfrom_SysControl()` routine.
- A hardware reset causes the processor to execute the reset sequence. For more information see [Chapter 26, “System Reset and Booting”](#).

If no operating mode transition is programmed, the PLL generates a wake-up signal, and the `bfrom_SysControl()` routine returns.

Dynamic Supply Voltage Control

In addition to clock frequency control, the processor's core is capable of running 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 multiple power domains. 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 product data sheet for details on the allowed voltage ranges for each power domain and power dissipation data.

Power Supply Management

V_{DDINT} is supplied by an external regulator and pin \overline{PG} is used to accept an active-low power-good indicator from the regulator. Note that the external regulator must comply with the V_{DDINT} specifications defined in the processor data sheet.

Changing Voltage

When changing the voltage using an external regulator, a specific programming sequence must be followed.

Unlike other Blackfin derivatives that feature an internal voltage regulator; the voltage level for the ADSP-BF51x cannot be changed by programming the `VR_CTL` register. With an internal voltage regulator, the PLL would automatically enter the active mode when the processor enters the `IDLE` state. At that point the voltage level would change and the PLL would re-lock to the new voltage. After the `PLL_LOCKCNT` has expired, the part returns to the full-on state.

With an external voltage regulator, this sequence must be reproduced in the program code by the user. The `PLL_LOCKCNT` register cannot be used in this case, but the value is still needed for calculating the required delay. A larger `PLL_LOCKCNT` value may be necessary for changing voltages than when changing just the PLL frequency. See the processor data sheet for details.

The processor must enter active mode before the user can access the external voltage regulator and program a new voltage level. See the data sheet of external voltage regulator for information on changing voltage levels. See the processor data sheet for more information about voltage tolerances and allowed rates of change.



Reducing the processor's 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 the recommended procedure is to bring the processor to the sleep operating mode before substantially varying the voltage.

The user must ensure a stable voltage and give the PLL time to re-lock at the new voltage level. This can be done by running the core in a loop for a certain amount of time before leaving active mode.

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 processor data sheet for the new operating voltage level.

Please see [“Changing Voltage Levels” on page 8-40](#) for more details on mode transitions and changing voltage levels.

The `VSTAT` bit in the `PLL_STAT` register can be used to indicate whether V_{DDINT} is stable and ready to use. The `VSTAT` bit works in conjunction with the \overline{PG} (Power Good) input signal of the ADSP-BF51x. The inverted version of a "power good" signal from the external regulator is fed to the ADSP-BF51x to indicate that the voltage has reached its programmed value. That in turn will set the `VSTAT` bit, which should be considered the end of your "wait" state for the voltage regulator to settle.

Powering Down the Core (Hibernate State)

The external regulator can be signaled to shut off V_{DDINT} using the `EXT_WAKE` signal. Writing `b#00` to the `FREQ` bits of the `VR_CTL` register, which disables `CCLK` and `SCLK`, will also make `EXT_WAKE` go low. `EXT_WAKE` will transition high if any wakeup sources occur, which will signal the external voltage regulator to turn V_{DDINT} on again. The wakeup sources are several user-selectable events, all of which are controlled in the `VR_CTL` register:

- Assertion of the \overline{RESET} pin always exits hibernate state and requires no modification to `VR_CTL`.
- RTC event. Set the wake-up enable control bit (`WAKE`) to enable wake-up upon an RTC interrupt. This can be any of the RTC interrupts (alarm, daily alarm, day, hour, minute, second, or stopwatch).

Dynamic Power Management Controller

- External GP event or Ethernet PHY event. Set the PHY wakeup enable control (PHYWE) bit to enable wakeup upon assertion of the PHY_INT/PF15 pin by an external PHY device. If no external PHY interrupt is needed, set this bit to enable a general-purpose external event via the PF15 pin.
- Pin EXT_WAKE is provided to indicate the occurrence of wakeup. EXT_WAKE is an output pin, which is a logical OR of the above wakeup sources, except hardware reset. The pin follows the wakeup signal of the various wakeup sources.



When the core is powered down, V_{DDINT} is set to 0 V, and the internal state of the processor is not maintained, with the exception of the VR_CTL register. 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. Be sure to set the drive SCKE low during reset control (SCKELOW) bit in VR_CTL to protect against the default reset state behavior of setting the EBIU pins to their inactive state. Failure to set the SCKELOW 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.

Powering down V_{DDINT} does not affect V_{DDEXT} or V_{DDMEM} . While V_{DDEXT} and V_{DDMEM} are still applied to the processor, external pins are maintained at a three-state level unless specified otherwise.



The SCKE pin will be three-stated during hibernate. In addition to setting the SCKELOW bit in VR_CTL prior to entering the hibernate state, an external pull-down resistor on the SCKE pin is required to also keep the pin low when the Blackfin processor is not driving it.

To signal the external regulator to power down V_{DDINT} :

1. Write 0 to the appropriate bits in the `SIC_IWRx` registers to prevent enabled peripheral resources from interrupting the hibernate process.
2. Call the `bfrom_SysControl()` routine; ensure that the `FREQ` bits in the `VR_CTL` variable are set to `b#00`, and the appropriate wake-up bit(s) to 1 (`WAKE` and/or `Ethernet Phy`). Optionally, set the `SCKELOW` bit if `SDRAM` data should be maintained.
3. The `bfrom_SysControl()` routine executes until V_{DDINT} transitions to 0 V. `bfrom_SysControl()` never returns.
4. When the processor is woken up, the PLL relocks and the boot sequence defined by the `BMODE[2:0]` pin settings takes effect.

The `WURESET` in the `SYCTRL` register is set and stays set until the next hardware reset. The `WURESET` bit may control a conditional boot process.

PLL and VR Registers

The user interface to the PLL and VR registers is through the system control ROM function (`bfrom_SysControl()`) described in “[System Control ROM Function](#)” on page 8-24. The memory-mapped registers (MMRs) are shown in [Table 8-7](#) and illustrated in [Figure 8-3](#) through [Figure 8-7](#).

[Table 8-7](#) shows the functions of the PLL/VR registers.

PLL and VR Registers

Table 8-7. PLL/VR Register Mapping

Register Name	Function	Notes	For More Information See:
PLL_CTL	PLL control register	Requires reprogramming sequence when written	Figure 8-4 on page 8-21
PLL_DIV	PLL divisor register	Can be written freely	Figure 8-3 on page 8-21
PLL_STAT	PLL status register	Monitors active modes of operation	Figure 8-5 on page 8-22
PLL_LOCKCNT	PLL lock count register	Number of SCLKs allowed for PLL to relock	Figure 8-6 on page 8-22
VR_CTL	Voltage regulator control register	Requires PLL reprogramming sequence when written	Figure 8-7 on page 8-23

PLL_DIV Register

PLL Divide Register (PLL_DIV)

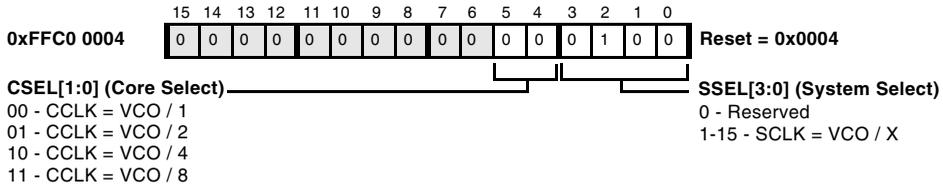


Figure 8-3. PLL Divide Register

PLL_CTL Register

PLL Control Register (PLL_CTL)

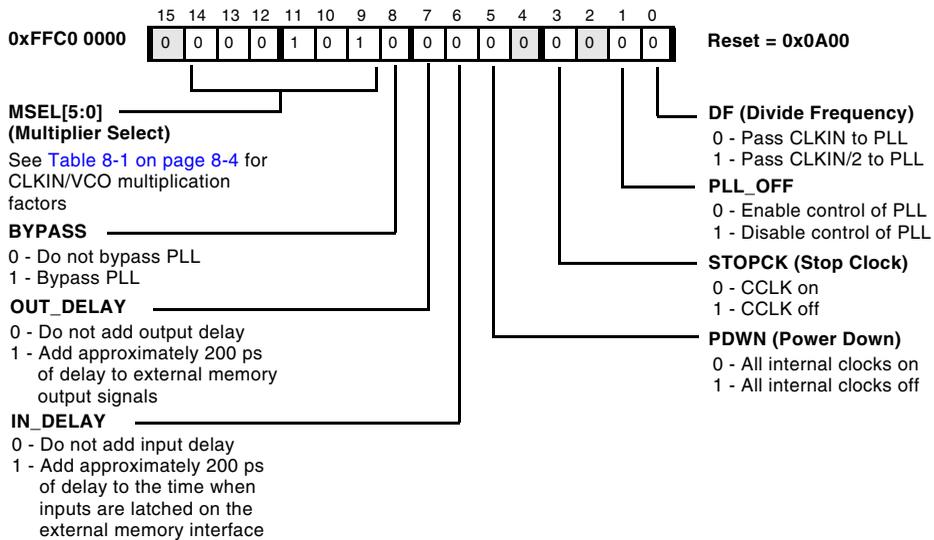


Figure 8-4. PLL Control Register

PLL_STAT Register

PLL Status Register (PLL_STAT)

Read only. Unless otherwise noted, 1 - Processor operating in this mode. For more information, see "Operating Modes" on page 8-7.

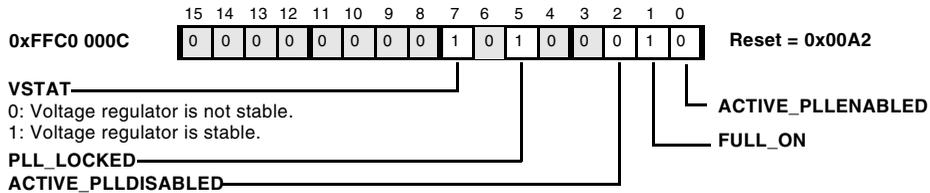


Figure 8-5. PLL Status Register

PLL_LOCKCNT Register

PLL Lock Count Register (PLL_LOCKCNT)

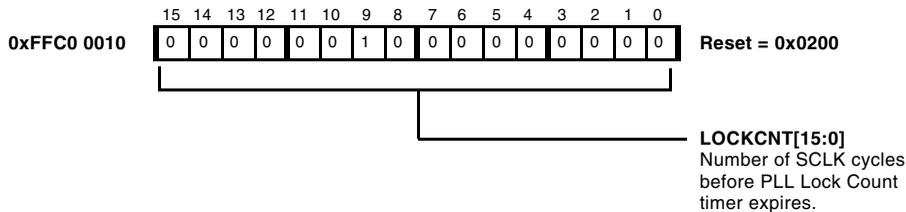


Figure 8-6. PLL Lock Count Register

VR_CTL Register

The CLKIN buffer output enable (CLKBUFOE) control bit allows another device, most likely the Ethernet PHY, and the Blackfin processor to run from a single crystal oscillator. Clearing this bit prevents the CLKBUF pin from driving a buffered version of the input clock CLKIN.

Voltage Regulator Control Register (VR_CTL)

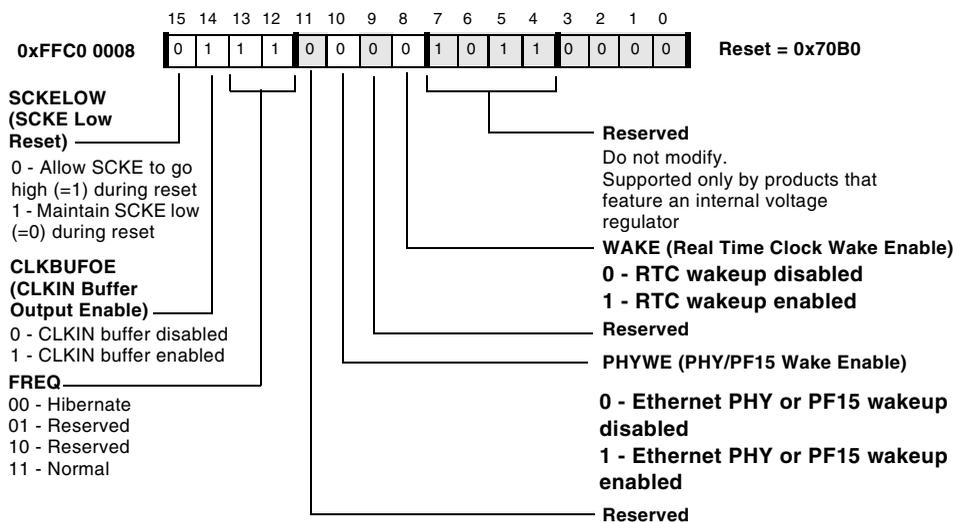


Figure 8-7. Voltage Regulator Control Register

System Control ROM Function

The PLL and voltage regulator registers should not be accessed directly. Instead, use the `bfrom_SysControl()` function to alter or read the register values. The function resides in the on-chip ROM and can be called by the user following C-language style calling conventions.

Entry address: 0xEF00 0038

Arguments:

- `dActionFlags` word in R0
- `pSysCtrlSettings` pointer in R1
- zero value in R2

A potential error message of internally called `bfrom_OtpRead()` function forwarded and returned in R0.



The system control ROM function does not verify the correctness of the forwarded arguments. Therefore, it is up to the programmer to choose the correct values.

C prototype: `u32 bfrom_SysControl(u32 dActionFlags, ADI_SYSCTRL_VALUES *pSysCtrlSettings, void *reserved);`

The first argument (`u32 dActionFlags`) to the system control ROM function holds the instruction flags. The following flags are supported.

```
#define SYSCTRL_READ           0x00000000
#define SYSCTRL_WRITE         0x00000001
#define SYSCTRL_SYSRESET      0x00000002
#define SYSCTRL_SOFTRESET     0x00000004
#define SYSCTRL_VRCTL         0x00000010
#define SYSCTRL_EXTVOLTAGE    0x00000020
#define SYSCTRL_OTPVOLTAGE    0x00000040
#define SYSCTRL_PLLCTL        0x00000100
```

```
#define SYSCTRL_PLLDIV      0x00000200
#define SYSCTRL_LOCKCNT    0x00000400
#define SYSCTRL_PLLSTAT    0x00000800
```

With `SYSCTRL_READ` and `SYSCTRL_WRITE`, a read or a write operation is initialized. The `SYSCTRL_SYSRESET` flag performs a system reset, while the `SYSCTRL_SOFTRESET` flag combines a core and system reset. The `SYSCTRL_EXTVOLTAGE` flag indicates that V_{DDINT} is supplied externally. The `SYSCTRL_OTPVOLTAGE` flag is for factory purposes only. The last five flags (`_VRCTL`, `_PLLCTL`, `_PLLDIV`, `_LOCKCNT`, `_PLLSTAT`) tells the system control ROM function which registers to be written to or read from. Note that `SYSCTRL_PLLSTAT` flag is read-only.

The second argument (`ADI_SYSCTRL_VALUES *pSysCtrlSettings`) to the system control ROM function passes a pointer to a special structure, which has entries for all PLL and voltage regulator registers. It is pre-defined in the `bfrom.h` header file as follows.

```
typedef struct
{
    u16 uwVrCtl;
    u16 uwPllCtl;
    u16 uwPllDiv;
    u16 uwPllLockCnt;
    u16 uwPllStat;
} ADI_SYSCTRL_VALUES;
```

The third argument to the system control ROM function is reserved and should be kept zero (NULL pointer).

The function's return value is described in the following `bfrom_OtpRead()` ROM routine descriptions; whereby a single-bit warning is suppressed.



The system control ROM function executes the correct steps and programming sequence for the Dynamic Power Management System of the Blackfin processor.

System Control ROM Function

Programming Model

The programming model for the system control ROM function in C/C++ and Assembly is described in the following sections.

Accessing the System Control ROM Function in C/C++

To read the `PLL_DIV` and `PLL_CTL` register values, for example, specify the `SYCTRL_READ` instruction flag along with the `SYCTRL_PLLCTL` and `SYCTRL_PLLDIV` register flags. The `bfrom_OtpRead()` function then only updates the `uwPllCtl` and `uwPllDiv` variables:

```
ADI_SYCTRL_VALUES read;  
bfrom_SysControl (SYCTRL_READ | SYCTRL_PLLCTL | SYCTRL_PLLDIV,  
&read, NULL);
```

The `read.uwPllCtl` and `read.uwPllDiv` variables access the `PLL_CTL` and `PLL_DIV` register values, respectively. To update the register values, specify the `SYCTRL_WRITE` instruction flag along with the register flags of those registers that should be modified and have valid data in the respective `ADI_SYCTRL_VALUES` variables:

```
ADI_SYCTRL_VALUES write;  
write.uwPllCtl = 0x0A00;  
write.uwPllDiv = 0x0004;  
bfrom_SysControl (SYCTRL_WRITE | SYCTRL_PLLCTL | SYCTRL_PLLDIV,  
&write, NULL);
```

Accessing the System Control ROM Function in Assembly

The assembler supports C structs, which is required to import the file `bfrom.h`:

```
#include <bfrom.h>
.IMPORT "bfrom.h";
.STRUCT ADI_SYSCTRL_VALUES dpm;
```

You can pre-load the struct:

```
.STRUCT ADI_SYSCTRL_VALUES dpm = { 0x70B0, 0x0A00, 0x0004,
0x0200, 0x00A2 };
```

or load the values dynamically inside the code:

```
P5.H = hi(dpm);
```

```
P5.L = lo(dpm->uwVrCtl);
R7 = 0x70B0 (z);
w[P5] = R7;
```

```
P5.L = lo(dpm->uwP11Ctl);
R7 = 0x0A00 (z);
w[P5] = R7;
```

```
P5.L = lo(dpm->uwP11Div);
R7 = 0x0004 (z);
w[P5] = R7;
```

```
P5.L = lo(dpm->uwP11LockCnt);
R7 = 0x0200 (z);
w[P5] = R0;
```

System Control ROM Function

The function `u32 bfrom_SysControl(u32 dActionFlags, ADI_SYSCTRL_VALUES *pSysCtrlSettings, void *reserved)`; can be accessed by `BFROM_SYSCONTROL`. Following the C/C++ run-time environment conventions, the parameters passed are held by the data registers R0, R1, and R2.

```
/* 10 = sizeof(ADI_SYSCTRL_VALUES). uimm18m4: 18-bit unsigned
field that must be a multiple of 4, with a range of 8 through
262,152 bytes (0x00000 through 0x3FFFC) */
link sizeof(ADI_SYSCTRL_VALUES)+2;
```

```
[--SP] = (R7:0,P5:0);
```

```
/* Allocate at least 12 bytes on the stack for outgoing argu-
ments, even if the function being called requires less than this.
*/
```

```
SP += -12;
```

```
R0 = SYSCTRL_WRITE      |
      SYSCTRL_VRCTL     |
      SYSCTRL_EXTVOLTAGE |
      SYSCTRL_PLLCTL    |
      SYSCTRL_PLLDIV    ;
```

```
R1.H = hi(dpm);
```

```
R1.L = lo(dpm);
```

```
R2 = 0 (z);
```

```
P5.H = hi(BFROM_SYSCONTROL);
```

```
P5.L = lo(BFROM_SYSCONTROL);
```

```
call(P5);
```

```
SP += 12;
```

```
(R7:0,P5:0) = [SP++];
```

```
unlink;
```

```
rts;
```

The processor's internal scratchpad memory can be used as an alternative for taking a C struct. Therefore, the stack/frame pointer must be loaded and passed.

```
/* 10 = sizeof(ADI_SYSCTRL_VALUES). uimm18m4: 18-bit unsigned
field that must be a multiple of 4, with a range of 8 through
262,152 bytes (0x00000 through 0x3FFFC) */
link sizeof(ADI_SYSCTRL_VALUES)+2;

[--SP] = (R7:0,P5:0);

/* Allocate at least 12 bytes on the stack for outgoing argu-
ments, even if the function being called requires less than this.
*/
SP += -12;

R7 = 0;
R7.L = 0x70B0;
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwVrCtl)] = R7;
R7.L = 0x0A00;
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwP11Ctl)] = R7;
R7.L = 0x0004;
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwP11Div)] = R7;
R7.L = 0x0200;
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwP11LockCnt)] = R7;

R0 = SYSCTRL_WRITE      |
      SYSCTRL_VRCTL     |
      SYSCTRL_EXTVOLTAGE |
      SYSCTRL_PLLCTL    |
      SYSCTRL_PLLDIV    ;
```

Programming Examples

```
R1 = FP;
R1 += -sizeof(ADI_SYSCTRL_VALUES);
R2 = 0;

P5.H = hi(BFROM_SYSCONTROL);
P5.L = lo(BFROM_SYSCONTROL);
call(P5);

SP += 12;
(R7:0,P5:0) = [SP++];
unlink;
rts;
```

Programming Examples

The following code examples illustrate how to use the system control ROM function to effect various operating mode transitions.

 The following examples are only meant to demonstrate how to program the PLL registers. Do not assume that the voltages and frequencies shown in the examples are supported by your processor. Instead, check your product's data sheet for supported voltages and frequencies.

Some setup code has been removed for clarity, and the following assumptions are made.

- PLL control (PLL_CTL) register setting: 0x0A00
- PLL divider (PLL_DIV) register setting: 0x0004
- PLL lock count (PLL_LOCKCNT) register setting: 0x0200
- Clock in (CLKIN) frequency: 25 MHz

VCO frequency is 125 MHz, core clock frequency is 125 MHz, and system clock frequency is 31.25 MHz.

- Voltage regulator control (VR_CTL) register setting: 0x70B0
- Logical voltage level (VDDINT) is at 1.20 V

For operating mode transition and voltage regulator examples:

- **C**
 - `#include <blackfin.h>`
 - `#include <bfrom.h>`
- **Assembly**
 - `#include <blackfin.h>`
 - `#include <bfrom.h>`
 - `.IMPORT "bfrom.h";`
 - `#define IMM32(reg,val) reg##.H=hi(val);`
 - `reg##.L=lo(val);`

Full-on Mode to Active Mode and Back

[Listing 8-1](#) and [Listing 8-2](#) provide code for transitioning from the full-on operating mode to active mode in C and Blackfin assembly code, respectively.

Programming Examples

Listing 8-1. Transitioning from Full-on Mode to Active Mode (C)

```
void active(void)
{
    ADI_SYSCTRL_VALUES active;
    bfrom_SysControl(SYSCTRL_READ | SYSCTRL_EXTVOLTAGE |
        SYSCTRL_PLLCTL, &active, NULL);
    active.uwPllCtl |= (BYPASS | PLL_OFF); /* PLL_OFF bit optional */
    bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_EXTVOLTAGE |
        SYSCTRL_PLLCTL, &active, NULL);
    return;
}
```

Listing 8-2. Transitioning from Full-on Mode to Active Mode (ASM)

```
__active:

    link sizeof(ADI_SYSCTRL_VALUES)+2;
    [--SP] = (R7:0,P5:0);
    SP += -12;

    R0 = (SYSCTRL_READ | SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL);
    R1 = FP;
    R1 += -sizeof(ADI_SYSCTRL_VALUES);
    R2 = 0 (z);
    IMM32(P4,BFROM_SYSCONTROL);
    call(P4);

    R0 = w[FP+sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwPllCtl)];
    bitset(R0,bitpos(BYPASS));
    bitset(R0,bitpos(PLL_OFF)); /* optional */
    w[FP+sizeof(ADI_SYSCTRL_VALUES)+offse-
tof(ADI_SYSCTRL_VALUES,uwPllCtl)] = R0;
```

```
R0 = (SYSCTRL_WRITE | SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL);
R1 = FP;
R1 += -sizeof(ADI_SYSCTRL_VALUES);
R2 = 0 (z);
IMM32(P4,BFROM_SYSCONTROL);
call(P4);

SP += 12;
(R7:0,P5:0) = [SP++];
unlink;
rts;

__active.end:
```

To return from active mode (go back to full-on mode), the `BYPASS` bit and the `PLL_OFF` bit must be cleared again, respectively.

Transition to Sleep Mode or Deep Sleep Mode

[Listing 8-3](#) and [Listing 8-4](#) provide code for transitioning from the full-on operating mode to sleep or deep sleep mode in C and Blackfin assembly code, respectively.

Listing 8-3. Transitioning to Sleep Mode or Deep Sleep Mode (C)

```
void sleep(void)
{
    ADI_SYSCTRL_VALUES sleep;
    bfrom_SysControl(SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL |
        SYSCTRL_READ, &sleep, NULL);
    sleep.uwPllCtl |= STOPCK;    /* either: Sleep Mode */
    sleep.uwPllCtl |= PDWN;     /* or: Deep Sleep Mode */
    bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_EXTVOLTAGE |
```

Programming Examples

```
SYSCtrl_PllCtl, &sleep, NULL);  
return;  
}
```

Listing 8-4. Transitioning to Sleep Mode or Deep Sleep Mode (ASM)

```
__sleep:  
  
link sizeof(ADI_SYSCtrl_VALUES)+2;  
[--SP] = (R7:0,P5:0);  
SP += -12;  
  
R0 = (SYSCtrl_READ | SYSCtrl_EXTVOLTAGE |  
SYSCtrl_PllCtl);  
R1 = FP;  
R1 += -sizeof(ADI_SYSCtrl_VALUES);  
R2 = 0 (z);  
IMM32(P4,BFROM_SYSCtrl);  
call(P4);  
  
R0 = w[FP+sizeof(ADI_SYSCtrl_VALUES)+offse-  
tof(ADI_SYSCtrl_VALUES,uwPllCtl)];  
bitset(R0,bitpos(STOPCK)); /* either: Sleep Mode */  
bitset(R0,bitpos(PDWN)); /* or: Deep Sleep Mode */  
w[FP+sizeof(ADI_SYSCtrl_VALUES)+offse-  
tof(ADI_SYSCtrl_VALUES,uwPllCtl)] = R0;  
  
R0 = (SYSCtrl_WRITE | SYSCtrl_EXTVOLTAGE | SYSCtrl_PllCtl);  
R1 = FP;  
R1 += -sizeof(ADI_SYSCtrl_VALUES);  
R2 = 0 (z);  
IMM32(P4,BFROM_SYSCtrl);  
call(P4);
```

```
SP += 12;
(R7:0,P5:0) = [SP++];
unlink;
rts;

__sleep.end:
```

Set Wakeups and Entering Hibernate State

[Listing 8-5](#) and [Listing 8-6](#) provide code for configuring the regulator wakeups (RTC wakeup) and placing the regulator in the hibernate state in C and Blackfin assembly code, respectively.

Listing 8-5. Configuring Regulator Wakeups and Entering Hibernate State (C)

```
void hibernate(void)
{
    ADI_SYSCTRL_VALUES hibernate;
    /* SCKELOW = 1: Enable Drive SCKE Low During Reset */
    /* Protect SDRAM contents during reset after wakeup */
    hibernate.uwVrCtl=SCKELOW |
                    WAKE | /* RTC/Reset Wake-Up Enable */
                    HIBERNATE; /* Powerdown */
    bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_VRCTL |
    SYSCTRL_EXTVOLTAGE, &hibernate, NULL);
    /* Hibernate State: no code executes until wakeup triggers reset
    */
}
```

Programming Examples

Listing 8-6. Configuring Regulator Wakeups and Entering Hibernate State (ASM)

```
__hibernate:

link sizeof(ADI_SYSCCTRL_VALUES)+2;
[--SP] = (R7:0,P5:0);
SP += -12;

cli R6; /* disable interrupts, copy IMASK to R6 */

/* SCKELOW = 1: Enable Drive SCKE Low During Reset */
/* Protect SDRAM contents during reset after wakeup */
R0.L = SCKELOW |
        WAKE      | /* RTC/Reset Wake-Up Enable */
        HIBERNATE ; /* Powerdown */
w[FP+-sizeof(ADI_SYSCCTRL_VALUES)+
offsetof(ADI_SYSCCTRL_VALUES,uwVrCtl)] = R0;

R0 = (SYSCTRL_WRITE | SYSCTRL_VRCTL | SYSCTRL_EXTVOLTAGE);
R1 = FP;
R1 += -sizeof(ADI_SYSCCTRL_VALUES);
R2 = 0 (z);
IMM32(P4,BFROM_SYSCONTROL);
call(P4);

/* Hibernate State: no code executes until wakeup triggers reset
*/

__hibernate.end;
```

Perform a System Reset or Soft-Reset

[Listing 8-7](#) and [Listing 8-8](#) provide code for executing a system reset *or* a soft-reset (system and core reset) in C and Blackfin assembly code, respectively.

Listing 8-7. Execute a System Reset or a Soft-Reset (C)

```
void reset(void)
{
    bfrom_SysControl(SYSCTRL_SYSRESET, NULL, NULL); /* either */
    bfrom_SysControl(SYSCTRL_SOFTRESET, NULL, NULL); /* or */
    return;
}
```

Listing 8-8. Execute a System Reset or a Soft-Reset (ASM)

```
__reset:

    link sizeof(ADI_SYSCTRL_VALUES)+2;
    [--SP] = (R7:0,P5:0);
    SP += -12;

    R0 = SYSCTRL_SYSRESET; /* either */
    R0 = SYSCTRL_SOFTRESET; /* or */
    R1 = 0 (z);
    R2 = 0 (z);
    IMM32(P4,BFROM_SYSCONTROL);
    call(P4);

    SP += 12;
    (R7:0,P5:0) = [SP++];
    unlink;
    rts;
```

```
__reset.end:
```

In Full-on Mode, Change VCO Frequency, Core Clock Frequency, and System Clock Frequency

[Listing 8-9](#) and [Listing 8-10](#) provide C and Blackfin assembly code for changing the CLKIN to VCO multiplier (from 10x to 21x), keeping the CSEL divider at 1, and changing the SSEL divider (from 5 to 4) in the full-on operating mode.

Listing 8-9. Transition of Frequencies (C)

```
void frequency(void)
{
    ADI_SYSCTRL_VALUES frequency;

    /* Set MSEL = 0-63 --> VCO = CLKIN*MSEL */
    frequency.uwPllCtl = SET_MSEL(21) ;

    /* Set SSEL = 1-15 --> SCLK = VCO/SSEL */
    /* CCLK = VCO / 1 */
    frequency.uwPllDiv = SET_SSEL(4) |
                        CSEL_DIV1    ;

    frequency.uwPllLockCnt = 0x0200;

    bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_EXTVOLTAGE |
                    SYSCTRL_PLLCTL | SYSCTRL_PLLDIV | SYSCTRL_LOCKCNT |, &frequency,
                    NULL);
    return;
}
```

Listing 8-10. Transition of Frequencies (ASM)

```

__frequency:

link sizeof(ADI_SYSCTRL_VALUES)+2;
[--SP] = (R7:0,P5:0);
SP += -12;

/* write the struct */
R0 = 0;

R0.L = SET_MSEL(21) ; /* Set MSEL = 0-63 --> VCO = CLKIN*MSEL */
w[FP+sizeof(ADI_SYSCTRL_VALUES)+
offsetof(ADI_SYSCTRL_VALUES,uwP11Ctl)] = R0;

R0.L = SET_SSEL(4) | /* Set SSEL = 1-15 --> SCLK = VCO/SSEL */
      CSEL_DIV1    ; /* CCLK = VCO / 1 */
w[FP+sizeof(ADI_SYSCTRL_VALUES)+
offsetof(ADI_SYSCTRL_VALUES,uwP11Div)] = R0;

R0.L = 0x0200;
w[FP+sizeof(ADI_SYSCTRL_VALUES)+
offsetof(ADI_SYSCTRL_VALUES,uwP11LockCnt)] = R0;

/* argument 1 in R0 */
R0 = (SYSCTRL_WRITE | SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL |
SYSCTRL_PLLDIV);

/* argument 2 in R1: structure lays on local stack */
R1 = FP;
R1 += -sizeof(ADI_SYSCTRL_VALUES);

/* argument 3 must always be NULL */
R2 = 0;

```

Programming Examples

```
/* call of SysControl function */
IMM32(P4,BFROM_SYSCONTROL);
call (P4); /* R0 contains the result from SysControl */

SP += 12;
(R7:0,P5:0) = [SP++];
unlink;
rts;

__frequency.end:
```

Changing Voltage Levels

[Listing 8-11](#) provides C code for changing the voltage level dynamically. The User must include his own code for accessing the external voltage regulator.

Listing 8-11. Changing Core Voltage (C)

```
void voltage(void)
{
    ADI_SYSCTRL_VALUES voltage;
    u32 ulCnt = 0;
    bfrom_SysControl( SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL |
    SYSCTRL_READ, &init, NULL );
    init.uwPllCtl |= BYPASS;
    init.uwPllLockCnt = 0x0200;
    bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_PLLCTL | SYSCTRL_LOCKCNT
    | SYSCTRL_EXTVOLTAGE, &voltage, NULL);
    /* Put your code for accessing the external voltage regulator
    here */
}
```

```
/* A delay loop is required to ensure VDDint is stable and the
PLL has re-locked. As this is depending on the external voltage
regulator circuitry the user must ensure timings are kept. The
compiler (no optimization enabled) will create a loop that takes
about 10 cycles. Time base is CLKIN as the PLL is bypassed. We
need 0x0200 CLKIN cycles that represent PLL_LOCKCNT and addition-
ally the time required by the circuitry */
ulCnt = 0x0200 + 0x0200;
while (ulCnt != 0) {ulCnt--;}
init.uwPl1Ctl &= ~BYPASS;
bfrom_SysControl(SYSCTRL_WRITE | SYSCTRL_PLLCTL |
SYSCTRL_EXTVOLTAGE, &voltage, NULL);
return;
}
```

Programming Examples

9 GENERAL-PURPOSE PORTS

This chapter describes the general-purpose ports. Following an overview and a list of key features is a block diagram of the interface and a description of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Overview

The ADSP-BF51x Blackfin processors feature a rich set of peripherals, which, through a powerful pin multiplexing scheme, provides great flexibility to the external application space.

Table 9-1 shows all the peripheral signals that can be accessed off-chip. ADSP-BF51x processors feature 42 peripheral pins through which all on-chip peripheral are multiplexed.

Table 9-1. General-Purpose and Special Function Signals

Peripheral	Signals
10/100 Ethernet MAC ¹ with IEEE-1588 ²	MII interface (18) or RMII (11), IEEE-1588(3)
RSI interface ³	Data (8), clock (1), command (1)
PWM	Channels (6), sync (1), trip (1)
PPI Interface	Data (16), frame sync (3), clock (1)
SPI Interface	Data (4), clock (2), slave select (2), slave enable (9)
SPORTs	Data (8), clock (4), frame sync (4)

Features

Table 9-1. General-Purpose and Special Function Signals (Continued)

Peripheral	Signals
UARTs	Data (4)
Timers	PWM/capture/clock (8), alternate clock input (4), alternate capture input (7)
General-Purpose I/O	GPIO (40)
Handshake MemDMA	MemDMA request (2)

1 ADSP-BF516 and ADSP-BF518 only.

2 ADSP-BF518 only.

3 ADSP-BF514, ADSP-BF516, and ADSP-BF518 only.

Features

The peripheral pins are functionally organized into general-purpose ports designated port F, port G, and port H.

Port F provides 16 pins:

- MII/RMII signals (ADSP-BF516 and ADSP-BF518 only)
- PPI data signals
- PWM signals
- Primary Timer signals
- Additional SPI0 and SPI1 slave selects
- GPIOs

Port G provides 16 pins:

- SPORT0 signals
- Primary SPI0 signals

- UART0 signals
- RSI signals (ADSP-BF514, ADSP-BF516, and ADSP-BF518 only)
- Handshake memDMA request signals
- PPI Clock and Frame Sync signals
- GPIOs

Port H provides 8 pins:

- SPORT1 signals
- UART1 signals
- SPI1 primary signals
- Up/Down Counter
- Primary Timer signals
- GPIOs

Interface Overview

By default, all port F, port G, and port H pins are in general-purpose I/O (GPIO) mode. Port J does not provide GPIO functionality. In this mode, a pin can function as a digital input, digital output, or interrupt input. See [“General-Purpose I/O Modules” on page 9-14](#) for details. Peripheral functionality must be explicitly enabled by the function enable registers (PORTE_FER, PORTG_FER, and PORTH_FER). The competing peripherals on port F, port G, and port H are controlled by the respective multiplexer control register (PORTE_MUX, PORTG_MUX, PORTH_MUX).

Interface Overview



In this chapter, the naming convention for registers and bits uses a lowercase x to represent F, G, or H. For example, the name `PORT x _FER` represents `PORTF_FER`, `PORTG_FER`, and `PORTH_FER`. The bit name `P x 0` represents `PF0`, `PG0`, and `PH0`. This convention is used to discuss registers common to these three ports.

External Interface

The external interface of the general-purpose ports are described in the following sections.

Port F Structure

[Table 9-2](#) shows the multiplexer scheme for port F. Port F is controlled by the `PORTF_MUX` and the `PORTF_FER` registers.

Port F consists of 16 pins, referred to as `PF0` to `PF15`, as shown in [Table 9-2](#). Besides the 16 GPIOs, this port houses all the PPI data signals (`PPID15-0`) and MII/RMII signals. The PPI signals are multiplexed with PWM signals. With an 8-bit PPI, there is no restriction to use the 4 channels of PWM. All the input signals in the “Additional Use” column are enabled by their module only, regardless of the state of the `PORT x _MUX` and `PORT x _FER` registers.

Any GPIO can be enabled individually and overrides the peripheral function if the respective bit in `PORTF_FER` is cleared.

Table 9-2. Port F Multiplexing Scheme

PORTF_MUX	00	01	10	11		
	1st Function	2nd Function	3rd Function	4th Function	Additional Use	GPIO
Bit[1:0]	MII ETxD2	PPI D0	SPI1SSEL2	-	TACLK6	PF0
Bit[3:2]	MII ERxD2 MII ETxD3 MII ERxD3 MII ERxCLK MII ERxDV MII COL	PPI D1 PPI D2 PPI D3 PPI D4 PPI D5 PPI D6	PWM_AH PWM_AL PWM_BH PWM_BL PWM_CH PWM_CL	- - - - - -	TACLK7 TACLK0 TACLK1 TACI0 TACI1	PF1 PF2 PF3 PF4 PF5 PF6
Bit[5:4]	SPI0SSEL1	PPI D7	PWM_SYNC	-		PF7
Bit[7:6]	RMII MDC RMII MDIO RMII ETxD0	PPI D8 PPI D9 PPI D10	SPI1SSEL4 TMR2 TMR3	- - -		PF8 PF9 PF10
Bit[9:8]	RMII ERxD0 RMII ETxD1 RMII ERxD1 RMII ETxEN	PPI D11 PPI D12 PPI D13 PPI D14	PWM_AH PWM_AL PWM_BH PWM_BL	- - - -	TACI3	PF11 PF12 PF13 PF14
Bit[11:10]	RMII PHYINT	PPI D15	PWM_SYNC	-		PF15



Bits 13-15 in the PORTF_MUX register are reserved. If TMRCLK is used as an input to a GP Timer but the PPI is disabled, then bit 12 of PORTF_MUX should be set.

Port G Structure

Table 9-3 shows the multiplexer scheme for port G. Port G is controlled by the PORTG_MUX and PORTG_FER registers.

Port G consists of 16 pins, referred to as PG0 to PG15, as shown in Table 9-3. Besides the 16 GPIOs, this port houses SPORT0 and SPI0 signals along with the RSI data, clock, and command signals. If a secondary channel on SPORT0 is not required, you can enable UART0 signals or an additional timer.

Interface Overview

Special attention is required for the use of the timers with PPI enabled. Timer0 and Timer1 are typically used for PPI frame sync generation.

Any GPIO can be enabled individually and overrides the peripheral function if the respective bit in the PORTG_FER register is cleared.

Table 9-3. Port G Multiplexing Scheme

PORTG_MUX	00	01	10	11		
	1st Function	2nd Function	3rd Function	4th Function	Additional Use	GPIO
Bit[1:0]	RMII CRS RMII ERxER RMII TxCLK	HWAIT DMAR1 DMAR0	SPI1SSEL3 PWM_CH PWM_CL	- - -		PG0 PG1 PG2
Bit[3:2]	DR0PRI	RSI_DATA0	SPI0SSEL5	-	TACLK3	PG3
Bit[5:4]	RSCLK0	RSI_DATA1	TMR5	-	TACI5	PG4
Bit[7:6]	RFS0 TFS0 DT0PRI	RSI_DATA2 RSI_DATA3 RSI_CMD	PPICLK/TMRCLK TMR0/PPIFS1 TMR1/PPIFS2	- - -		PG5 PG6 PG7
Bit[9:8]	TSCLK0	RSI_CLK	TMR6	-	TACI6	PG8
Bit[11:10]	DT0SEC DR0SEC	UART0 TX UART0 RX	TMR4 PWM_TRIPB	- -	TACI4	PG9 PG10
Bit[13:12]	SPI0 SS	AMS2	SPI1SSEL5	-	TACLK2	PG11
Bit[15:14]	SPI0 SCK SPI0 MISO SPI0 MOSI SPI0SSEL2	PPICLK/TMRCLK TMR0/PPIFS1 TMR1/PPIFS2 PPI FS3	PTP_PPS PTP_CLKOUT PWM_TRIPB AMS3	- - - -	PTP_AUXIN	PG12 PG13 PG14 PG15

Port H Structure

Table 9-4 shows the multiplexer scheme for port H. Port H is controlled by the PORTH_MUX and PORTH_FER registers.

Port H consists of 9 pins. PH0 to PH7 (shown in Table 9-4) are GPIO capable and operate in the same fashion as the Port F and Port G pins. PH8 has limited GPIO capability and connects to the chip enable of the optional internal SPI flash in the ADSP-BF51x package. Port H also houses the

SPORT1 and SPI signals. If a secondary channel on SPORT1 is not required, you can enable UART1 signals, or an additional SPI1 slave enable signal or an additional timer.

Any GPIO can be enabled individually and overrides the peripheral function if the respective bit in the PORTH_FER register is cleared.

Table 9-4. Port H Multiplexing Scheme

PORTH_MUX	00	01	10	11		
	1st Function	2nd Function	3rd Function	4th Function	Additional Use	GPIO
Bit[1:0]	DR1PRI RFS1 RSCLK1 DT1PRI	SPI1 SS SPI1 MISO SPI1 SCK SPI1 MOSI	RSI_DATA4 RSI_DATA5 RSI_DATA6 RSI_DATA7	- - - -		PH0 PH1 PH2 PH3
Bit[3:2]	TFS1 TSCLK1	AOE ARDY	SPI0SSEL3 ECLK	- -	CUD CDG	PH4 PH5
Bit[5:4]	DT1SEC	UART1 TX	SPI1SSEL1	-	CZM	PH6
Bit[7:6]	DR1SEC	UART1 RX	TMR7	-	TACI2	PH7



Bits 8-15 in the PORTH_MUX register are reserved.

Input Tap Considerations

Input taps are shown in [Table 9-2](#), [Table 9-3](#) and [Table 9-4](#) under the “Additional Use” column. When input taps (as well as GPIO based taps) are used with other functionality enabled on the GPIO pins, the signals seen by the input tap modules might be different from what is seen on the pins. This is because different pin functions have different signal requirements with respect to when the signal is latched, if at all. Because of this, input taps multiplexed on certain pins may behave differently than those

Interface Overview

on other pins, depending on which pin function is selected. The input taps will see different signals than at the pins in the following cases:

- All GPIO inputs except PG2, PG4, PG8, PG9, PG11, PG12, PH2, PH6 when GPIO is tapped with `PORTx_FER` set to 1.
- TACLK6 if `PORTF_FER[0] = 1` and `PORTF_MUX[1:0] = b#01`
- TACLK7 if `PORTF_FER[1] = 1` and `PORTF_MUX[3:2] = b#00, b#01`
- TACLK0 if `PORTF_FER[3] = 1` and `PORTF_MUX[3:2] = b#00, b#01`
- TACLK1 if `PORTF_FER[4] = 1` and `PORTF_MUX[3:2] = b#01`
- TACLK3 if `PORTG_FER[3] = 1` and `PORTG_MUX[3:2] = b#00`
- TACI0 if `PORTF_FER[5] = 1` and `PORTF_MUX[3:2] = b#00, b#01`
- TACI1 if `PORTF_FER[6] = 1` and `PORTF_MUX[3:2] = b#01`
- TACI3 if `PORTF_FER[11] = 1` and `PORTF_MUX[9:8] = b#00, b#01`
- TACI4 if `PORTG_FER[10] = 1` and `PORTG_MUX[11:10] = b#00`
- TACI2 if `PORTH_FER[7] = 1` and `PORTH_MUX[7:6] = b#00`
- PTP_AUXIN if `PORTG_FER[14] = 1` and `PORTG_MUX[15:14] = b#00, b#01`
- CUD if `PORTH_FER[4] = 1` and `PORTH_MUX[3:2] = b#00`
- CDG if `PORTH_FER[5] = 1` and `PORTH_MUX[3:2] = b#01`

PWM Unit Considerations

PWM signals that appear in multiple ports, if selected on both, will have inputs and outputs enabled only on PF1–PF7. `PWM_TRIPB` appears twice within Port G: on PG10 and PG14. If both are configured as `PWM_TRIPB` and selected, inputs will only be enabled on PG10.

If `PWM_TRIPB` is not selected on either `PG14` or `PG10`, then the internal `PWM_TRIPB` signal to the PWM module will be driven low. That is, the PWM unit will be tripped if neither of these `PWM_TRIPB` signals is selected via the `PORTG_MUX` register.

RSI Considerations

Pull up/pull down enabling for RSI:

- Pull down for `RSI_DATA[3]` will be enabled only if RSI is selected on `PG6` (that is, `PORTG_MUX[7:6] == b#01`) and the `PD_Dat3` bit is set in the `RSI_CONFIG` register.
- Pull up for `RSI_DATA[3]` will be enabled only if RSI is selected on `PG6` (that is, `PORTG_MUX[7:6] == b#01`) and the `PU_Dat3` bit is set in the `RSI_CONFIG` register.
- Pull up for `RSI_DATA[0]` will be enabled only if `RSI_DATA[0]` is selected on `PG3` (that is, `PORTG_MUX[3:2] == b#01`) and the `PU_Dat` bit is set in the `RSI_CONFIG` register.
- Pull up for `RSI_DATA[1]` will be enabled only if `RSI_DATA[1]` is selected on `PG4` (that is, `PORTG_MUX[5:4] == b#01`) and the `PU_Dat` bit is set in the `RSI_CONFIG` register.
- Pull up for `RSI_DATA[2]` will be enabled only if `RSI_DATA[2]` is selected on `PG5` (that is, `PORTG_MUX[7:6] == b#01`) and the `PU_Dat` bit is set in the `RSI_CONFIG` register.
- Pull up for `RSI_DATA[7:4]` will be enabled only if `RSI_DATA[7:4]` is selected on `PH[3:0]` (that is, `PORTH_MUX[1:0] == b#10`) and the `PU_Dat` bit is set in the `RSI_CONFIG` register.

If `RSI_DATA[3]` is not selected on `PG6` (that is, `PG_MUX[7:6] ≠ b#01`) then the `RSI_DATA[3]` signal to RSI module will be driven low. This is to prevent a spurious card detect interrupt generated by RSI due to data

Interface Overview

toggle on the PG6 pin when it is selected for SPORT/PPI/TMR/GPIO operation.

Internal Interfaces

Port control and GPIO registers are part of the system memory-mapped registers (MMRs). The addresses of the GPIO module MMRs appear in Appendix B. Core access to the GPIO configuration registers is through the system bus.

The PORT_x_MUX registers control the muxing schemes of port F, port G, and port H.

The function enable registers (PORTF_FER, PORTG_FER, PORTH_FER) enable the peripheral functionality for each individual pin of port x.

SPI0 and Internal Flash Usage

PH8 has limited GPIO capability and connects to the chip enable of the optional internal SPI flash in the ADSP-BF51x package.

 On ADSP-BF51x parts, with and without external flash, PORTH_FER[8] must always be set to b#1 for SPI0 to function properly—unless PH8 is used to select the internal SPI flash.

Table 9-5. SPI0 Usage Scenarios

Scenario	PORTH_FER[8]	PORTHIO_DIR[8]	Processors	Additional Information
SPI0 master external access using SPI_FLG	b#1	b#0	ADSP-BF51x and ADSP-BF51xF	
SPI0 master external access using GPIO slave selection	b#1	b#0	ADSP-BF51x and ADSP-BF51xF	

Table 9-5. SPI0 Usage Scenarios (Continued)

Scenario	PORTH_FER[8]	PORTHIO_DIR[8]	Processors	Additional Information
SPI0 slave access	b#1	b#0	ADSP-BF51x and ADSP-BF51xF	
SPI0 master internal flash access using SPI_FLG FLS4	b#1	b#0	ADSP-BF51xF only	
SPI0 master internal flash access using GPIO slave selection	b#0	b#1	ADSP-BF51xF only	Use PORTHIO_SET[8], PORTHIO_CLEAR[8], PORTHIO_TOGGLE[8] to access internal flash

Broadcast mode is restricted for SPI0. SPI0 can broadcast to all slaves except the internal SPI flash (FLS4 in the SPI0_FLG register). Therefore, in broadcast mode, FLS4 of SPI_FLG register should be set to b#0 and all other slave selects (FLS1–3 and FLS5) may be set to b#1.

GP Timer Interaction With Other Blocks

The TACLK_x and TACI_x inputs of the GP Timers connect to several different subsystems of the ADSP-BF51x processor. Following are the details of these connections.

Buffered CLKIN (CLKBUF)

TACLK5 and TACLK4 connect internally to the CLKBUF pin

GP Counter

TACI7 connects to the COUNTER0 TO output internally.

Interface Overview

PPI

TMR0 is internally looped back to PPI_FS1 (to be used as internally generated frame sync). In this case, PPI_CLK is the clock input for the Timer0 module.

TMR1 is internally looped back to PPI_FS2 (to be used as internally generated frame sync) In this case, PPI_CLK is the clock input for the Timer1 module.

PPI_CLK/TMRCLK can be used as a clock input for any of the timers. If TMRCLK is used as an input to a GP Timer but the PPI is disabled, then bit 12 of PORTF_MUX should be set.

PPI/TMR signals (PPICLK/TMRCLK, TMR0/PPIFS1, TMR1/PPIFS2) that appear in multiple ports, if selected on both, will have inputs and outputs enabled only on PG12–PG14.

UART

TACI4 can be used for autobaud detection of UART0 RX.

TACI2 can be used for autobaud detection of UART1 RX.

SPORT

If TMR5 is configured as an output and PORTG_MUX[5:4] == b#10 and SPORT0's RSCLK0 input enable is active, then TMR5 is the clock input for RSCLK0.

If TMR6 is configured as an output and PORTG_MUX[9:8] == b#10, and SPORT0's TSCLK0 input enable is active, then TMR6 is the clock input for TSCLK0.

If SPORT0's RSCLK0 is configured as an output and PORTG_MUX[5:4] == b#00 and TMR5 input enable is active, then RSCLK0 is the clock input for TMR5.

If SPORT0's TSCLK0 is configured as an output and PORTG_MUX[9:8] == b#00 and TMR6 input enable is active, then TSCLK0 is the clock input for TMR6.

If TAC15 is selected in the TMR5 module, then the signal from the PG4 pin is fed to both SPORT0's RSCLK0 and TAC15.

If TAC16 is selected in the TMR6 module, then the signal from the PG8 pin is fed to both SPORT0's TSCLK0 and TAC16.

Performance/Throughput

The PFX, PGx, and PHx pins are synchronized to the system clock (SCLK). When configured as outputs, the GPIOs 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 port 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 signal is asserted on the pin 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.

Description of Operation

The operation of the general-purpose ports is described in the following sections.

Description of Operation

Operation

The GPIO pins on port F, port G, and port H can be controlled individually by the function enable registers (`PORTx_FER`). With a control bit in these registers cleared, the peripheral function is fully decoupled from the pin. It functions as a GPIO pin only. To drive the pin in GPIO output mode, set the respective direction bit in the `PORTxIO_DIR` register. To make the pin a digital input or interrupt input, enable its input driver in the `PORTxIO_INEN` register.

 By default all peripheral pins are configured as inputs after reset. port F, port G, and port H pins are in GPIO mode. However, GPIO input drivers are disabled to minimize power consumption and any need of external pulling resistors.

When the control bit in the function enable registers (`PORTx_FER`) is set, the pin is set to its peripheral functionality and is no longer controlled by the GPIO module. However, the GPIO module can still sense the state of the pin. When using a particular peripheral interface, pins required for the peripheral must be individually enabled. Keep the related function enable bit cleared if a signal provided by the peripheral is not required by your application. This allows it to be used in GPIO mode.

General-Purpose I/O Modules

The processor supports 40 bidirectional or general-purpose I/O (GPIO) signals. These 40 GPIOs are managed by three different GPIO modules, which are functionally identical. One is associated with port F, one with port G, and one with port H. Port F and port G each consist of 16 GPIOs (`PF15-0` and `PG15-0`), respectively. Port H consists of eight GPIOs (`PH7-0`).

Each GPIO can be individually configured as either an input or an output by using the GPIO direction registers (`PORTxIO_DIR`).

When configured as output, the GPIO data registers (`PORTFIO`, `PORTGIO`, and `PORTHIO`) can be directly written to specify the state of the GPIOs.

The GPIO direction registers are read-write registers with each bit position corresponding to a particular GPIO. A logic 1 configures a GPIO as an output, driving the state contained in the GPIO data register if the peripheral function is not enabled by the function enable registers. A logic 0 configures a GPIO as an input.

 Note when using the GPIO as an input, the corresponding bit should also be set in the GPIO input enable register. Otherwise, changes at the input pins will not be recognized by the processor.

The GPIO input enable registers (`PORTFIO_INEN`, `PORTGIO_INEN`, and `PORTHIO_INEN`) are used to enable the input buffers on any GPIO that is being used as an input. Leaving the input buffer disabled eliminates the need for pull-ups and pull-downs when a particular `PFx`, `PGx`, or `PHx` pin is not used in the system. By default, the input buffers are disabled.

 Once the input driver of a GPIO pin is enabled, the GPIO is not allowed to operate as an output anymore. Never enable the input driver (by setting `PORTxIO_INEN` bits) and the output driver (by setting `PORTxIO_DIR` bits) for the same GPIO.

A write operation to any of the GPIO data registers sets the value of all GPIOs in this port that are configured as outputs. GPIOs configured as inputs ignore the written value. A read operation returns the state of the GPIOs defined as outputs and the sense of the inputs, based on the polarity and sensitivity settings, if their input buffers are enabled. [Table 9-6](#) helps to interpret read values in GPIO mode, based on the settings of the `PORTxIO_POLAR`, `PORTxIO_EDGE`, and `PORTxIO_BOTH` registers.

Description of Operation

Table 9-6. GPIO Value Register Pin Interpretation

POLAR	EDGE	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 GPIOs 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 GPIOs, the pin state is checked every cycle, so the readback value will change when the original level on the pin changes.

The state of the output is reflected on the associated pin only if the function enable bit in the `PORTx_FER` register is cleared.

Write operations to the GPIO data registers modify the state of all GPIOs of a port. In cases where only one or a few GPIOs need to be changed, the user may write to the GPIO set registers, `PORTxIO_SET`, the GPIO clear registers, `PORTxIO_CLEAR`, or to the GPIO toggle registers, `PORTxIO_TOGGLE` instead.

While a direct write to a GPIO data register alters all bits in the register, writes to a GPIO set register can be used to set a single or a few bits only. No read-modify-write operations are required. The GPIO set registers are write-1-to-set registers. All 1s contained in the value written to a GPIO set register sets the respective bits in the GPIO data register. The 0s have no effect. For example, assume that `PF0` is configured as an output. Writing `0x0001` to the GPIO set register drives a logic 1 on the `PF0` pin without

affecting the state of any other PF_x pins. The GPIO set registers are typically also used to generate GPIO interrupts by software. Read operations from the GPIO set registers return the content of the GPIO data registers.

The GPIO clear registers provide an alternative port to manipulate the GPIO data registers. While a direct write to a GPIO data register alters all bits in the register, writes to a GPIO clear register can be used to clear individual bits only. No read-modify-write operations are required. The clear registers are write-1-to-clear registers. All 1s contained in the value written to the GPIO clear register clears the respective bits in the GPIO data register. The 0s have no effect. For example, assume that PF_4 and PF_5 are configured as outputs. Writing 0x0030 to the `PORTFIO_CLEAR` register drives a logic 0 on the PF_4 and PF_5 pins without affecting the state of any other PF_x pins.



If an edge-sensitive pin generates an interrupt request, the service routine must acknowledge the request by clearing the respective GPIO latch. This is usually performed through the clear registers.

Read operations from the GPIO clear registers return the content of the GPIO data registers.

The GPIO toggle registers provide an alternative port to manipulate the GPIO data registers. While a direct write to a GPIO data register alters all bits in the register, writes to a toggle register can be used to toggle individual bits. No read-modify-write operations are required. The GPIO toggle registers are write-1-to-toggle registers. All 1s contained in the value written to a GPIO toggle register toggle the respective bits in the GPIO data register. The 0s have no effect. For example, assume that PG_1 is configured as an output. Writing 0x0002 to the `PORTGIO_TOGGLE` register changes the pin state (from logic 0 to logic 1, or from logic 1 to logic 0) on the PG_1 pin without affecting the state of any other PG_x pins. Read operations from the GPIO toggle registers return the content of the GPIO data registers.

The state of the GPIOs can be read through any of these data, set, clear, or toggle registers. However, the returned value reflects the state of the input

Description of Operation

pin only if the proper input enable bit in the `PORTxIO_INEN` register is set. Note that GPIOs can still sense the state of the pin when the function enable bits in the `PORTx_FER` registers are set.

Since function enable registers and GPIO input enable registers reset to zero, no external pull-ups or pull-downs are required on the unused pins of port F, port G, and port H.

GPIO Interrupt Processing

Each GPIO can be configured to generate an interrupt. The processor can sense up to 40 asynchronous off-chip signals, requesting interrupts through five interrupt channels. To make a pin function as an interrupt pin, the associated input enable bit in the `PORTxIO_INEN` register must be set. The function enable bit in the `PORTx_FER` register is typically cleared. Then, an interrupt request 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 GPIO polarity registers (`PORTFIO_POLAR`, `PORTGIO_POLAR`, and `PORTHIO_POLAR`), and the GPIO interrupt sensitivity registers (`PORTFIO_EDGE`, `PORTGIO_EDGE`, and `PORTHIO_EDGE`). If configured for edge sensitivity, the GPIO set on both edges registers (`PORTFIO_BOTH`, `PORTGIO_BOTH`, and `PORTHIO_BOTH`) let the interrupt request generate on both edges.

The GPIO polarity registers are used to configure the polarity of the GPIO input source. To select active high or rising edge, set the bits in the GPIO polarity register to 0. To select active low or falling edge, set the bits in the GPIO polarity register to 1. This register has no effect on GPIOs that are defined as outputs. The contents of the GPIO polarity registers are cleared at reset, defaulting to active high polarity.

The GPIO interrupt sensitivity registers are used to configure each of the inputs 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. The GPIO interrupt sensitivity register has no effect on GPIOs that are defined as outputs. The contents of the GPIO interrupt sensitivity registers are cleared at reset, defaulting to level sensitivity.

The GPIO set on both edges registers are used to enable interrupt generation on both rising and falling edges. When a given GPIO has been set to edge-sensitive in the GPIO interrupt sensitivity register, setting the respective bit in the GPIO 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 GPIOs that are defined as level-sensitive or as outputs. See [Table 9-6](#) for information on how the GPIO set on both edges register interacts with the GPIO polarity and GPIO interrupt sensitivity registers.

When the GPIO's input drivers are enabled while the GPIO direction registers configure it as an output, software can trigger a GPIO interrupt by writing to the data/set/toggle registers. The interrupt service routine should clear the GPIO to acknowledge the request.

Each of the three GPIO modules provides two independent interrupt channels. Identical in functionality, these are called interrupt A and interrupt B. Both interrupt channels have their own mask register which lets you assign the individual GPIOs to none, either, or both interrupt channels.

Since all mask registers reset to zero, none of the GPIOs is assigned any interrupt by default. Each GPIO represents a bit in each of these registers. Setting a bit means enabling the interrupt on this channel.

Interrupt A and interrupt B operate independently. For example, writing 1 to a bit in the mask interrupt A register does not affect interrupt channel B. This facility allows GPIOs to generate GPIO interrupt A, GPIO interrupt B, both GPIO interrupts A and B, or neither.

A GPIO interrupt is generated by a logical OR of all unmasked GPIOs for that interrupt. For example, if PF0 and PF1 are both unmasked for GPIO

Description of Operation

interrupt channel A, GPIO interrupt A will be generated when triggered by PF0 or PF1. The interrupt service routine must evaluate the GPIO data register to determine the signaling interrupt source. [Figure 9-1](#) illustrates the interrupt flow of any GPIO module's interrupt A channel.

 When using either rising or falling edge-triggered interrupts, the interrupt condition must be cleared each time a corresponding interrupt is serviced by writing 1 to the appropriate bit in the GPIO clear register.

At reset, all interrupts are masked and disabled.

Similarly to the GPIOs themselves, the mask register can either be written through the GPIO mask data registers (PORTxIO_MASKA, PORTxIO_MASKB) or be controlled by the mask A/mask B set, clear and toggle registers.

The GPIO mask interrupt set registers (PORTxIO_MASKA_SET, PORTxIO_MASKB_SET) provide an alternative port to manipulate the GPIO mask interrupt registers. While a direct write to a mask interrupt register alters all bits in the register, writes to a mask interrupt set register can be used to set a single or a few bits only. No read-modify-write operations are required.

The mask interrupt set registers are write-1-to-set registers. All ones contained in the value written to the mask interrupt set register set the respective bits in the mask interrupt register. The zeroes have no effect. Writing a one to any bit enables the interrupt for the respective GPIO.

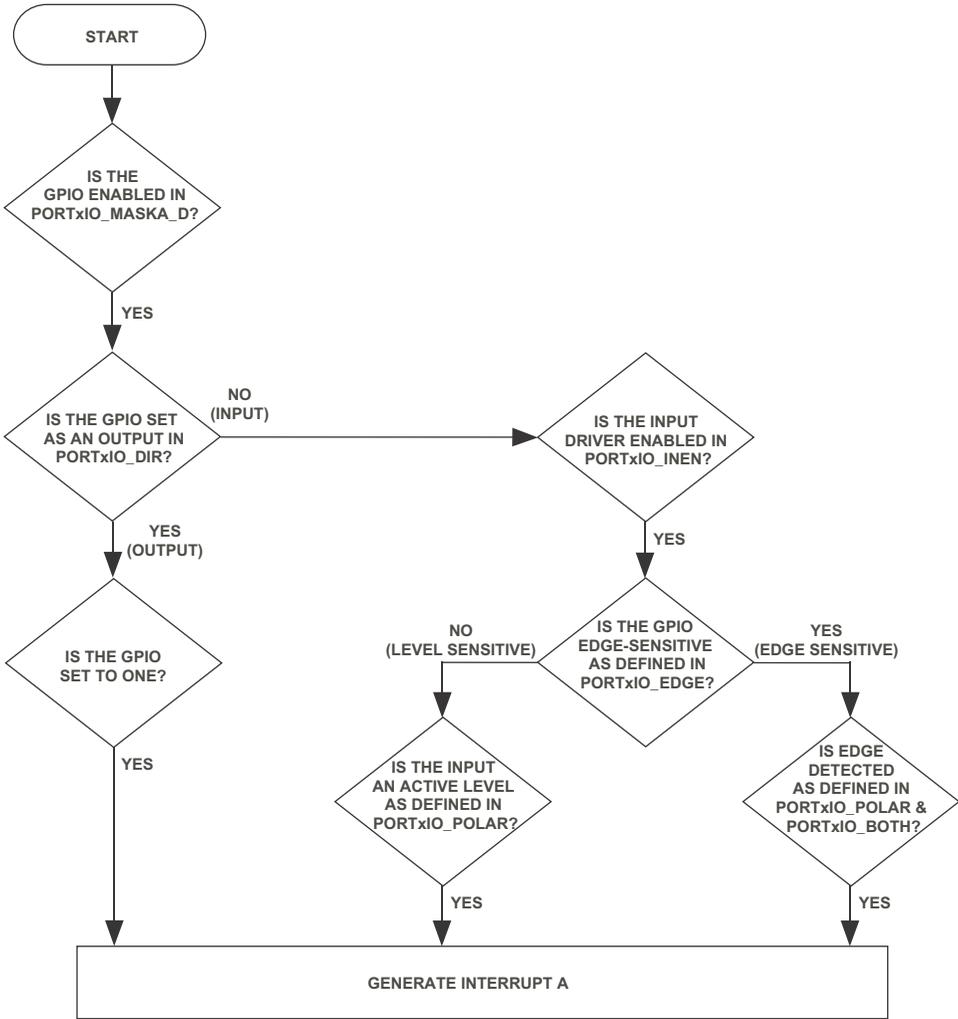


Figure 9-1. GPIO Interrupt Generation Flow for Interrupt Channel A

The GPIO mask interrupt clear registers (PORTxIO_MASKA_CLEAR, PORTxIO_MASKB_CLEAR) provide an alternative port to manipulate the GPIO mask interrupt registers. While a direct write to a mask interrupt register alters all bits in the register, writes to the mask interrupt clear

Description of Operation

register can be used to clear a single bit or a few bits only. No read-modify-write operations are required.

The mask interrupt clear registers are write-1-to-clear registers. All ones contained in the value written to the mask interrupt clear register clear the respective bits in the mask interrupt register. The zeroes have no effect. Writing a one to any bit disables the interrupt for the respective GPIO.

The GPIO mask interrupt toggle registers (`PORTxIO_MASKA_TOGGLE`, `PORTxIO_MASKB_TOGGLE`) provide an alternative port to manipulate the GPIO mask interrupt registers. While a direct write to a mask interrupt register alters all bits in the register, writes to a mask interrupt toggle register can be used to toggle a single bit or a few bits only. No read-modify-write operations are required.

The mask interrupt toggle registers are write-1-to-clear registers. All ones contained in the value written to the mask interrupt toggle register toggle the respective bits in the mask interrupt register. The zeroes have no effect. Writing a one to any bit toggles the interrupt for the respective GPIO.

[Figure 9-1](#) illustrates the interrupt flow of any GPIO module's interrupt A channel. The interrupt B channel behaves identically.

All GPIOs assigned to the same interrupt channel are OR'ed. If multiple GPIOs are assigned to the same interrupt channel, it is up to the interrupt service routine to evaluate the GPIO data registers to determine the signaling interrupt source.

Although each GPIO module provides two independent interrupt channels, the interrupt A channels of port F and port G are OR'ed as shown in [Figure 9-2](#). Therefore the total number of GPIO interrupt channels is five.

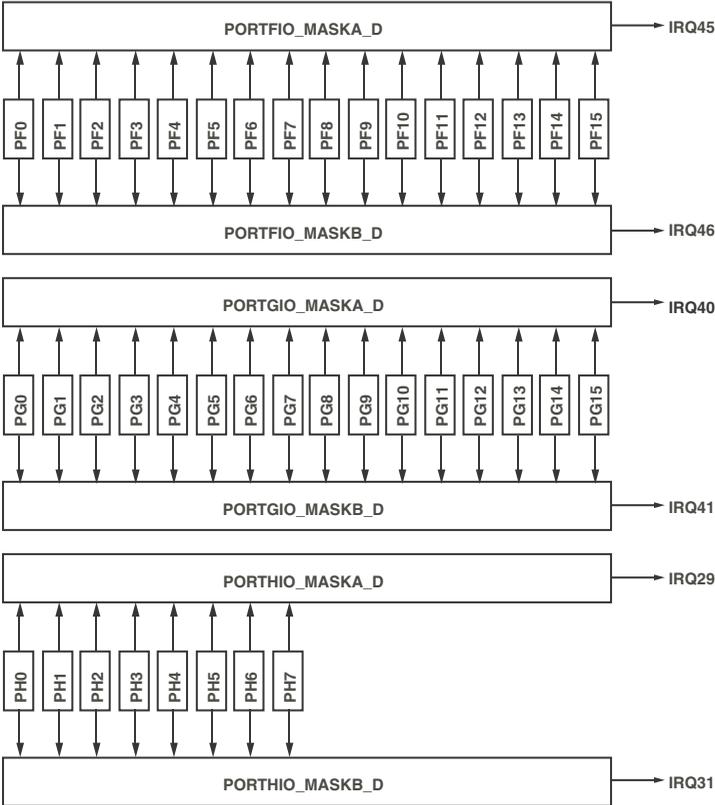


Figure 9-2. GPIO Interrupt Channels

Programming Model

Figure 9-3 and Figure 9-4 show the programming model for the general-purpose ports.

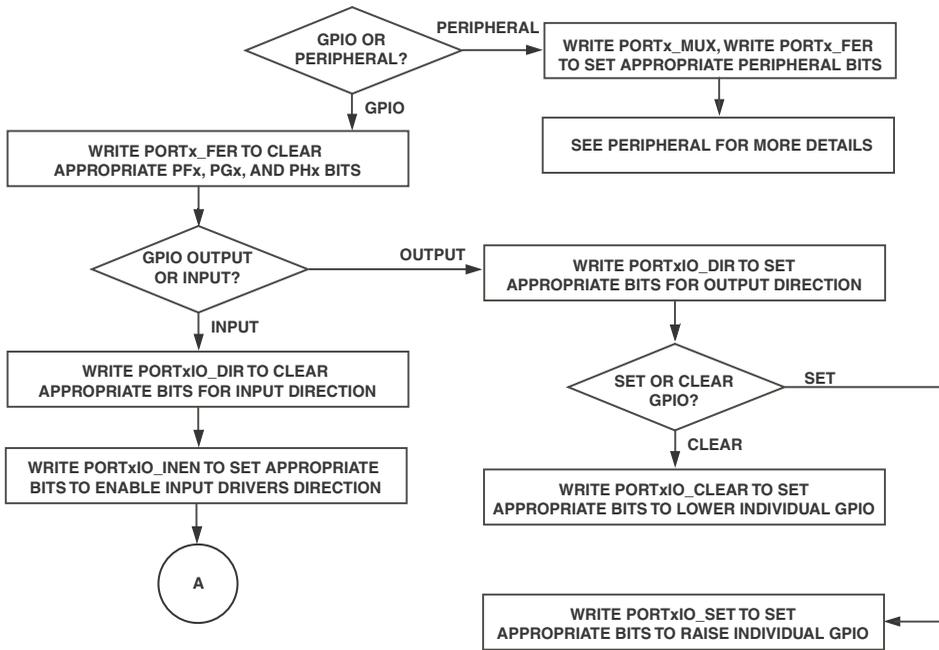


Figure 9-3. GPIO Flow Chart (Part 1 of 2)

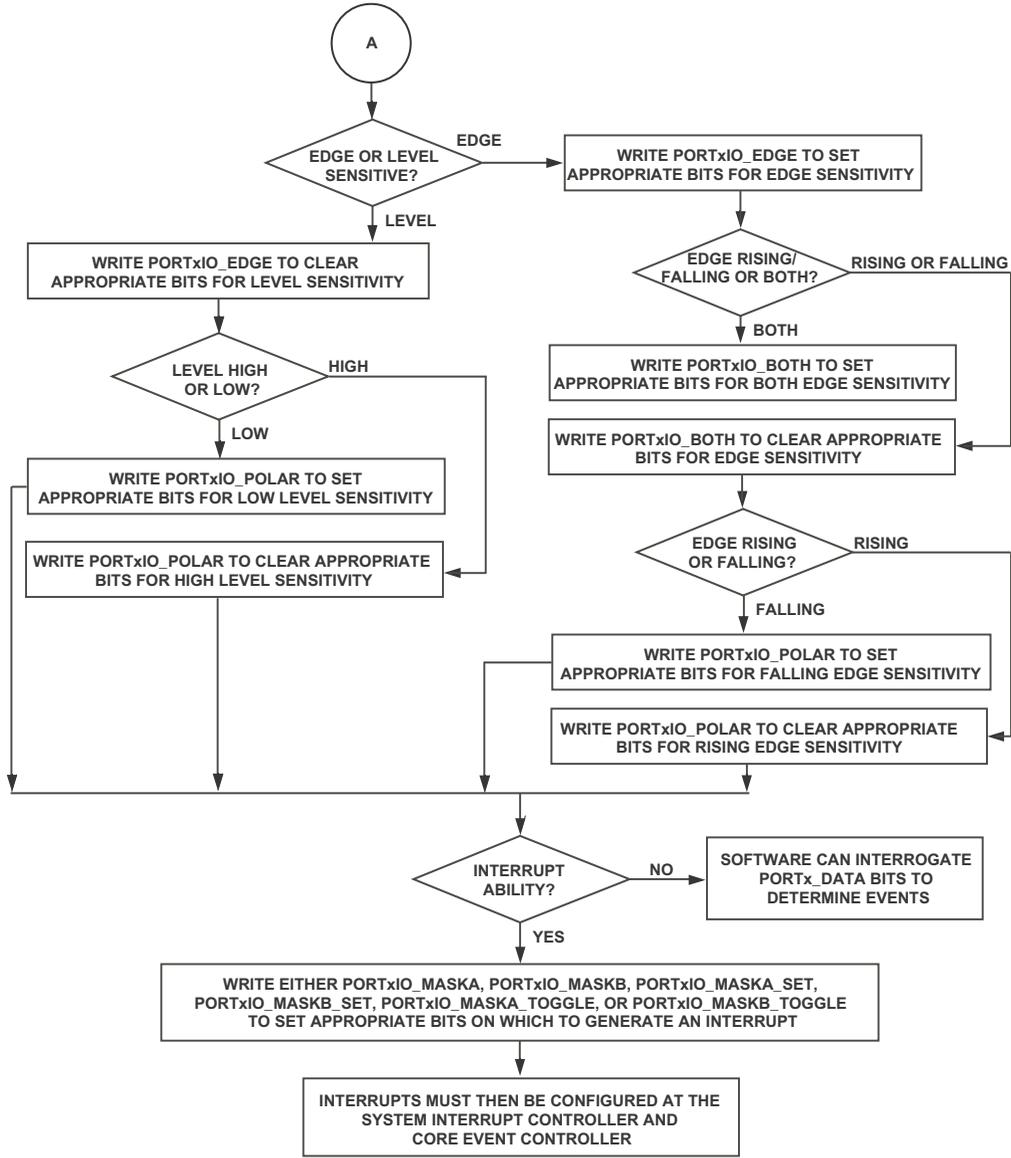


Figure 9-4. GPIO Flow Chart (Part 2 of 2)

Memory-Mapped GPIO Registers

The GPIO registers are part of the system memory-mapped registers (MMRs). [Figure 9-10](#) through [Figure 9-30](#) on [page 9-44](#) illustrate the GPIO registers. The addresses of the programmable flag MMRs appear in [Appendix B](#).

 In [Figure 9-10](#) through [Figure 9-30](#), bits 8-15 are reserved for Port H register descriptions.

PORTx Hysteresis Control (PORTx_HYSTERESIS) Register

The ADSP-BF51x contains additional registers controlling the hysteresis (via Schmitt triggering) for Port F, Port G and Port H. These are also included for several pins other than GPIOs. [Figure 9-5](#) to [Figure 9-7](#) show the bit descriptions of these registers.

This register configures Schmitt triggering (SE) for the PORTx inputs. The Schmitt trigger can be set only for pin groups, classified by the pin muxing controls. For each controlled group of pins, b#00 will disable Schmitt triggering, while b#01 will enable it. Combinations of b#1x are reserved.

Port F Hysteresis Register (PORTF_HYSTERESIS)

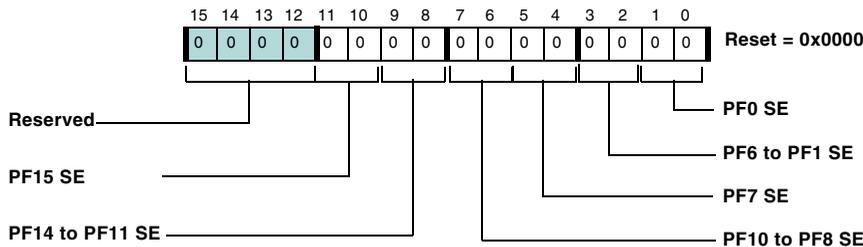


Figure 9-5. Port F Hysteresis Register

Port G Hysteresis Register (PORTG_HYSTERESIS)

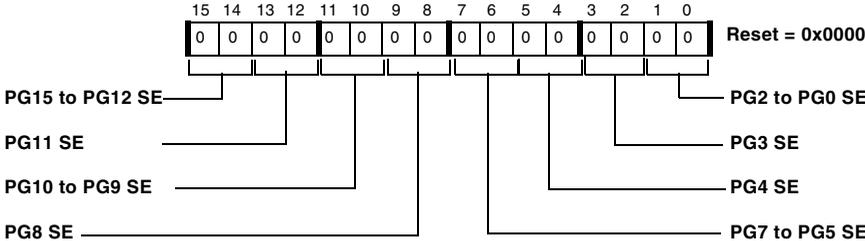


Figure 9-6. Port G Hysteresis Register

Port H Hysteresis Register (PORTH_HYSTERESIS)

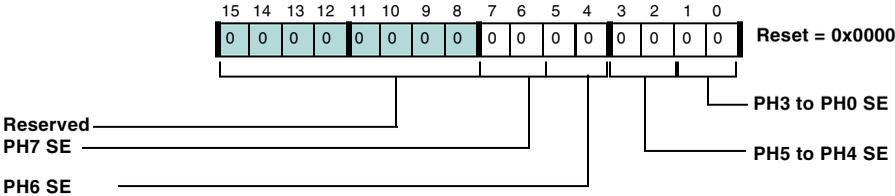


Figure 9-7. Port H Hysteresis Register

Memory-Mapped GPIO Registers

Non-GPIO Drive Strength Control Register

This register sets the drive strength and tolerance for the TWI signals on the ADSP-BF51x as specified in the diagram.

Non-GPIO Ports Drive Strength Control Register (NONGPIO_DRIVE)

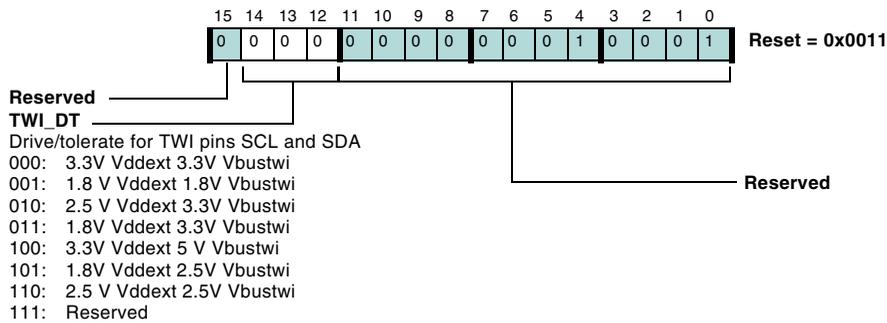


Figure 9-8. Non-GPIO Drive Strength Control Register

Non-GPIO Hysteresis (NONGPIO_HYSTERESIS) Register

This register sets the Schmitt trigger (SE) for various ADSP-BF51x signals. The bits relating to SPI flash (on parts with on-chip SPI flash) are intended for power conservation where possible.

Non-GPIO Hysteresis Register (NONGPIO_HYSTERESIS)

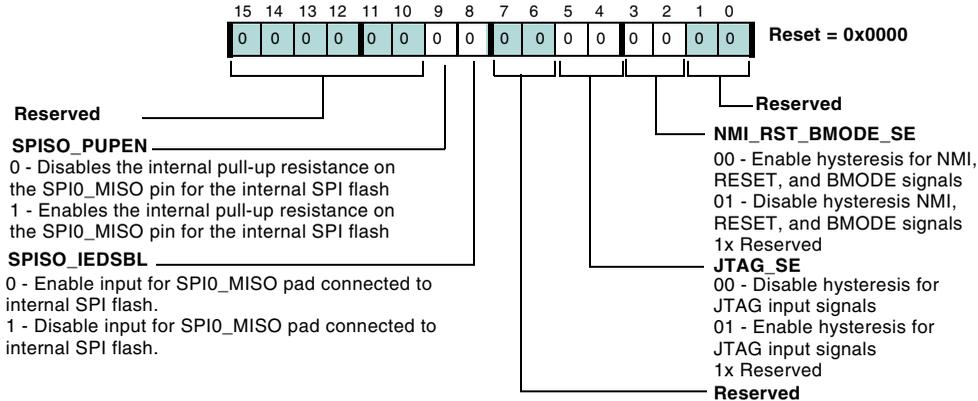
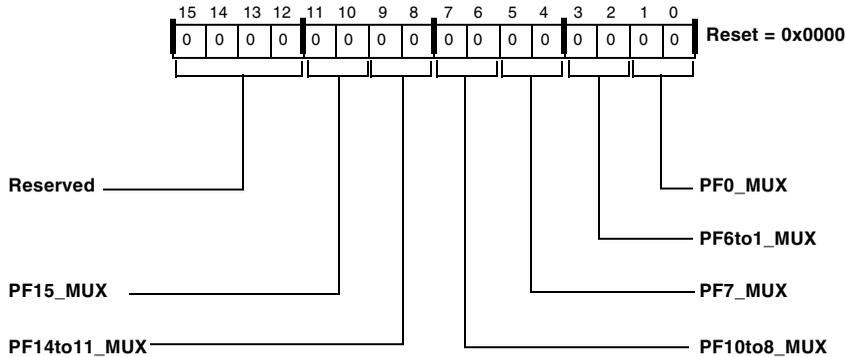


Figure 9-9. Non-GPIO Hysteresis Register

Port Multiplexer Control Register (PORTF_MUX)

Port F Multiplexer Control Register (PORTF_MUX)



For all bit fields:

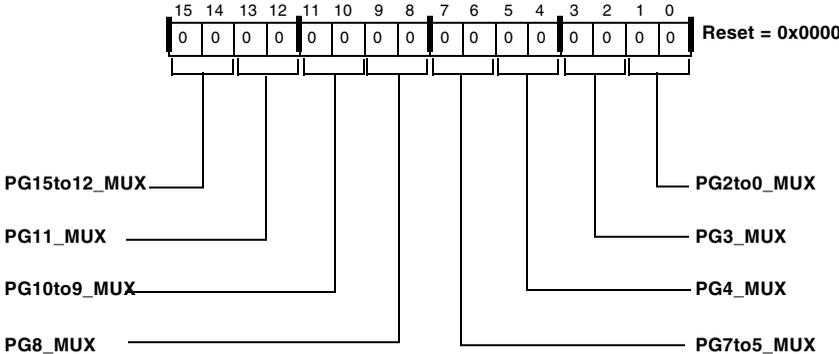
- 00 = 1st Peripheral function
- 01 = 1st alternate peripheral function
- 10 = 2nd alternate peripheral function
- 11 = Reserved

Refer to [Table 9-2 on page 9-5](#) to [Table 9-4 on page 9-7](#) for reserved bits in the PORTF_MUX register.

Figure 9-10. Port F Multiplexer Control Register

Port Multiplexer Control Register (PORTG_MUX)

Port G Multiplexer Control Register (PORTG_MUX)



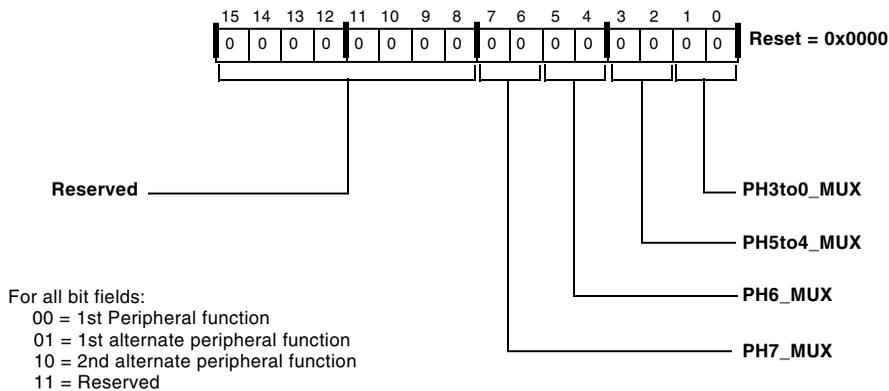
For all bit fields:
00 = 1st Peripheral function
01 = 1st alternate peripheral function
10 = 2nd alternate peripheral function
11 = Reserved

Refer to [Table 9-2 on page 9-5](#) to [Table 9-4 on page 9-7](#) for reserved bits in the PORTG_MUX register.

Figure 9-11. Port G Multiplexer Control Register

Port Multiplexer Control Register (PORTH_MUX)

Port H Multiplexer Control Register (PORTH_MUX)



Refer to [Table 9-2 on page 9-5](#) to [Table 9-4 on page 9-7](#) for reserved bits in the PORTH_MUX register.

Figure 9-12. Port H Multiplexer Control Register

Function Enable Registers (PORTx_FER)

Function Enable Registers (PORTx_FER)

For all bits, 0 - GPIO mode, 1 - Enable peripheral function

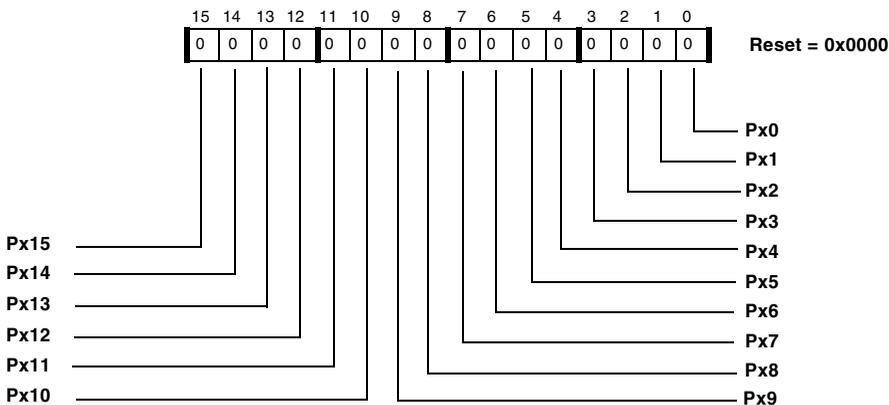


Figure 9-13. Function Enable Registers

GPIO Direction Registers (PORTxIO_DIR)

GPIO Direction Registers (PORTxIO_DIR)

For all bits, 0 - Input, 1 - Output

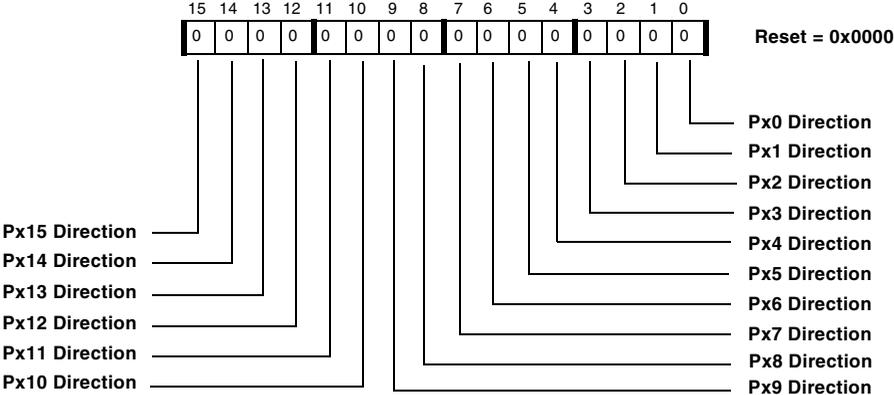


Figure 9-14. GPIO Direction Registers

GPIO Input Enable Registers (PORTxIO_INEN)

GPIO Input Enable Registers (PORTxIO_INEN)

For all bits, 0 - Input Buffer Disabled, 1 - Input Buffer Enabled

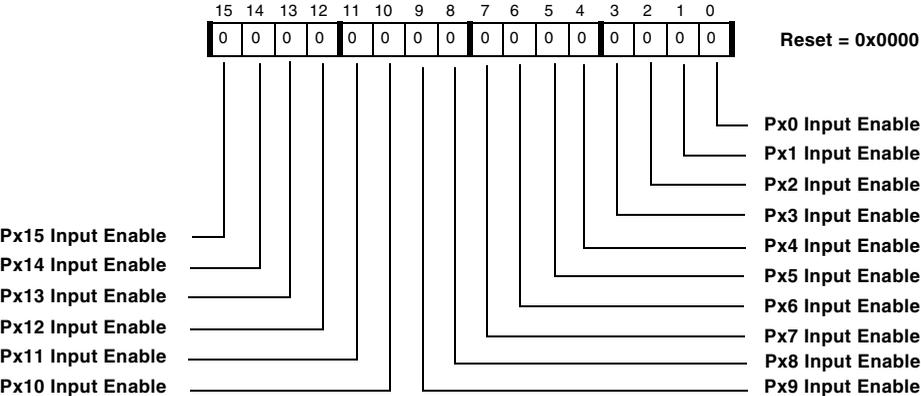


Figure 9-15. GPIO Input Enable Registers

GPIO Data Registers (PORTxIO)

GPIO Data Registers (PORTxIO)

1 - Set, 0 - Clear

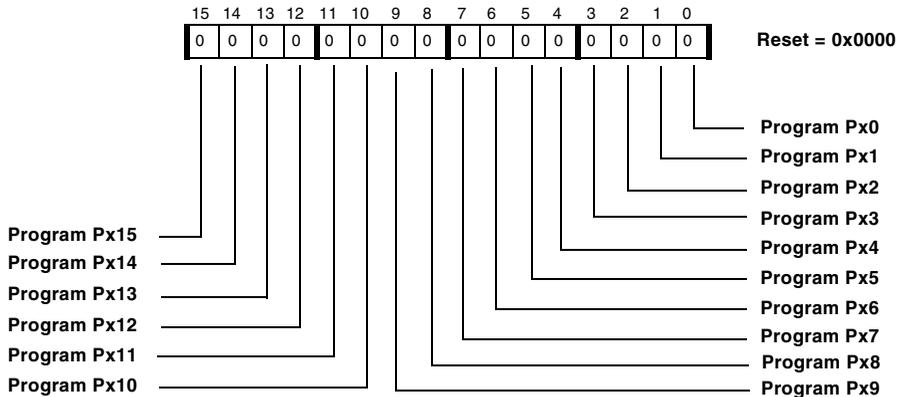


Figure 9-16. GPIO Data Registers

GPIO Set Registers (PORTxIO_SET)

GPIO Set Registers (PORTxIO_SET)

Write-1-to-set

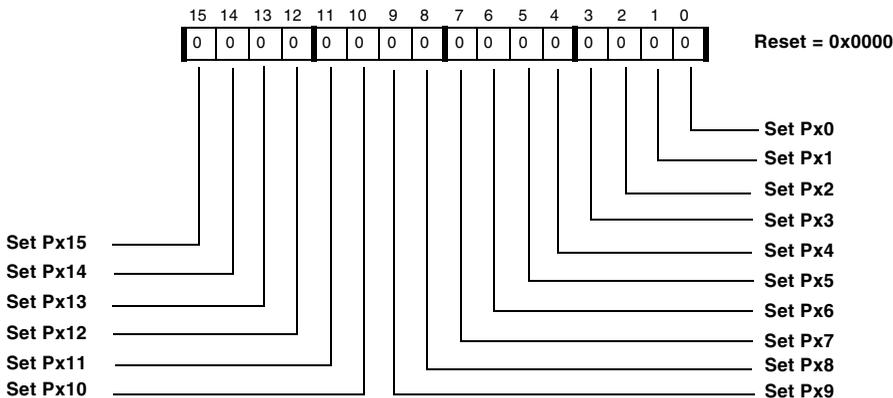


Figure 9-17. GPIO Set Registers

GPIO Clear Registers (PORTxIO_CLEAR)

GPIO Clear Registers (PORTxIO_CLEAR)

Write-1-to-clear

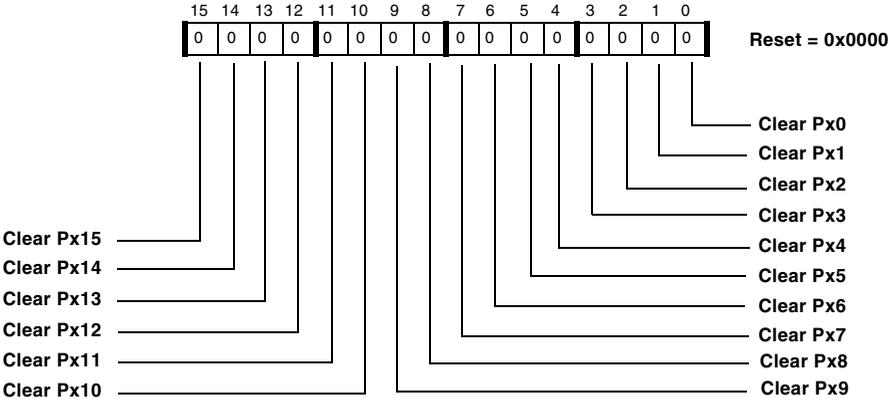


Figure 9-18. GPIO Clear Registers

GPIO Toggle Registers (PORTxIO_TOGGLE)

GPIO Toggle Registers (PORTxIO_TOGGLE)

Write-1-to-toggle

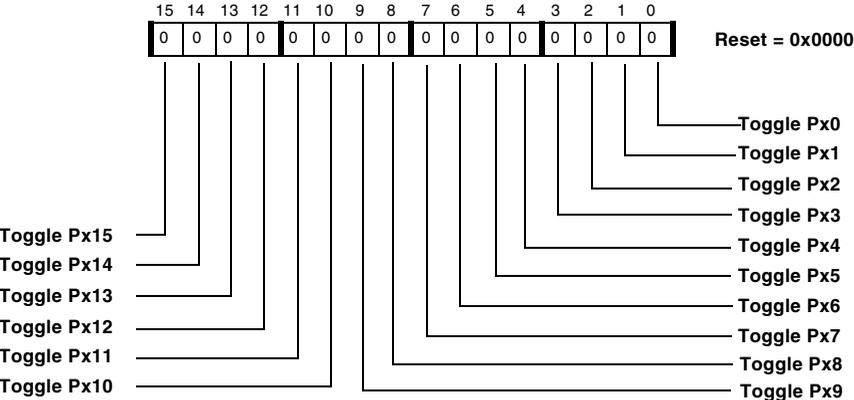


Figure 9-19. GPIO Toggle Registers

GPIO Polarity Registers (PORTxIO_POLAR)

GPIO Polarity Registers (PORTxIO_POLAR)

For all bits, 0 - Active high or rising edge, 1 - Active low or falling edge

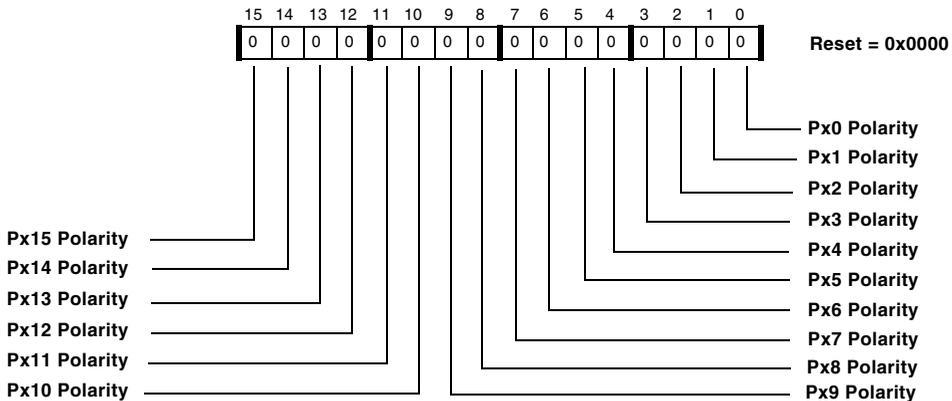


Figure 9-20. GPIO Polarity Registers

Interrupt Sensitivity Registers (PORTxIO_EDGE)

Interrupt Sensitivity Registers (PORTxIO_EDGE)

For all bits, 0 - Level, 1 - Edge

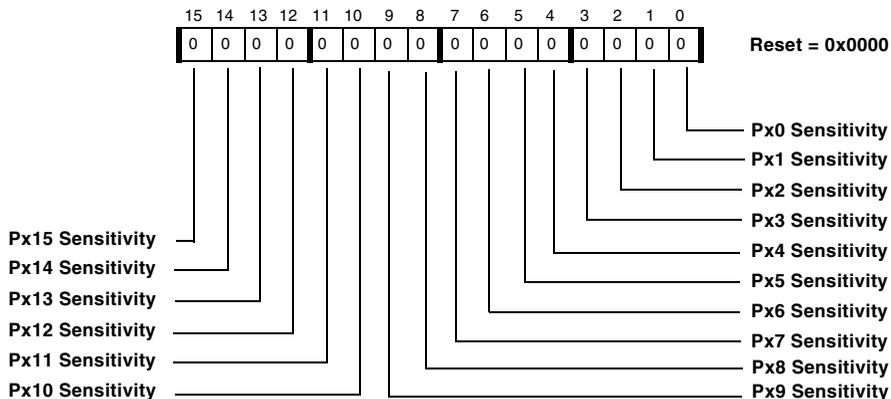


Figure 9-21. Interrupt Sensitivity Registers

GPIO Set on Both Edges Registers (PORTxIO_BOTH)

GPIO Set on Both Edges Registers (PORTxIO_BOTH)

For all bits when enabled for edge-sensitivity, 0 - Single edge, 1 - Both edges

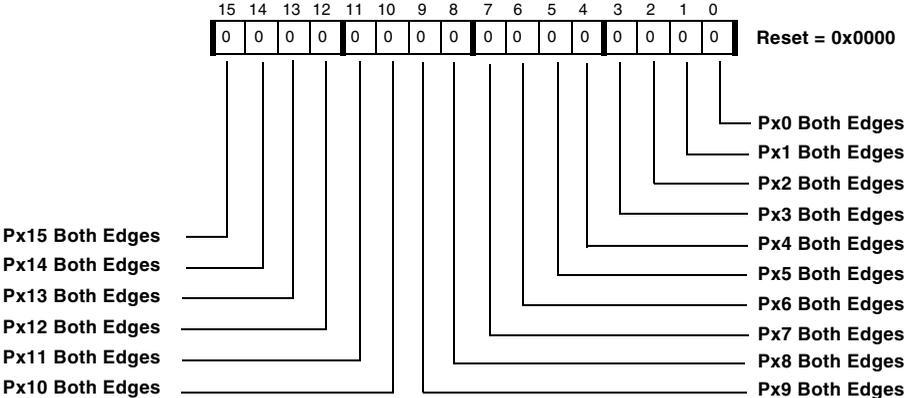


Figure 9-22. GPIO Set on Both Edges Registers

GPIO Mask Interrupt Registers (PORTxIO_MASKA/B)

GPIO Mask Interrupt A Registers (PORTxIO_MASKA)

For all bits, 1 - Enable, 0 - Disable

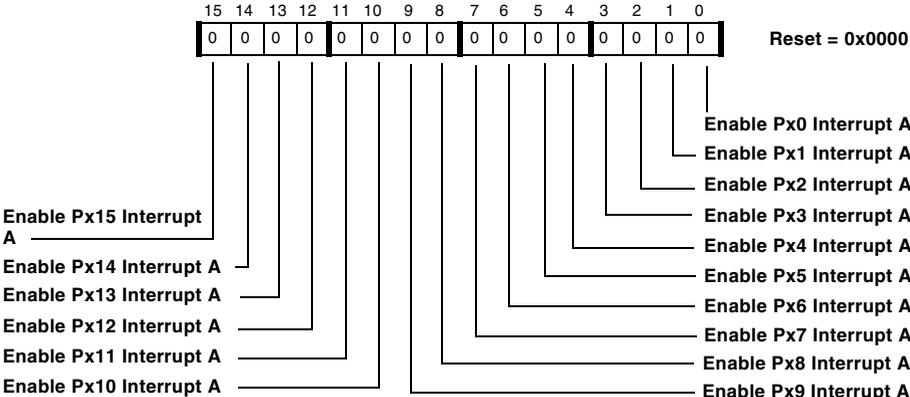


Figure 9-23. GPIO Mask Interrupt A Registers

Memory-Mapped GPIO Registers

GPIO Mask Interrupt B Registers (PORTxIO_MASKB)

For all bits, 1 - Enable

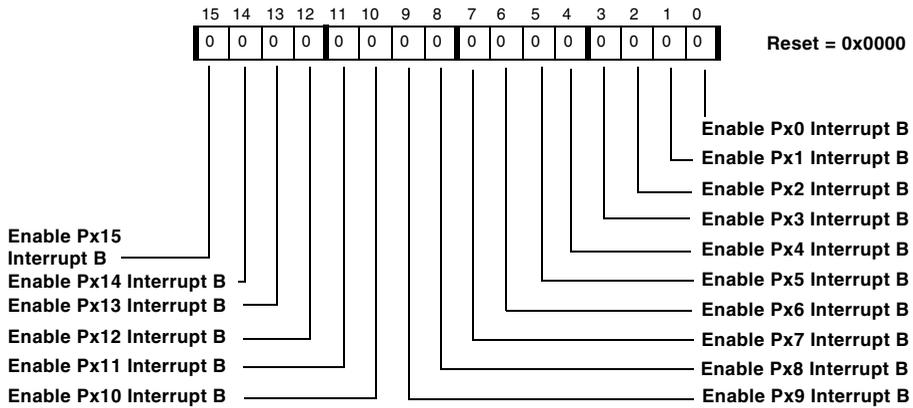


Figure 9-24. GPIO Mask Interrupt B Registers

GPIO Mask Interrupt Set Registers (PORTxIO_MASKA/B_SET)

GPIO Mask Interrupt A Set Registers (PORTxIO_MASKA_SET)

For all bits, 1 - Set

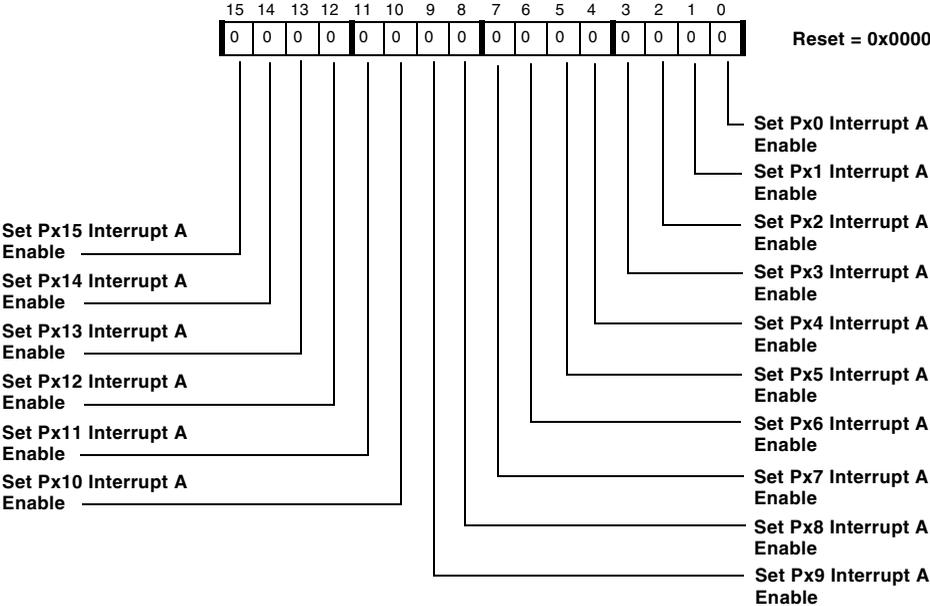


Figure 9-25. GPIO Mask Interrupt A Set Registers

Memory-Mapped GPIO Registers

GPIO Mask Interrupt B Set Registers (PORTxIO_MASKB_SET)

For all bits, 1 - Set

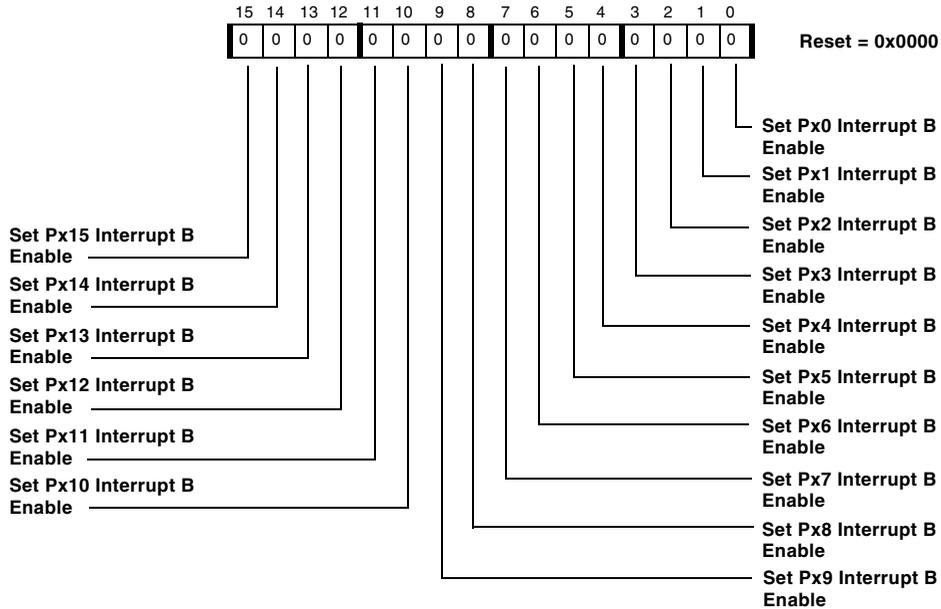


Figure 9-26. GPIO Mask Interrupt B Set Registers

GPIO Mask Interrupt Clear Registers (PORTxIO_MASKA/B_CLEAR)

GPIO Mask Interrupt A Clear Registers (PORTxIO_MASKA_CLEAR)

For all bits, 1 - Clear

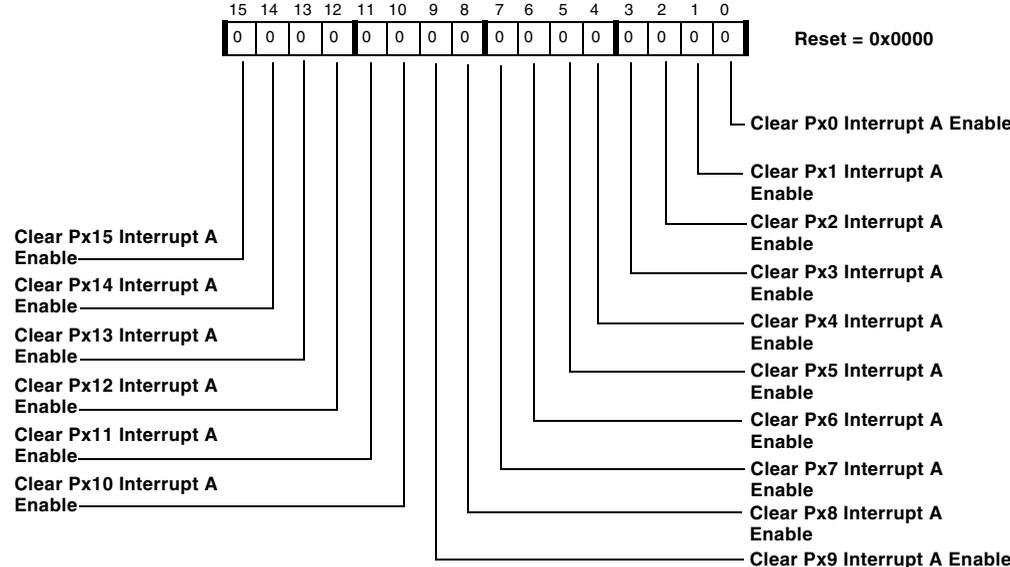


Figure 9-27. GPIO Mask Interrupt A Clear Registers

Memory-Mapped GPIO Registers

GPIO Mask Interrupt B Clear Registers (PORTxIO_MASKB_CLEAR)

For all bits, 1 - Clear

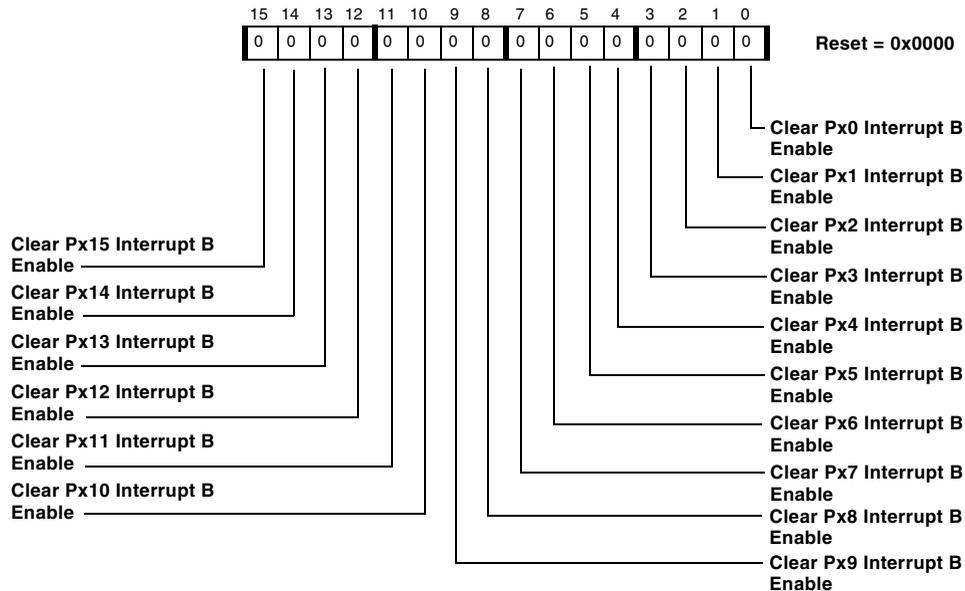


Figure 9-28. GPIO Mask Interrupt B Clear Registers

GPIO Mask Interrupt Toggle Registers (PORTxIO_MASKA/B_TOGGLE)

GPIO Mask Interrupt A Toggle Registers (PORTxIO_MASKA_TOGGLE)

For all bits, 1 - Toggle

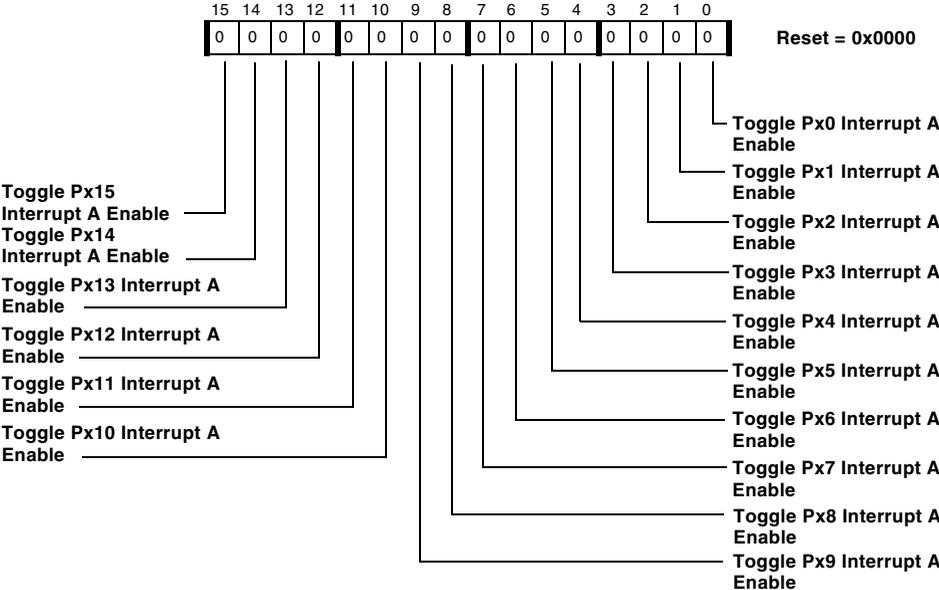


Figure 9-29. GPIO Mask Interrupt A Toggle Registers

Programming Examples

GPIO Mask Interrupt B Toggle Registers (PORTxIO_MASKB_TOGGLE)

For all bits, 1 - Toggle

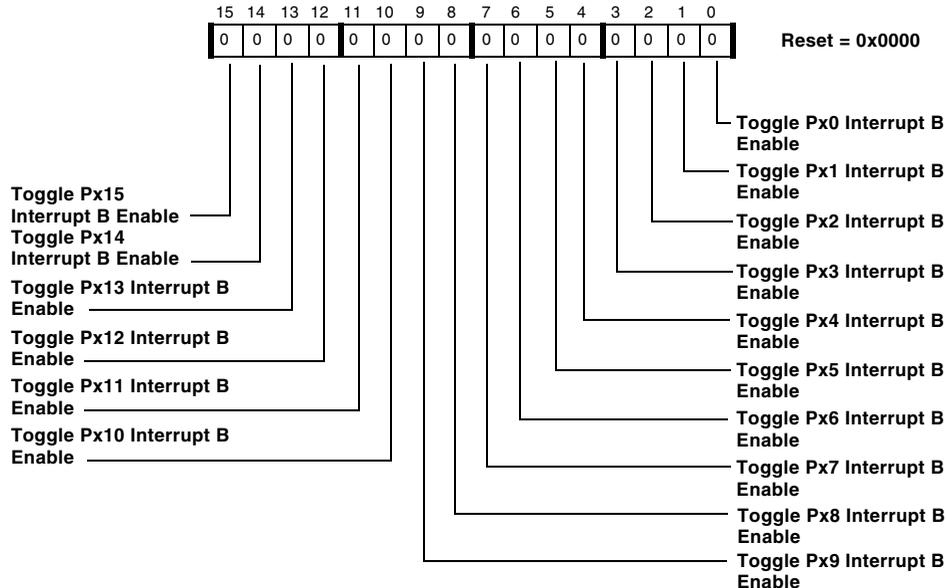


Figure 9-30. GPIO Mask Interrupt B Toggle Registers

Programming Examples

[Listing 9-1](#) provides examples for using the general-purpose ports.

Listing 9-1. General-Purpose Ports

```
/* set port f function enable register to GPIO (not peripheral)
*/
p0.l = lo(PORTF_FER);
p0.h = hi(PORTF_FER);

R0.h = 0x0000;
r0.l = 0x0000;
```

```
w[p0] = r0;

/* set port f direction register to enable some GPIO as output,
remaining are input */
p0.l = lo(PORTFIO_DIR);
p0.h = hi(PORTFIO_DIR);
r0.h = 0x0000;
r0.l = 0x0FC0;
w[p0] = r0;
ssync;

/* set port f clear register */
p0.l = lo(PORTFIO_CLEAR);
p0.h = hi(PORTFIO_CLEAR);
    r0.l = 0xFC0;
    w[p0] = r0;
    ssync;

/* set port f input enable register to enable input drivers of
some GPIOs */
p0.l = lo(PORTFIO_INEN);
p0.h = hi(PORTFIO_INEN);
r0.h = 0x0000;
r0.l = 0x003C;
w[p0] = r0;
ssync;

/* set port f polarity register */
p0.l = lo(PORTFIO_POLAR);
p0.h = hi(PORTFIO_POLAR);
r0 = 0x00000;
w[p0] = r0;
ssync;
```

Programming Examples

10 GENERAL-PURPOSE TIMERS

This chapter describes the general-purpose (GP) timer module. Following an overview and a list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of GP timers for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For GP Timer interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the GP Timers is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each GP Timer, refer to [Chapter A, “System MMR Assignments”](#).

GP timer behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 10-56](#).

Overview

The general-purpose timers support the following operating modes:

- Single-shot mode for interval timing and single pulse generation
- Pulse width modulation (PWM) generation with consistent update of period and pulse width values
- External signal capture mode with consistent update of period and pulse width values
- External event counter mode

Feature highlights are:

- Synchronous operation
- Consistent management of period and pulse width values
- Interaction with PPI module for video frame sync operation
- Autobaud detection for UART module
- Graceful bit pattern termination when stopping
- Support for center-aligned PWM patterns
- Error detection on implausible pattern values
- All read and write accesses to 32-bit registers are atomic
- Every timer has its dedicated interrupt request output
- Unused timers can function as edge-sensitive pin interrupts

The internal structure of the individual timers is illustrated by [Figure 10-1](#), which shows the details of timer 0 as a representative example. The other timers have identical structure.

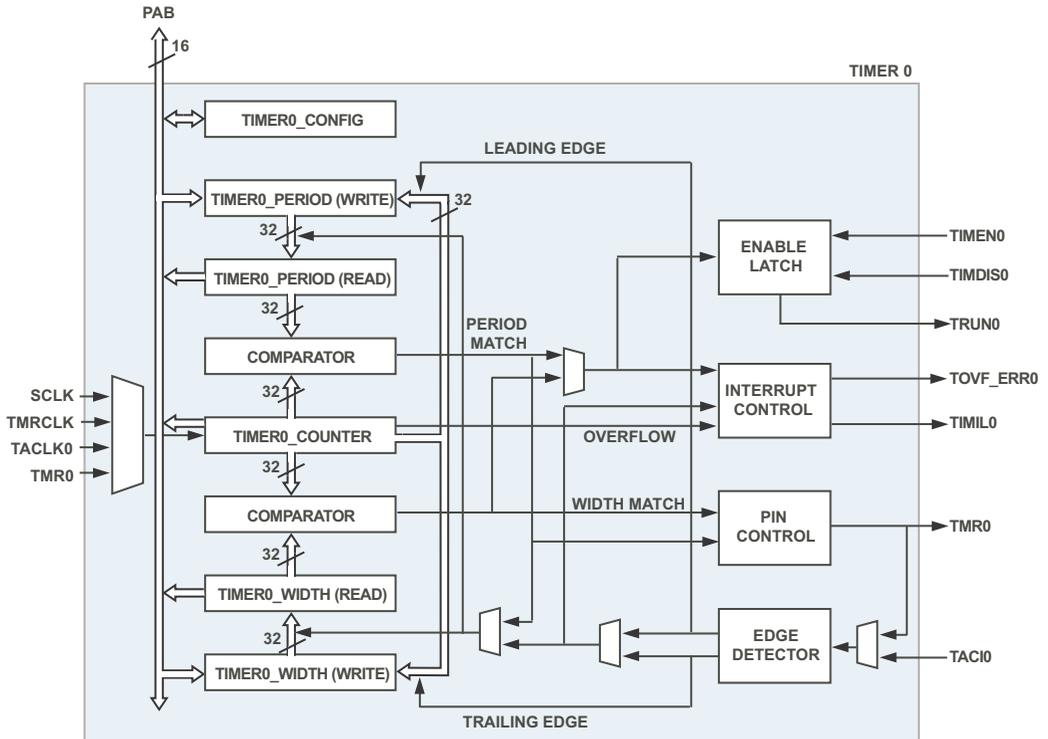


Figure 10-1. Internal Timer Structure

External Interface

Every timer has a dedicated TMR pin. If enabled, the TMR pins output the single-pulse or PWM signals generated by the timer. The TMR pins function as input in capture and counter modes. Polarity of the signals is programmable.

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.

Description of Operation

Clock and capture input pins are sampled every `SCLK` cycle. The duration of every low or high state must be at least one `SCLK`. Therefore, the maximum allowed frequency of timer input signals is `SCLK/2`.

Internal Interface

Timer registers are always accessed by the core through the 16-bit PAB bus. Hardware ensures that all read and write operations from and to 32-bit timer registers are atomic.

Every timer has a dedicated interrupt request output that connects to the system interrupt controller (SIC).

Description of Operation

The core of every timer is a 32-bit counter, that can be interrogated through the read-only `TIMER_COUNTER` register. Depending on the mode of operation, the counter is reset to either `0x0000 0000` or `0x0000 0001` when the timer is enabled. The counter always counts upward. Usually, it is clocked by `SCLK`. In PWM mode it can be clocked by the alternate clock input `TACLK` or, alternatively, the common timer clock input `TMRCLK`. In counter mode, the counter is clocked by edges on the `TMR` input pin. The significant edge is programmable.

After $2^{32}-1$ clocks, the counter overflows. This is reported by the overflow/error bit `TOVF_ERR` in the `TIMER_STATUS` register. In PWM and counter mode, the counter is reset by hardware when its content reaches the values stored in the `TIMER_PERIOD` register. In capture mode, the counter is reset by leading edges on the `TMR` or `TACI` input pin. If enabled, these events cause the interrupt latch `TIMIL` in the `TIMER_STATUS` register to be set and issue a system interrupt request. The `TOVF_ERR` and `TIMIL` latches are sticky and should be cleared by software using `W1C` (write-1-to-clear) operations to clear the interrupt request. The global `TIMER_STATUS` register

is 32-bits wide. A single atomic 32-bit read can report the status of all corresponding timers.

Before a timer can be enabled, its mode of operation is programmed in the individual timer-specific `TIMER_CONFIG` register. Then, the timers are started by writing a “1” to the representative bits in the global `TIMER_ENABLE` register.

The `TIMER_ENABLE` register can be used to enable all timers simultaneously. The register contains `W1S` (write-1-to-set) control bits, one for each timer. Correspondingly, the `TIMER_DISABLE` register contains `W1C` control bits to allow simultaneous or independent disabling of the timers. Either register can be read 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 `TIMEN` bit is set.

While the PWM mode is used to generate PWM patterns, the capture mode (`WDTH_CAP`) is designed to “receive” PWM signals. A PWM pattern is represented by a pulse width and a signal period. This is described by the `TIMER_WIDTH` and `TIMER_PERIOD` register pair. In capture mode these registers are read only. Hardware always captures both values. Regardless of whether in PWM or capture mode, shadow buffers always ensure consistency between the `TIMER_WIDTH` and `TIMER_PERIOD` values. In PWM mode, hardware performs a plausibility check by the time the timer is enabled. If there is an error, the type is reported by the `TIMER_CONFIG` register and signalled by the `TOVF_ERR` bit.

Interrupt Processing

Each timer can generate a single interrupt. The resulting interrupt signals are routed to the system interrupt controller block for prioritization and masking. The timer status (`TIMER_STATUS`) register latches the timer interrupts to provide a means for software to determine the interrupt source.

Description of Operation

Figure 10-2 shows the interrupt structure of the timers.

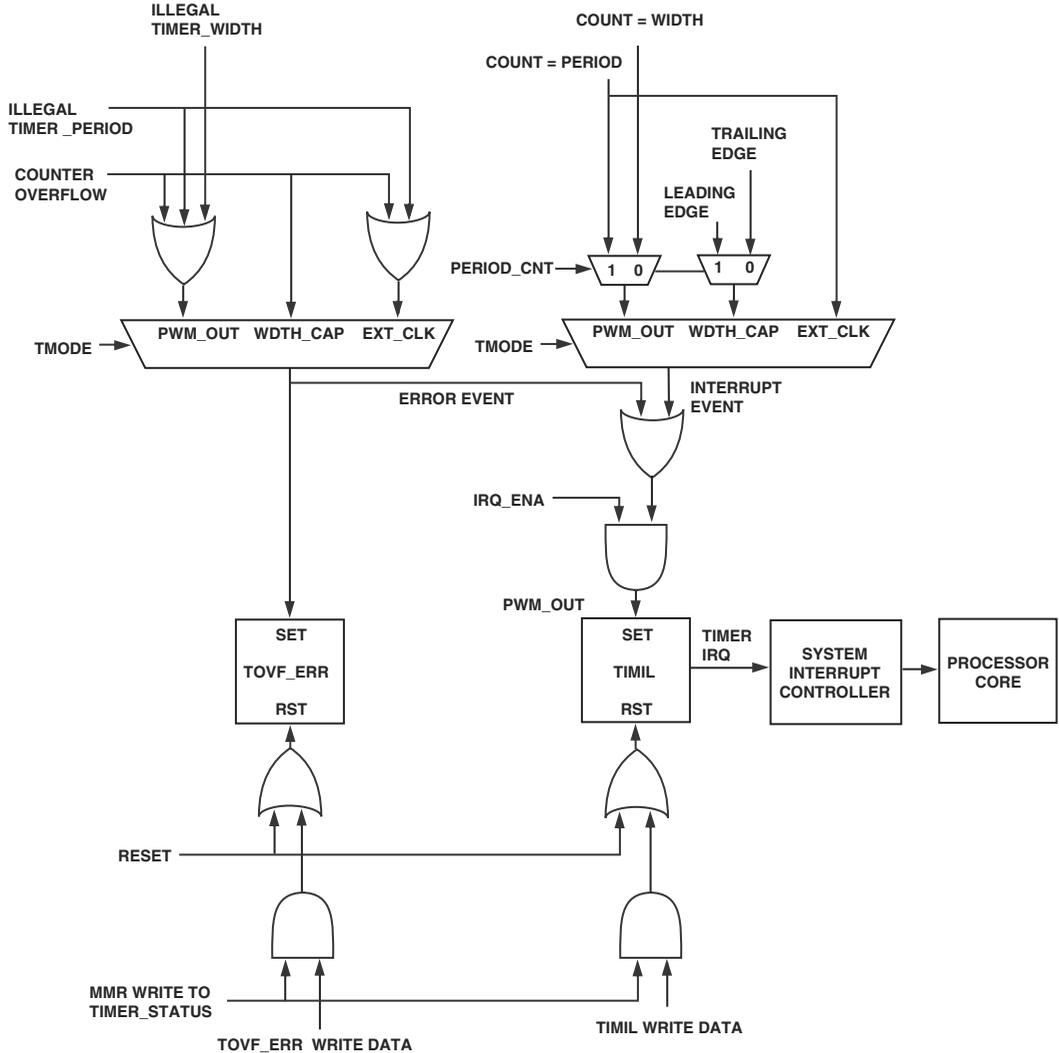


Figure 10-2. Timers Interrupt Structure

To enable interrupt generation, set the `IRQ_ENA` bit and unmask the interrupt source in the `IMASK` and `SIC_IMASK` registers. To poll the `TIMIL` bit

without interrupt generation, set `IRQ_ENA` but leave the interrupt masked at the system level. If enabled by `IRQ_ENA`, interrupt requests are also generated by error conditions as reported by the `TOVF_ERR` bits.

The system interrupt controller enables flexible interrupt handling. All timers may or may not share the same CEC interrupt channel, so that a single interrupt routine services more than one timer. In PWM mode, multiple timers may run with the same period settings and issue their interrupt requests simultaneously. In this case, the service routine might clear all `TIMIL` latch bits at once by writing `0x000F 000F` to the `TIMER_STATUS` register.

If interrupts are enabled, make sure that the interrupt service routine (ISR) clears the `TIMIL` bit in the `TIMER_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 `TIMIL` clear command from the `RTI` instruction, an extra `SSYNC` instruction may be inserted. In `EXT_CLK` mode, reset the `TIMIL` bit in the `TIMER_STATUS` register at the very beginning of the interrupt service routine to avoid missing any timer events.

Illegal States

Every timer features an error detection circuit. It handles overflow situations but also performs pulse width vs. period plausibility checks. Errors are reported by the `TOVF_ERR` bits in the `TIMER_STATUS` register and the `ERR_TYP` bit field in the individual `TIMER_CONFIG` registers. [Table 10-1](#) provides a summary of error conditions, using these terms:

- **Startup.** The first clock period during which the timer counter is running after the timer is enabled by writing `TIMER_ENABLE`.
- **Rollover.** The time when the current count matches the value in `TIMER_PERIOD` and the counter is reloaded with the value “1”.

Description of Operation

- **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 write “1” to clear `TOVF_ERR` to acknowledge the error.

The following table can be read as: “In mode __ at event __, if `TIMER_PERIOD` is __ and `TIMER_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 `TMR` pin.

Table 10-1. Overview of Illegal States

Mode	Event	TIMER_PERIOD	TIMER_WIDTH	ERR_TYP	TOVF_ERR	
PWM_OUT, PERIOD_CNT = 1	Startup (No boundary condition tests performed on TIMER_WIDTH)	== 0	Anything	b#10	Set	
		== 1	Anything	b#10	Set	
		≥ 2	Anything	No change	No change	
	Rollover	== 0	Anything	b#10	Set	
		== 1	Anything	b#11	Set	
		≥ 2	== 0	b#11	Set	
		≥ 2	< TIMER_PERIOD	No change	No change	
	Overflow, not possible unless there is also another error, such as TIMER_PERIOD == 0	Anything	Anything	b#01	Set	
	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	No change	No change		
Rollover		Rollover is not possible in this mode.				
Overflow, not possible unless there is also another error, such as TIMER_WIDTH == 0	Anything	Anything	b#01	Set		

Modes of Operation

Table 10-1. Overview of Illegal States (Continued)

Mode	Event	TIMER_PERIOD	TIMER_WIDTH	ERR_TYP	TOVF_ERR
WDTH_CAP	Startup	TIMER_PERIOD and TIMER_WIDTH are read-only in this mode, no error possible.			
	Rollover	TIMER_PERIOD and TIMER_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	No change	No change
	Rollover	== 0	Anything	b#10	Set
		≥ 1	Anything	No change	No change
	Overflow, not possible unless there is also another error, such as TIMER_PERIOD == 0	Anything	Anything	b#01	Set

Modes of Operation

The following sections provide a functional description of the general-purpose timers in various operating modes.

Pulse Width Modulation (PWM_OUT) Mode

Use the PWM_OUT mode for PWM signal or single-pulse generation, for interval timing or for periodic interrupt generation. [Figure 10-3](#) illustrates PWM_OUT mode.

Setting the `TMODE` field to `b#01` in the `TIMER_CONFIG` register enables `PWM_OUT` mode. Here, the `TMR` pin is an output, but it can be disabled by setting the `OUT_DIS` bit in the `TIMER_CONFIG` register.

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`).

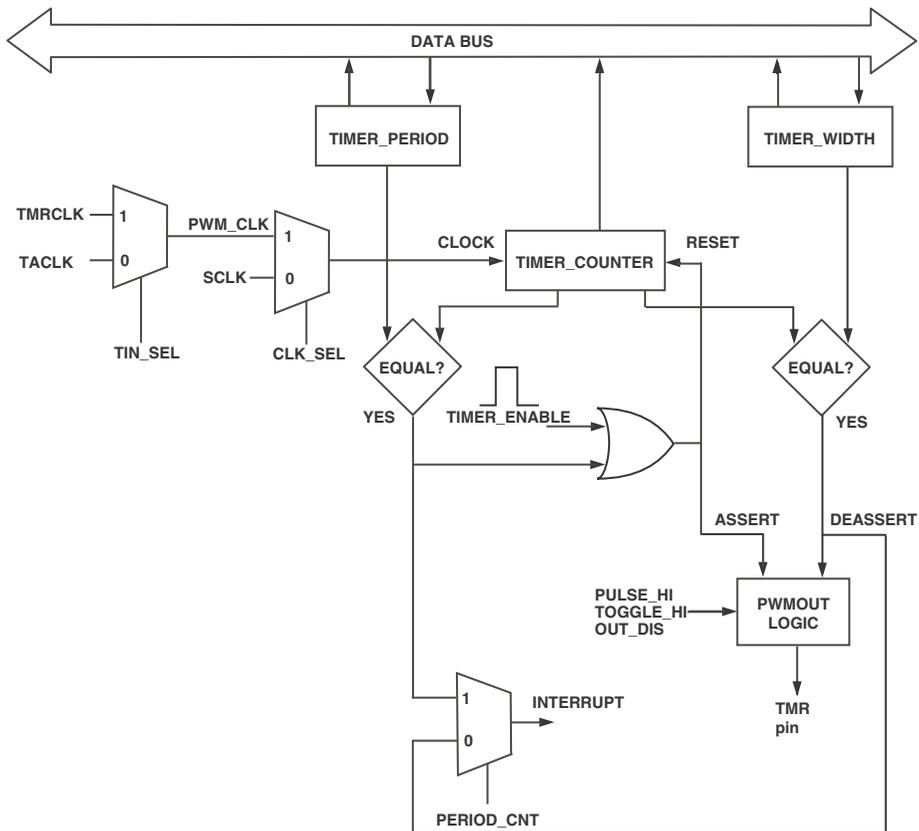


Figure 10-3. Timer Flow Diagram, PWM_OUT Mode

Modes of Operation

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 `0x0` 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 `TIMER_WIDTH` register, generates one asserting and one deasserting edge, then generates an interrupt (if enabled) and stops. When `PERIOD_CNT` is set (`PWM_OUT` continuous pulse mode), the timer uses both the `TIMER_PERIOD` and `TIMER_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.



The `TIMER_PERIOD` and `TIMER_WIDTH` registers are read-only in some operation modes. Be sure to set the `TMODE` field in the `TIMER_CONFIG` register to `b#01` before writing to these registers.

Output Pad Disable

The output pin can be disabled in `PWM_OUT` mode by setting the `OUT_DIS` bit in the `TIMER_CONFIG` register. The `TMR` pin is then three-stated regardless of the setting of `PULSE_HI` and `TOGGLE_HI`. This can reduce power consumption when the output signal is not being used. The `TMR` pin can also be disabled by the function enable and the multiplexer control registers.

Single Pulse Generation

If the `PERIOD_CNT` bit is cleared, the `PWM_OUT` mode generates a single pulse on the `TMR` pin. This mode can also be used to implement a precise delay.

The pulse width is defined by the `TIMER_WIDTH` register, and the `TIMER_PERIOD` register is not used. See [Figure 10-4](#).

At the end of the pulse, the timer interrupt latch bit `TIMIL` is set, and the timer is stopped automatically. No writes to the `TIMER_DISABLE` register are required in this mode. If the `PULSE_HI` bit is set, an active high pulse is generated on the `TMR` pin. If `PULSE_HI` is not set, the pulse is active low.

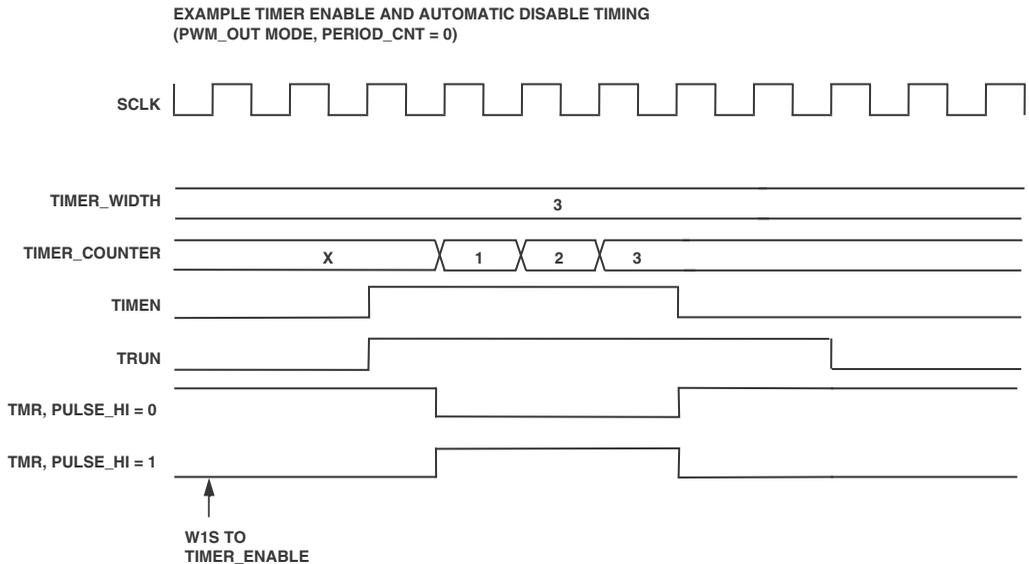


Figure 10-4. Timer Enable and Automatic Disable Timing

The pulse width may be programmed to any value from 1 to $(2^{32}-1)$, inclusive.

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 (PWM patterns). This mode also generates periodic interrupts for real-time signal processing.

Modes of Operation

The 32-bit `TIMER_PERIOD` and `TIMER_WIDTH` registers are programmed with the values required by the PWM signal.

When the timer is enabled in this mode, the `TMR` pin is pulled to a deasserted state each time the counter equals the value of the 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 `TMR` pin, the `PULSE_HI` bit in the corresponding `TIMER_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 `TMR` pin is driven to the deasserted level.

Figure 10-5 shows timing details.

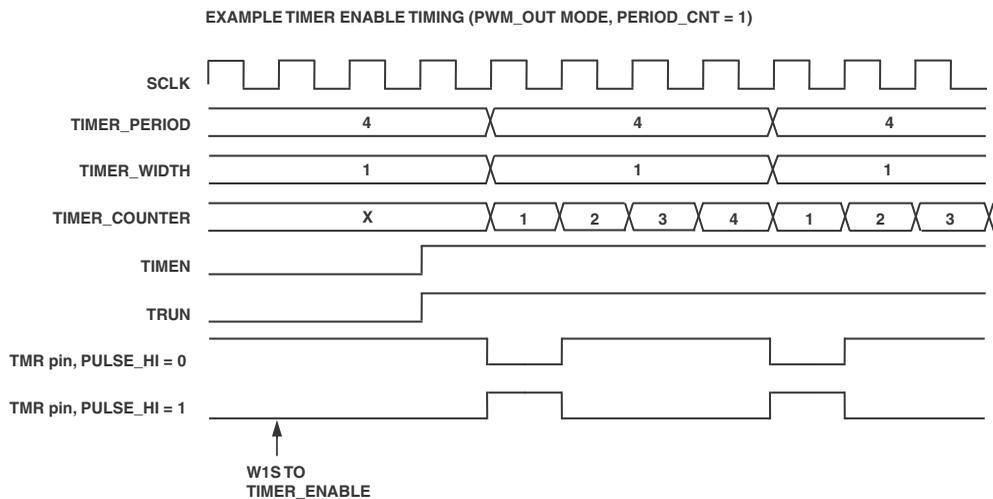


Figure 10-5. Timer Enable Timing

If enabled, a timer interrupt is generated at the end of each period. An interrupt service routine must clear the interrupt latch bit (`TIMIL`) and might alter period and/or width values. In PWM applications, the software needs to update period and pulse width values while the timer is

running. When software updates either the `TIMER_PERIOD` or `TIMER_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. Reads from `TIMER_PERIOD` and `TIMER_WIDTH` registers return the old values until the period expires.

The `TOVF_ERR` status bit signifies an error condition in `PWM_OUT` mode. The `TOVF_ERR` bit is set if `TIMER_PERIOD = 0` or `TIMER_PERIOD = 1` at startup, or when the timer counter register rolls over. It is also set if the timer pulse width register is greater than or equal to the timer period register by the time the counter rolls over. The `ERR_TYP` bits are set when the `TOVF_ERR` bit is set.

Although the hardware reports an error if the `TIMER_WIDTH` value equals the `TIMER_PERIOD` value, this is still a valid operation to implement PWM patterns with 100% duty cycle. If doing so, software must generally ignore the `TOVL_ERR` flags. Pulse width values greater than the period value are not recommended. Similarly, `TIMER_WIDTH = 0` is not a valid operation. Duty cycles of 0% are not supported.

To generate the maximum frequency on the TMR output pin, set the period value to “2” and the pulse width to “1”. This makes the pin 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 $(\text{period} - 1)$, inclusive.

PULSE_HI Toggle Mode

The waveform produced in `PWM_OUT` mode with `PERIOD_CNT = 1` normally has a fixed assertion time and a programmable deassertion time (via the `TIMER_WIDTH` register). When two or more timers are running synchronously by the same period settings, the pulses are aligned to the asserting edge as shown in [Figure 10-6](#).

Modes of Operation

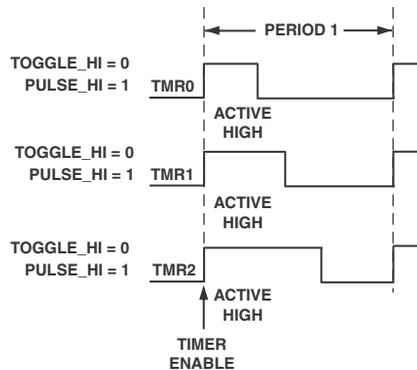


Figure 10-6. Example of Timers With Pulses Aligned to Asserting Edge

The `TOGGLE_HI` mode enables control of the timing of both the asserting and deasserting 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 and active high pulses, taken together, create two halves of a symmetrical rectangular waveform. The effective waveform is active high when `PULSE_HI` is set and active low when `PULSE_HI` is cleared. The value of the `TOGGLE_HI` bit 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 10-7 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 `TIMER_WIDTH` registers control the phase between the signals.

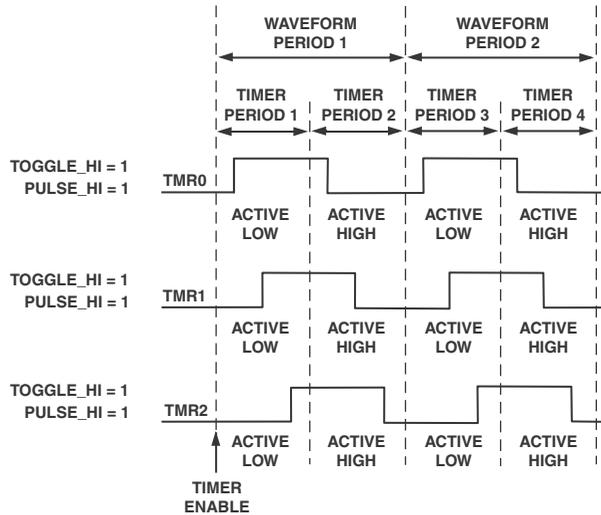


Figure 10-7. Three Timers With Same Period Settings

Similarly, two timers can generate non-overlapping clocks, by center-aligning the pulses while inverting the signal polarity for one of the timers (see Figure 10-8).

When `TOGGLE_HI = 0`, software updates the `TIMER_PERIOD` and `TIMER_WIDTH` registers once per waveform period. When `TOGGLE_HI = 1`, software updates the `TIMER_PERIOD` and `TIMER_WIDTH` registers twice per waveform. Period values are half as large. In odd-numbered periods, write $(\text{Period} - \text{Width})$ instead of `Width` to the `TIMER_WIDTH` register in order to obtain center-aligned pulses.

Modes of Operation

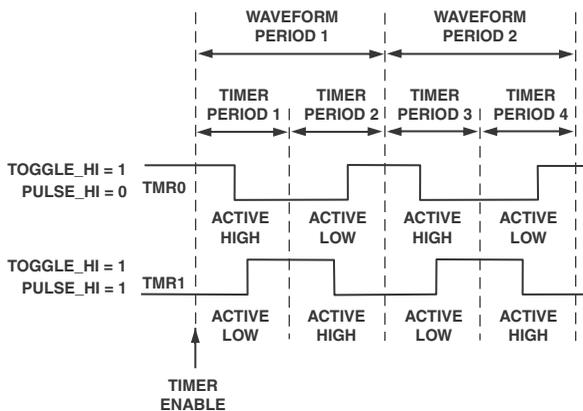


Figure 10-8. Two Timers With Non-overlapping Clocks

For example, if the pseudo-code when `TOGGLE_HI = 0` is:

```
int period, width;
for (;;) {
    period = generate_period(...);
    width = generate_width(...);

    waitfor (interrupt);

    write (TIMER_PERIOD, period);
    write (TIMER_WIDTH, width);
}
```

Then when `TOGGLE_HI = 1`, the pseudo-code 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 (TIMER_PERIOD, per1) ;
write (TIMER_WIDTH, per1 - wid1) ;

waitfor (interrupt) ;

write (TIMER_PERIOD, per2) ;
write (TIMER_WIDTH, wid2) ;

}
```

As shown in this example, the pulses produced do not need to be symmetric (*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 `TRUN` bit in the `TIMER_STATUS` register is updated only at the end of even-numbered periods in `TOGGLE_HI` mode. When `TIMER_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 the `TIMER_PERIOD` register is either set to "0" or "1", or when the width value is greater than or equal to the period value.

Externally Clocked PWM_OUT

By default, the timer is clocked internally by `SCLK`. Alternatively, if the `CLK_SEL` bit in the `TIMER_CONFIG` register is set, the timer is clocked by

Modes of Operation

PWM_CLK. The PWM_CLK is normally input from the TACLK pin, but may be taken from the common TMRCLK pin regardless of whether 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 TIMER_WIDTH register.

When CLK_SEL is set, the counter resets to 0x0 at startup and increments on each rising edge of PWM_CLK. The TMR 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 TMR 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 alternate timer clock inputs (TACLK) are enabled when a timer is in PWM_OUT mode with CLK_SEL = 1 and TIN_SEL = 0, without regard to the content of the multiplexer control and function enable registers.

Using PWM_OUT Mode With the PPI

Some timers may be 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 20-19](#).

Stopping the Timer in PWM_OUT Mode

In all PWM_OUT mode variants, the timer treats a disable operation (W1C to `TIMER_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 TMR pin. The processor can determine when the timer stops running by polling for the corresponding `TRUN` bit in the `TIMER_STATUS` register to read “0” or by waiting for the last interrupt (if enabled). Note the timer cannot be reconfigured (`TIMER_CONFIG` cannot be written to a new value) until after the timer stops and `TRUN` reads “0”.

In PWM_OUT single pulse mode (`PERIOD_CNT = 0`), it is not necessary to write `TIMER_DISABLE` to stop the timer. At the end of the pulse, the timer stops automatically, the corresponding bit in `TIMER_ENABLE` (and `TIMER_DISABLE`) is cleared, and the corresponding `TRUN` bit is cleared. See [Figure 10-4 on page 10-13](#). To generate multiple pulses, write a “1” to `TIMER_ENABLE`, wait for the timer to stop, then write another “1” to `TIMER_ENABLE`.

In continuous PWM generation mode (`PWM_OUT, PERIOD_CNT = 1`) software can stop the timer by writing to the `TIMER_DISABLE` register. To prevent the ongoing PWM pattern from being stopped in an unpredictable way, the timer does not stop immediately when the corresponding “1” has been written to the `TIMER_DISABLE` register. Rather, the write simply clears the enable latch and the timer still completes the ongoing PWM patterns gracefully. It stops cleanly at the end of the first period when the enable latch is cleared. During this final period the `TIMEN` bit returns “0”, but the `TRUN` bit still reads as a “1”.

If the `TRUN` bit is not cleared explicitly, and the enable latch can be cleared and 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.

Modes of Operation

Figure 10-9 shows detailed timing.

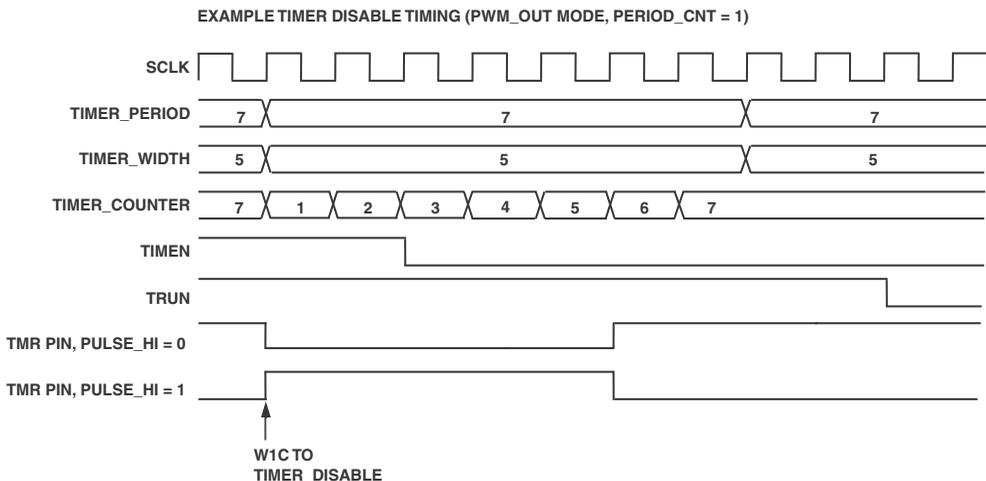


Figure 10-9. Timer Disable Timing

If necessary, the processor can force a timer in `PWM_OUT` mode to abort immediately. Do this by first writing a “1” to the corresponding bit in `TIMER_DISABLE`, and then writing a “1” to the corresponding `TRUN` bit in `TIMER_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 TMR pin.

When a timer is disabled, the `TIMER_COUNTER` register retains its state; when a timer is re-enabled, the timer counter is reinitialized based on the operating mode. The `TIMER_COUNTER` register is read-only. Software cannot overwrite or preset the timer counter value directly.

Pulse Width Count and Capture (WDTH_CAP) Mode

Use the `WDTH_CAP` mode, often simply called “capture mode,” to measure pulse widths on the `TMR` or `TACI` input pins, or to “receive” PWM signals. [Figure 10-10](#) shows a flow diagram for `WDTH_CAP` mode.

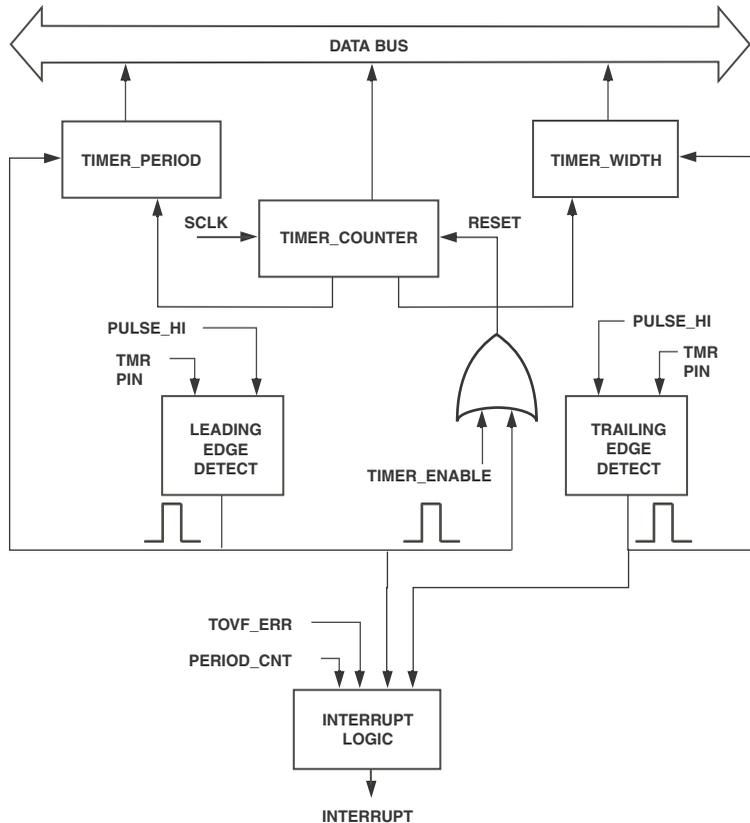


Figure 10-10. Timer Flow Diagram, `WDTH_CAP` Mode

In `WDTH_CAP` mode, the `TMR 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 `TIMER_CONFIG` register enables this mode.

Modes of Operation

When enabled in this mode, the timer resets the count in the `TIMER_COUNTER` register to `0x0000 0001` and does not start counting until it detects a leading edge on the `TMR` 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 `TIMER_COUNTER` register into the width buffer. At the next leading edge, the timer transfers the current 32-bit value of the `TIMER_COUNTER` register into the period buffer. 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 `TMR` pin, the `PULSE_HI` bit in the `TIMER_CONFIG` register is set or cleared. If the `PULSE_HI` bit is cleared, the measurement is initiated by a falling edge, the content of the counter register is captured to the pulse width buffer on the rising edge, and to the period buffer on the next falling edge. When the `PULSE_HI` bit is set, the measurement is initiated by a rising edge, the counter value is captured to the pulse width buffer on the falling edge, and to the period buffer on the next rising edge.

In `WDTH_CAP` mode, these three events always occur at the same time:

1. The `TIMER_PERIOD` register is updated from the period buffer.
2. The `TIMER_WIDTH` register is updated from the width buffer.
3. The `TIMIL` bit gets set (if enabled) but does not generate an error.

The `PERIOD_CNT` bit in the `TIMER_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 `TOVF_ERR` bit does not get set at a measurement report. A measurement report occurs, at most, once per input signal period.

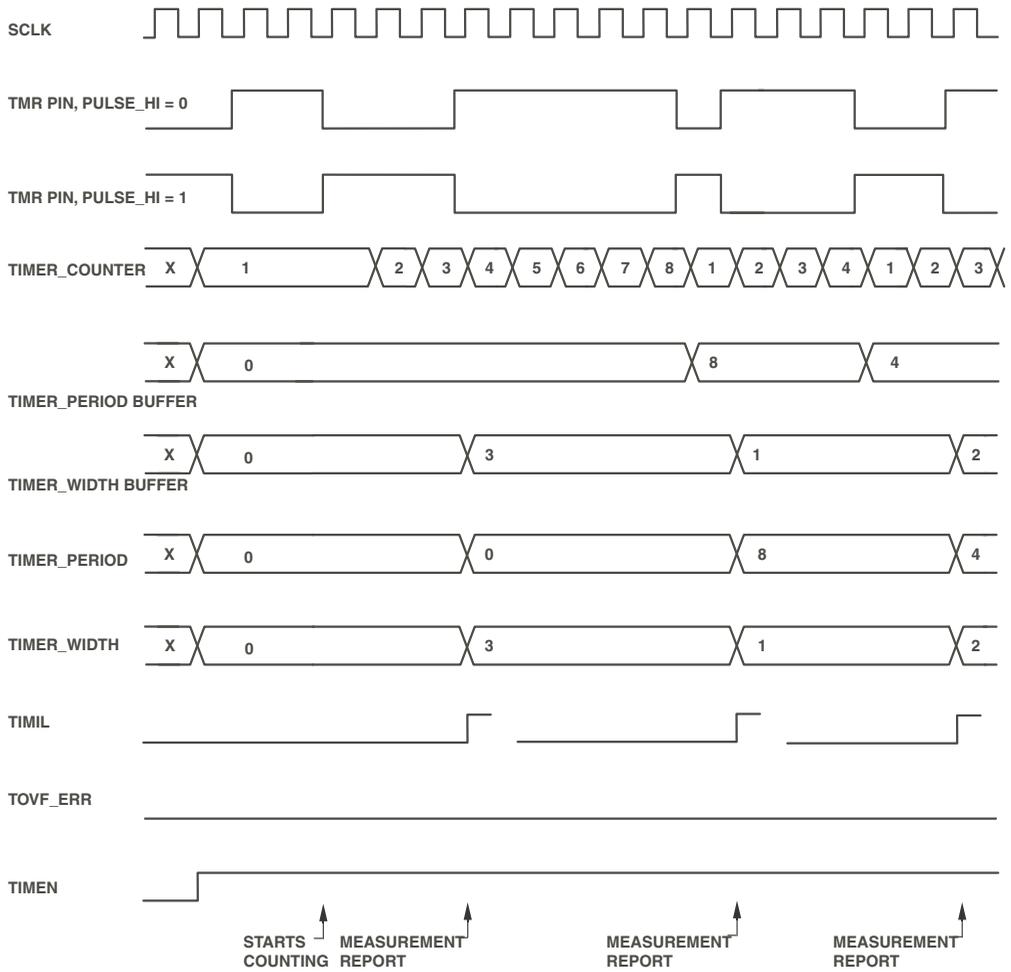
The current timer counter value is always copied to the width buffer and period buffer registers at the trailing and leading edges of the input signal, respectively, but these values are not visible to software. A measurement report event samples the captured values into visible registers and sets the timer interrupt to signal that `TIMER_PERIOD` and `TIMER_WIDTH` are ready to be read. When the `PERIOD_CNT` bit is set, the measurement report occurs just after the period buffer captures its value (at a leading edge). When the `PERIOD_CNT` bit is cleared, the measurement report occurs just after the width buffer captures its value (at a trailing edge).

If the `PERIOD_CNT` bit is set and a leading edge occurred (see [Figure 10-11](#)), then the `TIMER_PERIOD` and `TIMER_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 (see [Figure 10-12](#)), then the `TIMER_WIDTH` register reports the pulse width measured in the pulse that just ended, but the `TIMER_PERIOD` register reports the pulse period measured at the end of the previous period.

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 `TIMER_PERIOD` value in this case returns “0”, as shown in [Figure 10-12](#). 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. Instead, an error report interrupt is generated (if enabled) when the counter range is exceeded and the counter wraps around. In this case, both `TIMER_WIDTH` and `TIMER_PERIOD` read “0” (because no measurement report occurred to copy the value captured in the width buffer to `TIMER_WIDTH`). See the first interrupt in [Figure 10-13](#).



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 reenabled as appropriate in



NOTE: FOR SIMPLICITY, THE SYNCHRONIZATION DELAY BETWEEN TMR PIN EDGES AND BUFFER REGISTER UPDATES IS NOT SHOWN.

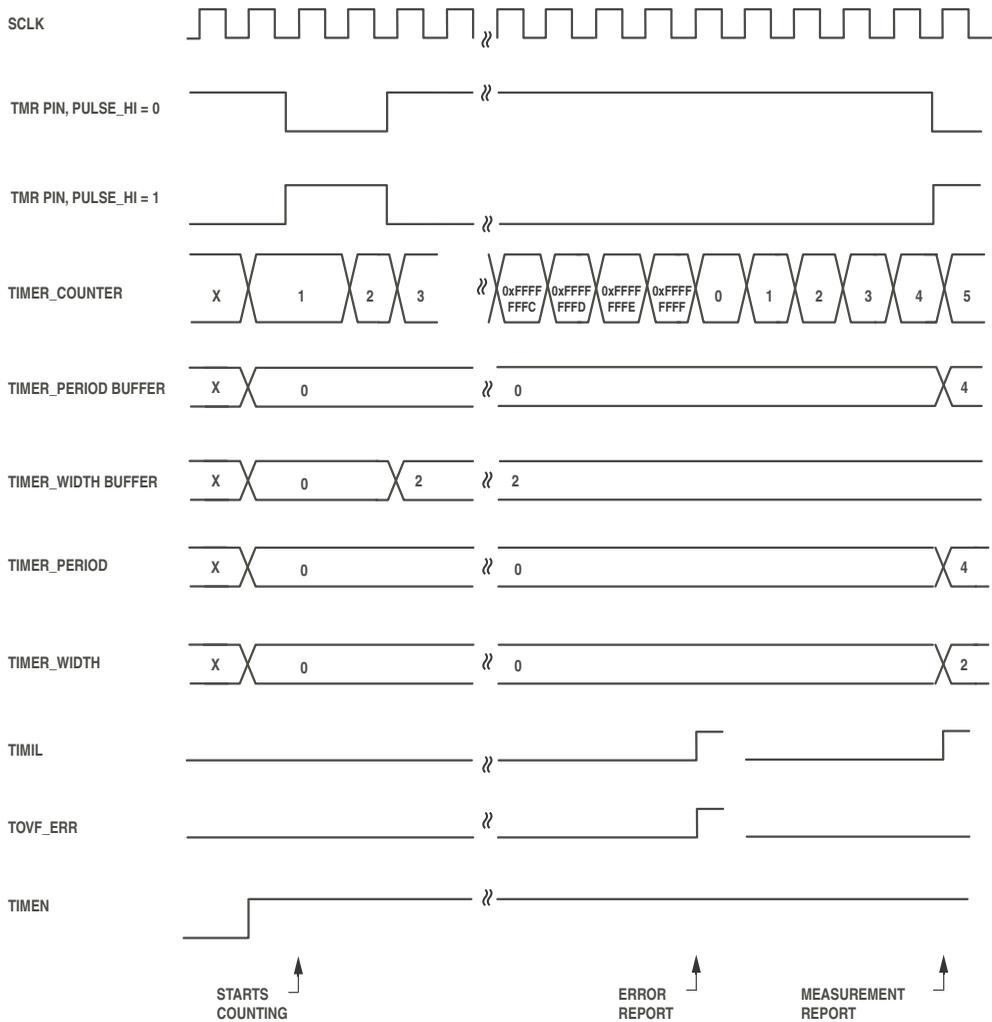
Figure 10-12. Example of Width Capture Measurement Report Timing (WDTH_CAP mode, PERIOD_CNT = 0)

preparation for another measurement. This procedure prevents the timer from free-running after the width measurement, and from logging errors generated by the timer count overflowing.

Modes of Operation

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_ERR` bit in the `TIMER_STATUS` register and the `ERR_TYP` bits in the `TIMER_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 `TIMER_PERIOD` and `TIMER_WIDTH` registers are never updated at the time an error is signaled.

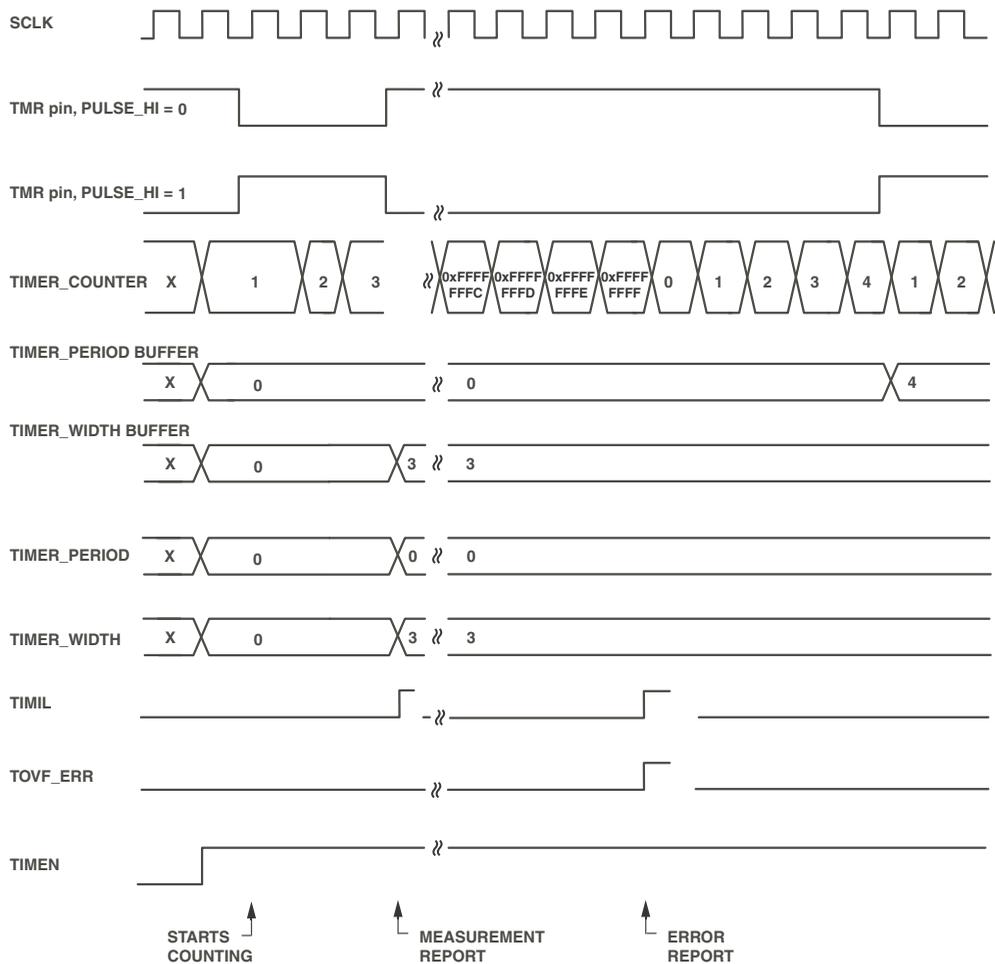
Refer to [Figure 10-13](#) and [Figure 10-14](#) for more information.



NOTE: FOR SIMPLICITY, THE SYNCHRONIZATION DELAY BETWEEN TMR PIN EDGES AND BUFFER REGISTER UPDATES IS NOT SHOWN.

Figure 10-13. Example Timing for Period Overflow Followed by Period Capture (WDTM_CAP mode, PERIOD_CNT = 1)

Modes of Operation



NOTE: FOR SIMPLICITY, THE SYNCHRONIZATION DELAY BETWEEN TMR PIN EDGES AND BUFFER REGISTER UPDATES IS NOT SHOWN.

Figure 10-14. Example Timing for Width Capture Followed by Period Overflow (WIDTH_CAP mode, PERIOD_CNT = 0)

Both `TIMIL` and `TOVF_ERR` are sticky bits, and software must explicitly clear them. If the timer overflowed and `PERIOD_CNT = 1`, neither the `TIMER_PERIOD` nor the `TIMER_WIDTH` register were updated. If the timer overflowed and `PERIOD_CNT = 0`, the `TIMER_PERIOD` and `TIMER_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 10-13](#) 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 `TMR` pin is not required to have a 50% duty cycle, but the minimum `TMR` pin low time is one `SCLK` period and the minimum `TMR` pin high time is one `SCLK` period. This implies the maximum `TMR` pin 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

On some devices, in `WDTH_CAP` mode, some of the timers can provide autobaud detection for the Universal Asynchronous Receiver/Transmitter (UART) interface(s). The `TIN_SEL` bit in the `TIMER_CONFIG` register causes the timer to sample the `TACI` pin instead of the `TMR` pin when enabled for `WDTH_CAP` mode. Autobaud detection can be used for initial bit rate negotiations as well as for detection of bit rate drifts while the interface is in operation.

External Event (EXT_CLK) Mode

Use the `EXT_CLK` mode (sometimes referred to as the counter mode) to count external events—that is, signal edges on the `TMR` pin (which is an

Modes of Operation

input in this mode). [Figure 10-15](#) shows a flow diagram for EXT_CLK mode.

The timer works as a counter clocked by an external source, which can also be asynchronous to the system clock. The current count in `TIMER_COUNTER` represents the number of leading edge events detected. Setting the `TMODE` field to `b#11` in the `TIMER_CONFIG` register enables this mode. The `TIMER_PERIOD` register is programmed with the value of the maximum timer external count.

The waveform applied to the `TMR` pin is not required to have a 50% duty cycle, but the minimum `TMR` low time is one `SCLK` period, and the minimum `TMR` high time is one `SCLK` period. This implies the maximum `TMR` pin input frequency is $SCLK/2$.

Period may be programmed to any value from 1 to $(2^{32} - 1)$, inclusive.

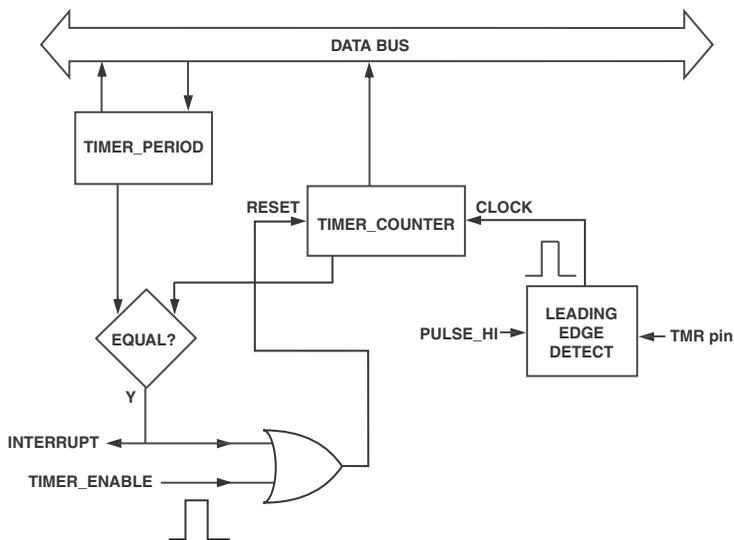


Figure 10-15. Timer Flow Diagram, EXT_CLK Mode

After the timer has been enabled, it resets the `TIMER_COUNTER` register to `0x0` and then waits for the first leading edge on the `TMR` 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 `TIMIL` 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_ERR` and `ERR_TYP` bits are set if the `TIMER_COUNTER` 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_WIDTH` register is unused.

Programming Model

The architecture of the timer block enables any of the timers within this block to work individually or synchronously along with others as a group of timers. Regardless of the operating mode, the programming model is always straightforward. Because of the error checking mechanism, always follow this order when enabling timers:

1. Set timer mode.
2. Write `TIMER_WIDTH` and `TIMER_PERIOD` registers as applicable.
3. Enable timer.

If this order is not followed, the plausibility check may fail because of undefined width and period values, or writes to `TIMER_WIDTH` and `TIMER_PERIOD` may result in an error condition, because the registers are read-only in some modes. The timer may not start as expected.

Timer Registers

If in `PWM_OUT` mode the PWM patterns of the second period differ from the patterns of the first one, the initialization sequence above might become:

1. Set timer mode to `PWM_OUT`.
2. Write first `TIMER_WIDTH` and `TIMER_PERIOD` value pair.
3. Enable timer.
4. Immediately write second `TIMER_WIDTH` and `TIMER_PERIOD` value pair.

Hardware ensures that the buffered width and period values become active when the first period expires.

Once started, timers require minimal interaction with software, which is usually performed by an interrupt service routine. In `PWM_OUT` mode software must update the pulse width and/or settings as required. In `WDTH_CAP` mode it must store captured values for further processing. In any case, the service routine should clear the `TIMIL` bits of the timers it controls.

Timer Registers

The timer peripheral module provides general-purpose timer functionality. It consists of multiple identical timer units.

Each timer provides four registers:

- `TIMER_CONFIG[15:0]` – timer configuration register
- `TIMER_WIDTH[31:0]` – timer pulse width register
- `TIMER_PERIOD[31:0]` – timer period register
- `TIMER_COUNTER[31:0]` – timer counter register

Additionally, three registers are shared between the timers within a block:

- `TIMER_ENABLE[15:0]` – timer enable register
- `TIMER_DISABLE[15:0]` – timer disable register
- `TIMER_STATUS[31:0]` – timer status register

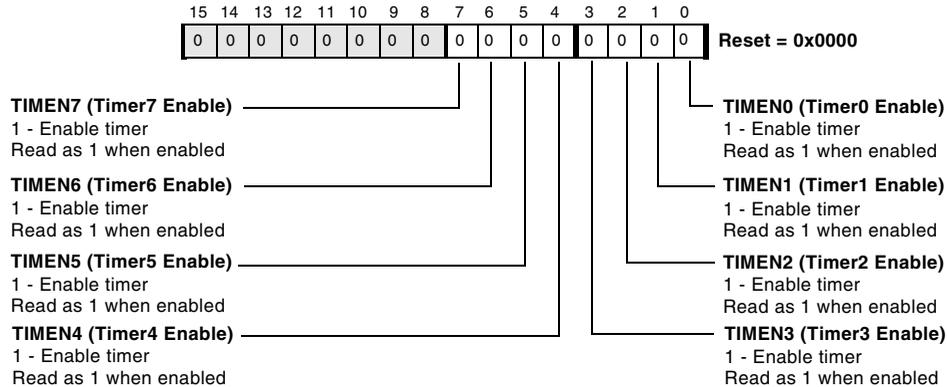
The size of accesses is enforced. A 32-bit access to a `TIMER_CONFIG` register or a 16-bit access to a `TIMER_WIDTH`, `TIMER_PERIOD`, or `TIMER_COUNTER` register results in a memory-mapped register (MMR) error. Both 16- and 32-bit accesses are allowed for the `TIMER_ENABLE`, `TIMER_DISABLE`, and `TIMER_STATUS` registers. On a 32-bit read of one of the 16-bit registers, the upper word returns all 0s.

Timer Enable Register (`TIMER_ENABLE`)

[Figure 10-16](#) shows an example of the `TIMER_ENABLE` register for a product with eight timers. The register allows simultaneous enabling of multiple timers so that they can 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 bits can be set individually or in any combination. A read of the `TIMER_ENABLE` register 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 Registers

Timer Enable Register (TIMER_ENABLE)



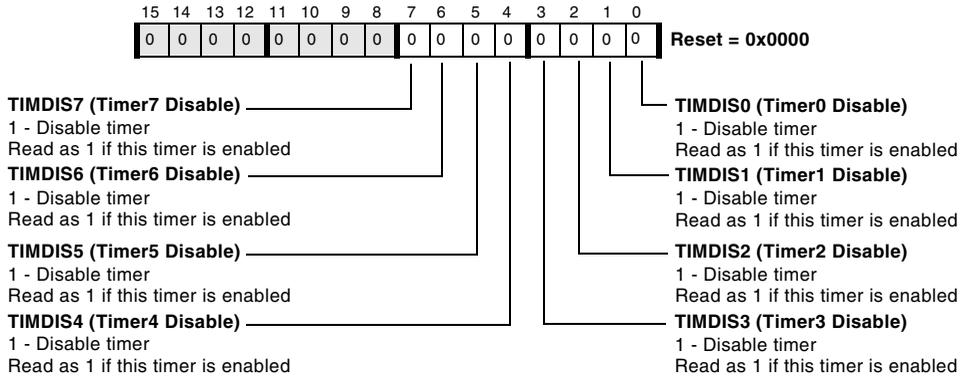
This diagram shows an example configuration for eight timers. Different products have different numbers of timers.

Figure 10-16. Timer Enable Register

Timer Disable Register (TIMER_DISABLE)

Figure 10-17 shows an example of the `TIMER_DISABLE` register for a product with eight timers. The register allows simultaneous disabling of multiple timers. For each timer there is a single `W1C` control bit. Writing a “1” disables the corresponding timer; writing a “0” has no effect. The bits can be cleared individually or in any combination. A read of the `TIMER_DISABLE` register returns a value identical to a read of the `TIMER_ENABLE` register. A “1” indicates that the timer is enabled. All unused bits return “0” when read.

Timer Disable Register (TIMER_DISABLE)



This diagram shows an example configuration for eight timers. Different products have different numbers of timers.

Figure 10-17. Timer Disable Register

In PWM_OUT mode, a write of a “1” to TIMER_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 TRUN bit in TIMER_STATUS. See [“Stopping the Timer in PWM_OUT Mode” on page 10-21](#).

In WDT_CAP and EXT_CLK modes, a write of a “1” to TIMER_DISABLE stops the corresponding timer immediately.

Timer Status Register (TIMER_STATUS)

The TIMER_STATUS register indicates the status of the timers and is used to check the status of multiple timers 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 TIMER_STATUS

Timer Registers

register read access, all reserved or unused bits return a “0”. [Figure 10-18](#) shows an example of the `TIMER_STATUS` register for a product with eight timers.

For detailed behavior and usage of the `TRUN` bit see [“Stopping the Timer in PWM_OUT Mode” on page 10-21](#). Writing the `TRUN` bits has no effect in other modes or when a timer has not been enabled. Writing the `TRUN` bits to “1” in `PWM_OUT` mode has no effect on a timer that has not first been disabled.

Error conditions are explained in [“Illegal States” on page 10-7](#).

Timer Configuration Register (TIMER_CONFIG)

The operating mode for each timer is specified by its `TIMER_CONFIG` register. The `TIMER_CONFIG` register, shown in [Figure 10-19](#), 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 `TIMER_STATUS` before attempting to reprogram `TIMER_CONFIG`. The `TIMER_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.

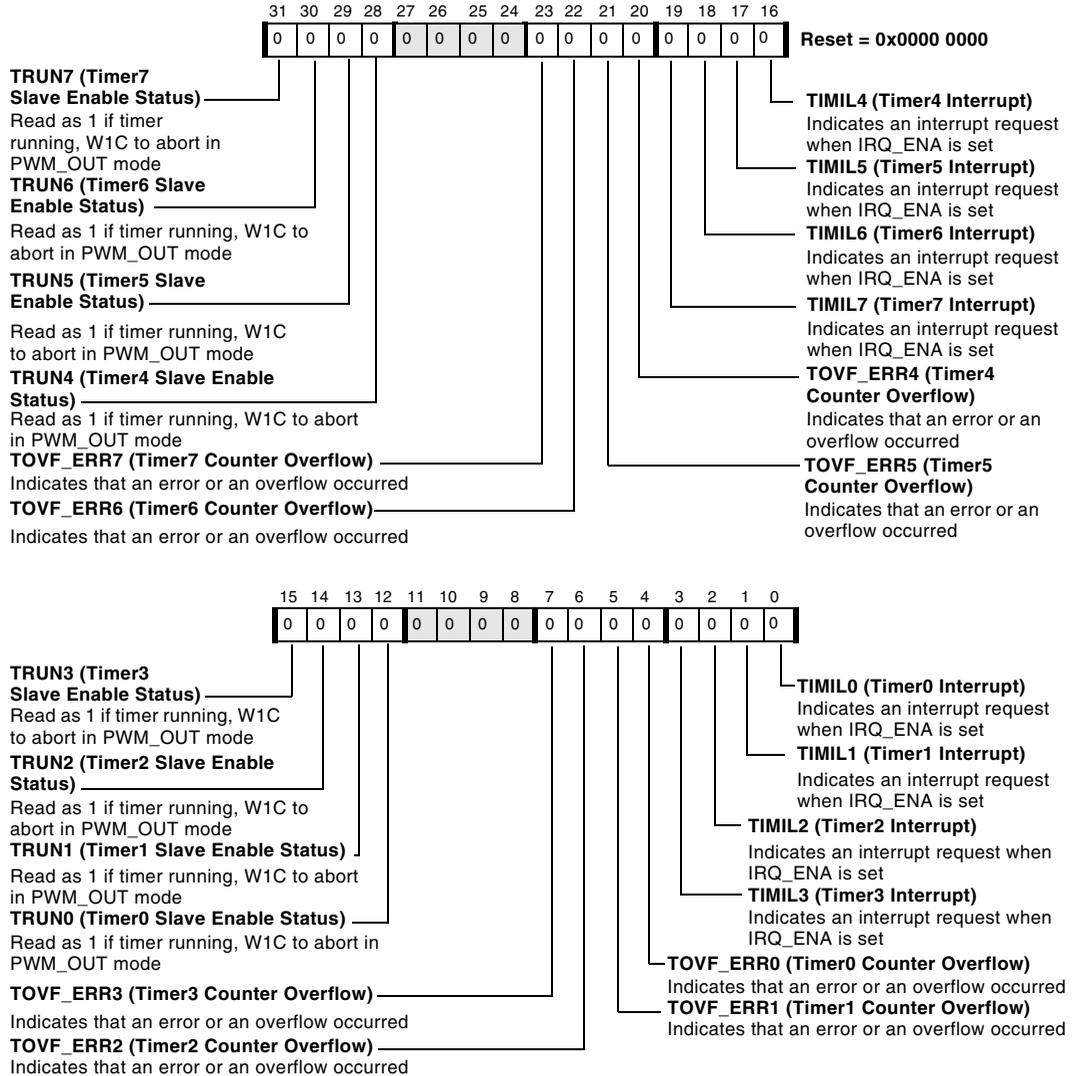
Each time `TOVF_ERR` is set, `ERR_TYP[1:0]` is loaded with a code that identifies the type of error that was detected. This value is held until the next error or timer enable occurs. For an overview of error conditions, see [Table 10-1 on page 10-9](#). The `TIMER_CONFIG` register also controls the behavior of the `TMR` pin, which becomes an output in `PWM_OUT` mode (`TMODE = 01`) when the `OUT_DIS` bit is cleared.



When operating the PPI in GP output modes with internal frame syncs, the `CLK_SEL` and the `TIN_SEL` bits for the timers involved must be set to “1”.

Timer Status Register (TIMER_STATUS)

All bits are W1C



This diagram shows an example configuration for eight timers. Different products have different numbers of timers, therefore some of the bits may not be applicable to your device.

Figure 10-18. Timer Status Register

Timer Registers

Timer Configuration Register (TIMER_CONFIG)

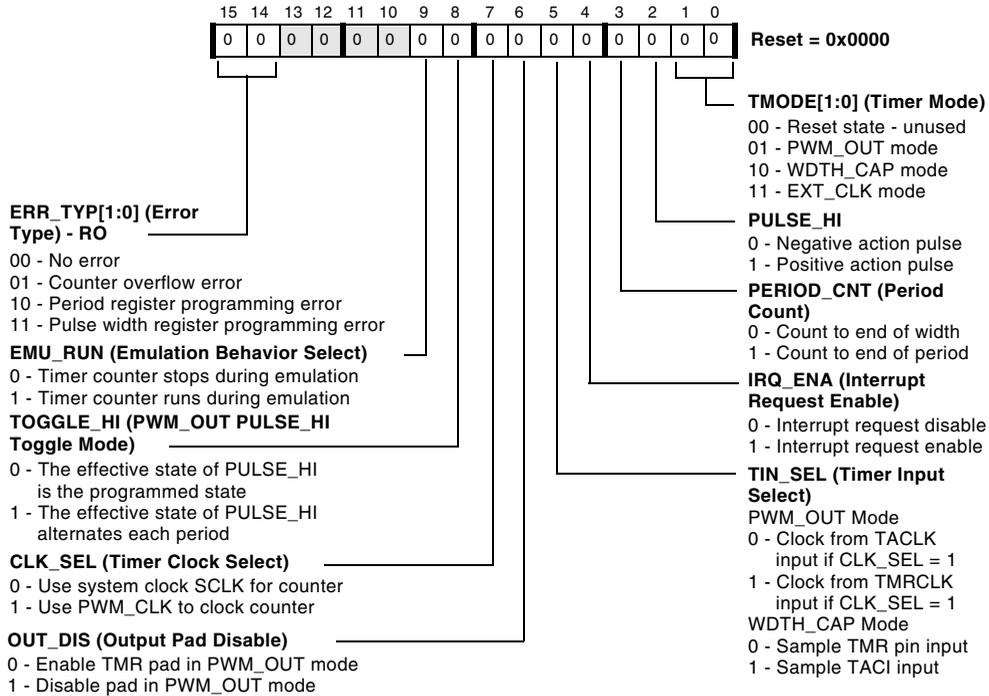


Figure 10-19. Timer Configuration Register

Timer Counter Register (TIMER_COUNTER)

This read-only register retains its state when disabled. When enabled, the `TIMER_COUNTER` register is reinitialized by hardware based on configuration and mode. The `TIMER_COUNTER` register, shown in [Figure 10-20](#), may be read at any time (whether the timer is running or stopped), and it returns an atomic 32-bit value. Depending on the operating mode, the incrementing counter can be clocked by four different sources: `SCLK`, the TMR pin, the alternative timer clock pin `TACLK`, or the common `TMRCLK` pin, which is most likely used as the PPI clock (`PPI_CLK`).

While the processor core is being accessed by an external emulator debugger, all code execution stops. By default, the `TIMER_COUNTER` register also halts its counting during an emulation access in order to remain synchronized with the software. While stopped, the count does not advance—in `PWM_OUT` mode, the `TMR` pin waveform is “stretched”; in `WDTH_CAP` mode, measured values are incorrect; in `EXT_CLK` mode, input events on the `TMR` pin may be missed. All other timer functions such as register reads and writes, interrupts previously asserted (unless cleared), and the loading of `TIMER_PERIOD` and `TIMER_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 `TIMER_CONFIG` to enable this behavior.

Timer Counter Register (`TIMER_COUNTER`)

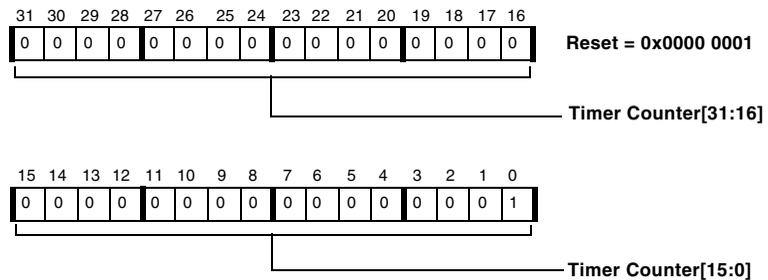


Figure 10-20. Timer Counter Register

Timer Period (`TIMER_PERIOD`) and Timer Width (`TIMER_WIDTH`) Registers

i When a timer is enabled and running, and the software writes new values to the `TIMER_PERIOD` register and the `TIMER_WIDTH` register, the writes are buffered and do not update the registers until the end of the current period (when `TIMER_COUNTER` equals `TIMER_WIDTH`).

Timer Registers

Usage of the `TIMER_PERIOD` register, shown in [Figure 10-21](#), and the `TIMER_WIDTH` register, shown in [Figure 10-22](#), varies depending on the mode of the timer:

- In `PWM_OUT` mode, both the `TIMER_PERIOD` and `TIMER_WIDTH` register values can be updated “on-the-fly” since the values change simultaneously.
- In `WIDTH_CAP` mode, the timer period and timer pulse width buffer values are captured at the appropriate time. The `TIMER_PERIOD` and `TIMER_WIDTH` registers are then updated simultaneously from their respective buffers. Both registers are read-only in this mode.
- In `EXT_CLK` mode, the `TIMER_PERIOD` register is writable and can be updated “on-the-fly.” The `TIMER_WIDTH` register is not used.

If new values are not written to the `TIMER_PERIOD` register or the `TIMER_WIDTH` register, the value from the previous period is reused. Writes to the 32-bit `TIMER_PERIOD` register and `TIMER_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_WIDTH` registers are always stored in the buffer registers. Reads from the `TIMER_PERIOD` or `TIMER_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, they do not become active until after the `TIMER_PERIOD` and `TIMER_WIDTH` registers are updated from their respective buffers at the end of the current period. See [Figure 10-1 on page 10-3](#).

When the timer is disabled, writes to the buffer registers are immediately copied through to the `TIMER_PERIOD` or `TIMER_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_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_WIDTH` register. The next period may use one old value and one new value. In order to prevent “pulse width \geq period” errors, write the `TIMER_WIDTH` register before the `TIMER_PERIOD` register when decreasing the values, and write the `TIMER_PERIOD` register before the `TIMER_WIDTH` register when increasing the value.

Timer Period Register (TIMER_PERIOD)

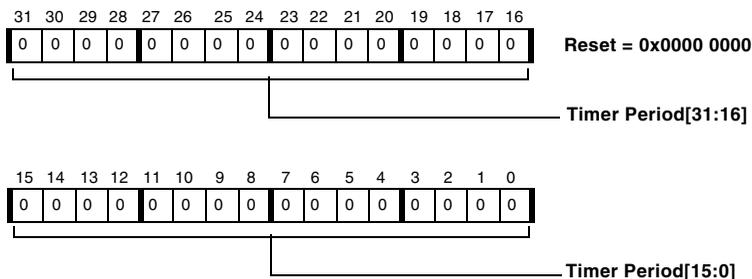


Figure 10-21. Timer Period Register

Timer Registers

Timer Width Register (TIMER_WIDTH)

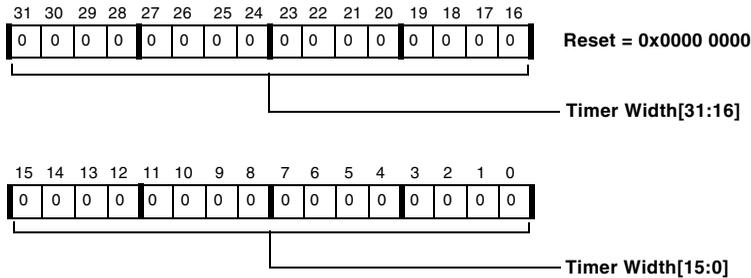


Figure 10-22. Timer Width Register

Summary

Table 10-2 summarizes control bit and register usage in each timer mode.

Table 10-2. Control Bit and Register Usage Chart

Bit / Register	PWM_OUT Mode	WDTH_CAP Mode	EXT_CLK Mode
TIMER_ENABLE	1 - Enable timer 0 - No effect	1 - Enable timer 0 - No effect	1 - Enable timer 0 - No effect
TIMER_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

Table 10-2. Control Bit and Register Usage Chart (Continued)

Bit / Register	PWM_OUT Mode	WDTH_CAP Mode	EXT_CLK Mode
TIN_SEL	Depends on CLK_SEL: If CLK_SEL = 1, 1 - Count TMRCLK clocks 0 - Count TACKL clocks If CLK_SEL = 0, Unused	1 - Select TACI input 0 - Select TMR pin input	Unused
OUT_DIS	1 - Disable TMR pin 0 - Enable TMR pin	Unused	Unused
CLK_SEL	1 - PWM_CLK clocks timer 0 - SCLK clocks timer	Unused	Unused
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

Timer Registers

Table 10-2. Control Bit and Register Usage Chart (Continued)

Bit / Register	PWM_OUT Mode	WIDTH_CAP Mode	EXT_CLK Mode
Counter	RO: Counts up on SCLK or PWM_CLK	RO: Counts up on SCLK	RO: Counts up on TMR pin 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
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

Programming Examples

Listing 10-1 configures the port control registers in a way that enables TMR pins associated with Port G. This example assumes TMR1-7 are connected to Port G bits 5–11.

Listing 10-1. Port Setup

```
timer_port_setup:
    [--sp] = (r7:7, p5:5);
    p5.h = hi(PORTG_FER);
    p5.l = lo(PORTG_FER);
    r7.l = PG5|PG6|PG7|PG8|PG9|PG10|PG11;
    w[p5] = r7;
    p5.l = lo(PORTG_MUX);
    r7.l = PFTE;
    w[p5] = r7;
    (r7:7, p5:5) = [sp++];
    rts;
timer_port_setup.end;
```

Listing 10-2 generates signals on the TMR4 and TMR5 outputs. By default, timer 5 generates a continuous PWM signal with a duty cycle of 50% (period = 0x40 SCLKs, width = 0x20 SCLKs) while the PWM signal generated by timer 4 has the same period but 25% duty cycle (width = 0x10 SCLKs).

If the preprocessor constant `SINGLE_PULSE` is defined, every TMR pin outputs only a single high pulse of 0x20 (timer 4) and 0x10 SCLKs (timer 5) duration.

In any case the timers are started synchronously and the rising edges are aligned. That is, the pulses are left aligned.

Programming Examples

Listing 10-2. Signal Generation

```
// #define SINGLE_PULSE
timer45_signal_generation:
    [--sp] = (r7:7, p5:5);
    p5.h = hi(TIMER_ENABLE);
    p5.l = lo(TIMER_ENABLE);
#ifdef SINGLE_PULSE
    r7.l = PULSE_HI | PWM_OUT;
#else
    r7.l = PERIOD_CNT | PULSE_HI | PWM_OUT;
#endif
    w[p5 + TIMER5_CONFIG - TIMER_ENABLE] = r7;
    w[p5 + TIMER4_CONFIG - TIMER_ENABLE] = r7;
    r7 = 0x10 (z);
    [p5 + TIMER5_WIDTH - TIMER_ENABLE] = r7;
    r7 = 0x20 (z);
    [p5 + TIMER4_WIDTH - TIMER_ENABLE] = r7;
#ifdef SINGLE_PULSE
    r7 = 0x40 (z);
    [p5 + TIMER5_PERIOD - TIMER_ENABLE] = r7;
    [p5 + TIMER4_PERIOD - TIMER_ENABLE] = r7;
#endif
    r7.l = TIMEN5 | TIMEN4;
    w[p5] = r7;
    (r7:7, p5:5) = [sp++];
    rts;
timer45_signal_generation.end;
```

All subsequent examples use interrupts. Therefore, [Listing 10-3](#) illustrates how interrupts are generated and how interrupt service routines can be registered. In this example, the timer 5 interrupt is assigned to the IVG12 interrupt channel of the CEC controller.

Listing 10-3. Interrupt Setup

```

timer5_interrupt_setup:
    [--sp] = (r7:7, p5:5);
    p5.h = hi(IMASK);
    p5.l = lo(IMASK);
/* register interrupt service routine */
    r7.h = hi(isr_timer5);
    r7.l = lo(isr_timer5);
    [p5 + EVT12 - IMASK] = r7;
/* unmask IVG12 in CEC */
    r7 = [p5];
    bitset(r7, bitpos(EVT_IVG12));
    [p5] = r7;
/* assign timer 5 IRQ (= IRQ37 in this example) to IVG12 */
    p5.h = hi(SIC_IAR4);
    p5.l = lo(SIC_IAR4);
/*SIC_IAR register mapping is processor dependent*/
    r7.h = 0xFF5F;
    r7.l = 0xFFFF;
    [p5] = r7;
/* enable timer 5 IRQ */
    p5.h = hi(SIC_IMASK1);
    p5.l = lo(SIC_IMASK1);
/*SIC_IMASK register mapping is processor dependent*/
    r7 = [p5];
    bitset(r7, 5);
    [p5] = r7;
/* enable interrupt nesting */
    (r7:7, p5:5) = [sp++];
    [--sp] = reti;
    rts;
timer5_interrupt_setup.end:

```

Programming Examples

The example shown in [Listing 10-4](#) does not drive the TMR pin. It generates periodic interrupt requests every 0x1000 SCLK cycles. If the preprocessor constant `SINGLE_PULSE` was defined, timer 5 requests an interrupt only once. Unlike in a real application, the purpose of the interrupt service routine shown in this example is just the clearing of the interrupt request and counting interrupt occurrences.

Listing 10-4. Periodic Interrupt Requests

```
// #define SINGLE_PULSE
timer5_interrupt_generation:
    [--sp] = (r7:7, p5:5);
    p5.h = hi(TIMER_ENABLE);
    p5.l = lo(TIMER_ENABLE);
#ifdef SINGLE_PULSE
    r7.l = EMU_RUN | IRQ_ENA | OUT_DIS | PWM_OUT;
#else
    r7.l = EMU_RUN | IRQ_ENA | PERIOD_CNT | OUT_DIS | PWM_OUT;
#endif
    w[p5 + TIMER5_CONFIG - TIMER_ENABLE] = r7;
    r7 = 0x1000 (z);
#ifdef SINGLE_PULSE
    [p5 + TIMER5_PERIOD - TIMER_ENABLE] = r7;
    r7 = 0x1 (z);
#endif
    [p5 + TIMER5_WIDTH - TIMER_ENABLE] = r7;
    r7.l = TIMEN5;
    w[p5] = r7;
    (r7:7, p5:5) = [sp++];
    r0 = 0 (z);
    rts;
timer5_interrupt_generation.end:
isr_timer5:
    [--sp] = astat;
```

```

[sp] = (r7:7, p5:5);
p5.h = hi(TIMER_STATUS);
p5.l = lo(TIMER_STATUS);
r7.h = hi(TIMIL5);
r7.l = lo(TIMIL5);
[p5] = r7;
r0 += 1;
ssync;
(r7:7, p5:5) = [sp++];
astat = [sp++];
rti;
isr_timer5.end;

```

Listing 10-5 illustrates how two timers can generate two non-overlapping clock pulses as typically required for break-before-make scenarios. Both timers are running in PWM_OUT mode with `PERIOD_CNT = 1` and `PULSE_HI = 1`.

Figure 10-23 explains how the signal waveform represented by the period `P` and the pulse width `W` translates to timer period and width values. **Table 10-3** summarizes the register writes.

Table 10-3. Register Writes for Non-Overlapping Clock Pulses

Register	Before Enable	After Enable	At IRQ1	At IRQ2
TIMER5_PERIOD	$P/2$			
TIMER5_WIDTH	$P/2 - W/2$	$W/2$	$P/2 - W/2$	$W/2$
TIMER4_PERIOD	P	$P/2$		
TIMER4_WIDTH	$P - W/2$		$W/2$	$P/2 - W/2$

Since hardware only updates the written period and width values at the end of periods, software can write new values immediately after the timers have been enabled. Note that both timers' period expires at exactly the

Programming Examples

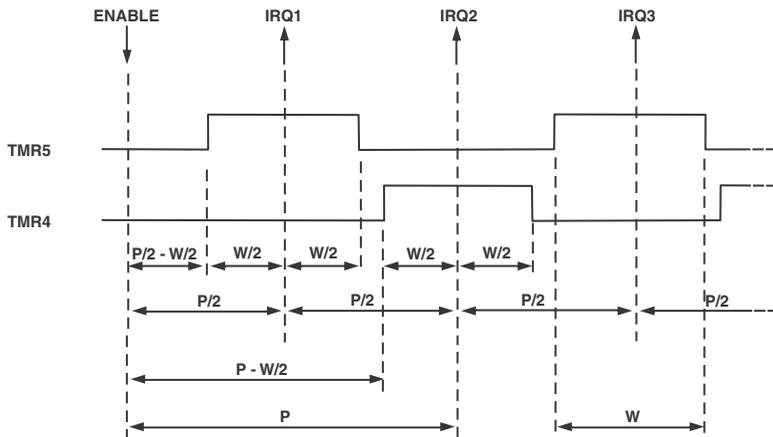


Figure 10-23. Non-Overlapping Clock Pulses

same times with the exception of the first timer 5 interrupt (at IRQ1) which is not visible to timer 4.

Listing 10-5. Non-Overlapping Clock Pulses

```
#define P 0x1000 /* signal period */
#define W 0x0600 /* signal pulse width */
#define N 4 /* number of pulses before disable */
timer45_toggle_hi:
    [--sp] = (r7:1, p5:5);
    p5.h = hi(TIMER_ENABLE);
    p5.l = lo(TIMER_ENABLE);
/* config timers */
    r7.l = IRQ_ENA | PERIOD_CNT | TOGGLE_HI | PULSE_HI | PWM_OUT;
    w[p5 + TIMER5_CONFIG - TIMER_ENABLE] = r7;
    r7.l = PERIOD_CNT | TOGGLE_HI | PULSE_HI | PWM_OUT;
    w[p5 + TIMER4_CONFIG - TIMER_ENABLE] = r7;
/* calculate timers widths and period */
    r0.l = lo(P);
```

```

    r0.h = hi(P);
    r1.l = lo(W);
    r1.h = hi(W);
    r2 = r1 >> 1;    /* W/2 */
    r3 = r0 >> 1;    /* P/2 */
    r4 = r3 - r2;    /* P/2 - W/2 */
    r5 = r0 - r2;    /* P - W/2 */
/* write values for initial period */
    [p5 + TIMER4_PERIOD - TIMER_ENABLE] = r0;
    [p5 + TIMER4_WIDTH - TIMER_ENABLE] = r5;
    [p5 + TIMER5_PERIOD - TIMER_ENABLE] = r3;
    [p5 + TIMER5_WIDTH - TIMER_ENABLE] = r4;
/* start timers */
    r7.l = TIMEN5 | TIMEN4 ;
    w[p5 + TIMER_ENABLE - TIMER_ENABLE] = r7;
/* write values for second period */
    [p5 + TIMER4_PERIOD - TIMER_ENABLE] = r3;
    [p5 + TIMER5_WIDTH - TIMER_ENABLE] = r2;
/* r0 functions as signal period counter */
    r0.h = hi(N * 2 - 1);
    r0.l = lo(N * 2 - 1);
    (r7:1, p5:5) = [sp++];
    rts;
timer45_toggle_hi.end:
isr_timer5:
    [--sp] = astat;
    [--sp] = (r7:5, p5:5);
    p5.h = hi(TIMER_ENABLE);
    p5.l = lo(TIMER_ENABLE);
/* clear interrupt request */
    r7.h = hi(TIMIL5);
    r7.l = lo(TIMIL5);
    [p5 + TIMER_STATUS - TIMER_ENABLE] = r7;
/* toggle width values (width = period - width) */

```

Programming Examples

```
r7 = [p5 + TIMER5_PERIOD - TIMER_ENABLE];
r6 = [p5 + TIMER5_WIDTH - TIMER_ENABLE];
r5 = r7 - r6;
[p5 + TIMER5_WIDTH - TIMER_ENABLE] = r5;
r5 = [p5 + TIMER4_WIDTH - TIMER_ENABLE];
r7 = r7 - r5;
CC = r7 < 0;
if CC r7 = r6;
[p5 + TIMER4_WIDTH - TIMER_ENABLE] = r7;
/* disable after a certain number of periods */
r0+= -1;
CC = r0 == 0;
r5.l = 0;
r7.l = TIMDIS5 | TIMDIS4;
if !CC r7 = r5;
w[p5 + TIMER_DISABLE - TIMER_ENABLE] = r7;
(r7:5, p5:5) = [sp++];
astat = [sp++];
rti;
isr_timer5.end;
```

Listing 10-5 generates N pulses on both timer output pins. Disabling the timers does not corrupt the generated pulse pattern anyhow.

Listing 10-6 configures timer 5 in WDTM_CAP mode. If looped back externally, this code might be used to receive N PWM patterns generated by one of the other timers. Ensure that the PWM generator and consumer both use the same PERIOD_CNT and PULSE_HI settings.

Listing 10-6. Timer Configured in WDTM_CAP Mode

```
.section L1_data_a;
.align 4;
#define N 1024
.var buffReceive[N*2];
```

```

.section L1_code;
timer5_capture:
    [--sp] = (r7:7, p5:5);
/* setup DAG2 */
    r7.h = hi(buffReceive);
    r7.l = lo(buffReceive);
    i2 = r7;
    b2 = r7;
    l2 = length(buffReceive)*4;
/* config timer for high pulses capture */
    p5.h = hi(TIMER_ENABLE);
    p5.l = lo(TIMER_ENABLE);
    r7.l = EMU_RUN|IRQ_ENA|PERIOD_CNT|PULSE_HI|WDTH_CAP;
    w[p5 + TIMER5_CONFIG - TIMER_ENABLE] = r7;
    r7.l = TIMEN5;
    w[p5 + TIMER_ENABLE - TIMER_ENABLE] = r7;
    (r7:7, p5:5) = [sp++];
    rts;
timer5_capture.end:
isr_timer5:
    [--sp] = astat;
    [--sp] = (r7:7, p5:5);
/* clear interrupt request first */
    p5.h = hi(TIMER_STATUS);
    p5.l = lo(TIMER_STATUS);
    r7.h = hi(TIMIL5);
    r7.l = lo(TIMIL5);
    [p5] = r7;
    r7 = [p5 + TIMER5_PERIOD - TIMER_STATUS];
    [i2++] = r7;
    r7 = [p5 + TIMER5_WIDTH - TIMER_STATUS];
    [i2++] = r7;
    ssync;
    (r7:7, p5:5) = [sp++];

```

Unique Information for the ADSP-BF51x Processor

```
    astat = [sp++];  
    rti;  
isr_timer5.end:
```

Unique Information for the ADSP-BF51x Processor

The ADSP-BF51x processor features one general-purpose timer module that contains eight identical 32-bit timers. Each timer can be individually configured to operate in various modes. Although the timers operate completely independently of each other, all of them can be started and stopped simultaneously for synchronous operation.

Interface Overview

[Figure 10-24](#) shows the ADSP-BF51x specific block diagram of the general-purpose timer module.

External Interface

The TMRCLK input is common to all eight timers. The PPI unit is clocked by the same pin; therefore any of the timers can be clocked by PPI_CLK. Since timer 0 and timer 1 are often used in conjunction with the PPI, they are internally looped back to the PPI module for frame sync generation.

The timer signals TMR0 and TMR1 are multiplexed with the PPI frame syncs when the frame syncs are applied externally. PPI modes requiring only one frame sync free up TMR1. For details, see the *Parallel Peripheral Interface* chapter.

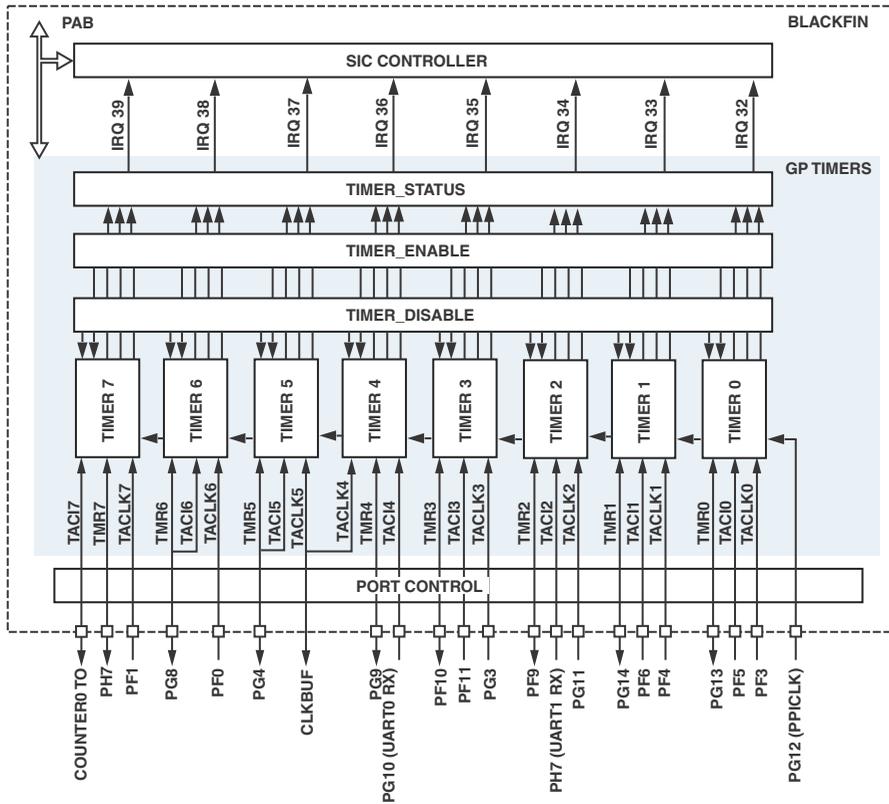


Figure 10-24. Timer Block Diagram



If the PPI frame syncs are applied externally, timer 0 and timer 1 are still fully functional and can be used for other purposes not involving the TMR_x pins. Timer 0 and timer 1 must not drive their TMR₀ and TMR₁ pins. If operating in PWM_OUT mode, the OUT_DIS bit in the TIMER₀_CONFIG and TIMER₁_CONFIG registers must be set.

Unique Information for the ADSP-BF51x Processor

11 CORE TIMER

This chapter describes the core timer. Following an overview, functional description, and consolidated register definitions, the chapter concludes with a programming example.

Specific Information for the ADSP-BF51x

For details regarding the number of core timers for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For Core Timer interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

For a list of MMR addresses for each Core Timer, refer to [Appendix A, “System MMR Assignments”](#).

Core timer behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 11-9](#).

Overview and Features

The core timer is a programmable 32-bit interval timer which can generate periodic interrupts. Unlike other peripherals, the core timer resides

Timer Overview

inside the Blackfin core and runs at the core clock (CCLK) rate. Core timer features include:

- 32-bit timer with 8-bit prescaler
- Operates at core clock (CCLK) rate
- Dedicated high-priority interrupt channel
- Single-shot or continuous operation

Timer Overview

Figure 11-1 provides a block diagram of the core timer.

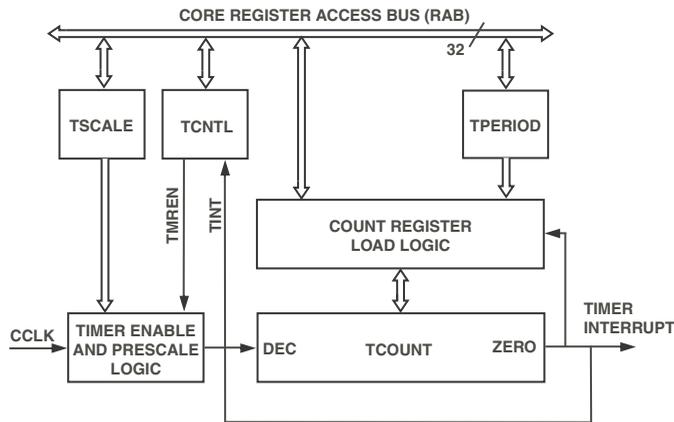


Figure 11-1. Core Timer Block Diagram

External Interfaces

The core timer does not directly interact with any pins of the chip.

Internal Interfaces

The core timer is accessed through the 32-bit register access bus (RAB). The module is clocked by the core clock CCLK. The timer's dedicated interrupt request is a higher priority than requests from all other peripherals.

Description of Operation

The software should initialize the TCOUNT register *before* the timer is enabled. The TCOUNT register can be written directly, but writes to the TPERIOD register are also passed through to TCOUNT.

When the timer is enabled by setting the TMREN bit in the core timer control register (TCNTL), the TCOUNT register is decremented once every time the prescaler TSCALE expires, that is, every TSCALE + 1 number of CCLK 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. If the TAUTORLD bit is not set, the timer stops operation.

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.

Interrupt Processing

The timer's dedicated interrupt request is a higher priority than requests from all other peripherals. The request goes directly to the core event controller (CEC) and does not pass through the system interrupt controller

Core Timer Registers

(SIC). Therefore, the interrupt processing is also completely in the CCLK domain.

 The core timer interrupt request is edge-sensitive and cleared by hardware automatically as soon as the interrupt is serviced.

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.

Core Timer Registers

The core timer includes four core memory-mapped registers, the timer control register (TCNTL), the timer count register (TCOUNT), the timer period register (TPERIOD), and the timer scale register (TSCALE). As with all core MMRs, these registers are always accessed by 32-bit read and write operations.

Core Timer Control Register (TCNTL)

The TCNTL register, shown in Figure 11-2, functions as control and status register.

Core Timer Control Register (TCNTL)

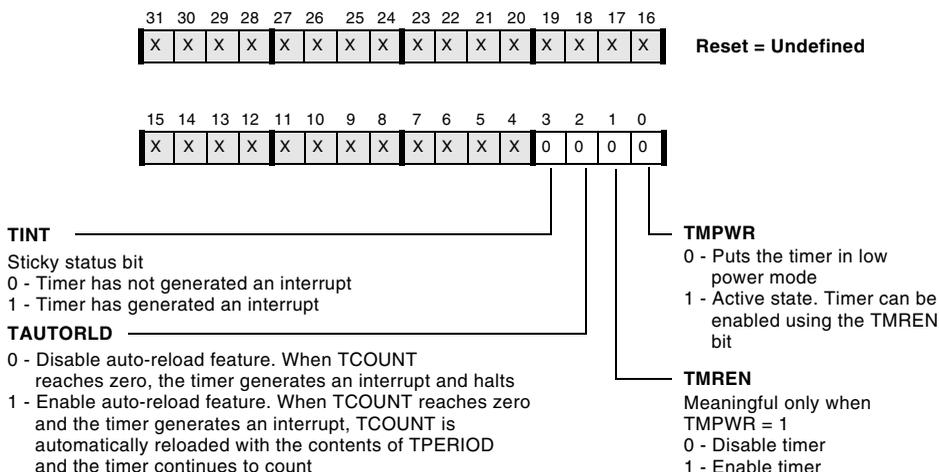


Figure 11-2. Core Timer Control Register

Core Timer Count Register (TCOUNT)

The TCOUNT register, shown in Figure 11-3, 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.

Values written to the TPERIOD register are automatically copied to the TCOUNT register. Nevertheless, the TCOUNT register can be written directly. In auto reload mode the value written to TCOUNT may differ from the TPERIOD value to let the initial period be shorter or longer than following periods. To do this, write to TPERIOD first and overwrite TCOUNT afterward.

Core Timer Registers

Writes to `TCOUNT` are ignored once the timer is running.

Core Timer Count Register (TCOUNT)

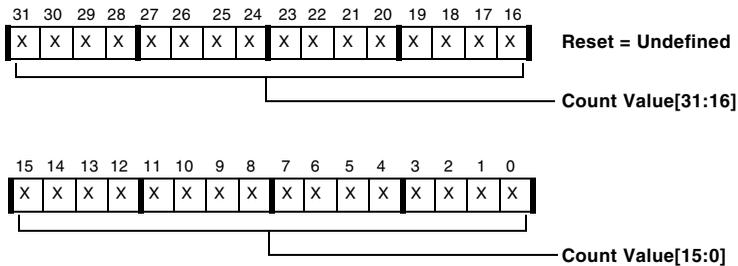


Figure 11-3. Core Timer Count Register

Core Timer Period Register (TPERIOD)

The `TPERIOD` register is shown in [Figure 11-4](#). When auto-reload is enabled, the `TCOUNT` register is reloaded with the value of the `TPERIOD` register whenever `TCOUNT` reaches 0. Writes to `TPERIOD` are ignored when the timer is running.

Core Timer Period Register (TPERIOD)

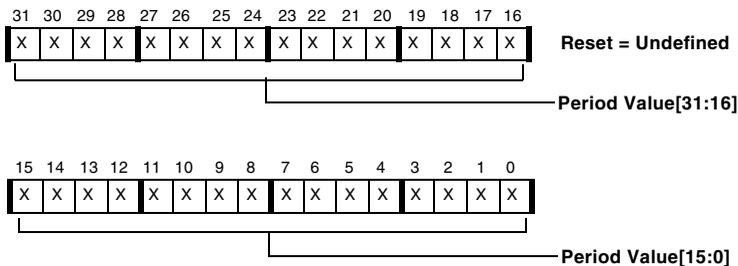


Figure 11-4. Core Timer Period Register

Core Timer Scale Register (TSCALE)

The `TSCALE` register is shown in [Figure 11-5](#). The register 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 0, the counter register decrements once every `CCLK` clock cycle. If `TSCALE` is 1, the counter decrements once every two cycles.

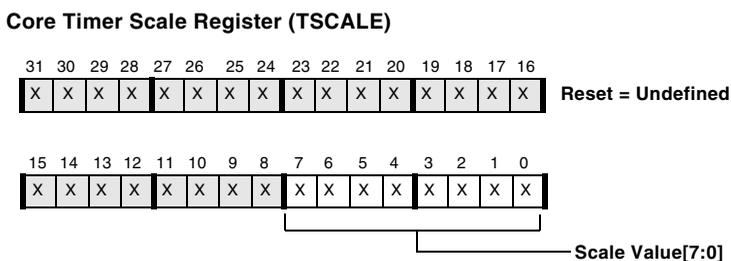


Figure 11-5. Core Timer Scale Register

Programming Examples

[Listing 11-1](#) configures the core timer in auto-reload mode. Assuming a `CCLK` of 500 MHz, the resulting period is 1 second. The initial period is twice as long as the others.

Listing 11-1. Core Timer Configuration

```
#include <defBF527.h> /*ADSP-BF527 product is used as an example*/
.section L1_code;
.global _main;
_main:
/* Register service routine at EVT6 and unmask interrupt */
    pl.l = lo(IMASK);
    pl.h = hi(IMASK);
```

Programming Examples

```
    r0.l = lo(isr_core_timer);
    r0.h = hi(isr_core_timer);
    [p1 + EVT6 - IMASK] = r0;
    r0 = [p1];
    bitset(r0, bitpos(EVT_IVTMR));
    [p1] = r0;
/* Prescaler = 50, Period = 10,000,000, First Period = 20,000,000
*/
    p1.l = lo(TCNTL);
    p1.h = hi(TCNTL);
    r0 = 50 (z);
    [p1 + TSCALE - TCNTL] = r0;
    r0.l = lo(10000000);
    r0.h = hi(10000000);
    [p1 + TPERIOD - TCNTL] = r0;
    r0 <<= 1;
    [p1 + TCOUNT - TCNTL] = r0;
/* R6 counts interrupts */
    r6 = 0 (z);
/* start in auto-reload mode */
    r0 = TAUTORLD | TMPWR | TMREN (z);
    [p1] = r0;
_main.forever:
    jump _main.forever;
_main.end:
/* interrupt service routine simple increments R6 */
_isr_core_timer:
    [--sp] = astat;
    r6+= 1;
    astat = [sp++];
    rti;
_isr_core_timer.end:
```

Unique Information for the ADSP-BF51x Processor

None.

Unique Information for the ADSP-BF51x Processor

12 WATCHDOG TIMER

This chapter describes the watchdog timer. Following an overview, functional description, and consolidated register definitions, the chapter concludes with programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of watchdog timers for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For Watchdog Timer interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

For a list of MMR addresses for each Watchdog Timer, refer to [Appendix A, “System MMR Assignments”](#).

Watchdog timer behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 12-11](#).

Overview and Features

The processor includes a 32-bit timer 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 watchdog expires before being updated by software.

Overview and Features

Watchdog timer key features include:

- 32-bit watchdog timer
- 8-bit disable bit pattern
- System reset on expire option
- NMI on expire option
- General-purpose interrupt option

Typically, the watchdog timer is used to supervise stability of the system software. When used in this way, software reloads the watchdog timer in a regular manner so that the downward counting timer never expires (never becomes 0). An expiring timer then indicates that system software might be out of control. At this point a special error handler may recover the system. For safety, however, it is often better to reset and reboot the system directly by hardware control.

Especially in slave boot configurations, a processor reset cannot automatically force the Blackfin device to be rebooted. In this case, the processor may reset without booting again and may negotiate with the host device by the time program execution starts. Alternatively, a watchdog event can cause an NMI event. The NMI service routine may request the host device reset and/or reboot the Blackfin processor.

The watchdog timer is often programmed to let the processor wake up from sleep mode after a programmable period of time.



For easier debugging, the watchdog timer does not decrement (even if enabled) when the processor is in emulation mode.

Interface Overview

Figure 12-1 provides a block diagram of the watchdog timer.

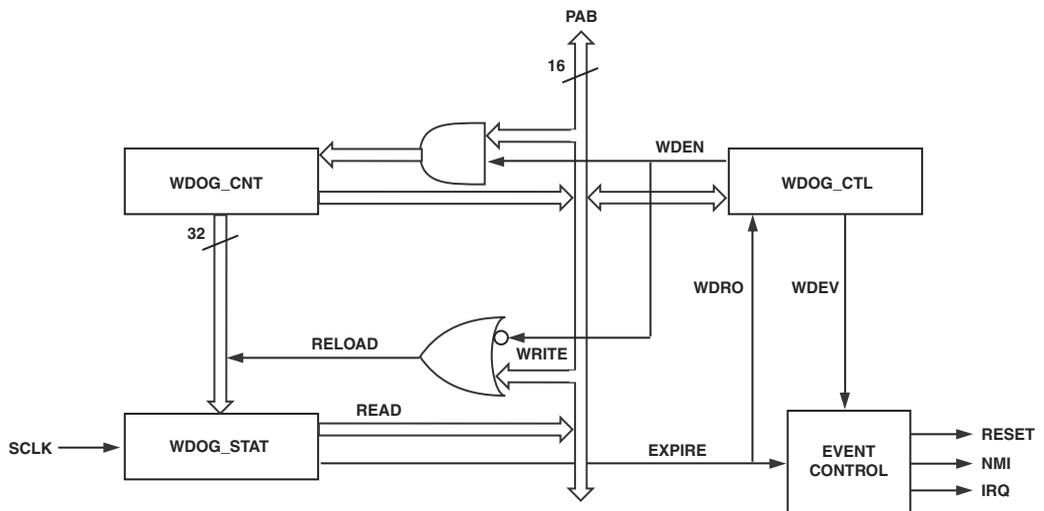


Figure 12-1. Watchdog Timer Block Diagram

External Interface

The watchdog timer does not directly interact with any pins of the chip.

Internal Interface

The watchdog timer is clocked by the system clock `SCLK`. Its registers are accessed through the 16-bit peripheral access bus (PAB). The 32-bit registers `WDOG_CNT` and `WDOG_STAT` must always be accessed by 32-bit read/write operations. Hardware ensures that those accesses are atomic.

When the counter expires, one of three event requests can be generated. Either a reset or an NMI request is issued to the core event controller

Description of Operation

(CEC) or a general-purpose interrupt request is passed to the system interrupt controller (SIC).

Description of Operation

If enabled, the 32-bit watchdog timer counts downward every `SCLK` cycle. If it becomes 0, one of three event requests can be issued to either the CEC or the SIC. Depending on how the `WDEV` bit field in the `WDOG_CTL` register is programmed, the event that is generated may be a reset, a non-maskable interrupt, or a general-purpose interrupt.

The counter value can be read through the 32-bit `WDOG_STAT` register. The `WDOG_STAT` register cannot, however, be written directly. Rather, software writes the watchdog period value into the 32-bit `WDOG_CNT` register *before* the watchdog is enabled. Once the watchdog is started, the period value cannot be altered.

To start the watchdog timer:

1. Set the count value for the watchdog timer by writing the count value into the watchdog count register (`WDOG_CNT`). Since the watchdog timer is not enabled yet, the write to the `WDOG_CNT` registers automatically pre-loads the `WDOG_STAT` register as well.
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.

If software does not service the watchdog in time, `WDOG_STAT` continues decrementing until it reaches 0. Then, the programmed event is generated. The counter stops decrementing and remains at zero. Additionally, the `WDRO` latch bit in the `WDOG_CTL` register is set and can be interrogated by software in case event generation is not enabled.

When the watchdog is programmed to generate a reset, it resets the processor core and peripherals. If the `NOBOOT` bit in the `SYSCR` register was set by the time the watchdog resets the part, the chip is not rebooted. This is recommended behavior in slave boot configurations. The reset handler may evaluate the `RESET_WDOG` bit in the software reset register `SWRST` to detect a reset caused by the watchdog. For details, see the *System Reset and Booting* chapter.

To prevent the watchdog from expiring, software services the watchdog by performing dummy writes to the `WDOG_STAT` register. The values written are ignored, but the write commands cause the `WDOG_STAT` register to be reloaded from the `WDOG_CNT` register.

If the watchdog is enabled with a zero value loaded to the counter and the `WDRO` bit was cleared, the `WDRO` bit of the watchdog control register is set immediately and the counter remains at zero without further decrements. If, however, the `WDRO` bit was set by the time the watchdog is enabled, the counter decrements to `0xFFFF FFFF` and continues operation.

Software can disable the watchdog timer only by writing a `0xAD` value to the `WDEN` field in the `WDOG_CTL` register.

Register Definitions

The watchdog timer is controlled by three registers.

Watchdog Count (`WDOG_CNT`) Register

The `WDOG_CNT` register, shown in [Figure 12-2](#), holds the 32-bit unsigned count value. The `WDOG_CNT` register must always be accessed with 32-bit read/writes.

A valid write to the `WDOG_CNT` register also preloads the watchdog counter. For added safety, the `WDOG_CNT` register can be updated only when the

Register Definitions

watchdog timer is disabled. A write to the `WDOG_CNT` register while the timer is enabled does not modify the contents of this register.

Watchdog Count Register (`WDOG_CNT`)

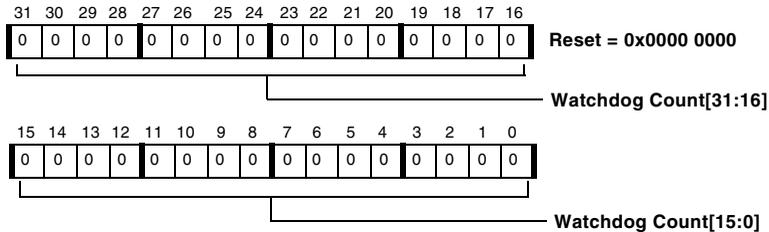


Figure 12-2. Watchdog Count Register

Watchdog Status (`WDOG_STAT`) Register

The 32-bit `WDOG_STAT` register, shown in [Figure 12-3](#), contains the current count value of the watchdog timer. Reads to `WDOG_STAT` return the current count value. 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, but not rolled over yet, writes to the `WDOG_STAT` register load it with the value in `WDOG_CNT`.



Enabling the watchdog timer does not automatically reload `WDOG_STAT` from `WDOG_CNT`.

The `WDOG_STAT` register is a 32-bit unsigned system MMR that must be accessed with 32-bit reads and writes.

Watchdog Status Register (WDOG_STAT)

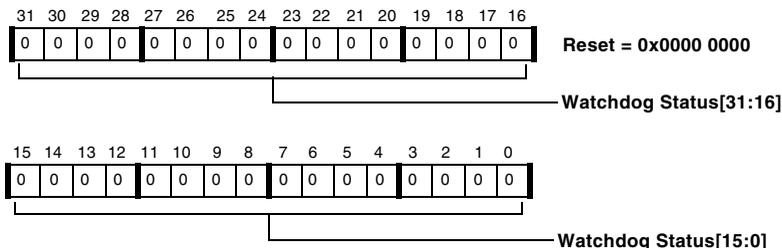


Figure 12-3. Watchdog Status Register

Watchdog Control (WDOG_CTL) Register

The `WDOG_CTL` register, shown in [Figure 12-4](#), is a 16-bit system MMR used to control the watchdog timer.

The watchdog event (`WDEV[1:0]`) bit 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 `SIC_IMASK` register that holds the watchdog timer mask bit 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 watchdog enable (`WDEN[7:0]`) bit field is used to enable and disable the watchdog timer. Writing any value other than the disable key (0xAD) 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 watchdog has expired by interrogating the `WDRO` status bit of the `WDOG_CTL` register. This is a sticky bit that is

Programming Examples

set whenever the watchdog timer count reaches 0. It can be cleared only by writing a “1” to the bit when the watchdog has been disabled first.

Watchdog Control Register (WDOG_CTL)

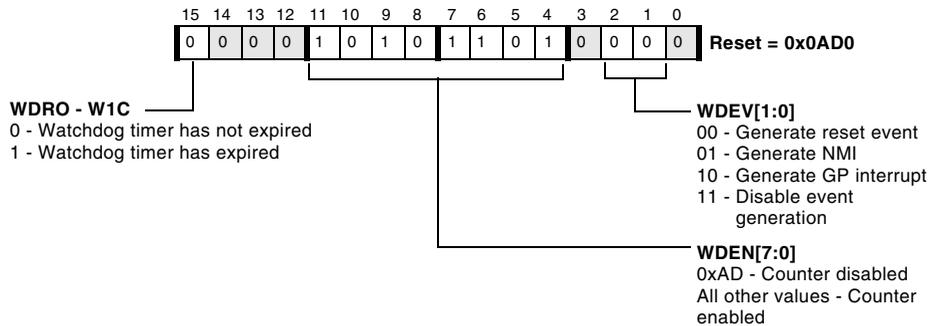


Figure 12-4. Watchdog Control Register

Programming Examples

[Listing 12-1](#) shows how to configure the watchdog timer so that it resets the chip when it expires. At startup, the code evaluates whether the recent reset event has been caused by the watchdog. Additionally, the example sets the NOBOOT bit to prevent the memory from being rebooted.

Listing 12-1. Watchdog Timer Configuration

```
#include <defBF527.h> /*ADSP-BF527 product is used as an example*/
#define WDOGPERIOD 0x00200000

.section L1_code;
.global _reset;
_reset:
    ...
    /* optionally, test whether reset was caused by watchdog */
```

```

    p0.h=hi(SWRST);
    p0.l=lo(SWRST);
    r6 = w[p0] (z);
    CC = bittst(r6, bitpos(RESET_WDOG));
    if !CC jump _reset.no_watchdog_reset;

/* optionally, warn at system level or host device here */

_reset.no_watchdog_reset:
/* optionally, set NOBOOT bit to avoid reboot in case */
    p0.h=hi(SYSCR);
    p0.l=lo(SYSCR);
    r0 = w[p0](z);
    bitset(r0,bitpos(NOBOOT));
    w[p0] = r0;

/* start watchdog timer, reset if expires */
    p0.h = hi(WDOG_CNT);
    p0.l = lo(WDOG_CNT);
    r0.h = hi(WDOGPERIOD);
    r0.l = lo(WDOGPERIOD);
    [p0] = r0;
    p0.l = lo(WDOG_CTL);
    r0.l = WDEN | WDEV_RESET;
    w[p0] = r0;
    ...
    jump _main;
_reset.end:

```

The subroutine shown in [Listing 12-2](#) can be called by software to service the watchdog. Note that the value written to the WDOG_STAT register does not matter.

Programming Examples

Listing 12-2. Service Watchdog

```
service_watchdog:
    [--sp] = p5;
    p5.h = hi(WDOG_STAT);
    p5.l = lo(WDOG_STAT);
    [p5] = r0;
    p5 = [sp++];
    rts;
service_watchdog.end:
```

[Listing 12-3](#) is an interrupt service routine that restarts the watchdog. Note that the watchdog must be disabled first.

Listing 12-3. Watchdog Restarted by Interrupt Service Routine

```
isr_watchdog:
    [--sp] = astat;
    [--sp] = (p5:5, r7:7);
    p5.h = hi(WDOG_CTL);
    p5.l = lo(WDOG_CTL);
    r7.l = WDDIS;
    w[p5] = r7;
    bitset(r7, bitpos(WDR0));
    w[p5] = r7;
    r7 = [p5 + WDOG_CNT - WDOG_CTL];
    [p5 + WDOG_CNT - WDOG_CTL] = r7;
    r7.l = WDEN | WDEV_GPI;
    w[p5] = r7;
    (p5:5, r7:7) = [sp++];
    astat = [sp++];
    rti;
isr_watchdog.end:
```

Unique Information for the ADSP-BF51x Processor

None.

Unique Information for the ADSP-BF51x Processor

13 GENERAL-PURPOSE COUNTER

This chapter describes the general-purpose up/down counter. The counter provides support for manually controlled rotary controllers, such as the volume wheel on a radio device. This unit also supports industrial encoders. Following the overview and list of key features is a description of the operating modes.

This chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of GP counters for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For GP counter interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the GP counters is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each GP counter, refer to [Appendix A, “System MMR Assignments”](#).

GP counter behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the

Overview

section [“Unique Information for the ADSP-BF51x Processor”](#) on page 13-38

Overview

The purpose of this interface is to convert pulses from incremental position encoders into data that is representative of the actual position. This is done by integrating (counting) pulses on one or two inputs. Since integration provides relative position, some devices also feature a zero position input (zero marker) that can be used to establish a reference point to verify that the acquired position does not drift over time.

In addition, the incremental position information can be used to determine speed, if the time intervals are measured.

The GP counter provides flexible ways to establish position information. When used in conjunction with the GP timer block, the GP counter allows for the acquisition of coherent position/time-stamp information that enables speed calculation.

Features

The GP counter includes the following features:

- 32-bit up/down counter
- Quadrature encoder mode (Gray code)
- Binary encoder mode
- Alternative frequency-direction mode
- Timed direction and up/down counting modes
- Zero marker/push button support

- Capture event timing in association with general purpose timer
- Boundary comparison and boundary setting features
- Input pin noise filtering (debouncing)
- Flexible error detection/signaling

Interface Overview

A block diagram of the GP counter is shown in [Figure 13-1](#). There are two input pins, the count up and direction (CUD) pin and the count down and gate (CDG) pin, that accept various forms of incremental inputs and are processed by the 32-bit counter. The third input, count zero marker (CZM), is the zero marker input. The module interfaces to the processor by way of the peripheral access bus (PAB) and can optionally generate an interrupt request through the IRQ line. There is also an output that can be used by the timer module to generate time-stamps on certain events.

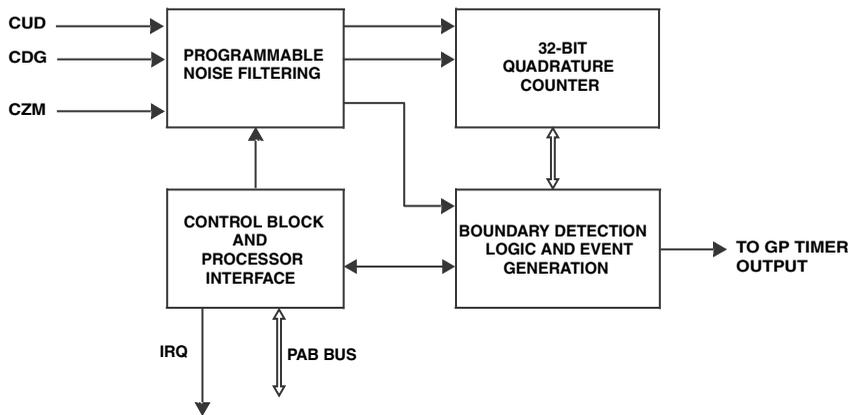


Figure 13-1. Block Diagram of the GP Counter Interface

Description of Operation

The GP counter has five modes of operation that are described in this section.

With the exception of the timed direction mode, the GP counter can operate with the GP timer block to capture additional timing information (time-stamps) associated with events detected by this block.

The third input (CZM) may be used as a zero marker or to sense the pressing of a push button. Refer to [“Zero Marker \(Push Button\) Operation” on page 13-9](#) for more details.

The three input pins may be filtered (debounced) before being evaluated by the GP counter. Refer to [“Input Noise Filtering \(Debouncing\)” on page 13-7](#) for more details.

The GP counter also features a flexible boundary comparison. In all of the operating modes, the counter can be compared to an upper and lower limit. A variety of actions can be taken when these limits are reached. Refer to [“Boundary Comparison Modes” on page 13-10](#) for more details.

Quadrature Encoder Mode

In this mode, the CUD and CDG inputs expect a quadrature-encoded signal that is interpreted as a 2-bit Gray code. The order of transitions of the CUD and CDG inputs determines whether the counter increments or decrements. The CNT_COUNTER register contains the number of transitions that have occurred. Refer to [Table 13-1](#) for more details.

Optionally, an interrupt is generated if both inputs change within one SCLK cycle. Such transitions are not allowed by Gray coding. Therefore, the CNT_COUNTER register remains unchanged and an error condition is signaled.

Table 13-1. Quadrature Events and Counting Mechanism

CNT_COUNTER Register Value	-4	-3	-2	-1	0	+1	+2	+3	+4
CDG:CUD Inputs	00	01	11	10	00	01	11	10	00

It is possible to reverse the count direction of the Gray coded signal. This can be achieved by enabling the polarity inverter of either the CUD pin or the CDG pin. Inverting both pins will not alter the behavior. This feature can be enabled with the CDGINV and CUDINV bits in the CNT_CONFIG register.

As an example, if the CDG:CUD inputs are 00 respectively and the next transition is to 01, this would normally increment the counter as is shown in [Table 13-1](#). If the CUD polarity is inverted this generates a received input of 01 followed by 00. This will result in a decrement of the counter, altering the behavior of the connected hardware.

Binary Encoder Mode

This mode is almost identical to the previous mode, with the exception that the CUD and CDG inputs expect a binary-encoded signal. The order of transitions of the CUD and CDG inputs determines whether the counter increments or decrements. The CNT_COUNTER register contains the number of transitions that have occurred. Refer to [Table 13-2](#).

Optionally, an interrupt is generated if the detected code steps by more than 1 (in binary arithmetic) within one SCLK cycle. Such transitions are considered erroneous. Therefore, the CNT_COUNTER register remains unchanged and an error condition is signaled.

Table 13-2. Binary Events and Counting Mechanism

CNT_COUNTER Register Value	-4	-3	-2	-1	0	+1	+2	+3	+4
CDG:CUD Inputs	00	01	10	11	00	01	10	11	00

Description of Operation

Reversing the `CUD` and `CDG` pin polarity has a different effect for the binary encoder mode than for the quadrature encoder mode. Inverting the polarity of the `CUD` pin only, or inverting both the `CUD` and `CDG` pins, will result in reversing the count direction.

Up/Down Counter Mode

In this mode, the counter is incremented or decremented at every active edge of the input pins.

If an active edge is detected at the `CUD` input, the counter increments. The active edge can be selected by the `CUDINV` bit in the `CNT_CONFIG` register. If this bit is cleared, a rising edge will increment the counter. If this bit is set, a falling edge will increment the counter.

If an active edge is detected at the `CDG` input, the counter decrements. The active edge can be selected by the `CDGINV` bit in the `CNT_CONFIG` register. If this bit is cleared, a rising edge will decrement the counter. If this bit is set, a falling edge will decrement the counter.

If simultaneous edges occur on pin `CDG` and pin `CUD`, the counter remains unchanged and both up-count and down-count events are signaled in the `CNT_STATUS` register.

Direction Counter Mode

In this mode, the counter is incremented or decremented at every active edge of the `CDG` input pin.

The state of the `CUD` input determines whether the counter increments or decrements. The polarity can be selected by the `CUDINV` bit in the `CNT_CONFIG` register. If this bit is cleared, a high `CUD` input will increment, a low input will decrement. If this bit is set, the polarity is inverted.

If an active edge is detected at the `CDG` input, the counter value changes by one in the selected direction. The active edge can be selected by the

CDGINV bit in the CNT_CONFIG register. If this bit is cleared, a rising edge will decrement the counter. If this bit is set, a falling edge will decrement the counter.

Timed Direction Mode

In this mode, the counter is incremented or decremented at each SCLK cycle.

The state of the CUD input determines whether the counter increments or decrements. The polarity can be selected by the CUDINV bit in the CNT_CONFIG register. If this bit is cleared, a high CUD input will increment the counter, a low input will decrement it. If this bit is set, the polarity is inverted.

The CDG pin can be used to gate the clock. The polarity can be selected by the CDGINV bit in the CNT_CONFIG register. If this bit is cleared, a high CDG input will enable the counter, a low input will stop it. If this bit is set, the polarity is inverted.

Functional Description

The following sections describe the various functions in more detail.

Input Noise Filtering (Debouncing)

In all modes, the three input pins can be filtered to present clean signals to the GP counter logic. This filtering can be enabled or disabled by the DEBE bit in the CNT_CONFIG register. [Figure 13-2](#) shows the filtering operation for the CUD pin.

Functional Description

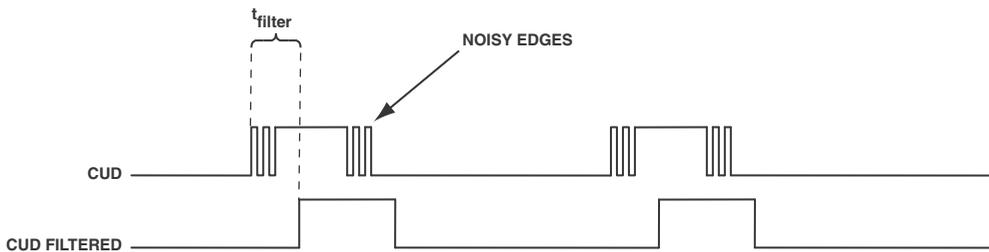


Figure 13-2. Programmable Noise Filtering

The filtering mechanism is implemented using counters for each pin. The counter for each pin is initialized from the `DPRESCALE` field of the `CNT_DEBOUNCE` register. When a transition is detected on a pin, the corresponding counter starts counting up to the programmed number of `SCLK` cycles. The state of the pin is latched after time t_{filter} and passed on to the GP counter logic.

The 5-bit `DPRESCALE` field in the `CNT_DEBOUNCE` register programs the desired number of cycles and therefore the debouncing time. The number of `SCLK` cycles for each pin can be selected in 18 steps ranging from 1×128 `SCLK` periods to 131072×128 `SCLK` periods (see [Figure 13-9 on page 13-25](#)).

The time t_{filter} is determined, given `SCLK` and the `DPRESCALE` value contained in the `CNT_DEBOUNCE` register, by the following formula:

$$t_{\text{filter}} = 128 \times (2^{\text{DPRESCALE}} \div \text{SCLK})$$

where `DPRESCALE` can contain values from 0 (minimum filtering) to 17 (maximum filtering).

Assuming an SCLK frequency of 133 MHz, the filter time range is shown by the following equations:

$$\text{DPRESCALE} = 0\text{b}0000$$

$$t_{\text{filter}} = 128 * 1 * 7.5\text{ns} = 960\text{ns} = (\text{approx.}) 1\mu\text{s}$$

$$\text{DPRESCALE} = 0\text{b}10001$$

$$t_{\text{filter}} = 128 * (131072) * 7.5\text{ns} = 125829\mu\text{s} = (\text{approx.}) 126\text{ms}$$

Zero Marker (Push Button) Operation

The CZM input pin can be used to sense the zero marker output of a rotary device or to detect the pressing of a push button. There are four programming schemes which are functional in all counter modes:

- **Push button mode**—This mode is enabled by setting the CZMIE bit in the CNT_IMASK register. An active edge at the CZM input will set the CZMII bit in the CNT_STATUS register. If enabled at the system interrupt controller, this will generate an interrupt request. The active edge is selected by the CZMINV bit in the CNT_CONFIG register (rising edge if cleared, falling edge if set to one).
- **Zero-marker-zeros-counter mode**—This mode is enabled by setting the ZMFC bit in the CNT_CONFIG register. An active level at the CZM input clears the CNT_COUNTER register and holds it until the CZM pin is deactivated. In addition, if enabled by the CZMZIE bit in the CNT_IMASK register, it will set the CZMZII bit in the CNT_STATUS register. If enabled by the peripheral interrupt controller, this will generate an interrupt request. The active level is selected by the CZMINV bit in the CNT_CONFIG register (active high if cleared, active low if set to one).
- **Zero-marker-error mode**—This mode is used to detect discrepancies between counter value and the zero marker output of certain rotary encoder devices. It is enabled by setting the CZMEIE bit in the

Functional Description

CNT_IMASK register. When an active edge is detected at the CZM input pin, the four LSBs of the CNT_COUNTER register are compared to zero. If they are not zero, a mismatch is signaled by way of the CZMERR bit in the CNT_STATUS register. If enabled by the peripheral interrupt controller, this will generate an interrupt request. The active edge is selected by the CZMINV bit in the CNT_CONFIG register: (rising edge if cleared, falling edge if set to one).

- **Zero-once mode**—This mode is used to perform an initial reset of the counter value when an active zero marker is detected. After that, the zero marker is ignored (the counter is not reset anymore). This mode is enabled by setting the W1ZMONCE bit in the CNT_COMMAND register. The CNT_COUNTER register and the W1ZMONCE bit are cleared on the next active edge on the CZM pin. Thus, the W1ZMONCE bit can be read to check whether the event has already occurred, if desired. The active edge of the CZM pin is selected by the CZMINV bit in the CNT_CONFIG register (rising edge if cleared, falling edge if set to one).

Boundary Comparison Modes

The GP counter includes two boundary registers, CNT_MIN (lower) and CNT_MAX (upper). The counter value is compared to the lower and upper boundary. Depending on which mode is selected, different actions are taken if the count value reaches either of the boundary values.

There are four boundary modes:

- **Boundary-compare mode**—The two boundary registers are simply compared to the CNT_COUNTER register. If, after incrementing, CNT_COUNTER equals CNT_MAX then the MAXCIE bit in the CNT_STATUS register is set. If the MAXCIE bit in the CNT_IMASK register is set, an interrupt request is generated. Similarly if, after decrementing, CNT_COUNTER equals CNT_MIN then the MINCIE status bit is set. If the

MINCII bit in the CNT_IMASK register is set, an interrupt request is generated. The MAXCII and MINCII bits are not set if the CNT_MAX and CNT_MIN registers are updated by software.

- **Boundary-zero mode**—This mode is similar to the boundary-compare mode. In addition to setting the status bits and requesting interrupts, the counter value in the CNT_COUNTER register is also set to zero.
- **Boundary auto-extend mode**—In this mode, the boundary registers are modified by hardware whenever the counter value reaches either of them. The CNT_MAX register is loaded with the current CNT_COUNTER value if the latter increments beyond the CNT_MAX value. Similarly, the CNT_MIN register is loaded with the CNT_COUNTER value if the latter decrements below the CNT_MIN value. This mode may be used to keep track of the widest angle the wheel ever reported, even if the software did not serve interrupts. At startup, the application software should set both boundary registers to the initial CNT_COUNTER value. The MAXCII and MINCII status bits are still set when the counter value matches the boundary register.
- **Boundary-capture mode**—In this mode, the CNT_COUNTER value is latched into the CNT_MIN register at one detected edge of the CZM input pin, and latched into CNT_MAX at the opposite edge. If the CZMINV bit in the CNT_CONFIG register is cleared, a rising edge captures into CNT_MIN and a falling edge into CNT_MAX. If the CZMINV bit is set, the edges are inverted. The MAXCII and MINCII status bits report the capture event.

The comparison is performed with signed arithmetic. The boundary registers and the counter value are all treated as signed integer values.

Functional Description

Control and Signaling Events

Eleven events can be signaled to the processor using status information and optional interrupt requests. The interrupts are enabled by the respective bits in the `CNT_IMASK` register. Dedicated bits in the `CNT_STATUS` register report events. When an interrupt from the GP counter is acknowledged, the application software is responsible for correct interpretation of the events. It is recommended to logically AND the content of the `CNT_IMASK` and `CNT_STATUS` registers to identify pending interrupts. Interrupt requests are cleared by write-one-to-clear (W1C) operations to the `CNT_STATUS` register. Hardware does not clear the status bits automatically, unless the counter module is disabled.

Illegal Gray/Binary Code Events

When the illegal transitions described in [“Quadrature Encoder Mode” on page 13-4](#) or [“Binary Encoder Mode” on page 13-5](#) occur, the `ICII` bit in the `CNT_STATUS` register is set. If enabled by the `ICIE` bit in the `CNT_IMASK` register, an interrupt request is generated. The `ICIE` bit should only be set in the quadrature encoder or binary encoder modes.

Up/Down Count Events

The `UCII` bit in the `CNT_STATUS` register indicates whether the counter has been incremented. Similarly, the `DCII` bit reports decrements. The two events are independent. For instance, if the counter first increments by one and then decrements by two, both bits remain set, even though the resulting counter value shows a decrement by one. In up/down counter mode, hardware may detect simultaneous active edges on the `CUD` and `CDG` inputs. In that case, the `CNT_COUNTER` remains unchanged, but both the `UCII` and `DCII` bits are set.

Interrupt requests for these events may be enabled through the `UCIE` and `DCIE` bits. This feature should be used carefully when the counter is

clocked at high rates. This is especially critical when the counter operates in `DIR_TMR` mode, as interrupts would be generated every `SCLK` cycle.

These events can also be used for additional push buttons, if GP counter features are not needed. When up/down counter mode is enabled, these count events can be used to report interrupts from push buttons that connect to the `CUD` and `CDG` inputs.

Zero-Count Events

The `CZEROII` status bit indicates that the `CNT_COUNTER` has reached a value equal to `0x0000 0000` after an increment or decrement. This bit is not set when the counter value is set to zero by a write to `CNT_COUNTER` or by setting the `W1LCNT_ZERO` bit in the `CNT_COMMAND` register. If enabled by the `CZEROIE` bit, an interrupt request is generated.

Overflow Events

There are two status bits that indicate whether the signed counter register has overflowed from a positive to a negative value or vice versa.

The `COV31II` bit reports that the 32-bit `CNT_COUNT` register has either incremented from `0x7FFF FFFF` to `0x8000 0000`, or decremented from `0x8000 0000` to `0x7FFF FFFF`. If enabled by the `COV31IE` bit, an interrupt request is generated.

Similarly, in applications where only the lower 16 bits of the counter are of interest, the `COV15II` status bit reports counter transitions from `0xFFFF 7FFF` to `0xFFFF 8000`, or from `0xFFFF 8000` to `0xFFFF 7FFF`. If enabled by the `COV15IE` bit, an interrupt request is generated.

Boundary Match Events

The `MINCII` and `MAXCII` status bits report boundary events as described in [“Boundary Comparison Modes” on page 13-10](#). These bits are not set if

Functional Description

the CNT_COUNTER, CNT_MAX or CNT_MIN registers are updated by software or the CNT_COMMAND register is written to.

The MINCIE and MAXCIE bits in the CNT_IMASK register enable interrupt generation on boundary events.

Zero Marker Events

There are three status bits CZMII, CZMEII and CZMZII associated with zero marker events, as described in [“Zero Marker \(Push Button\) Operation” on page 13-9](#). Each of these events can optionally generate an interrupt request, if enabled by the corresponding CZMIE, CZMEIE and CZMZIE bits in the CNT_IMASK register.

Capturing Timing Information

To calculate speed, many applications may wish to measure the time between two count events—in addition to accurately counting encoder pulses. For more accuracy, particularly at very low speeds, it is also necessary to obtain the time that has elapsed since the last count event. This additional information allows for estimating how much the GP counter has advanced since the last counter event.

For this purpose, the GP counter has an internal signal that connects to the alternate capture input (TACIX) of one of the GP timers. It is functional in all modes, with the exception of the timed direction mode. Refer to the “Internal Interfaces” section of [Chapter 9, “General-Purpose Ports”](#) for information regarding which GP timer(s) are associated with which GP counter module(s) for your device.

In order to use the timing measurements, the associated GP timer must be used in the WDTM_CAP mode. The alternate capture input is selected by setting the TIN_SEL bit in the GP timer’s TIMER_CONFIG register. For more information about the GP timers and their operating modes, refer to the *General-Purpose Timer* chapter.

Capturing Time Interval Between Successive Counter Events

When the only timing information of interest is the interval between successive count events, the associated timer should be programmed in `WDTH_CAP` mode with `PULSE_HI = 1`, `PERIOD_CNT = 1` and `TIN_SEL = 1`. Typically, this information is sufficient if the speed of GP counter events is known not to reach very low values. [Figure 13-3](#) shows the operation of the GP counter and the GP timer in this mode. TO generates a pulse every time a count event occurs. The GP timer will update its `TIMER_PERIOD` register with the period (measured from rising edge to rising edge) of the TO signal. The `TIMER_PERIOD` register is updated at every rising edge of the TO signal and contains the number of system clock (`SCLK`) cycles that have elapsed since the previous rising edge.

Incidentally, the `TIMER_WIDTH` register is also updated at the same time, but is generally of no interest in this mode of operation. If no reads of the `CNT_COUNTER` register occur between counter events, the `TIMER_WIDTH` register only contains the width of the TO pulse. If a read of `CNT_COUNTER` has occurred between events, the `TIMER_WIDTH` register will contain the time between the read of `CNT_COUNTER` and the next event.

This mode can also be used with `PULSE_HI = 0`. In this case, the period of TO is measured between falling edges. It will result in the same values as in the previous case, only the latching occurs one `SCLK` cycle later.

Functional Description

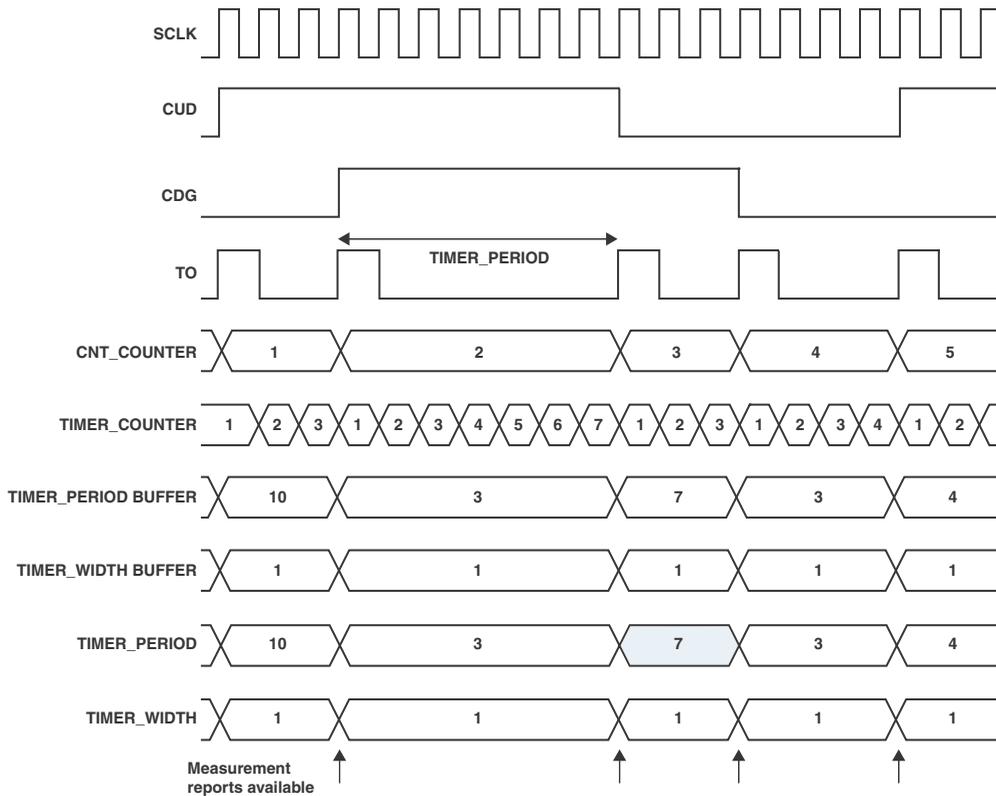


Figure 13-3. Operation With GP Timer Module

Capturing Counter Interval and CNT_COUNTER Read Timing

It is possible to also capture the time elapsed since the last count event. In this mode, the associated timer should be programmed in `WDTH_CAP` mode with `PULSE_HI = 0`, `PERIOD_CNT = 0` and `TIN_SEL = 1`. Typically, this additional information is used to estimate the advancement of the GP counter since the last count event, when the speed is very low. [Figure 13-4](#) shows the operation of the GP counter module and the GP timer module in this mode. TO generates a pulse every time a count event occurs.

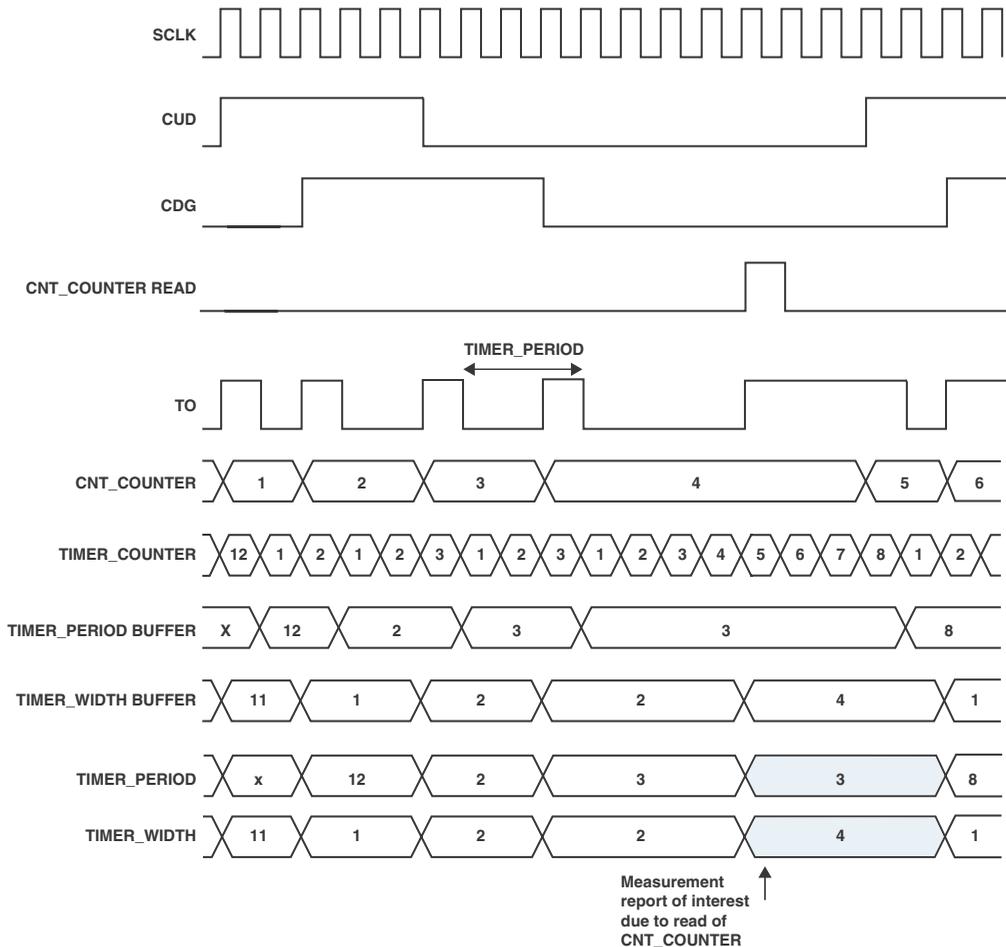


Figure 13-4. Capturing Counter Interval

In addition, when the processor reads the `CNT_COUNTER` register, the `TO` signal presents a pulse which is extended (high) until the next count event. The GP timer will update its `TIMER_PERIOD` register with the period (measured from falling edge to falling edge, because `PULSE_HI = 0`) of the `TO` signal. The `TIMER_WIDTH` register is updated with the pulse width (the portion where `TO` is low, again because `PULSE_HI = 0`). Both registers are

Programming Model

updated at every rising edge of the TO signal (because `PERIOD_CNT = 0`). Therefore, the `TIMER_PERIOD` register contains the period between the last two count events and the `TIMER_WIDTH` register contains the time since the last count event and the read of the `CNT_COUNTER` register, both measured in number of `SCLK` cycles.

The result is that when reading the `CNT_COUNTER` register, the two time measurements are also latched and the user has a coherent triplet of information to calculate speed and position.

-  Restrictions apply to the use of the TO signal in terms of speed. Therefore, the user must take care to not operate at very high count events. For instance, if `CNT_COUNTER` is incremented/decremented every `SCLK` cycle (timed direction mode), the TO signal is incorrect.

Programming Model

In a typical application, the user will initialize the GP counter for the desired mode, without enabling it. Normally the events of interest will be processed using interrupts rather than polling the status bit. In that case, clear all status bits and activate the generation of interrupt requests with the `CNT_IMASK` register. Set up the system interrupt controller and core interrupts. If timing information is required, set up the relevant GP timer in `WDTH_CAP` mode with the settings described in the “[Capturing Timing Information](#)” on page 13-14. Then, enable the interrupts and the peripheral itself.

Registers

The GP counter interface has eight memory-mapped registers (MMRs) that regulate its operation. Descriptions and bit diagrams for MMRs is provided in the sections that follow.

Counter Module Register Overview

Refer to [Table 13-3](#) for an overview of all MMRs associated with the GP counter interface.

Table 13-3. Counter Module Register Overview

Register Name	Width	PAB Operation	Reset Value
CNT_CONFIG	16 bits	R/W	0x0000
CNT_IMASK	16 bits	R/W	0x0000
CNT_STATUS	16 bits	R/W1C	0x0000
CNT_COMMAND	16 bits	R/W1A	0x0000
CNT_DEBOUNCE	16 bits	R/W	0x0000
CNT_COUNTER	32 bits	R/W (16/32 bits)	0x0000 0000
CNT_MAX	32 bits	R/W (16/32 bits)	0x0000 0000
CNT_MIN	32 bits	R/W (16/32 bits)	0x0000 0000

Registers

Counter Configuration Register (CNT_CONFIG)

This register is used to configure counter modes and input pins, as well as to enable the peripheral. It can be accessed at any time with 16-bit read and write operations.

Counter Configuration (CNT_CONFIG) Register

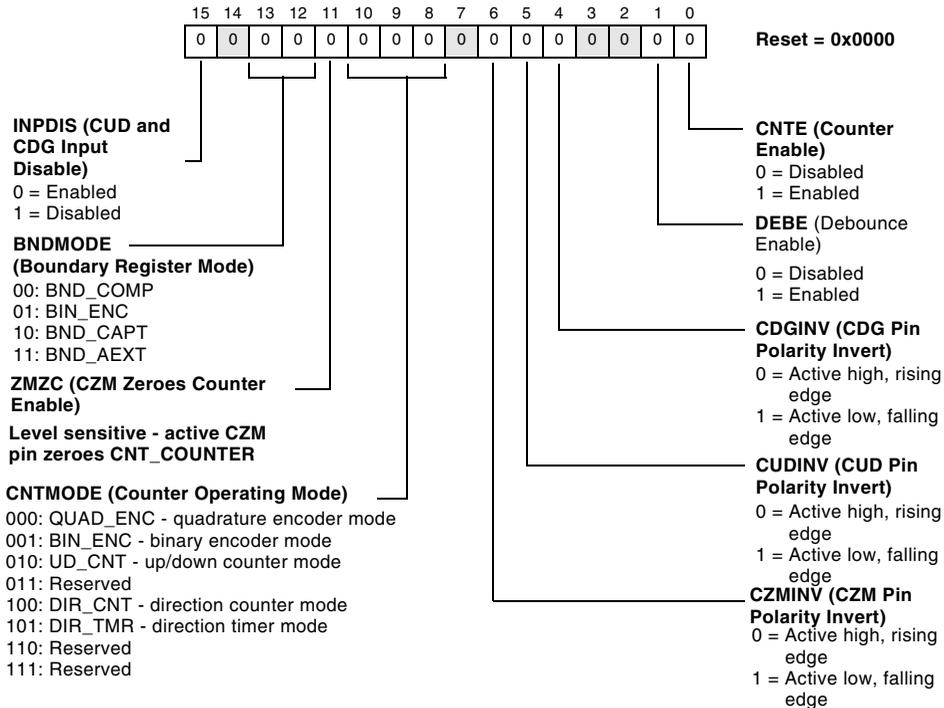


Figure 13-5. Counter Configuration Register

Counter Interrupt Mask Register (CNT_IMASK)

This register is used to enable interrupt request generation from each of the eleven events. It can be accessed at any time with 16-bit read and write operations. For explanations of the register bits, refer to the “Control and Signaling Events” on page 13-12.

Counter Interrupt Mask (CNT_IMASK) Register

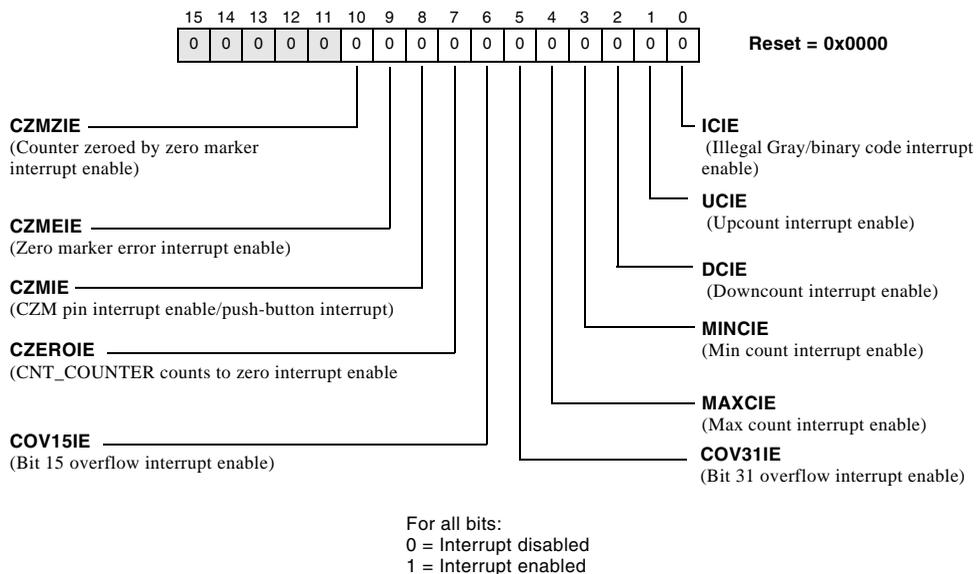


Figure 13-6. Counter Interrupt Mask Register

Counter Status Register (CNT_STATUS)

This register provides status information for each of the eleven events where 0 = no interrupt pending and 1 = interrupt pending. When an event is detected, the corresponding bit in this register is set. It remains set until either software writes a “1” to the bit (write-1-to-clear) or the GP counter is disabled. For explanations of the register bits, refer to the “Control and Signaling Events” on page 13-12.

Registers

Counter Status (CNT_STATUS) Register

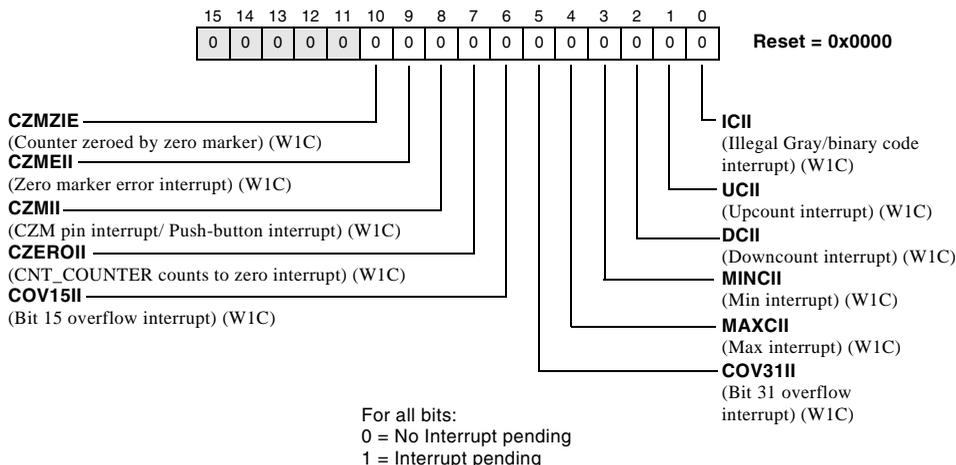


Figure 13-7. Counter Status Register

Counter Command Register (CNT_COMMAND)

The CNT_COMMAND register (shown in [Figure 13-8](#)) configures the GP counter, enabling operations such as zeroing a counter register and copying or swapping boundary registers. These actions are taken by writing a “one” to the appropriate bit.

Read operations from this register will not return meaningful values, with the exception of the W1ZONCE bit, where a “1” indicates that the bit has been set by software before, but no zero marker event has been detected on the CZM pin yet. Refer to the [“Zero Marker \(Push Button\) Operation” on page 13-9](#) for more details.

The CNT_COUNTER, CNT_MIN and CNT_MAX registers can be initialized to zero by writing a “one” to the W1LCNT_ZERO, W1LMIN_ZERO and W1LMAX_ZERO fields. In addition to clearing registers, CNT_COMMAND allows the boundary registers to be modified in a number of ways. The current counter value in CNT_COUNT can be captured and loaded into either of the two boundary

registers `CNT_MAX` and `CNT_MIN` to create new boundary limits. This is performed by setting the `W1LMAX_CNT` and `W1LMIN_CNT` bits. Alternatively, the counter can be loaded from `CNT_MAX` or `CNT_MIN` via the `W1LCNT_MAX` and `W1LCNT_MIN` bits. It is also possible to transfer the current `CNT_MAX` value into `CNT_MIN` (or vice versa) through the `W1LMIN_MAX` and `W1LMAX_MIN` bits. The final supported operation is the ability to only have the zero marker clear the `CNT_COUNT` register once, as described in “Zero Marker (Push Button) Operation” on page 13-9.

It is possible for multiple actions to be performed simultaneously by setting multiple bits in the `CNT_COMMAND` register. However, there are restrictions. The bits associated with each command have been grouped together such that all bits that involve a write to the `CNT_COUNTER` register are located within bits 3:0 of the `CNT_COMMAND` register. All commands that involve a write to the `CNT_MIN` register are located within bits 7:4 of the `CNT_COMMAND` register, and all commands that involve a write to the `CNT_MAX` register are located within bits 11:8 of the `CNT_COMMAND` register.

 A maximum of three commands can be issued at any one time, excluding the `W1ZMONCE` command. Note that (`W1LCNT_MIN`, `W1LCNT_MAX` and `W1LCNT_ZERO`) have to be used exclusively. Never set more than one of them at the same time. The same rule applies for (`W1LMAX_MIN`, `W1LMAX_CNT` and `W1LMAX_ZERO`) and for (`W1LMIN_MAX`, `W1LMIN_CNT`, and `W1LMIN_ZERO`).

Registers

Counter Command (CNT_COMMAND) Register

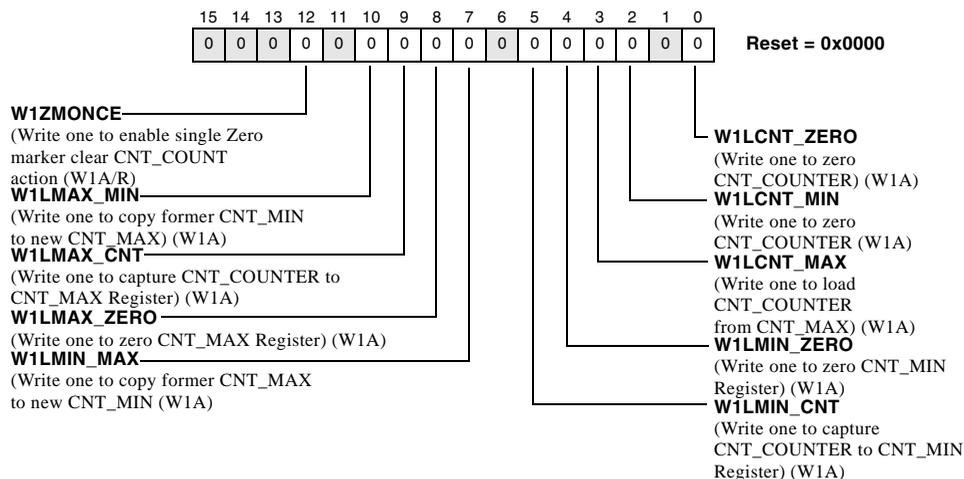


Figure 13-8. Counter Command Register

Counter Debounce Register (CNT_DEBOUNCE)

This register is used to select the noise filtering characteristic of the three input pins (see “[Input Noise Filtering \(Debouncing\)](#)” on page 13-7). Bits [4:0] determine the filter time. The register can be accessed at any time with 16-bit read and write operations.

$$t_{filter} = 128 \times (2^{DPRESCALE} \div SCLK)$$

Counter Debounce (CNT_DEBOUNCE) Register

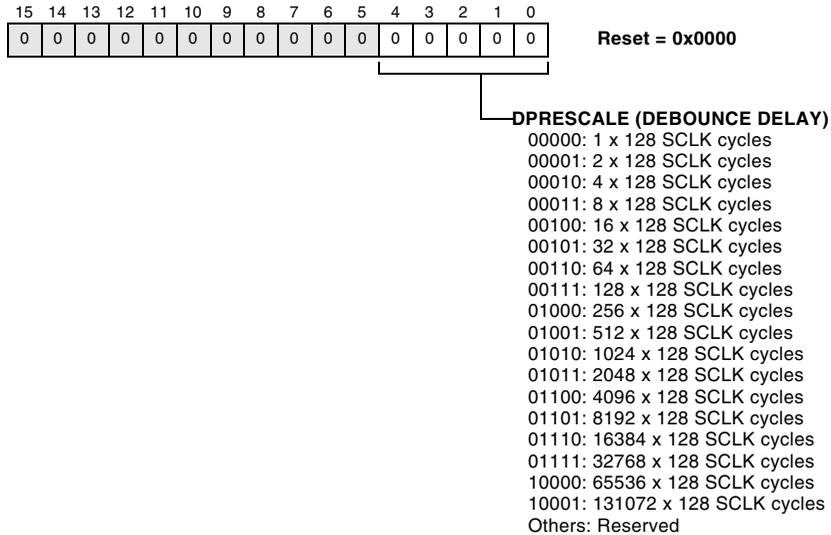


Figure 13-9. Counter Debounce Register

Counter Count Value Register (CNT_COUNTER)

This register holds the 32-bit, twos-complement, count value. It can be read and written at any time. Hardware ensures that reads and write are atomic, by providing respective shadow registers. This register can be accessed with either 32-bit or 16-bit operations. This allows use of the GP counter as a 16-bit counter if sufficient for the application.

Registers

Counter Count Value (CNT_COUNTER) Register

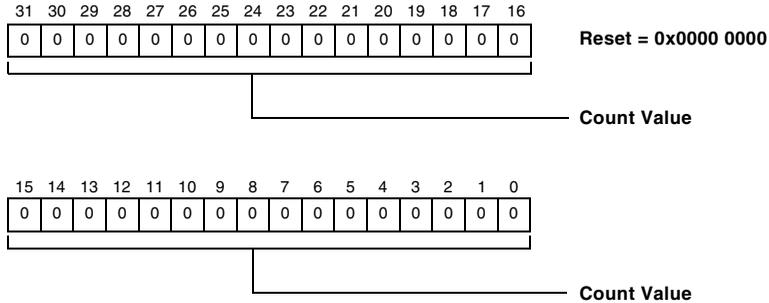


Figure 13-10. Counter Count Value Register

Counter Boundary Registers (CNT_MIN and CNT_MAX)

These registers hold the 32-bit, twos-complement, lower and upper boundary values. They can be read and written at any time. Hardware ensures that reads and write are atomic, by providing respective shadow registers. This register can be accessed with either 32-bit or 16-bit operations. This allows for using the GP counter as a 16-bit counter if sufficient for the application.

Counter Maximal Count (CNT_MAX) Register

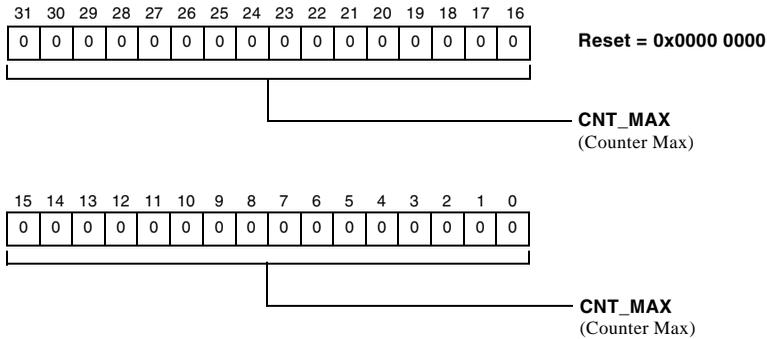


Figure 13-11. Counter Maximal Count Register

Counter Minimal Count (CNT_MIN) Register

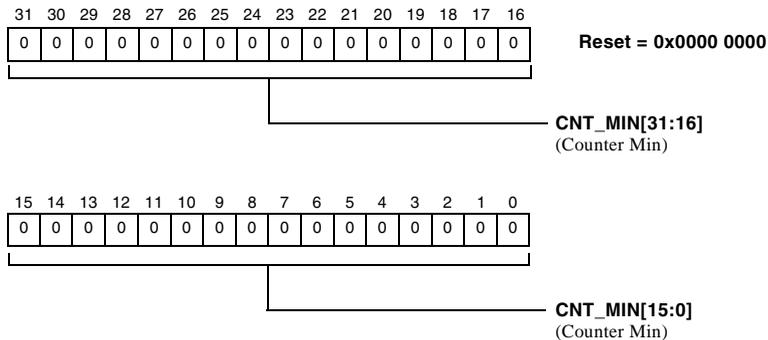


Figure 13-12. Counter Minimal Count Register

Programming Examples

[Listing 13-1](#) illustrates how to initialize the GP counter for various modes. The required interrupts are first unmasked. The GP counter is then configured for the required mode of operation. Note that at this point we do

Programming Examples

not yet enable the counter. Finally, some GP counter MMRs are cleared, as well as any interrupts that may be pending in the CNT_STATUS register.

Listing 13-1. Initializing the GP Counter

```
/* Setup Counter Interrupts */
P5.H = hi(CNT_IMASK);
P5.L = lo(CNT_IMASK);
R5 = nCZMZIE /* Counter zeroed by zero marker interrupt */
    | CZMEIE /* Zero marker error interrupt */
    | CZMIE /* CZM pin interrupt (push-button) */
    | CZEROIE /* Counts to zero interrupt */
    | nCOV15IE /* Counter bit 15 overflow interrupt */
    | nCOV31IE /* Counter bit 31 overflow interrupt */
    | MAXCIE /* Max count interrupt */
    | MINCIE /* Min count interrupt */
    | DCIE /* Downcount interrupt */
    | UCIE /* Upcount interrupt */
    | ICIE (z); /* Illegal gray/binary code interrupt */
w[P5] = R5;

/* Configure the GP Counter mode of operation */
P5.H = hi(CNT_CONFIG);
P5.L = lo(CNT_CONFIG);
R5 = nINPDIS /* Enable CUD and CDG inputs */
    | BNDMODE_COMP /* Boundary compare mode */
    | nZMZC /* Disable Zero Counter Enable */
    | CNTMODE_QUADENC /* Quadrature Encoder Mode */
    | CZMINV /* Polarity of CZM pin */
    | nCUDINV /* Polarity of CUD pin */
    | nCDGINV /* Polarity of CDG Pin */
    | nDEBE /* Disable the debounce */
    | nCNTE (z); /* Disable the counter */
w[P5] = R5;
```

```

/* Zero the CNT_COUNT, CNT_MIN and CNT_MAX registers
This is optional as after reset they are default to zero */
P5.H = hi(CNT_COMMAND);
P5.L = lo(CNT_COMMAND);
R5 = W1LCNT_ZERO | W1LMIN_ZERO | W1LMAX_ZERO (z);
w[P5] = R5;

/* Clear any identified interrupts */
P5.H = hi(CNT_STATUS);
P5.L = lo(CNT_STATUS);
R5.L = ICII      /* Illegal Gray/Binary Code Interrupt Identifier
*/
    | UCII      /* Up count Interrupt Identifier */
    | DCII      /* Down count Interrupt Identifier */
    | MINCII /* Min Count Interrupt Identifier */
    | MAXCII /* Max Count Interrupt Identifier */
    | COV31II /* Bit 31 Overflow Interrupt Identifier */
    | COV15II /* Bit 15 Overflow Interrupt Identifier */
    | CZEROII /* Count to Zero Interrupt Identifier */
    | CZMII    /* CZM Pin Interrupt Identifier */
    | CZMEII /* CZM Error Interrupt Identifier */
    | CZMZII; /* CZM Zeroes Counter Interrupt Identifier */
w[P5] = R5;

```

Listing 13-2 illustrates how to set up the peripheral and core interrupts for the GP counter. This example assumes the counter interrupts are generated on IRQ27, which is assumed to be mapped to the IVG11 interrupt. Finally, the system and peripheral interrupts are unmasked, and then the GP counter is enabled. This example can be easily tailored to processors with different SIC register mappings.

Programming Examples

Listing 13-2. Setting Up the Interrupts for the GP Counter

```
/* Assign the CNT interrupt to IVG11 */
P5.H = hi(SIC_IAR3);
P5.L = lo(SIC_IAR3);
R6.H = hi(0xFFFF4FFF);
R6.L = lo(0xFFFF4FFF);
R7.H = hi(0x00000000);
R7.L = lo(0x00000000);
R5 = [P5];
R5 = R5 & R6; /* zero the counter interrupt field */
R5 = R5 | R7; /* set Counter interrupt to required priority */
[P5] = R5;

/* Set up the interrupt vector for the counter */
R5.H = hi(_IVG11_handler);
R5.L = lo(_IVG11_handler);
P5.H = hi(EVT11);
P5.L = lo(EVT11);
[P5] = R5;

/* Unmask IVG11 interrupt in the IMASK register */
P5.H = hi(IMASK);
P5.L = lo(IMASK);
R5 = [P5];
bitset(R5, bitpos(EVT_IVG11));
[P5] = R5;

/* Unmask interrupt 27 generated by the counter */
P5.H = hi(SIC_IMASK0);
P5.L = lo(SIC_IMASK0);
R5 = [P5];
bitset(R5, bitpos(IRQ_CNT));
[P5] = R5;
```

```

/* Enable the counter */
P5.H = hi(CNT_CONFIG);
P5.L = lo(CNT_CONFIG);
R5 = w[P5](z);
bitset(R5, bitpos(CNTE));
w[P5] = R5.L;

```

Using the same assumptions from the previous example, [Listing 13-3](#) illustrates a sample interrupt handler that is responsible for servicing the GP counter interrupts. On entry to the handler, the SIC_ISR0 register is interrogated to determine if the counter is waiting for an interrupt to be serviced. If so, the handler responsible for processing all counter interrupts is called.

Listing 13-3. Sample Interrupt Handler for GP Counter Interrupts

```

_IVG11_handler:
    /* Stack management */
    [--SP] = RETS;
    [--SP] = ASTAT;
    [--SP] = (R7:0, P5:0);

    /* Was it a counter interrupt? */
    P5.H = hi(SIC_ISR0);
    P5.L = lo(SIC_ISR0);
    R5 = [P5];
    CC = bittst(R5, bitpos(IRQ_CNT));
    IF !CC JUMP _IVG11_handler.completed;
    CALL _IVG11_handler.counter;

_IVG11_handler.completed:

```

Programming Examples

```
SSYNC;
/* Restore from stack */
(R7:0, P5:0) = [SP++];
ASTAT = [SP++];
RETS = [SP++];
RTI; /* Exit the interrupt service routine */
_IVG11_handler.end:

_IVG11_handler.counter:
/* Stack management */
[--SP] = RETS;
[--SP] = (R7:0, P5:0);

/* Determine what counter interrupts we wish to service */
R5 = w[P5](z);
P5.H = hi(CNT_IMASK);
P5.L = lo(CNT_IMASK);
R5 = w[P5](z);

P5.H = hi(CNT_STATUS);
P5.L = lo(CNT_STATUS);
R6 = w[P5](z);
R5 = R5 & R6;

/* Interrupt handlers for all GP counter interrupts */
_IVG11_handler.counter.illegal_code:
CC = bittst(R5, bitpos(ICII));
IF !CC JUMP _IVG11_handler.counter.up_count;

/* Clear the serviced request */
R6 = ICII (z);
w[P5] = R6;

/* insert illegal code handler here */
```

```
_IVG11_handler.counter.illegal_code.end:

_IVG11_handler.counter.up_count:
    CC = bittst(R5, bitpos(UCII));
    IF !CC JUMP _IVG11_handler.counter.down_count;

    /* Clear the serviced request */
    R6 = UCII (z);
    w[P5] = R6;

    /* insert up count handler here */

_IVG11_handler.counter.up_count.end:

_IVG11_handler.counter.down_count:
    CC = bittst(R5, bitpos(DCII));
    IF !CC JUMP _IVG11_handler.counter.min_count;

    /* Clear the serviced request */
    R6 = DCII (z);
    w[P5] = R6;

    /* insert down count handler here */

_IVG11_handler.counter.down_count.end:

_IVG11_handler.counter.min_count:
    CC = bittst(R5, bitpos(MINCII));
    IF !CC JUMP _IVG11_handler.counter.max_count;

    /* Clear the serviced request */
    R6 = MINCII (z);
```

Programming Examples

```
w[P5] = R6;

/* insert min count handler here */

_IVG11_handler.counter.min_count.end:

_IVG11_handler.counter.max_count:
    CC = bittst(R5, bitpos(MAXCII));
    IF !CC JUMP _IVG11_handler.counter.b31_overflow;

/* Clear the serviced request */
R6 = MAXCII (z);
w[P5] = R6;

/* insert max count handler here */

_IVG11_handler.counter.max_count.end:

_IVG11_handler.counter.b31_overflow:
    CC = bittst(R5, bitpos(COV31II));
    IF !CC JUMP _IVG11_handler.counter.b15_overflow;

/* Clear the serviced request */
R6 = COV31II (z);
w[P5] = R6;

/* insert bit 31 overflow handler here */

_IVG11_handler.counter.b31_overflow.end:

_IVG11_handler.counter.b15_overflow:
    CC = bittst(R5, bitpos(COV15II));
    IF !CC JUMP _IVG11_handler.counter.count_to_zero;
```

```
/* Clear the serviced request */
R6 = COV15II (z);
w[P5] = R6;

/* insert bit 15 overflow handler here */

_IVG11_handler.counter.b15_overflow.end:

_IVG11_handler.counter.count_to_zero:
    CC = bittst(R5, bitpos(CZEROII));
    IF !CC JUMP _IVG11_handler.counter.czm;

/* Clear the serviced request */
R6 = CZEROII (z);
w[P5] = R6;

/* insert count to zero handler here */

_IVG11_handler.counter.count_to_zero.end:

_IVG11_handler.counter.czm:
    CC = bittst(R5, bitpos(CZMII));
    IF !CC JUMP _IVG11_handler.counter.czm_error;

/* Clear the serviced request */
R6 = CZMII (z);
w[P5] = R6;

/* insert czm handler here */

_IVG11_handler.counter.czm.end:

_IVG11_handler.counter.czm_error:
```

Programming Examples

```
CC = bittst(R5, bitpos(CZMEII));
IF !CC JUMP _IVG11_handler.counter.czm_zeroes_counter;

/* Clear the serviced request */
R6 = CZMEII (z);
w[P5] = R6;

/* insert czm error handler here */

_IVG11_handler.counter.czm_error.end:

_IVG11_handler.counter.czm_zeroes_counter:
CC = bittst(R5, bitpos(CZMZII));
IF !CC JUMP _IVG11_handler.counter.all_serviced;

/* Clear the serviced request */
R6 = CZMZII (z);
w[P5] = R6;

/* insert czm zeroes counter handler here */

_IVG11_handler.counter.czm_zeroes_counter.end:

_IVG11_handler.counter.all_serviced:

/* Restore from stack */
(R7:0, P5:0) = [SP++];
RETS = [SP++];
RTS;
_IVG11_handler.counter.end:
```

[Listing 13-4](#) shows how to set up timer 7 (as an example) to capture the period of counter events. Refer to the "Internal Interfaces" section of [Chapter 9, "General-Purpose Ports"](#) for information regarding which GP

timer(s) are associated with which GP counter module(s) for your device. The timer is configured for `WDTH_CAP` mode, and the period between the last two successive counter events is read from within the up count interrupt handler that was provided in [Listing 13-3 on page 13-31](#).

Listing 13-4. Setting Up Timer 7 for Counter Event Period Capture

```

/* configure the timer for WDTH_CAP mode */
P5.H = hi(TIMER7_CONFIG);
P5.L = lo(TIMER7_CONFIG);
R5 = PULSE_HI | PERIOD_CNT | TIN_SEL | WDTH_CAP (z);
w[P5] = R5.L;

/* Enable Timer 7
P5.H = hi(TIMER_ENABLE0);
P5.L = lo(TIMER_ENABLE0);
R5 = TIMEN7 (z);
w[P5] = R5.L;

...

_IVG11_handler.counter.up_count:
    CC = bittst(R5, bitpos(UCII));
    IF !CC JUMP _IVG11_handler.counter.down_count;

/* Clear the serviced request */
R6 = UCII (z);
w[P5] = R6;

/* insert up count handler here */

/* Read the period between the last two successive events */
P5.H = hi(TIMER7_PERIOD);
P5.L = lo(TIMER7_PERIOD);

```

Unique Information for the ADSP-BF51x Processor

```
R5 = [P5];  
  
P5.H = hi(_event_period);  
P5.L = lo(_event_period);  
[P5] = R5;  
_IVG11_handler.counter.up_count.end;
```

Unique Information for the ADSP-BF51x Processor

None.

14 PWM CONTROLLER

This chapter describes the PWM controller module. Following an overview and a list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model discussion and consolidated register definitions.

Specific Information for the ADSP-BF51x

For details regarding the number of PWMs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For PWM Controller interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how the PWM Controller is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for the PWM Controller, refer to [Appendix A, “System MMR Assignments”](#).

Overview

The PWM controller is a flexible, programmable, three-phase PWM waveform generator that can be programmed to generate the required switching patterns to drive a three-phase voltage source inverter for ac

Overview

induction motor (ACIM) or permanent magnet synchronous motor (PMSM) control.

In addition, the PWM block contains functions that considerably simplify the generation of the required PWM switching patterns for control of electronically commutated motors (ECMs) or brushless dc motors (BDCMs).

Programming the `PWM_SRMODE` bit of the `PWM_CTRL` register to 0 enables a special mode used for switched reluctance motors (SRMs). [Figure 14-1](#) shows a block diagram that represents the main functional blocks of the PWM Controller.

The following six blocks control the generation of the six output PWM signals (`PWM_AH`, `PWM_AL`, `PWM_BH`, `PWM_BL`, `PWM_CH`, and `PWM_CL`):

- **Three-Phase PWM Timing Unit.** As the core of the PWM Controller, this block generates three pairs of complemented, center-based PWM signals and `PWM_SYNC` coordination.
- **Dead Time Control Unit.** This block inserts emergency dead time after the “ideal” PWM output pair, including crossover, is generated.
- **Output Control Unit.** This block permits the redirection of the outputs of the Three-Phase Timing Unit for each channel to the high-side or the low-side output. In addition, the Output Control Unit allows individual enabling/disabling of each of the six PWM output signals.
- **Gate Drive Unit.** This block provides the correct polarity output PWM signals, based on the state of the `PWM_POLARITY` bit of the `PWM_CTRL` register. The Gate Drive Unit also permits the generation of the high-frequency chopping waveform and its subsequent mixing with the PWM output signals.

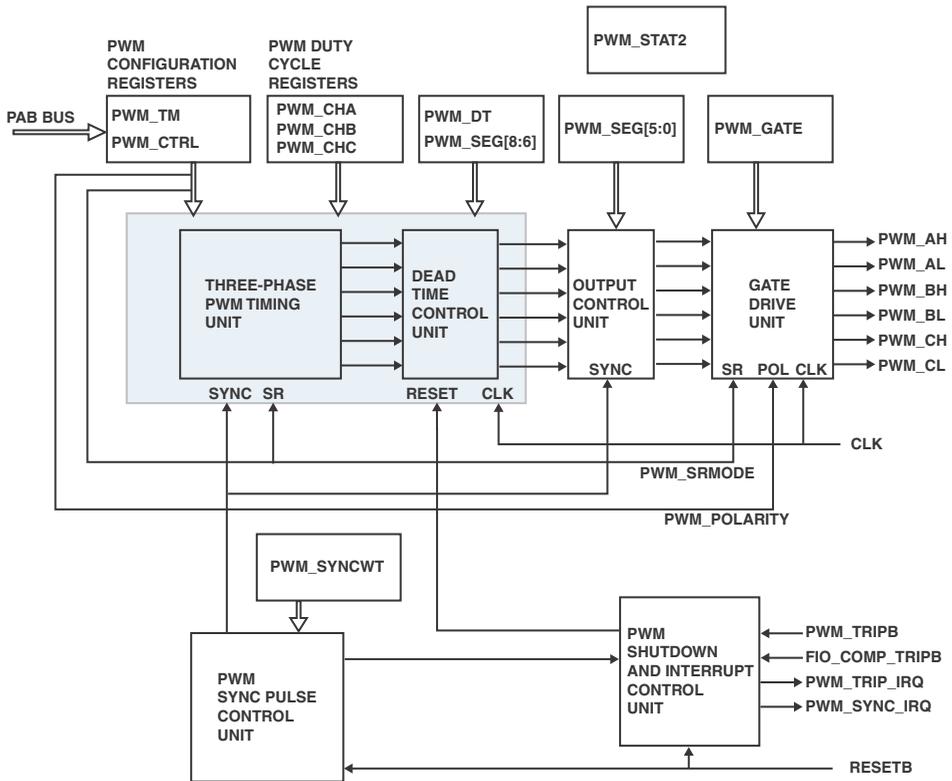


Figure 14-1. PWM Controller Block Diagram

- PWM Shutdown & Interrupt Control Unit.** This block takes care of the various PWM shutdown modes (via the `PWM_TRIPB` pin and the `PWM_CTRL` register). This unit generates the correct reset signal for the Three-Phase PWM Timing Unit and interrupt signals for the Interrupt Control Unit
- PWM Sync Pulse Control Unit.** This block generates the internal PWM synchronization pulse and also controls whether an external `PWM_SYNC` pulse is used.

Overview

The PWM Controller is driven by a clock, whose period is t_{SCLK} . The PWM generator produces three pairs (PWM_AH, PWM_AL, PWM_BH, PWM_BL, PWM_CH, and PWM_CL) of PWM signals on the six PWM output pins. There are three high-side drive signals (PWM_AH, PWM_BH, and PWM_CH) and three low-side drive signals (PWM_AL, PWM_BL, and PWM_CL). The polarity of the generated PWM signals may be programmed by the PWM_POLARITY bit of the PWM_CTRL register to generate active high or active low PWM patterns. The switching frequency and dead time of the generated PWM patterns are programmable via the PWM_TM and PWM_DT registers. In addition, three duty-cycle control registers (PWM_CHA, PWM_CHB, and PWM_CHC) directly control the duty cycles of the three pairs of PWM signals.

Each of the six PWM output signals can be enabled or disabled via separate output enable bits of the PWM_SEG register. In addition, three control bits of the PWM_SEG register permit independent crossover of the two signals of a PWM pair for easy control of ECMs or BDCMs. In crossover mode, the PWM signal destined for the high-side switch is diverted to the complementary low-side output, and the signal destined for the low-side switch is diverted to the corresponding high-side output signal for ECM or BDCM modes of operation. A typical configuration for these types of motors is shown in [Figure 14-2](#).

In common three-phase inverters, it is necessary to insert a so-called “dead time” between turning off one switch and turning on the other switch in the same leg, to prevent shoot-through. This dead time is inserted by an emergency dead-time insertion circuit, which enforces a dead time defined by the PWM_DT register between the high- and low-side drive signals of each PWM channel. This ensures that the correct dead time occurs at the power inverter.

In many applications, there is a need to provide an isolation barrier in the gate-drive circuits that turn on the power devices of the inverter. In general, there are two common isolation techniques: optical isolation using opto-isolators, and transformer isolation using pulse transformers. The PWM Controller permits the mixing of the output PWM signals with a

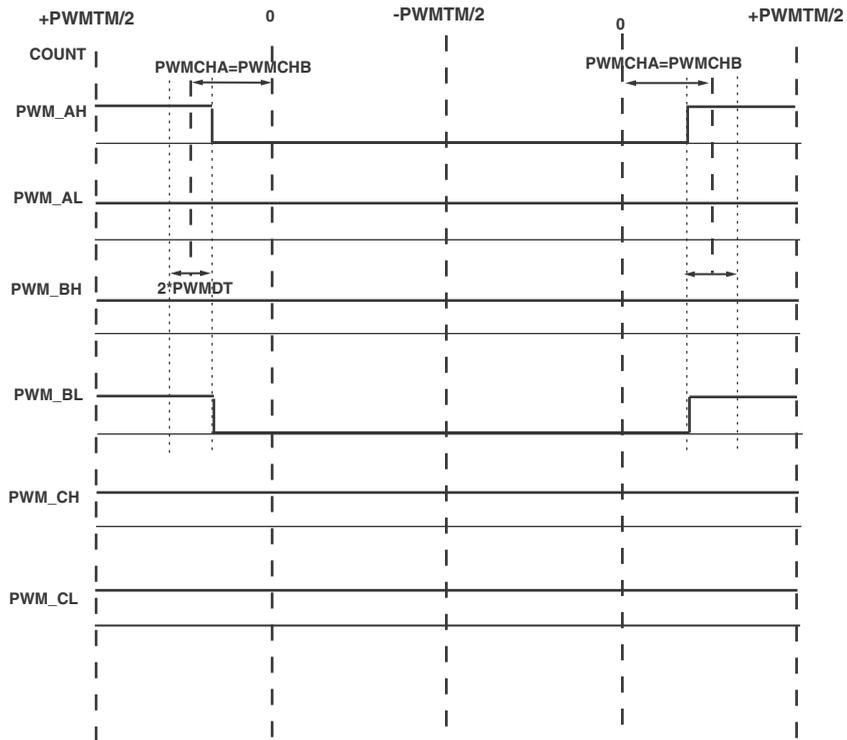


Figure 14-2. Active Low PWM Signals for ECM Control

high-frequency chopping signal, which provides an simple interface to pulse transformers. The features of gate-drive-chopping mode are controlled by the PWM_GATE register. An 8-bit value (GDCLK) within the PWM_GATE register directly controls the chopping frequency. In addition, high-frequency chopping can be independently enabled for the high- and low-side outputs using separate control bits in the PWM_GATE register. In addition, all PWM outputs require sufficient sink and source capability to directly drive most opto-isolators.

Overview

The PWM generator is capable of operating in two distinct modes:

- **Single-Update Mode.** In single-update mode, duty cycle values are programmable only once per PWM period; resultant PWM patterns are symmetrical about the mid-point of the PWM period.
- **Double-Update Mode.** In double-update mode, a second updating of the PWM registers is implemented at the midpoint of the PWM period. In double-update mode, it is possible to produce asymmetrical PWM patterns that produce lower harmonic distortion in three-phase PWM inverters. This technique also permits closed-loop controllers to change the average voltage applied to the machine windings at a faster rate, thus permitting faster closed-loop bandwidths to be achieved.

The operating mode of the PWM block (single- or double-update mode) is selected by the `PWM_DBL` bit in the `PWM_CTRL` register. Setting `PWM_DBL` to 1 selects double-update mode, and 0 selects single-update mode.

The PWM generator can provide an internal synchronization pulse on the `PWM_SYNC` pin that is synchronized to the PWM switching frequency. In single-update mode, a `PWM_SYNC` pulse is produced at the start of each PWM period. In double-update mode, an additional `PWM_SYNC` pulse is also produced at the midpoint of each PWM period. The width of the `PWM_SYNC` pulse is programmable through the `PWM_SYNCWT` register.

The PWM generator can also accept an external synchronization pulse on the `PWM_SYNC` pin. External synchronization is selected by setting the `PWM_EXTSYNC` bit in the `PWM_CTRL` register. The `PWM_SYNC` input timing can be synchronized to the internal system clock, which is selected by setting the `PWM_SYNCSEL` bit of the `PWM_CTRL` register. If the external synchronization pulse from the chip pin is asynchronous to the internal system clock (typical case), the external `PWM_SYNC` is considered asynchronous and should be synchronized. If the `PWM_SYNC` is actually received from another PWM on the same chip controlled by the same system clock, the `PWM_SYNC` can usually be considered synchronous. Synchronization logic will add

latency and jitter from the external sync pulse to the actual PWM outputs. If the same asynchronous external sync pulse is received by two independent PWM Controllers, synchronization of `PWM_SYNC` is also done independently and the jitter between the PWM Controllers will not be in unison. The size of the sync pulse on `PWM_SYNC` must be greater than two system clock periods.

The produced PWM output signals can be shut off via:

- **Hardware.** A dedicated asynchronous PWM shutdown pin (`PWM_TRIPB`) that when brought low (provided it is not disabled by the `PWMTRIP_DSBL` bit of the `PWM_CTRL` register) instantaneously places all six PWM outputs in the “off” state (as determined by the state of the `PWM_POLARITY` bit of the `PWM_CTRL` register). This hardware shutdown mechanism is asynchronous so that the associated PWM disable circuitry does not go through any clocked logic, thereby ensuring correct PWM shutdown even in the event of a loss of the processor system clock. A trip shutdown in hardware resets the `PWM_EN` bit in the `PWM_CTRL` register, but all the other programmable registers maintain their current state.
- **Software.** The PWM system may be shut down in software by disabling the `PWM_ENABLE` bit in the `PWM_CTRL` register.



On many processors, the PWM pins are multiplexed with other functionality. Because they can be in a high-impedance state before the `PORTx_MUX` registers are programmed to select the PWM functionality, there should be external pull-down logic for the `PWM_TRIPB` pin in these cases. For these and other questions about pin multiplexing, see the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

The PWM unit is capable of generating two different interrupt types. One interrupt (`PWM_SYNCINT`) is generated on the occurrence of a rising edge on the `PWM_SYNC` pulse, which is internally generated. The other interrupt (`PWM_TRIPINT`) is generated on the occurrence of `PWM_TRIPB`, the PWM

General Operation

shut-down action. Both interrupts are generated only when the corresponding enable bits (`PWMSYNCINT_EN` and `PWMTRIPINT_EN`) are set in the `PWM_CTRL` register.

The `PWM_STAT` register provides status information about the PWM system. In particular, the state of the `PWM_TRIPB` pin (`PWM_TRIP` bit), `PWM_POLARITY` (`PWM_POL` bit), and `PWM_SRMODE` (`PWM_SR` bit) are available, as well as a status bit (`PWM_PHASE`) that indicates whether operation is in the first half or the second half of the PWM period. The `PWM_STAT` register also reflects the status of the `PWM_SYNCINT` and `PWM_TRIPINT` interrupts, which are set if enabled in the `PWM_CTRL` register. The latter two bits are sticky; hence, the interrupt service routine must write-1-to-clear (W1C) these bits.

General Operation

Typically, the `PWM_SYNCINT` interrupt is used to periodically execute an interrupt service routine (ISR) to update the three PWM channel duties, according to a control algorithm based on expected motor operation and sampled data of the existing motor operation. `PWM_SYNC` can also trigger the ADC to sample data for use during the ISR. During processor boot, the PWM Controller is initialized and program flow enters a wait loop. When a `PWM_SYNCINT` interrupt occurs, the ADC samples data, the data is algorithmically interpreted, and then the new PWM channel duties are calculated and written to the PWM registers. More sophisticated implementations include different start-up, run-time, and shut-down algorithms to determine PWM channel duties, based on expected behavior and further features.

During initialization, the `PWM_TM` register is written to define the PWM period, and the `PWM_CHA`, `PWM_CHB` and `PWM_CHC` registers are written to define the initial channel pulse widths. The `PWM_SYNCWT`, `PWM_GATE`, `PWM_SEG`, `PWM_CHAL`, `PWM_CHBL` and `PWM_CHCL` registers are written, depending on the system configuration and modes. The `PWM_STAT` register can be

read to determine polarity, and whether switched reluctance (SR) mode (PWM_SR bit) is enabled, and whether an external trip situation is preventing the correct start-up of the PWM Controller. An active external trip event must be resolved prior to PWM startup. The PWM_CTRL register is then written to define the major operating mode and to enable the PWM outputs and PWM sync pulse.

During the PWM_SYNCINT interrupt-driven control loop, only the PWM_CHx duty values are updated typically. The PWM_SEG register may also be updated for other system implementations requiring output crossover.

During an external trip event (if not disabled), the PWM outputs will be turned off (that is, set to the opposite of the “on” polarity configured by the PWM_POLARITY bit of the PWM_CTRL register), and the PWM sync pulse will continue to operate if already enabled. A PWM_TRIPINT interrupt will occur if unmasked, notifying the software of this event. To handle cases where clock signal integrity is an issue, external trips will turn off the PWM outputs, with or without clocks.

Functional Description

This section describes the function of the following PWM features:

- “Three-Phase PWM Timing Unit and Dead Time Control Unit” on page 14-10
- “PWM Switching Frequency (PWM_TM) Register” on page 14-10
- “PWM Switching Dead Time (PWM_DT) Register” on page 14-11
- “PWM Operating Mode (PWM_CTRL and PWM_STAT) Registers” on page 14-12

Functional Description

- “PWM Duty Cycle (PWM_CHA, PWM_CHB, and PWM_CHC) Registers” on page 14-14
- “Special Consideration for PWM Operation in Over-Modulation” on page 14-19
- “Three-Phase PWM Timing Unit Operation” on page 14-22
- “Effective PWM Accuracy” on page 14-22
- “Switched Reluctance Mode” on page 14-24
- “Output Control Unit” on page 14-24
- “Switched Reluctance (SR) Mode” on page 14-30
- “PWM Sync Operation” on page 14-33
- “PWM Shutdown and Interrupt Control Unit” on page 14-35

Three-Phase PWM Timing Unit and Dead Time Control Unit

The 16-bit Three-Phase PWM Timing Unit is the core of the PWM Controller and produces three pairs of pulse-width modulated signals with high resolution and minimal processor overhead. The outputs of this unit are such that a low level is interpreted as a command to turn on (active-low) the associated power device. Three configuration registers (PWM_CTRL, PWM_TM, and PWM_DT) determine the fundamental characteristics of the PWM outputs. These registers, in conjunction with the three 16-bit duty cycle registers (PWM_CHA, PWM_CHB, and PWM_CHC), control the output of the Three-Phase PWM Timing Unit.

PWM Switching Frequency (PWM_TM) Register

The 16-bit read/write PWM period register (PWM_TM) controls the PWM switching frequency. The fundamental timing unit of the PWM

Controller is t_{SCLK} . Therefore, for a 100 MHz system clock ($SCLK$), f_{SCLK} , the fundamental time increment (t_{SCLK}) is 10 ns. The value written to the PWM_TM register is effectively the number of t_{SCLK} clock increments in half a PWM period. The required PWM_TM value as a function of the desired PWM switching frequency (f_{PWM}) is given by:

$$PWM_TM = \frac{f_{SCLK}}{2 \times f_{PWM}}$$

Therefore, the PWM switching period (T_s) can be written as:

$$T_s = 2 \times PWM_TM \times t_{SCLK}$$

For example, for an f_{SCLK} of 100 MHz and a desired PWM switching frequency (f_{PWM}) of 10 kHz ($T_s = 100 \mu s$), the correct value to load into the PWM_TM register is:

$$PWM_TM = \frac{100 \times 10^6}{2 \times 10 \times 10^3} = 5000$$

The largest value that can be written to the 16-bit PWM_TM register is $0xFFFF = 65,535$, which, at an f_{SCLK} of 100 MHz, corresponds to a minimum PWM switching frequency of:

$$f_{PWM(min)} = \frac{100 \times 10^6}{2 \times 65535} = 762 Hz$$



PWM_TM values of 0 and 1 are not defined and should not be used when the PWM outputs or PWM sync is enabled.

PWM Switching Dead Time (PWM_DT) Register

The second important parameter that must be set up in the initial configuration of the PWM Controller is the switching dead time. This is a short delay introduced between turning off one PWM signal (for example, AH)

Functional Description

and turning on the complementary signal (for example, AL). This short time delay permits the power switch being turned off (AH in this case) to completely recover its blocking capability before the complementary switch is turned on. This time delay prevents a potentially destructive short-circuit condition from developing across the dc link capacitor of a typical voltage source inverter.

The 10-bit, read/write `PWM_DT` register controls the dead time. This register controls the dead time inserted into the three pairs of PWM output signals. Dead time (T_d) is related to the value in the `PWM_DT` register by:

$$T_d = PWM_DT \times 2 \times t_{SCLK}$$

Therefore, a `PWM_DT` value of 0x00A introduces a 200 ns delay (for a `SCLK` of 100 MHz) between turning off any PWM signal (for example, AH) and then turning on its complementary signal (for example, AL). The length of the dead time can therefore be programmed in increments of $2 \times t_{SCLK}$ (or 20 ns for an `SCLK` of 100 MHz). The `PWM_DT` register is a 10-bit register whose maximum value of 0x3FF (1023 decimal) corresponds to a maximum programmed dead time of:

$$T_{d(max)} = 1023 \times 2 \times t_{SCLK} = 1023 \times 2 \times 10 \times 10^{-9} = 20.5 \mu s$$

for an `fSCLK` rate of 100 MHz. The dead time can be programmed to be zero by writing 0 to the `PWM_DT` register.

PWM Operating Mode (PWM_CTRL and PWM_STAT) Registers

The PWM Controller can operate in two distinct modes: single-update mode and double-update mode. The mode is determined by the state of `PWM_DBL` bit of the `PWM_CTRL` register. When this bit is cleared, the PWM Controller operates in single-update mode. Setting the `PWM_DBL` bit places the PWM Controller in double-update mode. Following a peripheral reset

or power on, the `PWM_DBL` bit is cleared; thus, PWM Controller operation defaults to single-update mode.

In single-update mode, a `PWM_SYNC` pulse is produced during each PWM period. The rising edge of this signal marks the start of a new PWM cycle and is used to latch new values from the PWM configuration registers (`PWM_TM`, `PWM_DT`, and `PWM_SYNCWT`), and the PWM duty cycle registers (`PWM_CHA`, `PWM_CHB`, `PWM_CHC`, `PWM_CHAL`, `PWM_CHBL`, and `PWM_CHCL`) into the Three-Phase PWM Timing Unit. In addition, the `PWM_SEG` register is also latched into the Output Control Unit on the rising edge of the `PWM_SYNC` pulse. In effect, this means that the characteristics and resultant duty cycles of the PWM signals can be updated only once per PWM period at the start of each cycle. This results in PWM patterns that are symmetrical about the midpoint of the switching period.

In double-update mode, an additional `PWM_SYNC` pulse is produced at the midpoint of each PWM period. The rising edge of this second `PWM_SYNC` pulse is again used to latch new values of the PWM configuration registers, duty cycle registers, and the `PWM_SEG` register. As a result, it is possible to alter both the characteristics (switching frequency, dead time, and `PWM_SYNC` pulse width) and the output duty cycles at the midpoint of each PWM cycle. Consequently, it is possible to produce PWM switching patterns that are no longer symmetrical about the midpoint of the period (asymmetrical PWM patterns).

In double-update mode, it may be necessary to know whether operation at any point in time is in the first or second half of the PWM cycle. This information is provided by the `PWM_PHASE` bit of the `PWM_STAT` register, which is cleared during operation in the first half of each PWM period (between the rising edge of the original `PWM_SYNC` pulse and the rising edge of the second `PWM_SYNC` pulse introduced in double-update mode). The `PWM_PHASE` bit is set during operation in the second half of each PWM period. This status bit allows determination of the particular half-cycle during implementation of the `PWM_SYNC` interrupt service routine.

Functional Description

The advantage of double-update mode is that the PWM process can produce lower harmonic voltages, and faster control bandwidths are possible. However, for a given PWM switching frequency, `PWM_SYNC` pulses occur at twice the rate in double-update mode. Since new duty cycle values are computed in each `PWM_SYNCINT` interrupt service routine, double-update mode places a larger computational burden on the processor. Alternatively, the same PWM update rate may be maintained at half the switching frequency, yielding lower switching losses.

The `PWM_STAT2` status register is provided for software simulation. This register contains the output values of all the three pairs of PWM signals (`PWM_AH`, `PWM_AL`, `PWM_BH`, `PWM_BL`, `PWM_CH`, and `PWM_CL`).

PWM Duty Cycle (`PWM_CHA`, `PWM_CHB`, and `PWM_CHC`) Registers

Three 16-bit read/write duty cycle registers (`PWM_CHA`, `PWM_CHB`, and `PWM_CHC`) control the duty cycles of the six PWM output signals on the `PWM_AH`, `PWM_AL`, `PWM_BH`, `PWM_BL`, `PWM_CH`, and `PWM_CL` pins when not in switched reluctance mode. The two's complement integer value in the `PWM_CHA` register controls the duty cycle of the signals on the `PWM_AH` and `PWM_AL` outputs; in `PWM_CHB`, it controls the duty cycle of the signals on `PWM_BH` and `PWM_BL`; in `PWM_CHC`, it controls the duty cycle of the signals on `PWM_CH` and `PWM_CL`. The duty cycle registers are programmed in two's complement integer counts of the fundamental time unit (t_{SCLK}) and define the desired on-time of the high-side PWM signal produced by the Three-Phase PWM Timing Unit over half the PWM period.

Each duty cycle register range is from $(-PWMTM/2 - PWMDT)$ to $(+PWMTM/2 + PWMDT)$, which, by definition, is scaled such that a value of 0 represents a 50% PWM duty cycle.

The switching signals produced by the Three-Phase PWM Timing Unit are also adjusted to incorporate the programmed dead time value in the `PWM_DT` register by programming active low polarity in `PWM_CTRL`. The

Three-Phase PWM Timing Unit produces active-low signals to turn on the associated power device.

Figure 14-3 shows a typical pair of PWM outputs (in this case, for PWM_AH and PWM_AL) from the Three-Phase PWM Timing Unit for operation in single-update mode. All illustrated time values indicate the integer value in the associated register and can be converted to time by multiplying by the fundamental time increment (t_{SCLK}) and comparing to the two's complement counter.

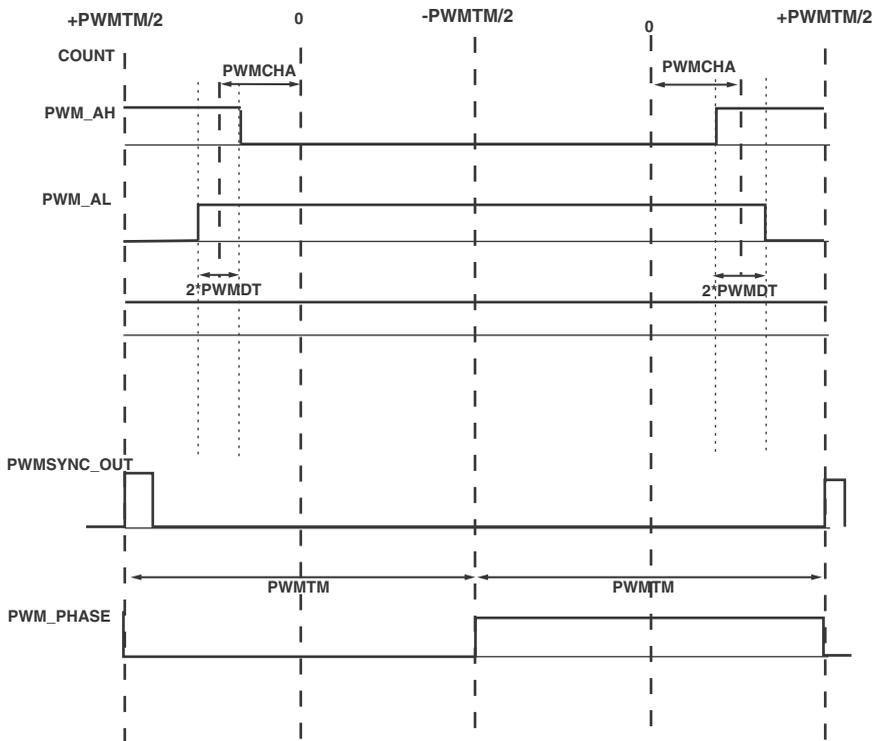


Figure 14-3. Typical PWM Outputs of Three-Phase Timing Unit in Single-Update Mode (Active-Low Waveforms)

Functional Description

Notice that the switching patterns are perfectly symmetrical about the midpoint of the switching period in single-update mode, since the same values of `PWM_CHA`, `PWM_TM`, and `PWM_DT` are used to define the signals in both half cycles of the period.

As implied by [Figure 14-3](#), the programmed duty cycles are adjusted to incorporate the desired dead time into the resultant pair of PWM signals by moving the switching instants of both PWM signals (`PWM_AH` and `PWM_AL`) away from the instant set by the `PWM_CHA` register. Both switching edges are moved by an equal amount ($PWMDT \times t_{SCLK}$) to preserve the symmetrical output patterns. [Figure 14-3](#) shows the `PWM_SYNC` pulse whose rising edge denotes the beginning of the switching period and whose width is set by the `PWM_SYNCWT` register. Also shown is the `PWM_PHASE` bit of the `PWM_STAT` register, which indicates whether operation is in the first half cycle or second half cycle of the PWM period.

The resultant on-times (active low) of the PWM signals over the full PWM period (two half periods) produced by the Three-Phase PWM Timing Unit and illustrated in [Figure 14-3](#), may be written as:

$$T_{AH} = (PWMTM + 2 \times (PWMCHA - PWMDT)) \times t_{SCLK}$$

$$\text{Range of } T_{AH} = [0, 2 \times PWMTM \times t_{SCLK}]$$

$$T_{AL} = (PWMTM - 2 \times (PWMCHA + PWMDT)) \times t_{SCLK}$$

$$\text{Range of } T_{AL} = [0, 2 \times PWMTM \times t_{SCLK}]$$

and the corresponding duty cycles are:

$$d_{AH} = \frac{T_{AH}}{T_S} = \frac{1}{2} + \frac{PWMCHA - PWMDT}{PWMTM}$$

$$d_{AL} = \frac{T_{AL}}{T_S} = \frac{1}{2} - \frac{PWMCHA + PWMDT}{PWMTM}$$

Obviously, negative values of T_{AH} and T_{AL} are not permitted, and the minimal permissible value is zero (corresponding to a 0% duty cycle). In a similar fashion, the maximal value is T_S , which is the PWM switching period that corresponds to a 100% duty cycle.

Figure 14-4 shows the output signals from the Three-Phase PWM Timing Unit in double-update mode. This figure illustrates a completely general case in which the switching frequency, dead time, and duty cycle are changed in the second half of the PWM period. Of course, the same value for any or all of these quantities may be used in both halves of the PWM cycle. However, it can be seen that there is no guarantee that a symmetrical PWM signal will be produced by the Three-Phase PWM Timing Unit in double-update mode. Additionally, it is seen that the dead time is inserted into the PWM signals similarly to single-update mode.

In general, the on-times (active low) of the PWM signals over the full PWM period in double-update mode can be defined as:

$$T_{AH} = \left(\frac{PWMTM_1}{2} + \frac{PWMTM_2}{2} + PWMCHA_1 + PWMCHA_2 - PWMDT_1 - PWMDT_2 \right) \times t_{SCLK}$$

$$T_{AL} = \left(\frac{PWMTM_1}{2} + \frac{PWMTM_2}{2} - PWMCHA_1 - PWMCHA_2 - PWMDT_1 - PWMDT_2 \right) \times t_{SCLK}$$

Functional Description

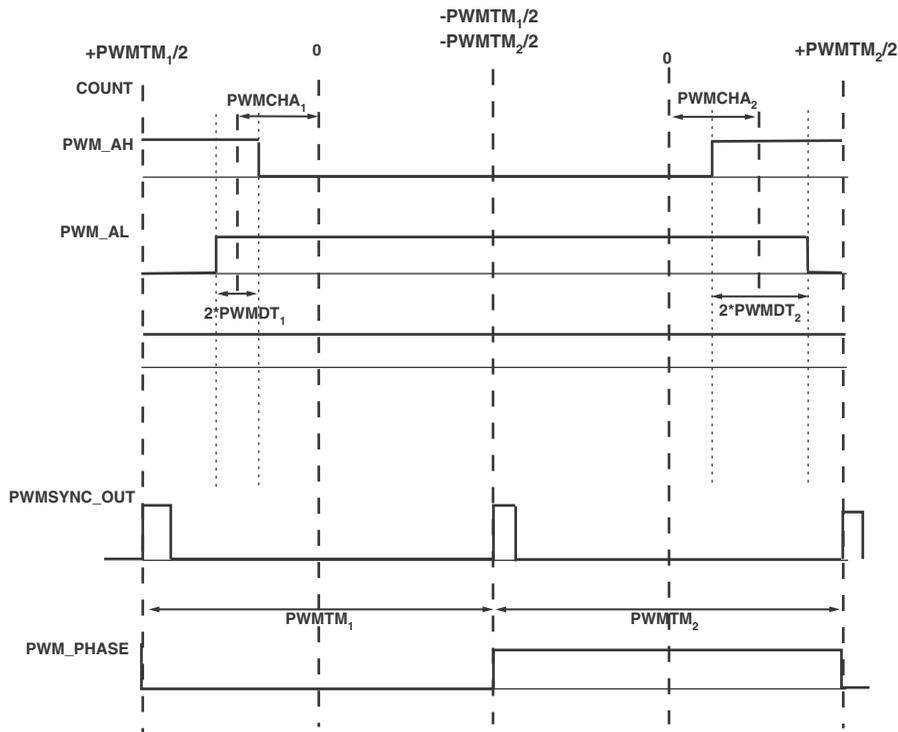


Figure 14-4. Typical PWM Outputs of Three-Phase Timing Unit in Double-Update Mode (Active Low Waveforms)

$$T_S = (PWMTM_1 + PWMTM_2) \times t_{SCLK}$$

where subscript 1 refers to the value of that register during the first half cycle and subscript 2 refers to the value during the second half cycle. The corresponding duty cycles are:

$$d_{AH} = \frac{T_{AH}}{T_S} = \frac{1}{2} + \frac{(PWMCHA_1 + PWMCHA_2 - PWMDT_1 - PWMDT_2)}{(PWMTM_1 + PWMTM_2)}$$

$$d_{AL} = \frac{T_{AL}}{T_S} = \frac{1}{2} - \frac{(PWMCHA_1 + PWMCHA_2 + PWMDT_1 + PWMDT_2)}{(PWMTM_1 + PWMTM_2)}$$

since for the completely general case in double-update mode, the switching period is given by:

$$T_S = (PWMTM_1 + PWMTM_2) \times t_{SCLK}$$

Again, the values of T_{AH} and T_{AL} are constrained to lie between zero and T_S . Similar PWM signals to those illustrated in [Figure 14-2 on page 14-5](#) and in [Figure 14-3 on page 14-15](#) can be produced on the BH, BL, CH, and CL outputs by programming the PWM_CHB and PWM_CHC registers in a manner identical to that described for PWM_CHA.

Special Consideration for PWM Operation in Over-Modulation

The Three-Phase PWM Timing Unit can produce PWM signals with variable duty-cycle values at the PWM output pins. At the extremities of the modulation process, both 0% and 100% modulation (termed *full off mode* and *full on mode*, respectively) are possible. In between, for other duty cycle values, the operation is termed *normal modulation*.

- Full On Mode.** The PWM for any pair of PWM signals is said to operate in full on mode when the desired high side output of the Three-Phase PWM Timing Unit is in the “on” state (low or high as specified by PWM_POLARITY bit of the PWM_CTRL register) between successive PWM_SYNC rising edges. This state may be entered by virtue of the commanded duty cycle values (in conjunction with the PWM_DT register).

Functional Description

- **Full Off Mode.** The PWM for any pair of PWM signals is said to operate in full off mode when the desired high side output of the Three-Phase PWM Timing Unit is in the “off” state (high or low as specified by the `PWM_POLARITY` bit of the `PWM_CTRL` register) between successive `PWM_SYNC` pulses. This state may be entered by virtue of the commanded duty cycle values (in conjunction with the `PWM_DT` register).
- **Normal Modulation.** The PWM for any pair of PWM signals is said to operate in normal modulation when the desired output duty cycle is other than 0% or 100% between successive `PWM_SYNC` pulses.

Certain situations exist whereby it is necessary to transition into or out of full on mode or full off mode in order to insert additional “emergency dead time” delays to prevent potential shoot-through conditions in the inverter. Crossover usage also can potentially cause outputs to violate shoot-through condition criteria, as described in [“Crossover Feature” on page 14-24](#). These transitions are detected automatically and, if appropriate for safety, emergency dead-time is inserted to prevent shoot-through conditions.

The insertion of additional emergency dead time into one of the PWM signals of a given pair during these transitions is necessary only when both PWM signals are required to toggle within a dead time of each other. The additional emergency dead time delay is inserted into the PWM signal that is toggling into the “on” state. In effect, the turn on (if turning on during this dead time region) of this signal is delayed by an amount ($2 * \text{PWM_DT} * t_{\text{SCLK}}$) from the rising edge of the opposite output. After this delay, the PWM signal is allowed to turn on, provided the desired output is still scheduled to be in the on state after the emergency dead time delay.

[Figure 14-5](#) illustrates two examples of such transitions. In the top half (marked A) of [Figure 14-5](#), no special action (dead time) is needed when transitioning from normal modulation to full on mode at the half cycle boundary in double-update mode. However, in the bottom half

(marked B) of Figure 14-5, when transitioning from normal modulation into full off mode at the same boundary, it can be seen that an additional emergency dead time is necessary (inserted by the PWM Controller). Clearly, this inserted dead time is different from the normal dead time as it is impossible to move one of the switching events back in time, because this would take it into the previous modulation cycle. Therefore, the entire emergency dead time is inserted by delaying the turn on of the appropriate signal by the full amount.

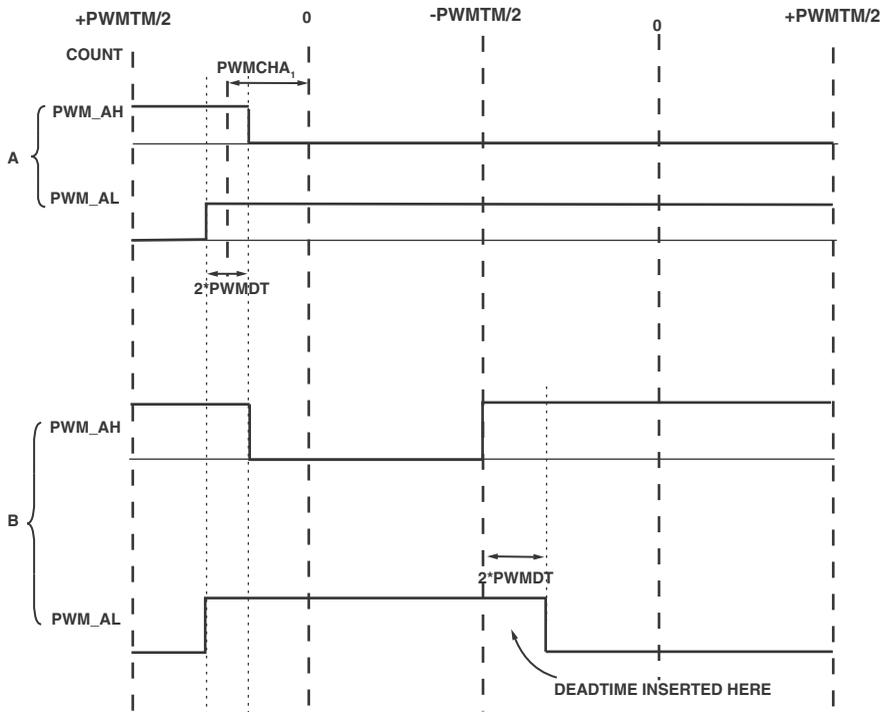


Figure 14-5. Examples of Transitioning from Normal Modulation to Full On Mode (A) or Full Off Mode (B)

Functional Description

Three-Phase PWM Timing Unit Operation

The internal operation of the PWM Controller is controlled by the Three-Phase PWM Timing Unit, which is clocked at the system clock rate with period t_{SCLK} . The operation of the Three-Phase PWM Timing Unit over one full PWM period is illustrated in [Figure 14-6](#).

During the first half cycle (when the `PWM_PHASE` bit of the `PWM_STAT` register is cleared), the Three-Phase PWM Timing Unit decrements from $+PWMTM/2$ to $-PWMTM/2$ using a two's complement count. Then the count direction changes, and the unit increments from $-PWMTM/2$ to the $+PWMTM/2$ value.

[Figure 14-6](#) also shows the PWM SYNC pulses during single-update mode and double-update mode. Clearly, an additional PWM SYNC pulse is generated at the midpoint of the PWM cycle in double-update mode. If the value of the `PWM_TM` register is altered at the midpoint in double-update mode, the duration of the second half period (when the `PWM_PHASE` bit of the `PWM_STAT` register is set) may differ from that of the first half cycle. `PWM_TM` is double-buffered; a change in one half of the PWM switching period will only take effect in the next half period.

Effective PWM Accuracy

The PWM Controller has 16-bit resolution, but accuracy depends on the PWM period. In single-update mode, the same values of `PWM_CHA`, `PWM_CHB`, and `PWM_CHC` define the on-times in both half cycles of the PWM period. As a result, the effective accuracy of the PWM generation process is $2 \times t_{\text{SCLK}}$ (20 ns for a 100 MHz f_{SCLK}). Incrementing one of the duty cycle registers by 1 changes the resultant on-time of the associated PWM signals by t_{SCLK} in each half period ($2 \times t_{\text{SCLK}}$ for the full period). In double-update mode, improved accuracy is possible since different values of the duty cycles registers are used to specify the on-times in both the first half and second half of the PWM period. As a result, it is possible to adjust the on-time over the entire period in increments of t_{SCLK} . This

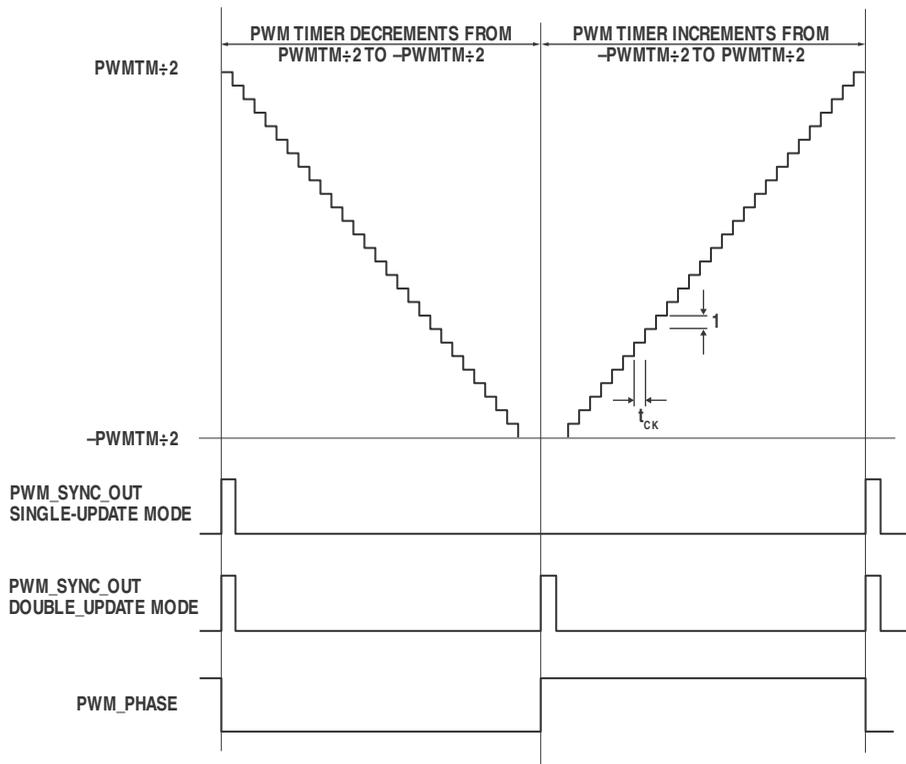


Figure 14-6. Operation of Internal PWM Timer

corresponds to an effective PWM accuracy of t_{SCLK} in double-update mode (10 ns for a 100 MHz f_{SCLK}). The minimum achievable PWM switching frequency at a given PWM accuracy is shown in [Table 14-1](#) for $SCLK = 100$ MHz.

Functional Description

Table 14-1. Minimum Achievable PWM Frequency Versus Bit Resolution for SCLK = 100 MHz

Resolution (bits)	PWM Frequency (kHz) in Single-Update Mode	PWM Frequency (kHz) in Double-Update Mode
8	195.3	390.6
9	97.7	195.3
10	48.8	97.7
11	24.4	48.8
12	12.2	24.4
13	6.1	12.2
14	3.05	6.1

Switched Reluctance Mode

A general-purpose mode utilizing independent edge placement of upper and lower signals of each of the three PWM channels is incorporated into the Three-Phase PWM Timing Unit. This mode is provided for SR motor operation and is described in detail in [“Switched Reluctance \(SR\) Mode” on page 14-30](#).

Output Control Unit

The operation of the Output Control Unit is controlled by the 9-bit read/write `PWM_SEG` register ([on page 14-44](#)) that controls two distinct features that are useful in the control of ECMs or BDCMs.

Crossover Feature

The `PWM_SEG` register contains three crossover bits—one for each pair of PWM outputs. Setting the `AHAL_XOVR` bit of the `PWM_SEG` register enables crossover mode for the `AH/AL` pair of PWM signals, setting `BHBL_XOVR` enables crossover on the `BH/BL` pair, and setting `CHCL_XOVR` enables

crossover on the CH/CL pair. If crossover mode is enabled for any pair of PWM signals, the high-side PWM signal (for example, AH) from the Three-Phase PWM Timing Unit is diverted to the associated low-side output of the Output Control Unit so that the signal ultimately appears at the AL pin. The corresponding low-side output of the Three-Phase PWM Timing Unit is also diverted to the complementary high-side output of the Output Control Unit so that the signal appears at the AH pin. Following a reset, the three crossover bits are cleared, disabling crossover mode on all three pairs of PWM signals. Even though crossover is considered an output control feature, dead time insertion occurs after crossover transitions (as necessary to eliminate shoot-through safety issues).

Mode Bits (POLARITY and SRMODE)

PWM_POLARITY and PWM_SRMODE are programmable bits of the PWM_CTRL register.

- ⊘ The incorrect programming of these two mode-select signals can have destructive consequences on the external power inverter connected to the PWM unit. Since PWM_POLARITY and PWM_SRMODE are software programmable bits, accidental power inverter shoot-through current may occur from incorrect programming.

Output Enable Function

The PWM_SEG register also contains six bits (bits 0 to 5) that can be used to individually enable or disable each of the six PWM outputs. The PWM signal of the AL pin is enabled by clearing the AL_EN bit of the PWM_SEG register, the AH_EN bit controls AH, the BL_EN bit controls BL, the BH_EN bit controls BH, the CL_EN bit controls CL, and the CH_EN bit controls the CH output. If the associated bit of the PWM_SEG register is set, the corresponding PWM output is disabled irrespective of the value of the corresponding duty cycle register. This PWM output signal will remain in the off state as long as the corresponding enable/disable bit of the PWM_SEG register is set. This output enable function is implemented after the crossover function.

Functional Description

Following a reset, all six enable bits of the `PWM_SEG` register are cleared so that all PWM outputs are enabled by default. In a manner identical to the duty cycle registers, the `PWM_SEG` register is latched on the rising edge of the `PWM_SYNC` signal so that changes to this register only become effective at the start of each PWM cycle in single-update mode. In double-update mode, the `PWM_SEG` register can also be updated at the midpoint of the PWM cycle.

Brushless DC Motor (Electronically Commutated Motor) Control

In the control of an electronically commutated motor (ECM), only two inverter legs are switched at any time. Often, the high-side device in one leg must be switched on at the same time as the low-side driver in a second leg. Therefore, by programming identical duty cycles values for two PWM channels (for example, `PWM_CHA = PWM_CHB`) and setting the `BHBL_XOVR` bit of the `PWM_SEG` register to crossover the `BH/BL` pair if PWM signals, it is possible to turn on the high-side switch of phase A and the low-side switch of phase B at the same time.

In ECM control, usually the third inverter leg (phase C in this example) is disabled for a number of PWM cycles. This is implemented by disabling the `CH` and `CL` outputs by setting the `CH_EN` and `CL_EN` bits of the `PWM_SEG` register.

This is illustrated in [Figure 14-7](#) where it can be seen that both the `AH` and `BL` signals are identical (since `PWM_CHA = PWM_CHB` and the crossover bit for phase B is set). In addition, the other four signals (`AL`, `BH`, `CH`, and `CL`) are disabled by setting the appropriate enable/disable bits of the `PWM_SEG` register.

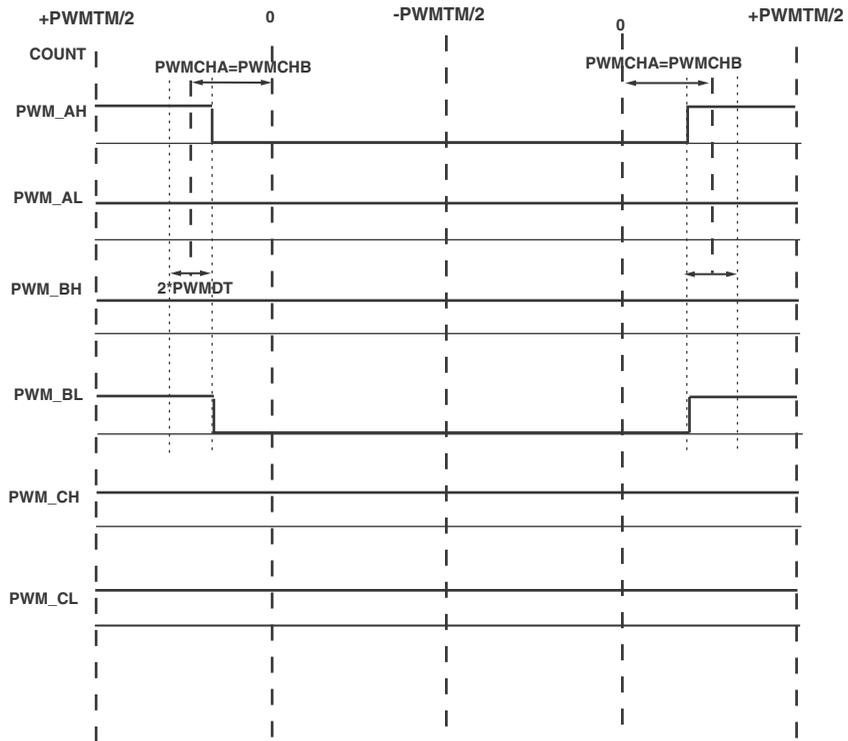


Figure 14-7. Example of Active Low Signals for ECM Control

For the situation illustrated in [Figure 14-7](#), an appropriate value for the `PWM_SEG` register is `0x00A7`. In normal ECM operation, each inverter leg is disabled for certain lengths of time, such that the `PWM_SEG` register is changed, based upon the position of the rotor shaft (motor commutation).

Gate Drive Unit

The Gate Drive Unit is described in the following sections:

- [“High-Frequency Chopping” on page 14-28](#)
- [“PWM Polarity Control” on page 14-29](#)

Functional Description

High-Frequency Chopping

The Gate Drive Unit of the PWM Controller simplifies the design of isolated gate drive circuits for PWM inverters. If a transformer-coupled power device gate drive amplifier is used, the active PWM signal must be chopped at a high frequency. The 10-bit read/write `PWM_GATE` register allows you to specify this high-frequency chopping mode.

Chopped active PWM signals may be required for high-side drivers only, for low-side drivers only, or for both high-side and low-side switches. Therefore, independent control of this mode for both high- and low-side switches is included with two separate control bits (`CHOPHI` and `CHOPLO`) in the `PWM_GATE` register.

Typical PWM output signals with high-frequency chopping enabled on both high- and low-side signals are shown in [Figure 14-8](#). Chopping the high-side PWM outputs (`AH`, `BH`, and `CH`) is enabled by setting the `CHOPHI` bit of the `PWM_GATE` register; chopping the low-side PWM outputs (`AL`, `BL`, and `CL`) is enabled by setting the `CHOPLO` bit of the `PWM_GATE` register. The high-frequency chopping frequency is controlled by the 8-bit word placed in bits 0 to 7 (`GDCLK`) of the `PWM_GATE` register. The period of this high-frequency carrier is:

$$T_{chop} = [4 \times (GDCLK + 1)] \times t_{SCLK}$$

and the chopping frequency is therefore an integral subdivision of the system clock frequency:

$$f_{chop} = \frac{f_{SCLK}}{[4 \times (GDCLK + 1)]}$$

The `GDCLK` value may range from 0 to 255, which corresponds to a programmable chopping frequency rate from 97.7 kHz to 25 MHz for a 100 MHz f_{SCLK} rate. The gate drive features must be programmed before

operation of the PWM Controller and typically are not changed during normal operation of the PWM Controller. Following a reset, all bits of the PWM_GATE register are cleared so that high-frequency chopping is disabled, by default.

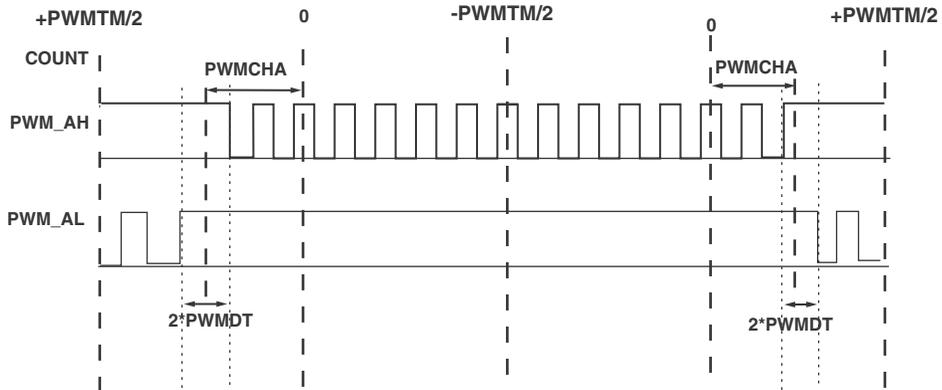


Figure 14-8. Example of Active Low PWM Signals for Gate Chopping

PWM Polarity Control

The polarity of the PWM signals produced at output pins AH to CL can be programmed via the PWM_POLARITY bit of the PWM_CTRL register. Setting this bit to 0 selects active low PWM outputs, such that a low level is interpreted as a command to turn on the associated power device. Conversely, setting the PWM_POLARITY bit to 1 selects active high PWM outputs, such that a high level at the PWM outputs turns on the associated power devices. The status of the polarity may be read from the PWM_POL bit of the PWM_STAT register, where a zero indicates a measured low level at the PWM_POLARITY bit.

Output Control Feature Precedence

It is important to understand the order in which output control features are applied to the PWM signal. The following hierarchy indicates the order (from most important to least important) in which signal features are applied to the PWM output signal.

1. Channel duty generation
2. Channel crossover
3. Low-side invert
4. Output enable
5. Emergency dead time insertion
6. Active signal chopping
7. Polarity

Switched Reluctance (SR) Mode

The PWM Controller provides a switched reluctance (SR) mode that is enabled by setting the `PWM_SRMODE` bit in the `PWM_CTRL` register to 0. This mode is not enabled by default. The state of this switched reluctance mode may be read from the `PWM_SR` bit of the `PWM_STAT` register. If the `PWM_SRMODE` bit is high (such that SR mode is disabled) the `PWM_SR` bit of the `PWM_STAT` register is set (indicating that the mode is disabled). Conversely, if the `PWM_SRMODE` bit is low and SR mode is enabled, the `PWM_SR` bit of `PWM_STAT` register is cleared.



Since this is a software programmable bit, be careful not to write it to an active state in a non-SR mode system and cause shoot-through at the power inverters, possibly leading to an unsafe situation.

In the typical power converter configuration for switched or variable reluctance motors, the motor winding is connected between the two power switches of a given inverter leg. Therefore, to allow for a complete circuit in the motor winding, it is necessary to turn on both switches at the same time.

SR mode provides four mode types: hard chop, alternate chop, soft chop-bottom on, and soft chop-top on (see [Table 14-2 on page 14-33](#)). Three registers (`PWM_CHAL`, `PWM_CHBL`, and `PWM_CHCL`) are used to define edge placement of the low side of the channel. The `PWM_DT` register, which is not used, is internally forced to 0 by hardware when SR mode is active. The four switched reluctance (SR) chop modes are specified via three bits (`PWM_SR_LSI_A`, `PWM_SR_LSI_B`, and `PWM_SR_LSI_C`) of the `PWM_LSI` register, full on mode, and full off mode.

The `PWM_CHA` and `PWM_CHAL` registers are programmed independently; `PWM_CHA` specifies edge placement for the high side of the channel, and `PWM_CHAL` specifies edge placement for the low side of the channel. Similarly, the `PWM_CHB` and `PWM_CHBL` pair, and the `PWM_CHC` and `PWM_CHCL` pair, respectively, specify high-side and low-side edge placement.

Functional Description

Figure 14-9 shows the four SR mode types as active-high PWM output signals, and Table 14-2 describes the four mode types.

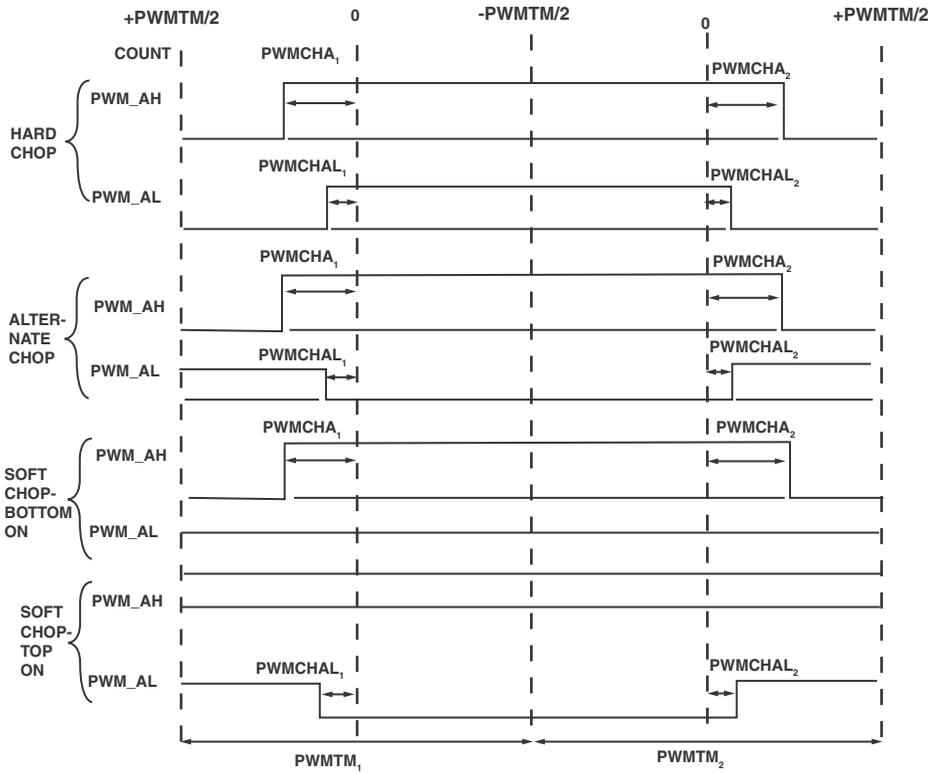


Figure 14-9. Four SR Mode Types

Table 14-2. Switched Reduction Mode (SR Mode) Types

Mode	Description
Hard chop	Contains independently programmed rising edges of channels' high and low signals in the same PWM half cycle, and both contain independently programmed falling edges in the next PWM half cycle. The <code>PWM_CHA</code> duty register is used for the high channel, and the <code>PWM_CHAL</code> duty register is used for the low channel. A similar structure is present for the B and C channels.
Alternate chop	Similar to normal PWM operation, but the PWM channel high and low signal edges are opposite and are independently programmed. The <code>PWM_CHA</code> duty register is used for the high channel, and the <code>PWM_CHAL</code> duty register is used for the low channel. A similar structure is present for the B and C channels. The <code>PWM_CTRL</code> and <code>PWM_LSI</code> registers are used to independently invert the low side of each PWM channel. The low-side invert is the only difference between hard chop mode and alternate chop mode
Soft chop-bottom on	Utilizes a 100% duty on the low side of the channel. Similar to hard chop mode, the <code>PWM_CHA</code> duty register is used for the high channel and the <code>PWM_CHAL</code> duty register is used for the low channel. A similar structure is present for the B and C channels.
Soft chop-top on	Utilizes a 100% duty on the high side of the channel. Similar to hard chop mode, the <code>PWM_CHA</code> duty register is used for the high channel and the <code>PWM_CHAL</code> duty register is used for the low channel. A similar structure is present for the B and C channels.

PWM Sync Operation

The PWM sync can be internally generated as a function of the `PWM_TM` and `PWM_SYNCWT` register values, or the PWM sync can be input externally. Multiple PWM configurations can be established, each of which can operate with its own independent PWM sync (or from its own external PWM sync signal or a shared external PWM sync signal). The external PWM sync can be synchronous to the internal clock, as in the case of a primary PWM Controller generating an internal `PWM_SYNC` signal that drives a secondary PWM Controller's `PWM_SYNC` pin. The external PWM sync can also be asynchronous to the internal clock, as is typically the case of an off-chip `PWM_SYNC` signal used to drive each PWM Controller's `PWM_SYNC` pin.

Functional Description

Internal PWM SYNC Generation

The PWM Controller produces an output PWM synchronization pulse at a rate equal to the PWM switching frequency in single-update mode and at twice the PWM frequency in double-update mode. This pulse is available for external use at the `PWM_SYNC` pin. The width of this PWM SYNC pulse is programmable by the 10-bit read/write `PWM_SYNCW` register. The width of the PWM SYNC pulse (T_{PWM_SYNC}) is given by:

$$T_{PWM_SYNC} = t_{SCLK} \times (PWMSYNCW + 1)$$

so that the width of the pulse is programmable from t_{SCLK} to $1024 \times t_{SCLK}$ (corresponding to 10 ns to 10.24 μ s for an f_{SCLK} rate of 100 MHz). Following a reset, the `PWM_SYNCW` register contains 0x3FF (1023 decimal) so that the default `PWM_SYNC` width is 10.24 μ s, again for an f_{SCLK} of 100 MHz.

External PWM SYNC Generation

By setting the `PWM_EXTSYNC` bit of the `PWM_CTRL` register, the PWM is set up in a mode to expect an external PWM SYNC on the `PWM_SYNC` pin. The external sync should be synchronized by setting the `PWM_SYNCSEL` bit of the `PWM_CTRL` register to 0, which assumes the selected external PWM SYNC is asynchronous.

The external PWM SYNC period is expected to be an integer multiple of the internal PWM SYNC period. When the rising edge of the external `PWM_SYNC` is detected, the PWM Controller is restarted at the beginning of the PWM cycle. If the external PWM SYNC period is not an integer multiple of the internal PWM SYNC, the behavior of the PWM channel outputs will be clipping. Note that a small amount of jitter inherent in the synchronization logic cannot be avoided when the external PWM SYNC is synchronized.

The latency from `PWM_SYNC` to the effect in PWM outputs is 3 `SCLK` cycles in synchronous mode and 5 `SCLK` cycles in asynchronous mode.

- ❗ In external sync pulse mode, do not allow changes in `PWM_SYNCSEL` (which selects between asynchronous/synchronous external sync pulse) ± 10 `SCLK` cycles of the toggling of an external sync pulse. If this rule is not followed, unexpected behavior may occur.

PWM Shutdown and Interrupt Control Unit

In the event of an external fault condition, it is essential that the PWM Controller be shut down instantaneously in a safe fashion. A falling edge on the `PWM_TRIPB` pin (assuming it is not disabled by the `PWM_TRIP_DSB` bit of the `PWM_CTRL` register) provides an instantaneous, asynchronous (independent of the processor clock) shutdown of the PWM controller. All six PWM outputs are placed in the off state (as defined by the `PWM_POLARITY` bit of the `PWM_CTRL` register). However, the `PWM_SYNC` pulse occurs if it was previously enabled, and the associated interrupt is also not stopped.

- ❗ The processor's `PWM_TRIPB` signal should have an external pull-down resistor; if the pin becomes disconnected, the PWM Controller will be disabled. The state of the `PWM_TRIPB` pin can be read from the `PWM_TRIP` bit of the `PWMSTAT` register.

On the occurrence of a PWM shutdown command (or from a signal on the `PWM_TRIPB` pin), a `PWM_TRIP` interrupt will be generated if enabled. In addition, if `PWM_SYNC_EN` is enabled in the `PWM_CTRL` register, the `PWM_SYNC` pulse will continue to appear at the output pin. Following a PWM shutdown, the PWM can be re-enabled (by a `PWM_TRIP` interrupt service routine, for example) by writing to the `PWM_EN` bit of the `PWM_CTRL` register. The PWM Controller will restart in a manner identical to that prior to the PWM shutdown, provided that the external fault has been cleared and `PWM_TRIPB` returned to a high level. That is, except for the

PWM Registers

PWM_EN bit in the PWM_CTRL register, all PWM registers retain their values during the PWM shutdown.

-  The dead time counters will be reset when a trip occurs, and the user is expected to restart the PWM only after waiting the required dead time. If restarting a PWM immediately after trip, for high dead time period cases, the dead time will not be met.
-  Do not allow changes in the PWM_TRIP_DSBL bit of the PWM_CTRL register (which is to select between trip enable and disable) ± 10 SCLK cycles of the toggling of an external trip pulse. If this rule is not followed, unexpected behavior may occur.

Between the time that the PWM_EN bit is written to 0 and the time the waveforms are disabled, the latency is 2 SCLK cycles. After enabling the PWM_EN bit, output waveforms will begin to appear from the next PWM pulse.

PWM Registers

Descriptions and bit diagrams for each of the PWM memory-mapped registers (MMRs) are provided in the following sections.

Table 14-3. PWM Registers

Name	Description
PWM_CTRL	PWM control register on page 14-37
PWM_STAT	PWM status register on page 14-39
PWM_TM	PWM period register on page 14-40
PWM_DT	PWM dead time register on page 14-41
PWM_GATE	PWM chopping control on page 14-41
PWM_CHA	PWM channel A duty control on page 14-42
PWM_CHB	PWM channel B duty control on page 14-42

Table 14-3. PWM Registers (Continued)

Name	Description
PWM_CHC	PWM channel C duty control on page 14-42
PWM_SEG	PWM crossover and output enable on page 14-44
PWM_SYNCWT	PWM sync pulse width control on page 14-46
PWM_CHAL	PWM channel AL duty control (SR mode only) on page 14-46
PWM_CHBL	PWM channel BL duty control (SR mode only) on page 14-46
PWM_CHCL	PWM channel CL duty control (SR mode only) on page 14-46
PWM_LSI	PWM low side invert (SR mode only) on page 14-48
PWM_STAT2	PWM simulation status register on page 14-49

PWM Control (PWM_CTRL) Register

The PWM_CTRL register is used for configuration of the PWM block. Bit diagrams and descriptions are provided in [Figure 14-10](#) and [Table 14-4](#).

PWM Control Register (PWM_CTRL)

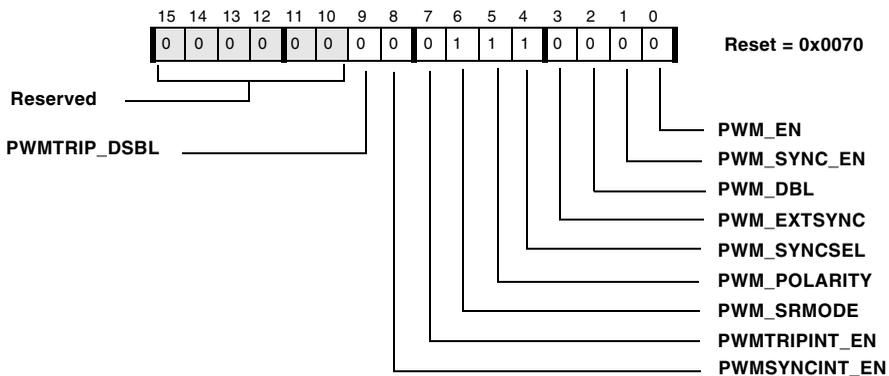


Figure 14-10. PWM Control Register

PWM Registers

Table 14-4. PWM_CTRL Register

Bit	Name	Function	Type	Default
0	PWM_EN Hardware modifiable bit.	PWM enable 0 = disabled 1 = enabled reset by PWM_TRIPB	RW	0
1	PWM_SYNC_EN	PWM sync enable 0 = disabled 1 = enabled	RW	0
2	PWM_DBL	Double-update mode 0 = single-update mode 1 = double-update mode	RW	0
3	PWM_EXTSYNC	External sync 0 = internal sync 1 = external sync	RW	0
4	PWM_SYNCSEL	External sync select 0 = asynchronous 1 = synchronous	RW	1
5	PWM_POLARITY	PWM output polarity 1 = active high 0 = active low	RW	1
6	PWM_SRMODE	PWM SR Mode 0 = enabled 1 = disabled	RW	1
7	PWMTRIPINT_EN	Interrupt enable for trip 1 = enabled 0 = disabled	RW	0
8	PWMSYNCINT_EN	Interrupt enable for sync 1 = enabled 0 = disabled	RW	0
9	PWMTRIP_DSBL	Disable for trip input 1 = disabled 0 = enabled	RW	0
15:10	Reserved			0

PWM Status (PWM_STAT) Register

The PWM_STAT register provides status information regarding PWM operation. Bit diagrams and descriptions are provided in [Figure 14-11](#) and [Table 14-5](#).

PWM Status Register (PWM_STAT)

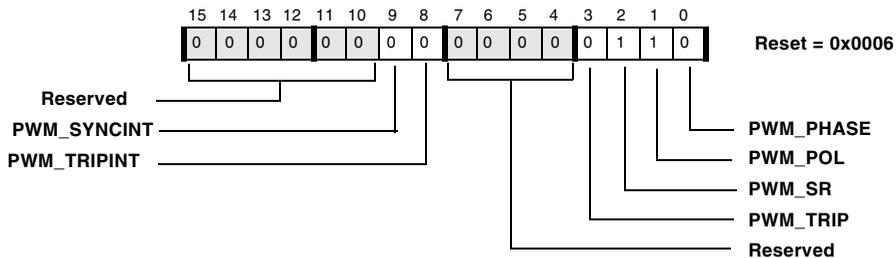


Figure 14-11. PWM Status Register

Table 14-5. PWM_STAT Register

Bit	Name	Function	Type	Default
0	PWM_PHASE	PWM phase 0 = first half 1 = second half	RO	0
1	PWM_POL	PWM polarity 1 = active high 0 = active low	RO	1
2	PWM_SR	PWM SR mode 0 = active 1 = inactive	RO	1
3	PWM_TRIP	PWM trip	RO	0
7:4	Reserved			0
8	PWM_TRIPINT	PWM trip interrupt (via hardware pin or software)	R/W1C	0

PWM Registers

Table 14-5. PWM_STAT Register (Continued)

Bit	Name	Function	Type	Default
9	PWM_SYNCINT	PWM sync interrupt	R/W1C	0
15:10	Reserved			0

PWM Period (PWM_TM) Register

The PWM_TM register controls the switching frequency of the generated PWM patterns. Bit diagrams and descriptions are provided in [Figure 14-12](#) and [Table 14-6](#).

PWM Period Register (PWM_TM)

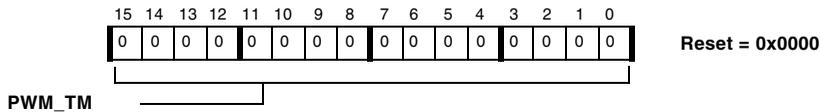


Figure 14-12. PWM Period Register

Table 14-6. PWM_TM Register

Bit	Name	Function	Type	Default
15:0	PWM_TM	PWM period (unsigned)	RW	0

PWM Dead Time (PWM_DT) Register

The `PWM_DT` register controls the dead time interval of the generated PWM patterns. Bit diagrams and descriptions are provided in [Figure 14-13](#) and [Table 14-7](#).

PWM Dead Time Register (PWM_DT)

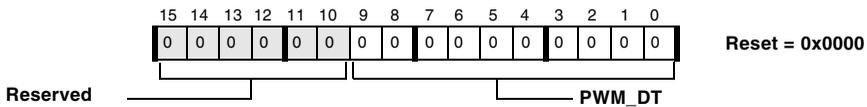


Figure 14-13. PWM Dead Time Register

Table 14-7. PWM_DT Register

Bit	Name	Function	Type	Default
9:0	PWM_DT	PWM dead time (unsigned)	RW	0
15:10	Reserved			0

PWM Chopping Control (PWM_GATE) Register

The PWM controller permits the mixing of the output PWM signals with a high-frequency chopping signal. The features of gate-drive-chopping mode are controlled by the `PWM_GATE` register. Bit diagrams and descriptions are provided in [Figure 14-14](#) and [Table 14-8](#).

PWM Registers

PWM Chopping Control Register (PWM_GATE)

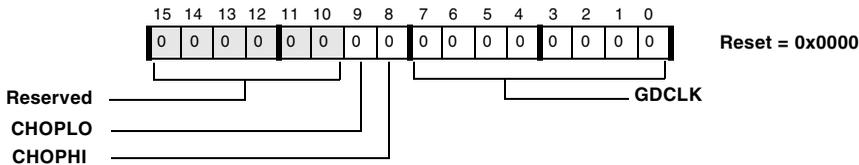


Figure 14-14. PWM Chopping Control Register

Table 14-8. PWM_GATE Register

Bit	Name	Function	Type	Default
7:0	GDCLK	PWM gate chopping period (unsigned)	RW	0
8	CHOPHI	Gate chopping enable high side	RW	0
9	CHOPLO	Gate chopping enable low side	RW	0
15:10	Reserved			0

PWM Channel A, B, C Duty Control (PWM_CHA, PWM_CHB, PWM_CHC) Registers

The three duty-cycle control registers (PWM_CHA , PWM_CHB, and PWM_CHC) directly control the duty cycles of the three pairs of PWM signals. Bit diagrams and descriptions for each are provided in [Figure 14-15](#) through [Figure 14-17](#), and [Table 14-9](#) through [Table 14-11](#).

PWM Channel A Duty Control Register (PWM_CHA)

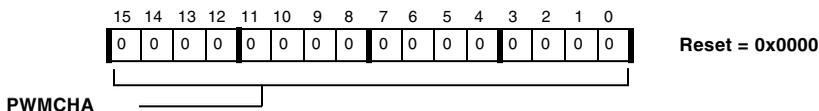


Figure 14-15. PWM Channel A Duty Control Register

Table 14-9. PWM_CHA Register

Bit	Name	Function	Type	Default
15:0	PWMCHA	Channel A duty (two's complement)	RW	0

PWM Channel B Duty Control Register (PWM_CHB)

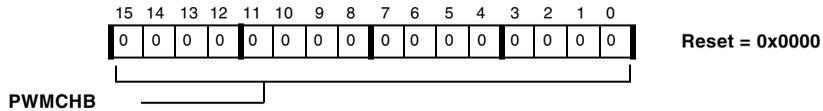


Figure 14-16. PWM Channel B Duty Control Register

Table 14-10. PWM_CHB Register

Bit	Name	Function	Type	Default
15:0	PWMCHB	Channel B duty (two's complement)	RW	0

PWM Channel C Duty Control Register (PWM_CHC)

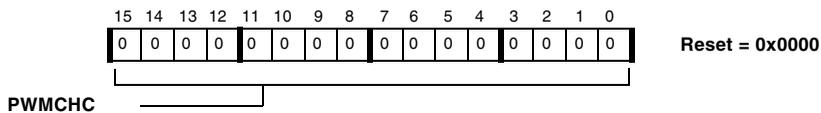


Figure 14-17. PWM Channel C Duty Control Register

Table 14-11. PWM_CHC Register

Bit	Name	Function	Type	Default
15:0	PWMCHC	Channel C duty (two's complement)	RW	0

PWM Crossover and Output Enable (PWM_SEG) Register

The PWM_SEG register controls output enabling of the high-side and low-side PWM outputs, and it also permits configuration of crossover mode for each output pair. Bit diagrams and descriptions are provided in [Figure 14-18](#) and [Table 14-12](#).

PWM Crossover and Output Enable Register (PWM_SEG)

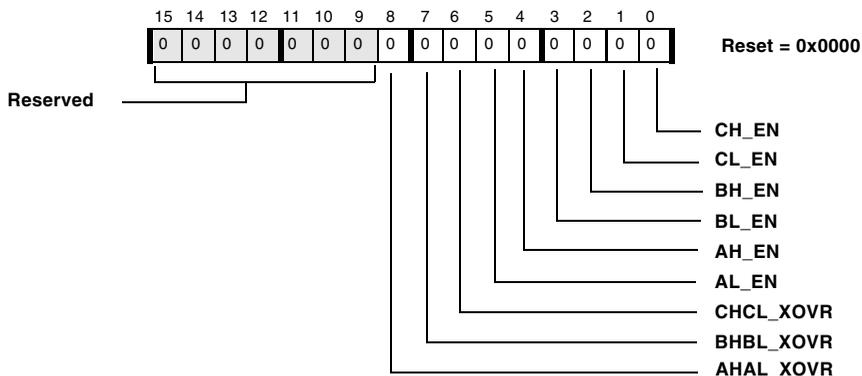


Figure 14-18. PWM Crossover and Output Enable Register

Table 14-12. PWM_SEG Register

Bit	Name	Function	Type	Default
0	CH_EN	CH output enable 1 = disabled 0 = enabled	RW	0
1	CL_EN	CL output enable 1 = disabled 0 = enabled	RW	0
2	BH_EN	BH output enable 1 = disabled 0 = enabled	RW	0

Table 14-12. PWM_SEG Register (Continued)

Bit	Name	Function	Type	Default
3	BL_EN	BL output enable 1 = disabled 0 = enabled	RW	0
4	AH_EN	AH output enable 1 = disabled 0 = enabled	RW	0
5	AL_EN	AL output enable 1 = disabled 0 = enabled	RW	0
6	CHCL_XOVR	Channel C output crossover 1 = XOVR 0 = not XOVR	RW	0
7	BHBL_XOVR	Channel B output crossover 1 = XOVR 0 = not XOVR	RW	0
8	AHAL_XOVR	Channel A output crossover 1 = XOVR 0 = not XOVR	RW	0
15:9	Reserved			0

PWM Registers

PWM Sync Pulse Width Control (PWM_SYNCWT) Register

The PWM_SYNCWT register allows programming of the PWM_SYNC pulse width. Bit diagrams and descriptions are provided in [Figure 14-19](#) and [Table 14-13](#).

PWM Sync Pulse Width Control Register (PWM_SYNCWT)

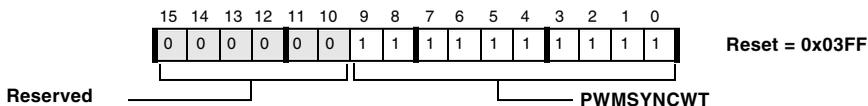


Figure 14-19. PWM Sync Pulse Width Control Register

Table 14-13. PWM_SYNCWT Register

Bit	Name	Function	Type	Default
9:0	PWMSYNCWT	PWM sync pulse width (unsigned)	RW	0x03FF
15:10	Reserved			0

PWM Channel AL, BL, CL Duty Control (PWM_CHAL, PWM_CHBL, PWM_CHCL) Registers

These registers are used to program duty cycle for a low-side channel in SR (switched reluctance) mode only. Bit diagrams and descriptions for each register are provided in [Figure 14-20](#) through [Figure 14-22](#), and [Table 14-14](#) through [Table 14-16](#).

PWM Channel AL Duty Control Register (PWM_CHAL)

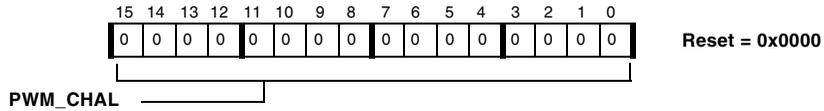


Figure 14-20. PWM Channel AL Duty Control Register

Table 14-14. PWM_CHAL Register

Bit	Name	Function	Type	Default
15:0	PWM_CHAL	Channel A duty (two's complement)	RW	0

PWM Channel BL Duty Control Register (PWM_CHBL)

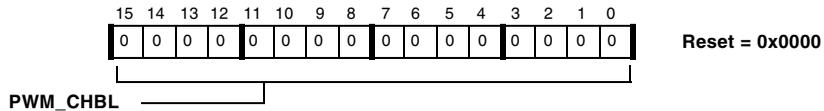


Figure 14-21. PWM Channel BL Duty Control Register

Table 14-15. PWM_CHBL Register

Bit	Name	Function	Type	Default
15:0	PWM_CHBL	Channel B duty (two's complement)	RW	0

PWM Channel CL Duty Control Register (PWM_CHCL)

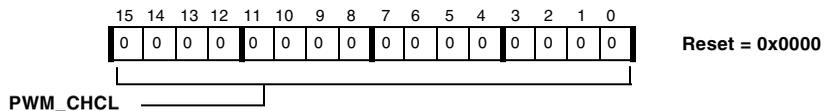


Figure 14-22. PWM Channel CL Duty Control Register

PWM Registers

Table 14-16. PWM_CHCL Register

Bit	Name	Function	Type	Default
15:0	PWM_CHCL	Channel C duty (two's complement)	RW	0

PWM Low Side Invert (PWM_LSI) Register

The PWM_LSI register is used for specifying switched reluctance (SR) chop modes. Bit diagrams and descriptions are provided in [Figure 14-23](#) and [Table 14-17](#).

PWM Low Side Invert Register (PWM_LSI)

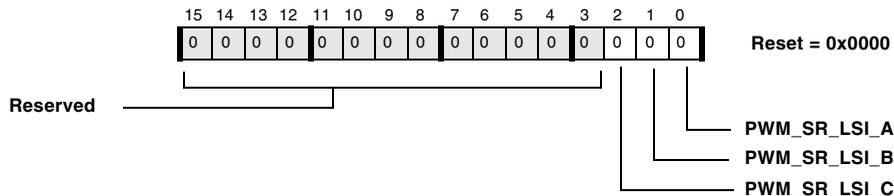


Figure 14-23. PWM Low Side Invert Register

Table 14-17. PWM_LSI Register

Bit	Name	Function	Type	Default
0	PWM_SR_LSI_A	PWM SR mode low side invert channel A	RW	0
1	PWM_SR_LSI_B	PWM SR mode low side invert channel B	RW	0
2	PWM_SR_LSI_C	PWM SR mode low side invert channel C	RW	0
15:3	Reserved			0

PWM Simulation Status (PWM_STAT2) Register

The PWM_STAT2 register provides a way to observe the status of the PWM high-side and low-side output channels via software. This can be useful for debug operation. Bit diagrams and descriptions are provided in [Figure 14-24](#) and [Table 14-18](#).

PWM Simulation Status Register (PWM_STAT2)

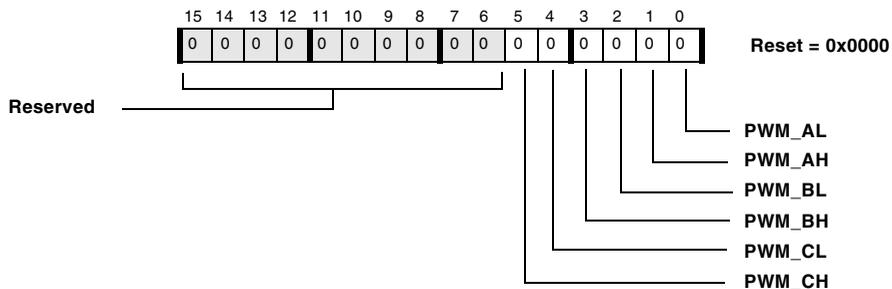


Figure 14-24. PWM Simulation Status Register

Table 14-18. PWM_STAT2 Register

Bit	Name	Function	Type	Default
0	PWM_AL	PWM_AL output signal for S/W observation	RO	0
1	PWM_AH	PWM_AH output signal for S/W observation	RO	0
2	PWM_BL	PWM_BL output signal for S/W observation	RO	0
3	PWM_BH	PWM_BH output signal for S/W observation	RO	0
4	PWM_CL	PWM_CL output signal for S/W observation	RO	0
5	PWM_CH	PWM_CH output signal for S/W observation	RO	0
15:6	Reserved			0

Unique Information for the ADSP-BF51x Processor

None

15 UART PORT CONTROLLERS

This chapter describes the universal asynchronous receiver/transmitter (UART) module. Following an overview and a list of key features is a description of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of UARTs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For UART DMA channel assignments, refer to [Table 6-7 on page 6-104](#) in [Chapter 6, “Direct Memory Access”](#).

For UART interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the UARTs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each UART, refer to [Appendix A, “System MMR Assignments”](#).

UART behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 15-42](#).

Overview

The UART module is a full-duplex peripheral compatible with PC-style industry-standard UARTs, sometimes called serial controller interfaces (SCI). UARTs convert data between serial and parallel formats. The serial communication follows an asynchronous protocol that supports various word length, stop bits, bit rate, and parity generation options.

Features

Each UART includes these features:

- 5 – 8 data bits
- 1 or 2 stop bits (1½ in 5-bit mode)
- Even, odd, and sticky parity bit options
- 3 interrupt outputs for reception, transmission, and status
- Independent DMA operation for receive and transmit
- SIR IrDA operation mode
- Internal loop back

The UART is logically compliant to EIA-232E, EIA-422, EIA-485 and LIN standards, but usually requires an external transceiver device to meet electrical requirements. In IrDA® (Infrared Data Association) mode, the UART meets the half-duplex IrDA SIR (9.6/115.2 Kbps rate) protocol.

Interface Overview

Figure 15-1 shows a simplified block diagram of a UART module and how it interconnects to the Blackfin architecture and to the outside world.

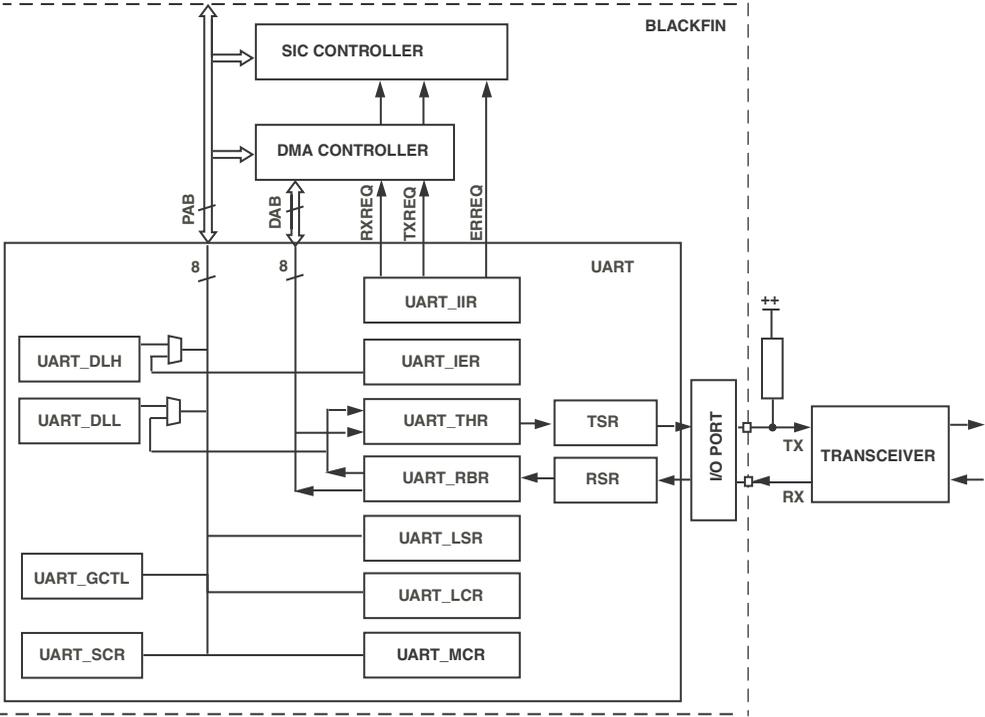


Figure 15-1. UART Block Diagram

External Interface

Each UART features an RX and a TX pin. These two pins usually connect to an external transceiver device that meets the electrical requirements of full duplex (for example, EIA-232, EIA-422, 4-wire EIA-485) or half duplex (for example, 2-wire EIA-485, LIN) standards.

Description of Operation

The RX and TX pins do not need to be used together. If only receive or transmit functionality of a UART module is needed, the unused pin may be used for an alternate function, depending on the port multiplexing scheme of a specific processor. For more details on functionality multiplexed with the UART pins, see [Chapter 9, “General-Purpose Ports”](#).

 Modem status and control functionality is not supported by the UART modules, but may be implemented using GPIO pins.

Internal Interface

UARTs are DMA-capable peripherals with support for separate TX and RX DMA master channels. They 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. Each UART has its own separate transmit and receive DMA channels. For more information on DMA, see the *Direct Memory Access* chapter.

All UART registers are eight bits wide. They connect to the peripheral bus. However, some registers share their address as controlled by the `DLAB` bit in the `UART_LCR` register. The `UART_RBR` and `UART_THR` registers also connect to the DAB bus

A hardware-assisted autobaud detection mechanism is accomplished by coupling a specific GP Timer with a specific UART. For information on GP Timer - UART pairings for autobaud detection, see [General-Purpose Ports](#) chapter.

Description of Operation

The following sections describe the operation of the UART.

UART Transfer Protocol

UART communication follows an asynchronous serial protocol, consisting of individual data words. A word has 5 to 8 data bits.

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 15-2 shows a typical physical bitstream measured on one of the TX pins.

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 16× 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 16× clock is used to determine an IrDA pulse sample window, from which the RZI-modulated NRZ code is recovered.

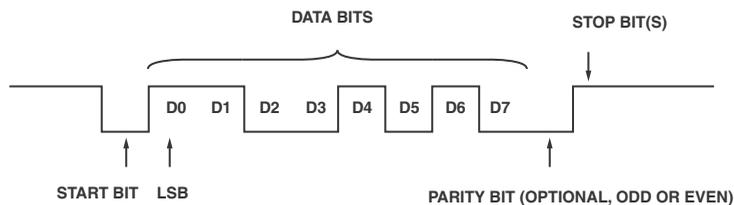


Figure 15-2. Bitstream on a TX Pin Transmitting an “S” Character (0x53)

Description of Operation

IrDA support is enabled by setting the `IREN` bit in the `UART_GCTL` register. The IrDA application requires external transceivers.

UART Transmit Operation

Receive and transmit paths operate independently except that the bit rate and the frame format are identical for both transfer directions.

Transmission is initiated by writes to the `UART_THR` register. If no former operation is pending, the data is immediately passed from the `UART_THR` register to the internal `TSR` register where it is shifted out at a bit rate equal to $SCLK/(16 \times \text{Divisor})$ (see “[Bit Rate Generation](#)” on page 15-12 for information about the divisor) with start, stop, and parity bits appended as defined the `UART_LCR` register. The least significant bit (LSB) is always transmitted first. This is bit 0 of the value written to `UART_THR`.

Writes to the `UART_THR` register clear the `THRE` flag. Transfers of data from `UART_THR` to the transmit shift registers (`TSR`) set this status flag in `UART_LSR` again.

When enabled by the `ETBEI` bit in the `UART_IER` register, a 0 to 1 transition of the `THRE` flag requests an interrupt on the dedicated `TXREQ` output. This signal is routed through the DMA controller. If the associated DMA channel is enabled, the `TXREQ` signal functions as a DMA request, otherwise the DMA controller simply forwards it to the system interrupt controller (SIC).

The `UART_THR` register and the internal `TSR` register can be seen as a two-stage transmit buffer. When data is pending in either one of these registers, the `TEMT` flag is low. As soon as the data has left the `TSR` register, the `TEMT` bit goes high again and indicates that all pending transmit operation has finished. At that time it is safe to disable the `UCEN` bit or to three-state off-chip line drivers.

UART Receive Operation

The receive operation uses the same data format as the transmit configuration, except that one valid stop bit is always sufficient. That is, the STB bit has no impact to the receiver.

After detection of the start bit, the received word is shifted into the internal shift register (RSR) at a bit rate of $SCLK/(16 \times \text{Divisor})$. Once the appropriate number of bits (including one stop bit) is received, the content of the RSR register is transferred to the UART_RBR registers, shown in [Figure 15-11 on page 15-26](#). Finally, the data ready (DR) bit and the status flags are updated in the UART_LSR register, to signal data reception, parity, and also error conditions, if required.

The RSR and the UART_RBR registers can be seen as almost a two-stage receive buffer. If the stop bit of a second byte is received before software reads the first byte from the UART_RBR register, an overrun error is reported and the first byte is overwritten.

If enabled by the ERBFI bit in the UART_IER register, a 0 to 1 transition of the DR flag requests an interrupt on the dedicated RXREQ output. This signal is routed through the DMA controller. If the associated DMA channel is enabled, the RXREQ signal functions as a DMA request, otherwise the DMA controller simply forwards it to the system interrupt controller.

If errors are detected during reception, an interrupt can be requested to a separate error interrupt output. This error request goes directly to the system interrupt controller. However, it is hard-wired with the error requests of other modules. The error handler routine may need to interrogate multiple modules as to whether they requested the event. Error requests must be enabled by the ELSI bit in the UART_IER register. The following error situations are detected. Every error has an indicating bit in the UART_LSR register.

- Overrun error (OE bit)
- Parity error (PE bit)

Description of Operation

- Framing error/Invalid stop bit (FE bit)
- Break indicator (BI bit)

Reception is started when a falling edge is detected on the RX input pin. The receiver attempts to see a start bit. For better immunity against noise and hazards on the line, the receiver oversamples every bit 16 times and does a majority decision based on the middle three samples. The data is shifted immediately into the internal `RSR` register. After the 9th sample of the first stop bit is processed, the received data is copied to the `UART_RBR` register and the receiver recovers itself for further data.

The sampling clock, equal to 16 times the bit rate, samples the data bits close to their midpoint. Because the receiver clock is usually asynchronous to the transmitter's data rate, the sampling point may drift relative to the center of the data bits. The sampling point is synchronized again 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.

IrDA Transmit Operation

To generate the IrDA pulse transmitted by the UART, the normal NRZ output of the transmitter is first inverted if the `TPOLC` bit is cleared, 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 clock periods out of 16 clock periods in the cycle. The pulse is centered around the middle of the bit time, as shown in [Figure 15-3](#). The final IrDA pulse is fed to the off-chip infrared driver.

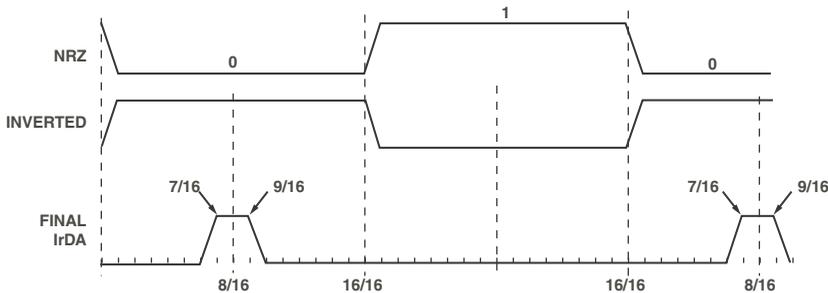


Figure 15-3. IrDA Transmit Pulse

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 15-1 on page 15-13](#), the error terms associated with the bit rate generator are very small and well within the tolerance of most infrared transceiver specifications.

IrDA Receive Operation

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 that because the system clock can

Description of Operation

change across systems, the longest glitch tolerated is inversely proportional to the system clock frequency.

The receive sampling window is determined by a counter that is clocked at the $16\times$ bit-time sample clock. The sampling window is re-synchronized 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 15-4 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.

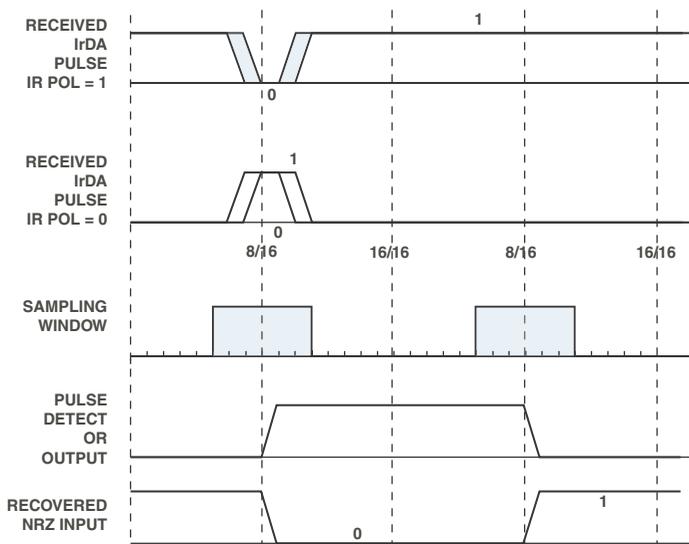


Figure 15-4. IrDA Receiver Pulse Detection

Interrupt Processing

Each UART module has three interrupt outputs. One is dedicated for transmission, one for reception, and the third is used to report line status. As shown in [Figure 15-1 on page 15-3](#), the transmit and receive requests are routed through the DMA controller. The status request goes directly to the system interrupt controller after being ORed with interrupt signals from other modules.

If the associated DMA channel is enabled, the request functions as a DMA request. If the DMA channel is disabled, it simply forwards the request to the system interrupt controller. Note that a DMA channel must be associated with the UART module to enable TX and RX interrupts. Otherwise, the transmit and receive requests cannot be forwarded. Refer to the description of the peripheral map registers in the *Direct Memory Access* chapter.

Transmit interrupts are enabled by the `ETBEI` bit in the `UART_IER` register. If set, the transmit request is asserted when the `THRE` bit in the `UART_LSR` register transitions from 0 to 1, indicating that the TX buffer is ready for new data.

Note that the `THRE` bit resets to 1. When the `ETBEI` bit is set in the `UART_IER` register, the UART module immediately issues an interrupt or DMA request. In this way, no special handling of the first character is required when transmission of a string is initiated. Simply 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 can be cleared if the string transmission has completed. For more information, see [“DMA Mode” on page 15-17](#).

The `THRE` bit is cleared by hardware when new data is written to the `UART_THR` register. These writes also clear the TX interrupt request. However, they also initiate further transmission. If software doesn't want to continue transmission, the TX request can alternatively be cleared by either clearing the `ETBEI` bit or by reading the `UART_IIR` register.

Description of Operation

Receive interrupts are enabled by the `ERBFI` bit in the `UART_IER` register. If set, the receive request is asserted when the `DR` bit in the `UART_LSR` register transitions from 0 to 1, indicating that new data is available in the `UART_RBR` register. When software reads the `UART_RBR`, hardware clears the `DR` bit again. Reading `UART_RBR` also clears the RX interrupt request.

Status interrupts are enabled by the `ELSI` bit in the `UART_IER` register. If set, the status interrupt request is asserted when any error bit in the `UART_LSR` register transitions from 0 to 1. Refer to [“UART Line Status \(UART_LSR\) Register” on page 15-24](#) for details. Reading the `UART_LSR` register clears the error bits destructively. These reads also clear the status interrupt request.

For legacy reasons, the `UART_IIR` registers still reflect the UART interrupt status. Legacy operation may require bundling all UART interrupt sources to a single interrupt channel and 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 controller.

 If either the line status interrupt or the receive data interrupt has been assigned a lower interrupt priority by the system interrupt controller, a deadlock condition can occur. To avoid this, always assign the lowest priority of the enabled UART interrupts to the `UART_THR` empty event.

Bit Rate Generation

The UART clock is enabled by the `UCEN` bit in the `UART_GCTL` register.

The bit rate is characterized by the system clock (`SCLK`) and the 16-bit divisor. The divisor is split into the `UART_DLL` and the `UART_DLH` registers. These registers form a 16-bit divisor. The bit clock is divided by 16 so that:

$$\text{bit rate} = \text{SCLK} / (16 \times \text{divisor})$$

$$\text{divisor} = 65536 \text{ when } \text{UART_DLL} = \text{UART_DLH} = 0$$

Table 15-1 provides example divide factors required to support most standard baud rates.

Table 15-1. UART Bit Rate Examples With 100 MHz SCLK

Bit Rate	DL	Actual	% Error
2400	2604	2400.15	0.006
4800	1302	4800.31	0.007
9600	651	9600.61	0.006
19200	326	19171.78	0.147
38400	163	38343.56	0.147
57600	109	57339.45	0.452
115200	54	115740.74	0.469
921600	7	892857.14	3.119
6250000	1	6250000	–

 Careful selection of SCLK frequencies, that is, even multiples of desired bit rates, can result in lower error percentages.

Note that the UART module is clocked 16 times faster than the bit clock. This is required to oversample bits on reception and to generate RZI code in IrDA mode.

Autobaud Detection

At the chip level, the UART RX pin is routed to the alternate capture input (TACIx) of a general purpose timer. When working in WDTM_CAP mode this timer can be used to automatically detect the bit rate applied to the RX pin by an external device. For more information, see [Chapter 9, “General-Purpose Ports”](#).

The capture capabilities of the timers are often used to supervise the bit rate at runtime. If the Blackfin UART talks to a device supplied by a weak

Description of Operation

clock oscillator that drifts over time, the Blackfin can re-adjust its UART bit rate dynamically.

Often, autobaud detection is used for initial bit rate negotiations. In this case, the Blackfin processor is most likely a slave device waiting for the host to send a predefined autobaud character (see below). This is the scenario used for UART booting. In this scenario, the UART clock enable bit `UCEN` should not be enabled while autobaud detection is performed. This prevents the UART from starting reception with incorrect bit rate matching. Alternatively, the UART can be disconnected from the `RX` pin by setting the `LOOP_ENA` bit.

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 `SCLK`—the pulse widths can be used to calculate the baud rate divisor for the UART.

$$\text{divisor} = \text{TIMER_WIDTH}/(16 \times \text{number of captured UART bits})$$

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. Traditionally, a NULL character (ASCII 0x00) was used in autobaud detection, as shown in [Figure 15-5](#).



Figure 15-5. Autobaud Detection Character 0x00

Because the example frame in [Figure 15-5](#) encloses 8 data bits and 1 start bit, apply the formula:

$$\text{divisor} = \text{TIMER_WIDTH}/(16 \times 9)$$

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 “@” (0x40) as the autobaud detection character and measure the period between two subsequent falling edges. As shown in [Figure 15-6](#), measure the period between the falling edge of the start bit and the falling edge after bit 6. Since this period encloses eight bits, apply the formula:

$$\text{divisor} = \text{TIMER_PERIOD} >> 7$$

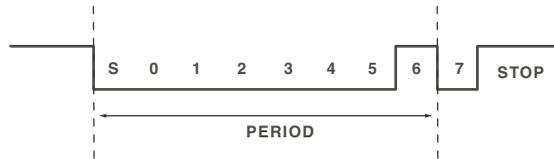


Figure 15-6. Autobaud Detection Character 0x40

An example is provided in [Listing 15-2 on page 15-33](#).

Programming Model

The following sections describe a programming model for the UART.

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 time.

Programming Model

To prevent any loss of data and misalignments of the serial datastream, the `UART_LSR` register 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 `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. Be careful if transmit and receive are served by different software threads, because read operations on the `UART_LSR` and `UART_IIR` registers are destructive. Polling the `SIC_ISR` register without enabling the interrupts by `SIC_MASK` is an alternate method of operation to consider. 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. Separate interrupt lines are provided for UART TX, UART RX, and UART error/status. The independent interrupts can be enabled individually by the `UART_IER` register.

The ISRs can evaluate the status bit field within the `UART_IIR` register 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's interrupt controller. The ISRs must clear the interrupt latches explicitly. See [Figure 15-13 on page 15-29](#).

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.

DMA channels provide a 4-deep FIFO, resulting in total buffer capabilities of 6 words at both the transmit and receive sides. In DMA mode, the latency is determined by the bus activity and arbitration mechanism and not by the processor loading and interrupt priorities.

DMA interrupt routines must explicitly write “1” to the corresponding `DMA_IRQ_STATUS` registers 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/or `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. The UART’s error interrupt goes directly to the system interrupt handling unit, bypassing the DMA unit completely.

For transmit DMA, it is recommended that the `SYNC` bit in the `DMA_CONFIG` register be set. With this bit set, the interrupt generation is delayed until the entire DMA FIFO has been drained to the UART module. The UART TX DMA interrupt service routine is allowed to start another DMA sequence or to clear the `ETBEI` control bit only when the `SYNC` bit is set.

If another DMA is started while data is still pending in the UART transmitter, there is no need to pulse the `ETBEI` bit to initiate the second DMA. If, however, the recent byte has already been loaded into the `TSR` registers

Programming Model

(that is, the `THRE` bit is set), then the `ETBEI` bit must be cleared and set again to let the second DMA start.

In DMA transmit mode, the `ETBEI` bit enables the peripheral request to the DMA FIFO. The strobe on the memory side is still enabled by the `DMAEN` bit. If the DMA count is less than the DMA FIFO depth, which is 4, then the DMA interrupt might be requested before the `ETBEI` bit is set. If this is not wanted, set the `SYNC` bit in the `DMA_CONFIG` register.

 Regardless of the `SYNC` setting, the DMA stream has not left the UART transmitter completely at the time the interrupt is generated. If the UART clock was disabled without additional polling of the `TEMT` bit, transmission may abort in the middle of the stream—causing data loss.

The UART's DMA supports 8-bit and 16-bit operation, but not 32-bit operation. Sign extension is not supported.

Mixing Modes

Especially on the transmit side, switching from DMA mode to non-DMA operation on the fly requires some thought. By default, the interrupt timing of the DMA is synchronized with the memory side of the DMA FIFOs. The TX DMA completion interrupt is generated after the last byte has been copied from the memory into the DMA FIFO. The TX DMA interrupt service routine is not yet permitted to start other DMA sequences or to switch to non-DMA transmission. The interrupt is requested by the time the `DMA_DONE` bit is set. The `DMA_RUN` bit, however, remains set until the data has completely left the TX DMA FIFO.

Therefore, when planning to switch from DMA to non-DMA of operation, always set the `SYNC` bit in the `DMA_CONFIG` word of the last descriptor or work unit before handing over control to non-DMA mode. Then, after the interrupt occurs, software can write new data into the `UART_THR` register as soon as the `THRE` bit permits. If the `SYNC` bit cannot be set, software can poll the `DMA_RUN` bit instead.

When switching from non-DMA to DMA operation, take care that the very first DMA request is issued properly. If the DMA is enabled while the UART is still transmitting, no precaution is required. If, however, the DMA is enabled after the `TEMT` bit became high, the `ETBEI` bit should be pulsed to initiate DMA transmission.

UART Registers

The processor provides a set of PC-style industry-standard control and status registers for each UART. These memory-mapped registers (MMRs) are byte-wide registers that are mapped as half words with the most significant byte zero filled. [Table 15-2](#) provides an overview of the UART registers.

Consistent with industry-standard devices, multiple registers are mapped to the same address location. The `UART_DLH` and `UART_DLL` registers share their addresses with the `UART_THR` registers, the `UART_RBR` registers, and the `UART_IER` registers. The `DLAB` bit in the `UART_LCR` register 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` registers buffer the transmit shift register (`TSR`) and the `UART_RBR` registers buffer the receive shift register (`LSR`). The shift registers are not directly accessible by software.

Table 15-2. UART Register Overview

Name	Address Offset	DLAB Bit Setting	Operation	Reset Value	Function
<code>UART_RBR</code>	0x0000	0	R	0x00	Receive buffer register
<code>UART_THR</code>	0x0000	0	W	0x00	Transmit holding register
<code>UART_DLL</code>	0x0000	1	R/W	0x01	Divisor latch low byte

UART Registers

Table 15-2. UART Register Overview (Continued)

Name	Address Offset	DLAB Bit Setting	Operation	Reset Value	Function
UART_IER	0x0004	0	R/W	0x00	Interrupt enable register
UART_DLH	0x0004	1	R/W	0x00	Divisor latch high byte
UART_IIR	0x0008	X	R Read operations are destructive	0x01	Interrupt identification register
UART_LCR	0x000C	X	R/W	0x00	Line control register
UART_MCR	0x0010	X	R/W	0x00	Modem control register
UART_LSR	0x0014	X	R Read operations are destructive	0x60	Line status register
UART_SCR	0x001C	X	R/W	0x00	Scratch register
UART_GCTL	0x0024	X	R/W	0x00	Global control register

UART Line Control (UART_LCR) Register

The `UART_LCR` register, shown in [Figure 15-7](#), controls the format of received and transmitted character frames.

UART Line Control Register (UART_LCR)

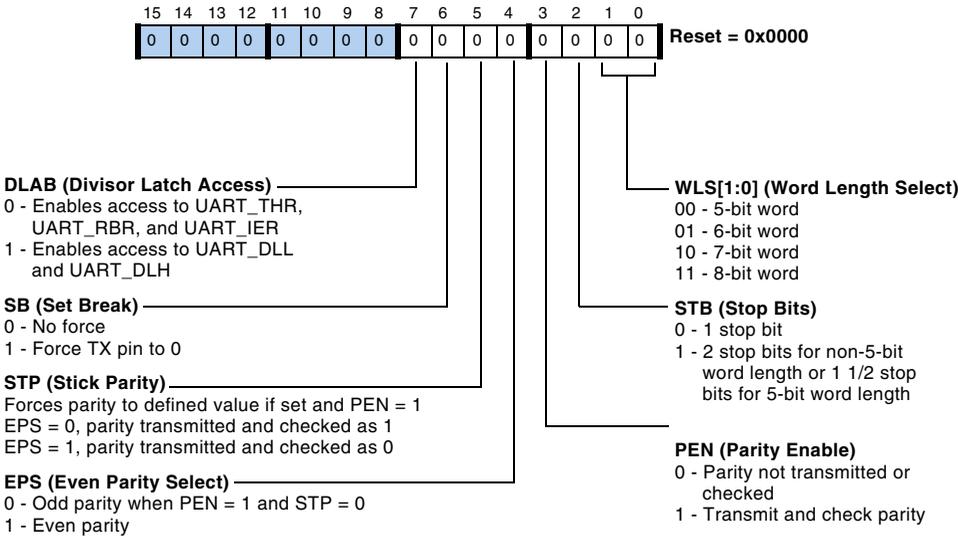


Figure 15-7. UART Line Control Register

The 2-bit `WLS` field determines whether the transmitted and received UART word consists of 5, 6, 7 or 8 data bits.

The `STB` bit controls how many stop bits are appended to transmitted data. When `STB = 0`, one stop bit is transmitted. If `WLS` is non zero, `STB = 1` instructs the transmitter to add one additional stop bit, two stop bits in total. If `WLS = 0` and 5-bit operation is chosen, `STB = 1` forces the transmitter to append one additional half bit, 1½ stop bits in total. Note that this bit does not impact data reception—the receiver is always satisfied with one stop bit.

UART Registers

The `PEN` bit inserts one additional bit between the most significant data bit and the first stop bit. The polarity of this so-called parity bit depends on data and the `STP` and `EPS` control bits. Both transmitter and receiver calculate the parity value. The receiver compares the received parity bit with the expected value and issues a parity error if they don't match. If `PEN` is cleared, the `STP` and the `EPS` bits are ignored.

The `STP` bit controls whether the parity is generated by hardware based on the data bits or whether it is set to a fixed value. If `STP = 0` the hardware calculates the parity bit value based on the data bits. Then, the `EPS` bit determines whether odd or even parity mode is chosen. If `EPS = 0`, odd parity is used. That means that the total count of `logical-1` data bits including the parity bit must be an odd value. Even parity is chosen by `STP = 0` and `EPS = 1`. Then, the count of `logical-1` bits must be a even value. If the `STP` bit is set, then hardware parity calculation is disabled. In this case, the sent and received parity equals the inverted `EPS` bit. The example in [Table 15-3](#) summarizes polarity behavior assuming 8-bit data words (`WLS = 3`).

Table 15-3. UART Parity

<code>PEN</code>	<code>STP</code>	<code>EPS</code>	Data (hex)	Data (binary, LSB first)	Parity
0	x	x	x	x	None
1	0	0	0x60	0000 0110	1
1	0	0	0x57	1110 1010	0
1	0	1	0x60	0000 0110	0
1	0	1	0x57	1110 1010	1
1	1	0	x	x	1
1	1	1	x	x	0

If set, the `SB` bit forces the TX pin to low asynchronously, regardless of whether or not data is currently transmitted. It 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.

The `DLAB` bit controls whether the `UART_RBR`, `UART_THR` and `UART_IER` registers are accessible by the peripheral bus (`DLAB = 0`) or the divisor latch registers `UART_DLH` and `UART_DLL` alternatively (`DLAB = 1`).

UART Modem Control (UART_MCR) Register

The `UART_MCR` register controls the UART port, as shown in [Figure 15-8](#). Even if modem functionality is not supported, the `UART_MCR` register is available in order to support the loopback mode.

UART Modem Control Register (UART_MCR)

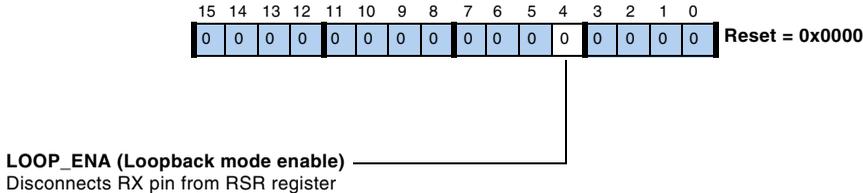


Figure 15-8. UART Modem Control Registers

Loopback mode disconnects the receiver’s input from the RX pin, but redirects it to the transmit output internally.

UART Registers

UART Line Status (UART_LSR) Register

The `UART_LSR` register contains UART status information as shown in Figure 15-9.

UART Line Status Register (UART_LSR)

read only

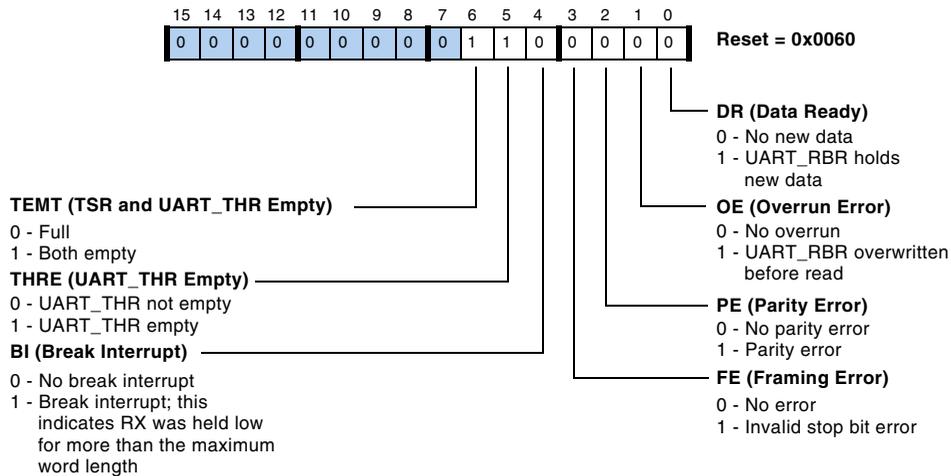


Figure 15-9. UART Line Status Register

The `DR` bit indicates that data is available in the receiver and can be read from the `UART_RBR` register. The bit is set by hardware when the receiver detects the first valid stop bit. It is cleared by hardware when the `UART_RBR` register is read.

The `OE` bit indicates that a start bit condition has been detected, but the internal receive shift register (`RSR`) and the receive buffer (`UART_RBR`) already contain data. New data overwrites the content of the buffers. To avoid overruns read the `UART_RBR` register in time. The `OE` bit cleared when the `UART_LSR` register is read.

The `PE` bit indicates that the received parity bit does not match the expected value. The `PE` bit is set simultaneously with the `DR` bit. The `PE` bit

cleared when the `UART_LSR` register is read. Invalid parity bits can be simulated by setting the `FPE` bit in the `UART_GCTL` register.

The `FE` bit indicates that the first stop bit has been sampled low. It is cleared by hardware when the `UART_RBR` register is read. Invalid stop bits can be simulated by setting the `FFE` bit in the `UART_GCTL` register.

The `BI` bit indicates that the first stop bit has been sampled low and the entire data word, including parity bit, consists of low bits only. It is cleared by hardware when the `UART_RBR` register is read.

 Because of the destructive nature of these read operations, special care should be taken. For more information, see the *Memory* chapter of the *ADSP-BF51x Blackfin Processor Hardware Reference*.

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 passed from `UART_THR` to the internal `TSR` register.

The `TEMT` bit indicates that both the `UART_THR` register and the internal `TSR` register are empty. In this case the program is permitted to write to the `UART_THR` register twice without losing data. The `TEMT` bit can also be used as an indicator that pending UART transmission has been completed. At that time it is safe to disable the `UCEN` bit or to three-state the off-chip line driver.

UART Transmit Holding (UART_THR) Register

The write-only `UART_THR` register, shown in [Figure 15-10](#), is mapped to the same address as the read-only `UART_RBR` and `UART_DLL` registers. 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.

UART Registers

UART Transmit Holding Register (UART_THR)

write only

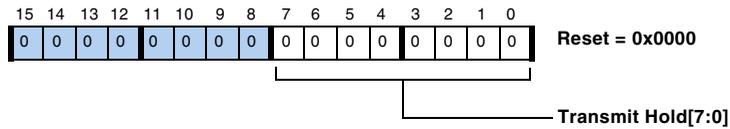


Figure 15-10. UART Transmit Holding Register

UART Receive Buffer (UART_RBR) Register

The read-only UART_RBR register, shown in [Figure 15-11](#), is mapped to the same address as the write-only UART_THR and UART_DLL registers. 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 value in the UART_RBR register.

UART Receive Buffer Register (UART_RBR)

read only

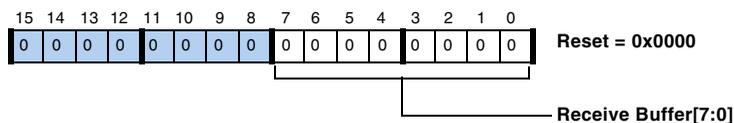


Figure 15-11. UART Receive Buffer Register

UART Interrupt Enable (UART_IER) Register

The UART_IER register, shown in [Figure 15-12](#), is used to enable requests for system handling of empty or full states of UART data registers. 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.

i Each UART features three separate interrupt channels to handle data transmit, data receive, and line status events independently, regardless of whether DMA is enabled or not. On some processors, the status interrupt channels from multiple UARTs may be ORED prior to being connected to the system interrupt controller. See [Chapter 5, “System Interrupts”](#) for more information.

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` registers are mapped to the same address as the `UART_DLH` registers. To access `UART_IER`, the `DLAB` bit in `UART_LCR` must be cleared.

UART Interrupt Enable Register (UART_IER)

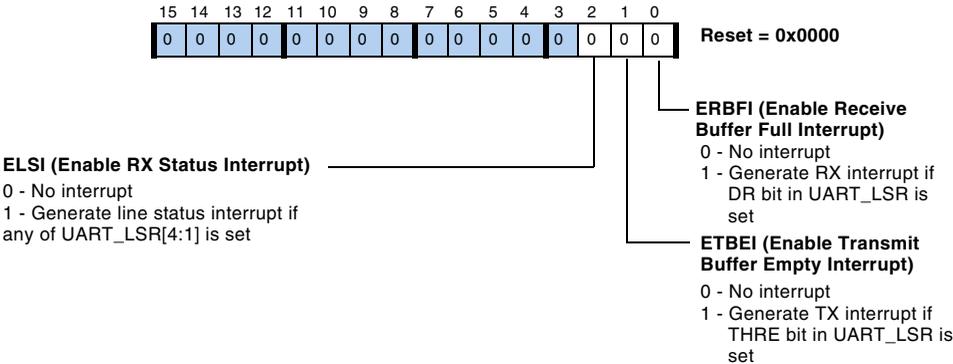


Figure 15-12. UART Interrupt Enable Register

The UART’s DMA is enabled by first setting up the system DMA control registers and then enabling the UART `ERBFI` and/or `ETBEI` interrupts in the `UART_IER` register. This is because the interrupt request lines double as

UART Registers

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_LSR` register:

- Receive overrun error (`OE`)
- Receive parity error (`PE`)
- Receive framing error (`FE`)
- Break interrupt (`BI`)

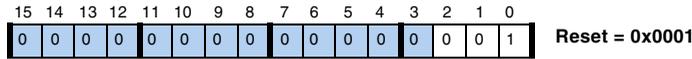
UART Interrupt Identification (`UART_IIR`) Register

The `UART_IIR` register conveys interrupt status within the UART. When cleared, the `NINT` bit 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 signaling, the `UART_IIR` reads `0x06`.

When a UART interrupt is pending, the interrupt service routine needs to clear the interrupt latch explicitly. [Figure 15-13](#) shows how to clear any of the three latches.

UART Interrupt Identification Register (UART_IIR)

read only



NINT (Pending interrupt)
 0 - Interrupt pending
 1 - No interrupt pending

STATUS[1:0]

- 00 - Reserved
- 01 - UART_THR empty. Write UART_THR or read UART_IIR to clear interrupt request.
- 10 - Receive data ready. Read UART_RBR to clear interrupt request.
- 11 - Receive line status. Read UART_LSR to clear interrupt request.

Figure 15-13. UART Interrupt Identification Register

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.



Because of the destructive nature of these read operations, special care should be taken. For more information, see the *Memory* chapter of the *ADSP-BF51x Blackfin Processor Hardware Reference*.

UART Divisor Latch (UART_DLL and UART_DLH) Registers

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 `UART_IER` register. The `DLAB` bit in `UART_LCR` must be set before the `UART_DLL` and `UART_DLH` registers, shown in [Figure 15-14](#), can be accessed.

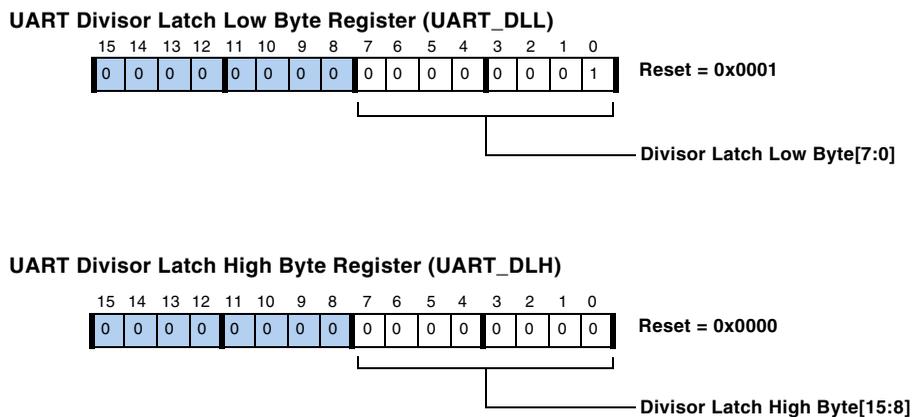


Figure 15-14. UART Divisor Latch Registers

i 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 saves power. The `UART_DLH` and `UART_DLL` registers can be programmed by software before or after setting the `UCEN` bit.

UART Scratch (UART_SCR) Register

The 8-bit UART_SCR register, shown in Figure 15-15, is used for general-purpose data storage and does not control the UART hardware in any way. The contents are reset to 0x00.

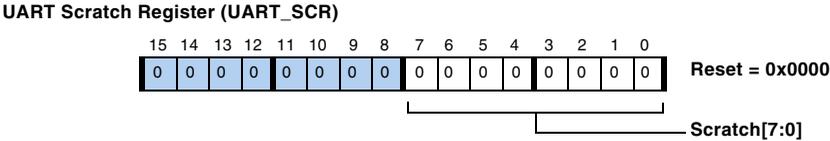


Figure 15-15. UART Scratch Register

UART Global Control (UART_GCTL) Register

The UART_GCTL register, shown in Figure 15-16, contains the enable bit for internal UART clocks and for the IrDA mode of operation of the UART.

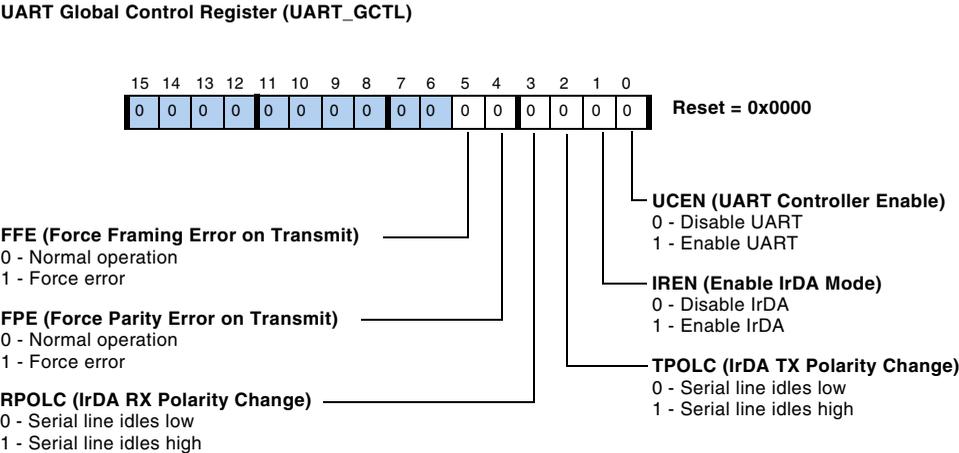


Figure 15-16. UART Global Control Register

Programming Examples

The `UCEN` bit enables the UART clocks. It also resets the state machine and control registers when cleared.

This bit 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.

Programming Examples

The subroutine in [Listing 15-1](#) shows a typical UART initialization sequence.

Listing 15-1. UART Initialization

```
/*
*****
* Configures UART in 8 data bits, no parity, 1 stop bit mode.
* Input parameters: r0 holds divisor latch value to be
*                   written into
*                   DLH:DLL registers.
*                   p0 contains the UART_GCTL register address
* Return values:   none
*****/
uart_init:
    [--sp] = r7;
    r7 = UCEN (z); /* First of all, enable UART clock */
    w[p0+UART0_GCTL-UART0_GCTL] = r7;

    r7 = DLAB (z); /* to set bit rate */
    w[p0+UART0_LCR-UART0_GCTL] = r7; /* set DLAB bit first */
    w[p0+UART0_DLL-UART0_GCTL] = r0; /* write lower byte to DLL */
```

```

r7 = r0 >> 8;
w[p0+UART0_DLH-UART0_GCTL] = r7; /* write upper byte to DLH */

r7 = STB | WLS(8) (z); /* clear DLAB again and config to */
w[p0+UART0_LCR-UART0_GCTL] = r7;
/* 8 bits, no parity, 2 stop bits */

r7 = [sp++];
rts;
uart_init.end:

```

The subroutine in [Listing 15-2](#) performs autobaud detection similarly to UART boot.

Listing 15-2. UART Autobaud Detection Subroutine

```

/*****
* Assuming 8 data bits, this functions expects a '@'
* (ASCII 0x40) character
* on the UART RX pin. A Timer performs the autobaud detection.
* Input parameters: p0 contains the UART_GCTL register address
*                   p1 contains the TIMER_CONFIG register
*                   address
* Return values:    r0 holds timer period value (equals 8 bits)
*****/
uart_autobaud:
    [--sp] = (r7:5,p5:5);
    r5.h = hi(TIMER0_CONFIG); /* for generic timer use calculate
*/
    r5.l = lo(TIMER0_CONFIG); /* specific bits first */
    r7 = p1;
    r7 = r7 - r5;
    r7 >>= 4; /* r7 holds the 'x' of TIMERx_CONFIG now */
    r5 = TIMEN0 (z);
    r5 <<= r7; /* r5 holds TIMENx/TIMDISx now */

```

Programming Examples

```
r6 = TRUN0 | TOVL_ERR0 | TIMILO (z);
r6 <<= r7;
CC = r7 <= 3;
r7 = r6 << 12;
if !CC r6 = r7; /* r6 holds TRUNx | TOVL_ERRx | TIMILx */

p5.h = hi(TIMER_STATUS);
p5.l = lo(TIMER_STATUS);
w[p5 + TIMER_DISABLE - TIMER_STATUS] = r5; /* disable Timer x
*/
[p5 + TIMER_STATUS - TIMER_STATUS] = r6;
/* clear pending latches */
/* period capture, falling edge to falling edge */
r7 = TIN_SEL | IRQ_ENA | PERIOD_CNT | WDTM_CAP (z);
w[p1 + TIMERO_CONFIG - TIMERO_CONFIG] = r7;
w[p5+TIMER_ENABLE-TIMER_STATUS] = r5;

uart_autobaud.wait: /* wait for timer event */
    r7 = w[p5 + TIMER_STATUS - TIMER_STATUS] (z);
    r7 = r7 & r5;
    CC = r7 == 0;
    if CC jump uart_autobaud.wait;
w[p5 + TIMER_DISABLE - TIMER_STATUS] = r5; /* disable Timer x
*/
[p5 + TIMER_STATUS - TIMER_STATUS] = r6;
/* clear pending latches */
/* Save period value to R0 */
r0 = [p1 + TIMERO_PERIOD - TIMERO_CONFIG];

/* delay processing as autobaud character is still ongoing */
r7 = OUT_DIS | IRQ_ENA | PERIOD_CNT | PWM_OUT (z);
w[p1 + TIMERO_CONFIG - TIMERO_CONFIG] = r7;
w[p5 + TIMER_ENABLE - TIMER_STATUS] = r5;
```

```

uart_autobaud.delay:
    r7 = w[p5 + TIMER_STATUS - TIMER_STATUS] (z);
    r7 = r7 & r5;
    CC = r7 == 0;
    if CC jump uart_autobaud.delay;
w[p5 + TIMER_DISABLE - TIMER_STATUS] = r5;
[p5 + TIMER_STATUS - TIMER_STATUS] = r6;
(r7:5,p5:5) = [sp++];
rts;
uart_autobaud.end:

```

The parent routine in [Listing 15-3](#) performs autobaud detection, using as an example a processor whose `TIMER4` is mapped to `UART0` for this purpose. Note also that this example assumes the processor's `UART0` pins are mapped to `PORT G` (`PG7` and `PG8`).

Listing 15-3. UART Autobaud Detection Parent Routine

```

p0.l = lo(PORTG_FER);
    /* function enable on UART0 pins PG7 and PG8 */
p0.h = hi(PORTG_FER);
r0 = PG7 | PG8 (z)
w[p0] = r0;
p0.l = lo(PORTG_MUX);
p0.h = hi(PORTG_MUX);
r0.l = 0x0020;
r0.h = 0x0000;
w[p0] = r0;
p0.l = lo(UART0_GCTL);    /* select UART 0 */
p0.h = hi(UART0_GCTL);
p1.l = lo(TIMER4_CONFIG); /* select TIMER 4 */
p1.h = hi(TIMER4_CONFIG);
call uart_autobaud;
r0 >>= 7;    /* divide PERIOD value by (16 x 8) */

```

Programming Examples

```
    call uart_init;
    ...
```

The subroutine in [Listing 15-4 on page 15-36](#) transmits a character by polling operation.

Listing 15-4. UART Character Transmission

```
/******
 * Transmit a single byte by polling the THRE bit.
 * Input parameters: r0 holds the character to be transmitted
 *                   p0 contains UART_GCTL register address
 * Return values: none
*****/
uart_putc:
    [--sp] = r7;
uart_putc.wait:
    r7 = w[p0+UART0_LSR-UART0_GCTL] (z);
    CC = bittst(r7, bitpos(THRE));
    if !CC jump uart_putc.wait;
    w[p0+UART0_THR-UART0_GCTL] = r0; /* write initiates transfer
*/
    r7 = [sp++];
    rts;
uart_putc.end:
```

Use the routine shown in [Listing 15-5](#) to transmit a C-style string that is terminated by a null character.

Listing 15-5. UART String Transmission

```
/******
 * Transmit a null-terminated string.
 * Input parameters: p1 points to the string
 *                   p0 contains UART_GCTL register address
*****
```

```

*   Return values: none
*****/
uart_puts:
    [--sp] = rets;
    [--sp] = r0;
uart_puts.loop:
    r0 = b[p1++] (z);
    CC = r0 == 0;
    if CC jump uart_puts.exit;
    call uart_putc;
    jump uart_puts.loop;
uart_puts.exit:
    r0 = [sp++];
    rets = [sp++];
    rts;
uart_puts.end:

```

Note that polling the `UART0_LSR` register for transmit purposes may clear the receive error latch bits. It is, therefore, not recommended to poll `UART0_LSR` for transmission this way while data is being received. In that case, write a polling loop that reads `UART_LSR` once and then evaluates *all* status bits of interest, as shown in [Listing 15-6](#).

Listing 15-6. UART Polling Loop

```

uart_loop:
    r7 = w[p0+UART0_LSR-UART0_GCTL] (z);
    CC = bittst(r7, bitpos(DR));
    if !CC jump uart_loop.transmit;
    r6 = w[p0+UART0_RBR-UART0_GCTL] (z);
    r5 = BI | OE | FE | PE (z);
    r5 = r5 & r7;
    CC = r5 == 0;
    if !CC jump uart_loop.error;
    b[p1++] = r6;          /* store byte */

```

Programming Examples

```
uart_loop.transmit:
    CC = bittst(r7, bitpos(THRE));
    if !CC jump uart_loop;
    r5 = b[p2++] (z);      /* load next byte */
    w[p0+UART0_THR-UART0_GCTL] = r5;
    jump uart_loop;
uart_loop.error:
    ...
    jump uart_loop;
```

In non-DMA interrupt operation, the three UART interrupt request lines may or may not be ORed together in the system interrupt controller. If they had three different service routines, they may look as shown in [Listing 15-7](#).

Listing 15-7. UART Non-DMA Interrupt Operation

```
isr_uart_rx:
    [--sp] = astat;
    [--sp] = r7;
    r7 = w[p0+UART0_RBR-UART0_GCTL] (z);
    b[p4++] = r7;
    ssync;
    r7 = [sp++];
    astat = [sp++];
    rti;
isr_uart_rx.end:

isr_uart_tx:
    [--sp] = astat;
    [--sp] = r7;
    r7 = b[p3++] (z);
    CC = r7 == 0;
    if CC jump isr_uart_tx.final;
    w[p0+UART0_THR-UART0_GCTL] = r7;
```

```

    r7 = [sp++];
    astat = [sp++];
    ssync;
    rti;
isr_uart_tx.final:
    r7 = w[p0+UART0_IER-UART0_GCTL] (z);
        /* clear TX interrupt enable */
    bitclr(r7, bitpos(ETBEI)); /* ensure this sequence is not */
    w[p0+UART0_IER-UART0_GCTL] = r7;
        /* interrupted by other IER accesses */

    ssync;
    r7 = [sp++];
    astat = [sp++];
    rti;
isr_uart_tx.end:

isr_uart_error:
    [--sp] = astat;
    [--sp] = r7;
    r7 = w[p0+UART0_LSR-UART0_GCTL] (z);
        /* read clears interrupt request */
        /* do something with the error */

    r7 = [sp++];
    astat = [sp++];
    ssync;
    rti;
isr_uart_error.end:

```

Programming Examples

[Listing 15-8](#) transmits a string by DMA operation, waits until DMA completes and sends an additional string by polling. Note the importance of the SYNC bit.

Listing 15-8. UART Transmission SYNC Bit Use

```
.section data;
.byte sHello[] = 'Hello Blackfin User',13,10,0;
.byte sWorld[] = 'How is life?',13,10,0;

.section program;
    ...
    p1.l = lo(IMASK);
    p1.h = hi(IMASK);
    r0.l = lo(isr_uart_tx);      /* register service routine */
    r0.h = hi(isr_uart_tx); /*Assume UART0 TX defaults to IVG10*/
    r0 = [p1 + IMASK - IMASK]; /* unmask interrupt in CEC */
    bitset(r0, bitpos(EVT_IVG10));
    [p1] = r0;
    p1.l = lo(SIC_IMASK0);
    p1.h = hi(SIC_IMASK0);
/* unmask interrupt in SIC */
/* (assume SIC_IMASK0 for this example)*/
    r0.l = 0x0080;
    r0.h = 0x0000;
    [p1] = r0;
    [--sp] = reti;      /* enable nesting of interrupts */

    p5.l = lo(DMA9_CONFIG);
/* setup DMA in STOP mode */
/* (assume DMA channel 9 for this example)*/
    p5.h = hi(DMA9_CONFIG);
    r7.l = lo(sHello);
    r7.h = hi(sHello);
```

```

[p5+DMA9_START_ADDR-DMA9_CONFIG] = r7;
r7 = length(sHello) (z);
r7+= -1;          /* do not send trailing null character */
w[p5+DMA9_X_COUNT-DMA9_CONFIG] = r7;
r7 = 1;
w[p5+DMA9_X_MODIFY-DMA9_CONFIG] = r7;
r7 = FLOW_STOP | WDSIZE_8 | DI_EN | SYNC | DMAEN (z);
w[p5] = r7;

p0.l = lo(UART0_GCTL); /* select UART 0 */
p0.h = hi(UART0_GCTL);
r0 = ETBEI (z);      /* enable and issue first request */
w[p0+UART0_IER-UART0_GCTL] = r0;

wait4dma: /* just one way to synchronize with the service rou-
tine */
    r0 = w[p5+DMA9_IRQ_STATUS-DMA9_CONFIG] (z);
    CC = bittst(r0,bitpos(DMA_RUN));
    if CC jump wait4dma;
    p1.l=lo(sWorld);
    p1.h=hi(sWorld);
    call uart_puts;

forever: jump forever;

isr_uart_tx:
    [--sp] = astat;
    [--sp] = r7;
    r7 = DMA_DONE (z); /* W1C interrupt request */
    w[p5+DMA9_IRQ_STATUS-DMA9_CONFIG] = r7;
    r7 = 0;          /* pulse ETBEI for general case */
    w[p0+UART0_IER-UART0_GCTL] = r7;
    ssync;
    r7 = [sp++];

```

Unique Information for the ADSP-BF51x Processor

```
    astat = [sp++];  
    rti;  
isr_uart_tx.end:
```

Unique Information for the ADSP-BF51x Processor

None.

16 TWO WIRE INTERFACE CONTROLLER

This chapter describes the two wire interface (TWI) port. Following an overview and a list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of TWIs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For TWI interrupt vector assignments, refer to [Table 5-3 on page 5-19 in Chapter 5, “System Interrupts”](#).

To determine how each of the TWIs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5 through Table 9-4 on page 9-7 in Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each TWI, refer to [Appendix A, “System MMR Assignments”](#).

TWI behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 16-58](#).

Overview

The TWI controller allows a device to interface to an inter IC bus as specified by the *Philips I²C Bus Specification version 2.1* dated January 2000.

The TWI is fully compatible with the widely used I²C bus standard. It was designed with a high level of functionality and is compatible with multi-master, multi-slave bus configurations. To preserve processor bandwidth the TWI controller can be set up with transfer initiated interrupts only to service FIFO buffer data reads and writes. Protocol related interrupts are optional.

The TWI externally moves 8-bit data while maintaining compliance with the I²C bus protocol. The TWI controller includes these features:

- Simultaneous master and slave operation on multiple device systems
- Support for multi-master bus arbitration
- 7-bit addressing
- 100K bits/second and 400K bits/second data rates
- General call address support
- Master clock synchronization and support for clock low extension
- Separate multiple-byte receive and transmit FIFOs
- Low interrupt rate
- Individual override control of data and clock lines in the event of bus lock-up

- Input filter for spike suppression
- Serial camera control bus support as specified in the *OmniVision Serial Camera Control Bus (SCCB) Functional Specification* version 2.1.

Interface Overview

Figure 16-1 provides a block diagram of the TWI controller. The interface is essentially a shift register that serially transmits and receives data bits, one bit at a time at the SCL rate, to and from other TWI devices. The SCL signal synchronizes the shifting and sampling of the data on the serial data pin.

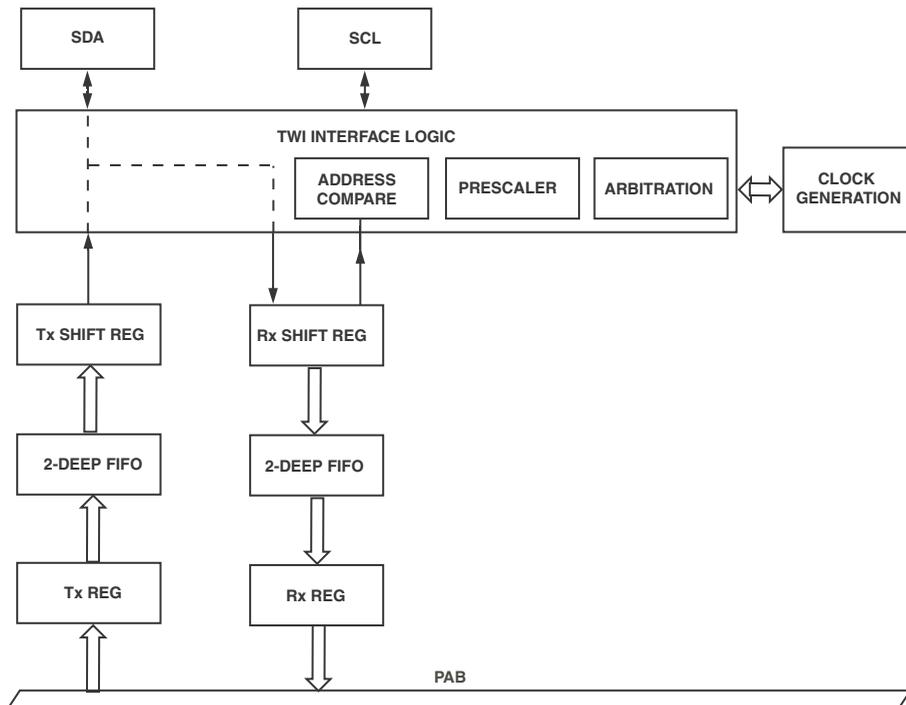


Figure 16-1. TWI Block Diagram

Interface Overview

External Interface

The SDA (serial data) and SCL (serial clock) signals are open drain and as such require pull-up resistors.

Serial Clock Signal (SCL)

In slave mode this signal is an input and an external master is responsible for providing the clock.

In master mode the TWI controller must set this signal to the desired frequency. The TWI controller supports the standard mode of operation (up to 100 KHz) or fast mode (up to 400 KHz).

The TWI control register (TWI_CONTROL) is used to set the PRESCALE value which gives the relationship between the system clock (SCLK) and the TWI controller's internally timed events. The internal time reference is derived from SCLK using a prescaled value.

$$\text{PRESCALE} = f_{\text{SCLK}}/10\text{MHz}$$

The PRESCALE value is the number of system clock (SCLK) periods used in the generation of one internal time reference. The value of PRESCALE must be set to create an internal time reference with a period of 10 MHz. It is represented as a 7-bit binary value.

Serial Data Signal (SDA)

This is a bidirectional signal on which serial data is transmitted or received depending on the direction of the transfer.

TWI Pins

Table 16-1 shows the pins for the TWI. Two bidirectional pins externally interface the TWI controller to the I²C bus. The interface is simple and no other external connections or logic are required.

Table 16-1. TWI Pins

Pin	Description
SDA	In/Out TWI serial data, high impedance reset value.
SCL	In/Out TWI serial clock, high impedance reset value.

Internal Interfaces

The peripheral bus interface supports the transfer of 16-bit wide data and is used by the processor in the support of register and FIFO buffer reads and writes.

The register block contains all control and status bits and reflects what can be written or read as outlined by the programmer's model. Status bits can be updated by their respective functional blocks.

The FIFO buffer is configured as a 1-byte-wide 2-deep transmit FIFO buffer and a 1-byte-wide 2-deep receive FIFO buffer.

The transmit shift register serially shifts its data out externally off chip. The output can be controlled for generation of acknowledgements or it can be manually overwritten.

The receive shift register receives its data serially from off chip. The receive shift register is 1 byte wide and data received can either be transferred to the FIFO buffer or used in an address comparison.

The address compare block supports address comparison in the event the TWI controller module is accessed as a slave.

Description of Operation

The prescaler block must be programmed to generate a 10 MHz time reference relative to the system clock. This time base is used for filtering of data and timing events specified by the electrical data sheet (See the Philips Specification), as well as for SCL clock generation.

The clock generation module is used to generate an external SCL clock when in master mode. It includes the logic necessary for synchronization in a multi-master clock configuration and clock stretching when configured in slave mode.

Description of Operation

The following sections describe the operation of the TWI interface.

TWI Transfer Protocols

The TWI controller follows the transfer protocol of the *Philips I²C Bus Specification version 2.1* dated January 2000. A simple complete transfer is diagrammed in [Figure 16-2](#).

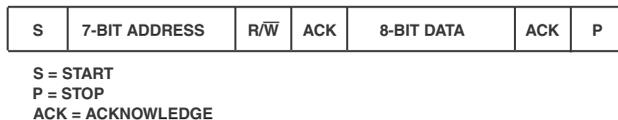


Figure 16-2. Basic Data Transfer

To better understand the mapping of TWI controller register contents to a basic transfer, [Figure 16-3](#) details the same transfer as above noting the corresponding TWI controller bit names. In this illustration, the TWI controller successfully transmits one byte of data. The slave has acknowledged both address and data.

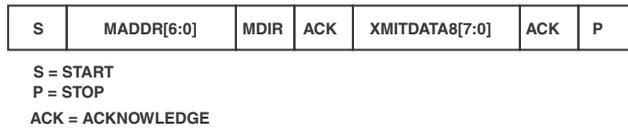


Figure 16-3. Data Transfer With Bit Illustration

Clock Generation and Synchronization

The TWI controller implementation only issues a clock during master mode operation and only at the time a transfer has been initiated. If arbitration for the bus is lost, the serial clock output immediately three-states. If multiple clocks attempt to drive the serial clock line, the TWI controller synchronizes its clock with the other remaining clocks. This is shown in [Figure 16-4](#).

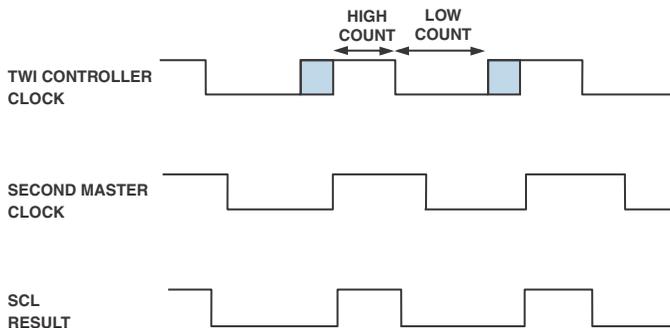


Figure 16-4. TWI Clock Synchronization

The TWI controller's serial clock (SCL) output follows these rules:

- Once the clock high (CLKHI) count is complete, the serial clock output is driven low and the clock low (CLKLOW) count begins.
- Once the clock low count is complete, the serial clock line is three-stated and the clock synchronization logic enters into a delay mode (shaded area) until the SCL line is detected at a logic 1 level. At this time the clock high count begins.

Description of Operation

Bus Arbitration

The TWI controller initiates a master mode transmission (MEN) only when the bus is idle. If the bus is idle and two masters initiate a transfer, arbitration for the bus begins. This is shown in [Figure 16-5](#).

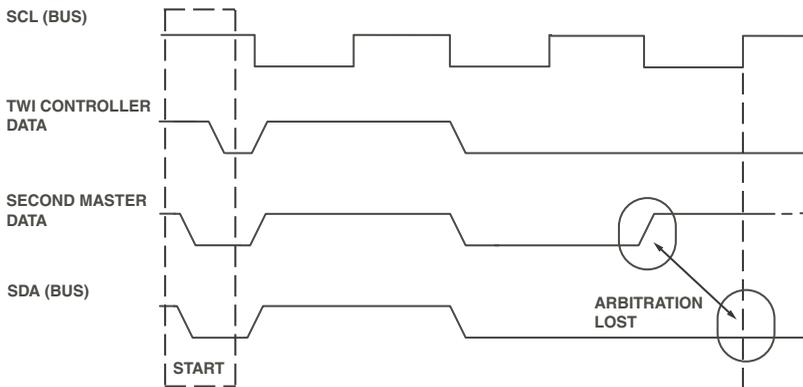


Figure 16-5. TWI Bus Arbitration

The TWI controller monitors the serial data bus (SDA) while SCL is high and if SDA is determined to be an active logic 0 level while the TWI controller's data is a logic 1 level, the TWI controller has lost arbitration and ends generation of clock and data. Note arbitration is not performed only at serial clock edges, but also during the entire time SCL is high.

Start and Stop Conditions

Start and stop conditions involve serial data transitions while the serial clock is a logic 1 level. The TWI controller generates and recognizes these transitions. Typically start and stop conditions occur at the beginning and at the conclusion of a transmission with the exception repeated start “combined” transfers, as shown in [Figure 16-6](#).

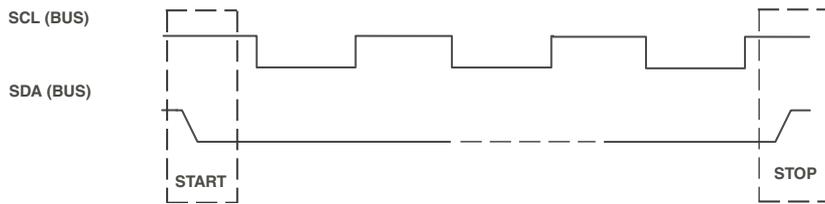


Figure 16-6. TWI Start and Stop Conditions

The TWI controller's special case start and stop conditions include:

- TWI controller addressed as a slave-receiver

If the master asserts a stop condition during the data phase of a transfer, the TWI controller concludes the transfer (SCOMP).

- TWI controller addressed as a slave-transmitter

If the master asserts a stop condition during the data phase of a transfer, the TWI controller concludes the transfer (SCOMP) and indicates a slave transfer error (SERR).

- TWI controller as a master-transmitter or master-receiver

If the stop bit is set during an active master transfer, the TWI controller issues a stop condition as soon as possible avoiding any error conditions (as if data transfer count had been reached).

General Call Support

The TWI controller always decodes and acknowledges a general call address if it is enabled as a slave (SEN) and if general call is enabled (GEN). general call addressing (0x00) is indicated by the GCALL bit being set and by nature of the transfer the TWI controller is a slave-receiver. If the data associated with the transfer is to be NAK'ed, the NAK bit can be set.

Functional Description

If the TWI controller is to issue a general call as a master-transmitter the appropriate address and transfer direction can be set along with loading transmit FIFO data.

Fast Mode

Fast mode essentially uses the same mechanics as standard mode of operation. It is the electrical specifications and timing that are most affected. When fast mode is enabled (*FAST*) the following timings are modified to meet the electrical requirements.

- Serial data rise times before arbitration evaluation (t_r)
- Stop condition set-up time from serial clock to serial data ($t_{SU;STO}$)
- Bus free time between a stop and start condition (t_{BUF})

Functional Description

The following sections describe the functional operation of the TWI.

General Setup

General setup refers to register writes that are required for both slave mode operation and master mode operation. General setup should be performed before either the master or slave enable bits are set.

- Program the `TWI_CONTROL` register to enable the TWI controller and set the prescale value. Program the prescale value to the binary representation of $f_{SCLK}/10\text{MHz}$

All values should be rounded up to the next whole number. The `TWI_ENA` bit enable must be set. Note once the TWI controller is enabled a bus

busy condition may be detected. This condition should clear after t_{BUF} has expired assuming no additional bus activity has been detected.

Slave Mode

When enabled, slave mode operation supports both receive and transmit data transfers. It is not possible to enable only one data transfer direction and not acknowledge (NAK) the other. This is reflected in the following setup.

1. Program `TWI_SLAVE_ADDR`. The appropriate 7 bits are used in determining a match during the address phase of the transfer.
2. Program `TWI_XMT_DATA8` or `TWI_XMT_DATA16`. These are the initial data values to be transmitted in the event the slave is addressed and a transmit is required. This is an optional step. If no data is written and the slave is addressed and a transmit is required, the serial clock (SCL) is stretched and an interrupt is generated until data is written to the transmit FIFO.
3. Program `TWI_INT_MASK`. Enable bits are associated with the desired interrupt sources. As an example, programming the value `0x000F` results in an interrupt output to the processor in the event that a valid address match is detected, a valid slave transfer completes, a slave transfer has an error, a subsequent transfer has begun yet the previous transfer has not been serviced.
4. Program `TWI_SLAVE_CTL`. Ultimately this prepares and enables slave mode operation. As an example, programming the value `0x0005` enables slave mode operation, requires 7-bit addressing, and indicates that data in the transmit FIFO buffer is intended for slave mode transmission.

Functional Description

Table 16-2 shows what the interaction between the TWI controller and the processor might look like using this example.

Table 16-2. Slave Mode Setup Interaction

TWI Controller Master	Processor
Interrupt: SINIT – Slave transfer in progress.	Acknowledge: Clear interrupt source bits.
Interrupt: RCVFULL – Receive buffer is full.	Read receive FIFO buffer. Acknowledge: Clear interrupt source bits.
...	...
Interrupt: SCOMP – Slave transfer complete.	Read receive FIFO buffer. Acknowledge: Clear interrupt source bits.

Master Mode Clock Setup

Master mode operation is set up and executed on a per-transfer basis. An example of programming steps for a receive and for a transmit are given separately in following sections. The clock setup programming step listed here is common to both transfer types.

- Program `TWI_CLKDIV`. This defines the clock high duration and clock low duration.

Master Mode Transmit

Follow these programming steps for a single master mode transmit:

1. Program `TWI_MASTER_ADDR`. This defines the address transmitted during the address phase of the transfer.
2. Program `TWI_XMT_DATA8` or `TWI_XMT_DATA16`. This is the initial data transmitted. It is considered an error to complete the address phase of the transfer and not have data available in the transmit FIFO buffer.

3. Program `TWI_FIFO_CTL`. Indicate if transmit FIFO buffer interrupts should occur with each byte transmitted (8-bits) or with each two bytes transmitted (16-bits).
4. Program `TWI_INT_MASK`. Enable bits associated with the desired interrupt sources. As an example, programming the value `0x0030` results in an interrupt output to the processor in the event that the master transfer completes, and the master transfer has an error.
5. Program `TWI_MASTER_CTL`. Ultimately this prepares and enables master mode operation. As an example, programming the value `0x0201` enables master mode operation, generates a 7-bit address, sets the direction to master-transmit, uses standard mode timing, and transmits 8 data bytes before generating a Stop condition.

Table 16-3 shows what the interaction between the TWI controller and the processor might look like using this example.

Table 16-3. Master Mode Transmit Setup Interaction

TWI Controller Master	Processor
Interrupt: <code>XMEMPTY</code> – Transmit buffer is empty.	Write transmit FIFO buffer. Acknowledge: Clear interrupt source bits.
...	...
Interrupt: <code>MCOMP</code> – Master transfer complete.	Acknowledge: Clear interrupt source bits.

Master Mode Receive

Follow these programming steps for a single master mode receive:

1. Program `TWI_MASTER_ADDR`. This defines the address transmitted during the address phase of the transfer.
2. Program `TWI_FIFO_CTL`. Indicate if receive FIFO buffer interrupts should occur with each byte received (8-bits) or with each two bytes received (16-bits).
3. Program `TWI_INT_MASK`. Enable bits associated with the desired interrupt sources. For example, programming the value `0x0030` results in an interrupt output to the processor in the event that the master transfer completes, and the master transfer has an error.
4. Program `TWI_MASTER_CTL`. Ultimately this prepares and enables master mode operation. As an example, programming the value `0x0205` enables master mode operation, generates a 7-bit address, sets the direction to master-receive, uses standard mode timing, and receives 8 data bytes before generating a Stop condition.

[Table 16-4](#) shows what the interaction between the TWI controller and the processor might look like using this example.

Table 16-4. Master Mode Receive Setup Interaction

TWI Controller Master	Processor
Interrupt: <code>RCVFULL</code> – Receive buffer is full.	Read receive FIFO buffer. Acknowledge: Clear interrupt source bits.
...	...
Interrupt: <code>MCOMP</code> – Master transfer complete.	Acknowledge: Clear interrupt source bits. Read receive FIFO buffer.

Repeated Start Condition

In general, a repeated start condition is the absence of a stop condition between two transfers. The two transfers can be of any direction type. Examples include a transmit followed by a receive, or a receive followed by a transmit. The following sections guide the programmer in developing a service routine.

Transmit/Receive Repeated Start Sequence

Figure 16-7 shows a repeated start data transmit followed by a data receive sequence.

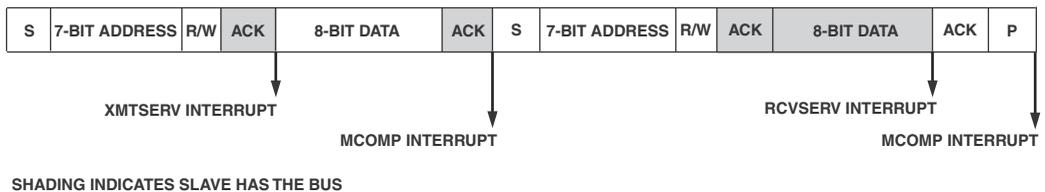


Figure 16-7. Transmit/Receive Data Repeated Start

The following tasks are performed at each interrupt.

- XMTSERV interrupt

This interrupt was generated due to a FIFO access. Since this is the last byte of this transfer, `FIFO_STATUS` indicates the transmit FIFO is empty. When read, `DCNT` would be zero. Set the `RSTART` bit to indicate a repeated start and set the `MDIR` bit if the following transfer will be a data receive.

Functional Description

- MCOMP interrupt

This interrupt was generated because all data has been transferred (DCNT = 0). If no errors were generated, a start condition is initiated. Clear the RSTART bit and program the DCNT with the desired number of bytes to receive.

- RCVSERV interrupt

This interrupt is generated due to the arrival of a byte in the receive FIFO. Simple data handling is all that is required.

- MCOMP interrupt

The transfer is complete.

Receive/Transmit Repeated Start Sequence

Figure 16-8 on page 16-16 illustrates a repeated start data receive followed by a data transmit sequence.

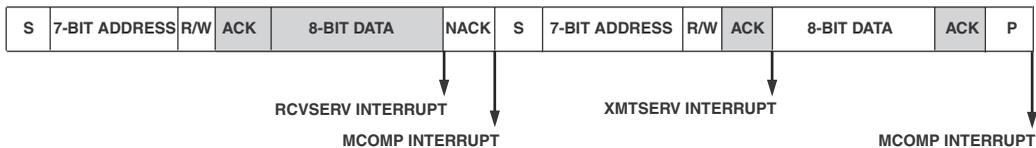


Figure 16-8. Receive/Transmit Data Repeated Start

The tasks performed at each interrupt are:

- RCVSERV interrupt

This interrupt is generated due to the arrival of a data byte in the receive FIFO. Set the RSTART bit to indicate a repeated start and clear the MDIR bit if the following transfer will be a data transmit.

- MCOMP interrupt

This interrupt has occurred due to the completion of the data receive transfer. If no errors were generated, a start condition is initiated. Clear the RSTART bit and program the DCNT with the desired number of bytes to transmit.

- XMTSERV interrupt

This interrupt is generated due to a FIFO access. Simple data handling is all that is required.

- MCOMP interrupt

The transfer is complete.



There is no timing constraint to meet the above conditions—the user can program the bits as required. Refer to [“Clock Stretching During Repeated Start Condition” on page 16-20](#) for more on how the controller stretches the clock during Repeated Start transfers.

Clock Stretching

Clock stretching is an added functionality of the TWI controller in Master Mode operation. This new behavior utilizes self-induced stretching of the I²C clock while waiting on servicing interrupts. Stretching is done automatically by the hardware and no programming is required for this.

The TWI Controller as Master supports three modes of clock stretching: [“Clock Stretching During FIFO Underflow” on page 16-17](#), [“Clock Stretching During FIFO Overflow” on page 16-19](#) and [“Clock Stretching During Repeated Start Condition” on page 16-20](#).

Clock Stretching During FIFO Underflow

During a master mode transmit, an interrupt is generated at the instant the transmit FIFO becomes empty. At this time, the most recent byte

Functional Description

begins transmission. If the XMTSERV interrupt is not serviced, the concluding “acknowledge” phase of the transfer will be stretched. Stretching of the clock continues until new data bytes are written to the transmit FIFO (TWI_XMT_DATA8 or TWI_XMT_DATA16). No other action is required to release the clock and continue the transmission. This behavior continues until the transmission is complete (DCNT = 0) at which time the transmission is concluded (MCOMP) as shown in Figure 16-9 and described in Table 16-5.

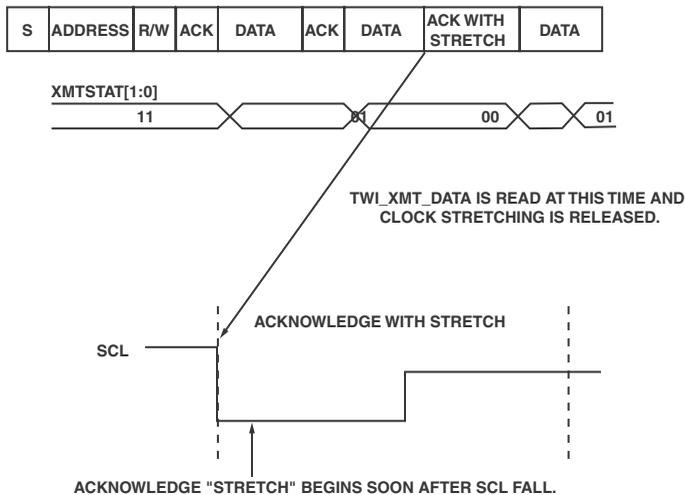


Figure 16-9. Clock Stretching During FIFO Underflow

Table 16-5. FIFO Underflow Case

TWI Controller	Processor
Interrupt: XMTSERV – Transmit FIFO buffer is empty.	Acknowledge: Clear interrupt source bits. Write transmit FIFO buffer.
...	...
Interrupt: MCOMP – Master transmit complete (DCNT= 0x00).	Acknowledge: Clear interrupt source bits.

Clock Stretching During FIFO Overflow

During a master mode receive, an interrupt is generated at the instant the receive FIFO becomes full. It is during the acknowledge phase of this received byte that clock stretching begins. No attempt is made to initiate the reception of an additional byte. Stretching of the clock continues until the data bytes previously received are read from the receive FIFO buffer (TWI_RCV_DATA8, TWI_RCV_DATA16). No other action is required to release the clock and continue the reception of data. This behavior continues until the reception is complete (DCNT = 0x00) at which time the reception is concluded (MCOMP) as shown in Figure 16-10 and described in Table 16-6.

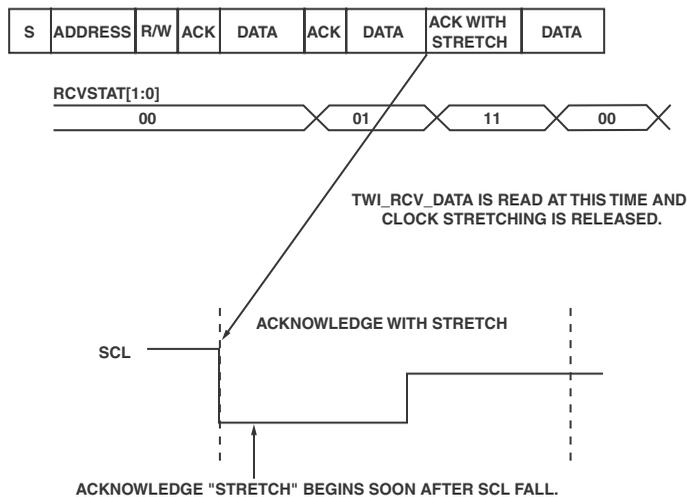


Figure 16-10. Clock Stretching During FIFO Overflow

Functional Description

Table 16-6. FIFO Overflow Case

TWI Controller	Processor
Interrupt: RCVSERV – Receive FIFO buffer is full.	Acknowledge: Clear interrupt source bits. Read receive FIFO buffer.
...	...
Interrupt: MCOMP – Master receive complete.	Acknowledge: Clear interrupt source bits.

Clock Stretching During Repeated Start Condition

The repeated start feature in I²C protocol requires transitioning between two subsequent transfers. With the use of clock stretching, the task of managing transitions becomes simpler and becomes common to all transfer types.

Once an initial TWI master transfer has completed (transmit or receive) the clock will initiate a stretch during the repeated start phase between transfers. Concurrent with this event the initial transfer will generate a transfer complete interrupt (MCOMP) to signify the initial transfer has completed (DCNT = 0). This initial transfer is handled without any special bit setting sequences or timings. The clock stretching logic described above applies here. With no system related timing constraints the subsequent transfer (receive or transmit) is setup and activated. This sequence can be repeated as many times as required to string a series of repeated start transfers together. This is shown in [Figure 16-11](#) and described in [Table 16-7](#).

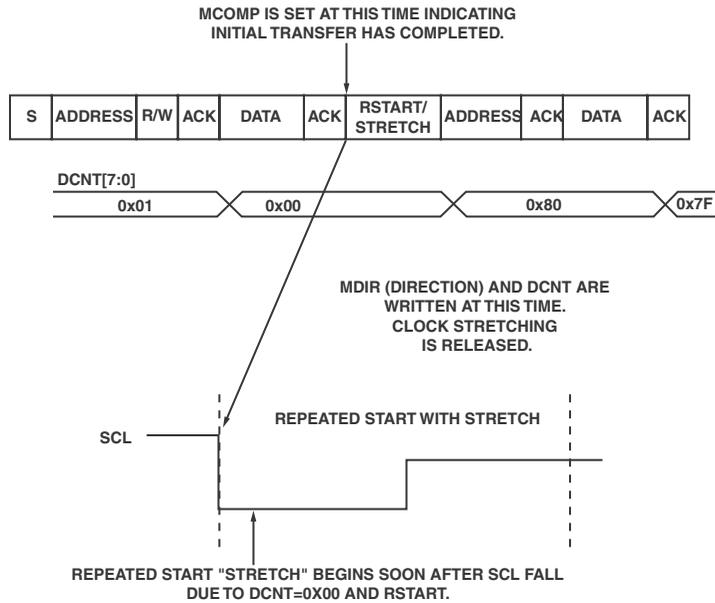


Figure 16-11. Clock Stretching during Repeated Start Condition

Table 16-7. Repeated Start Case

TWI Controller	Processor
Interrupt: MCOMP – Initial transmit has completed and DCNT = 0x00. Note: transfer in progress, RSTART previously set.	Acknowledge: Clear interrupt source bits. Write TWI_MASTER_CTL, setting MDIR (receive), clearing RSTART, and setting new DCNT value (nonzero).
Interrupt: RCVSERV – Receive FIFO is full.	Acknowledge: Clear interrupt source bits. Read receive FIFO buffer.
...	...
Interrupt: MCOMP – Master receive complete.	Acknowledge: Clear interrupt source bits.

Programming Model

Figure 16-12 and Figure 16-13 illustrate the programming model for the TWI.

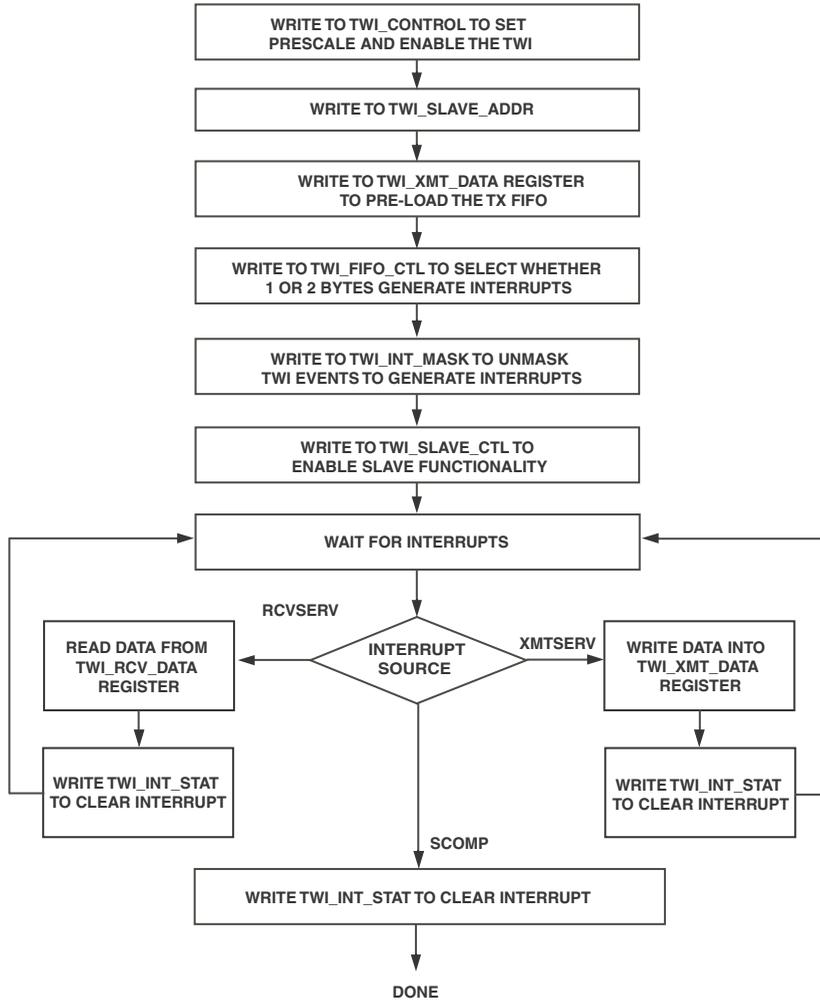


Figure 16-12. TWI Slave Mode

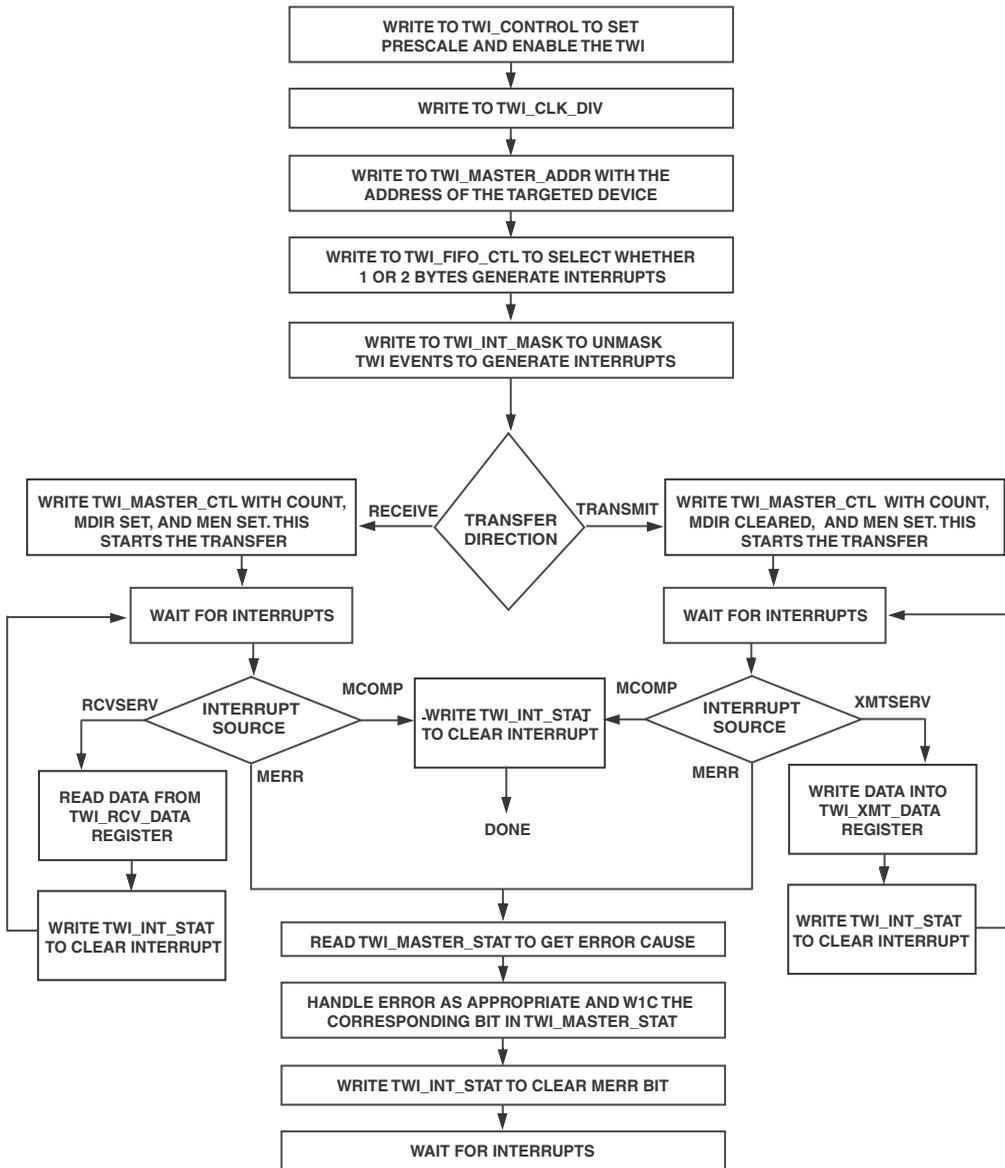


Figure 16-13. TWI Master Mode

Register Descriptions

The TWI controller has 16 registers described in the following sections. [Figure 16-14](#) through [Figure 16-31](#) on page 16-46 illustrate the registers.

TWI CONTROL Register (TWI_CONTROL)

The `TWI_CONTROL` register is used to enable the TWI module as well as to establish a relationship between the system clock (`SCLK`) and the TWI controller's internally timed events. The internal time reference is derived from `SCLK` using a prescaled value.

$$\text{PRESCALE} = f_{\text{SCLK}}/10\text{MHz}$$

SCCB compatibility is an optional feature and should not be used in an I²C bus system. This feature is turned on by setting the `SCCB` bit in the `TWI_CONTROL` register. When this feature is set all slave asserted acknowledgement bits are ignored by this master. This feature is valid only during transfers where the TWI is mastering an SCCB bus. Slave mode transfers should be avoided when this feature is enabled because the TWI controller always generates an acknowledge in slave mode.

For either master and/or slave mode of operation, the TWI controller is enabled by setting the `TWI_ENA` bit in the `TWI_CONTROL` register. It is recommended that this bit be set at the time `PRESCALE` is initialized and remain set. This guarantees accurate operation of bus busy detection logic.

The `PRESCALE` field of the `TWI_CONTROL` register specifies the number of system clock (`SCLK`) periods used in the generation of one internal time reference. The value of `PRESCALE` must be set to create an internal time reference with a period of 10 MHz. It is represented as a 7-bit binary value.

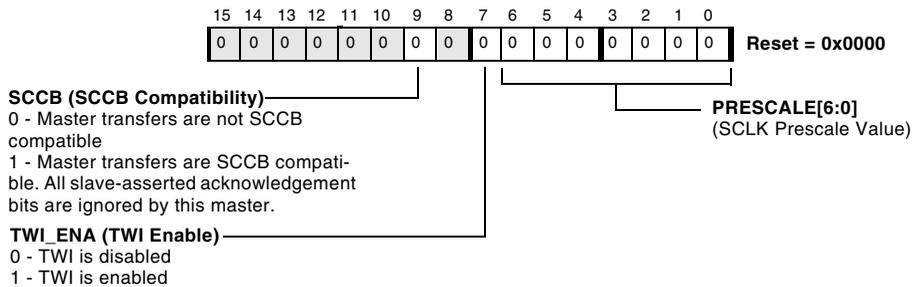
TWI Control Register (TWI_CONTROL)

Figure 16-14. TWI Control Register

SCL Clock Divider Register (TWI_CLKDIV)

The clock signal SCL is an output in master mode and an input in slave mode.

During master mode operation, the TWI_CLKDIV register values are used to create the high and low durations of the serial clock (SCL). Serial clock frequencies can vary from 400 KHz to less than 20 KHz. The resolution of the clock generated is 1/10 MHz or 100 ns.

$$\text{CLKDIV} = \text{TWI SCL period} / 10 \text{ MHz time reference}$$

For example, for an SCL of 400 KHz (period = 1/400 KHz = 2500 ns) and an internal time reference of 10 MHz (period = 100 ns):

$$\text{CLKDIV} = 2500 \text{ ns} / 100 \text{ ns} = 25$$

For an SCL with a 30% duty cycle, then CLKLOW = 17 and CLKHI = 8. Note that CLKLOW and CLKHI add up to CLKDIV.

The CLKHI field of the TWI_CLKDIV register specifies the number of 10 MHz time reference periods the serial clock (SCL) waits before a new clock

Additional information for the `TWI_SLAVE_CTL` register bits includes:

- **General call enable** (`GEN`)

General call address detection is available only when slave mode is enabled.

[0] General call address matching is not enabled.

[1] General call address matching is enabled. A general call slave receive transfer is accepted. All status and interrupt source bits associated with transfers are updated.

- **NAK** (`NAK`)

[0] Slave receive transfers generate an ACK at the conclusion of a data transfer.

[1] Slave receive transfers generate a data NAK (not acknowledge) at the conclusion of a data transfer. The slave is still considered to be addressed.

- **Slave transmit data valid** (`STDVAL`)

[0] Data in the transmit FIFO is for master mode transmits and is not allowed to be used during a slave transmit, and the transmit FIFO is treated as if it is empty.

[1] Data in the transmit FIFO is available for a slave transmission.

- **Slave enable** (`SEN`)

[0] The slave is not enabled. No attempt is made to identify a valid address. If cleared during a valid transfer, clock stretching ceases, the serial data line is released, and the current byte is not acknowledged.

[1] The slave is enabled. Enabling slave and master modes of operation concurrently is allowed.

TWI Slave Mode Address Register (TWI_SLAVE_ADDR)

The TWI_SLAVE_ADDR register holds the slave mode address, which is the valid address that the slave-enabled TWI controller responds to. The TWI controller compares this value with the received address during the addressing phase of a transfer.

TWI Slave Mode Address Register (TWI_SLAVE_ADDR)

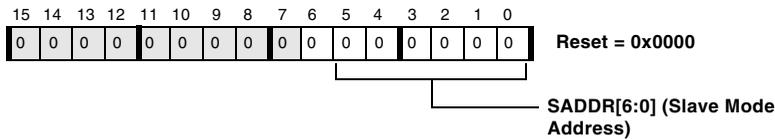


Figure 16-17. TWI Slave Mode Address Register

TWI Slave Mode Status Register (TWI_SLAVE_STAT)

TWI Slave Mode Status Register (TWI_SLAVE_STAT)

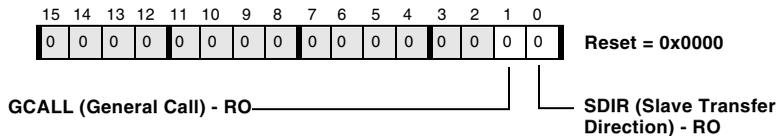


Figure 16-18. TWI Slave Mode Status Register

During and at the conclusion of register slave mode transfers, the `TWI_SLAVE_STAT` register holds information on the current transfer. Generally slave mode status bits are not associated with the generation of interrupts. Master mode operation does not affect slave mode status bits.

- **General call** (`GCALL`)

This bit self clears if slave mode is disabled (`SEN = 0`).

[0] At the time of addressing, the address was not determined to be a general call.

[1] At the time of addressing, the address was determined to be a general call.

- **Slave transfer direction** (`SDIR`)

This bit self clears if slave mode is disabled (`SEN = 0`).

[0] At the time of addressing, the transfer direction was determined to be slave receive.

[1] At the time of addressing, the transfer direction was determined to be slave transmit.

TWI Master Mode Control Register (`TWI_MASTER_CTL`)

The `TWI_MASTER_CTL` register controls the logic associated with master mode operation. Bits in this register do not affect slave mode operation and should not be modified to control slave mode functionality.

Register Descriptions

TWI Master Mode Control Register (TWI_MASTER_CTL)

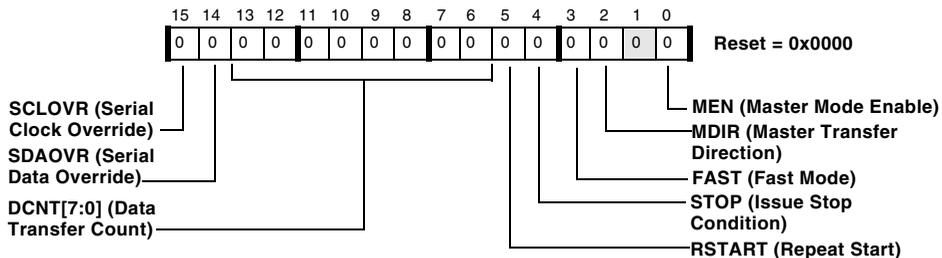


Figure 16-19. TWI Master Mode Control Register

Additional information for the TWI_MASTER_CTL register bits includes:

- **Serial clock override (SCLOVR)**

This bit can be used when direct control of the serial clock line is required. Normal master and slave mode operation should not require override operation.

[0] Normal serial clock operation under the control of master mode clock generation and slave mode clock stretching logic.

[1] Serial clock output is driven to an active 0 level overriding all other logic. This state is held until this bit is cleared.

- **Serial data (SDA) override (SDAOVR)**

This bit can be used when direct control of the serial data line is required. Normal master and slave mode operation should not require override operation.

[0] Normal serial data operation under the control of the transmit shift register and acknowledge logic.

[1] Serial data output is driven to an active 0 level overriding all other logic. This state is held until this bit is cleared.

- **Data transfer count** (DCNT[7:0])

Indicates the number of data bytes to transfer. As each data word is transferred, DCNT is decremented. When DCNT is 0, a stop condition is generated. Setting DCNT to 0xFF disables the counter. In this transfer mode, data continues to be transferred until it is concluded by setting the STOP bit.

- **Repeat start** (RSTART)

[0] Transfer concludes with a stop condition.

[1] Issue a repeat start condition at the conclusion of the current transfer (DCNT = 0) and begin the next transfer. The current transfer concludes with updates to the appropriate status and interrupt bits. If errors occurred during the previous transfer, a repeat start does not occur. In the absence of any errors, master enable (MEN) does not self clear on a repeat start.

- **Issue stop condition** (STOP)

[0] Normal transfer operation.

[1] The transfer concludes as soon as possible avoiding any error conditions (as if data transfer count had been reached) and at that time the TWI interrupt mask register (TWI_INT_MASK) is updated along with any associated status bits.

- **Fast mode** (FAST)

[0] Standard mode (up to 100K bits/s) timing specifications in use.

[1] Fast mode (up to 400K bits/s) timing specifications in use.

- **Master transfer direction** (MDIR)

[0] The initiated transfer is master transmit.

[1] The initiated transfer is master receive.

Register Descriptions

- **Master mode enable** (MEN)

This bit self clears at the completion of a transfer. This includes transfers terminated due to errors.

[0] Master mode functionality is disabled. If this bit is cleared during operation, the transfer is aborted and all logic associated with master mode transfers are reset. Serial data and serial clock (SDA, SCL) are no longer driven. Write-1-to-clear status bits are not affected.

[1] Master mode functionality is enabled. A start condition is generated if the bus is idle.

TWI Master Mode Address Register (TWI_MASTER_ADDR)

During the addressing phase of a transfer, the TWI controller, with its master enabled, transmits the contents of the TWI_MASTER_ADDR register. When programming this register, omit the read/write bit. That is, only the upper 7 bits that make up the slave address should be written to this register. For example, if the slave address is b#1010000X, where X is the read/write bit, then TWI_MASTER_ADDR is programmed with b#1010000, which corresponds to 0x50. When sending out the address on the bus, the TWI controller appends the read/write bit as appropriate based on the state of the MDIR bit in the master mode control register.

TWI Master Mode Address Register (TWI_MASTER_ADDR)

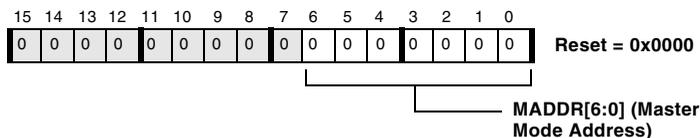


Figure 16-20. TWI Master Mode Address Register

TWI Master Mode Status Register (TWI_MASTER_STAT)

TWI Master Mode Status Register (TWI_MASTER_STAT)

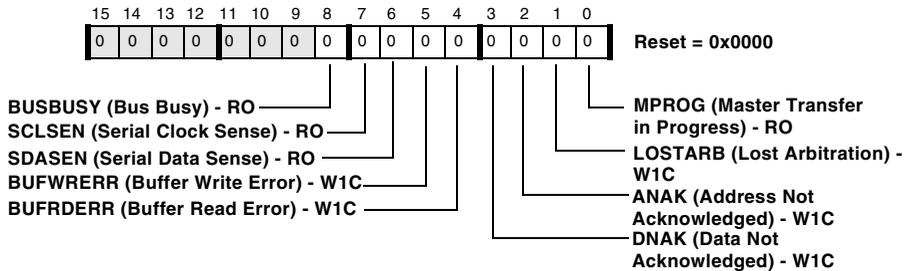


Figure 16-21. TWI Master Mode Status Register

The TWI_MASTER_STAT register holds information during master mode transfers and at their conclusion. Generally, master mode status bits are not directly associated with the generation of interrupts but offer information on the current transfer. Slave mode operation does not affect master mode status bits.

Note that—while the SCLSEN bit is set (this could be due to having no pull-up resistor on SCL or another agent is driving SCL low)—the acknowledge bits (ANAK and DNAK) do not update. This result occurs because the acknowledge conditions are sampled during the high phase of SCL.

- **Bus busy** (BUSBUSY)

Indicates whether the bus is currently busy or free. This indication is not limited to only this device but is for all devices. Upon a start condition, the setting of the register value is delayed due to the input filtering. Upon a stop condition the clearing of the register value occurs after t_{BUF} .

Register Descriptions

[0] The bus is free. The clock and data bus signals have been inactive for the appropriate bus free time.

[1] The bus is busy. Clock or data activity has been detected.

- **Serial clock sense** (SCLSEN)

This status bit can be used when direct sensing of the serial clock line is required. The register value is delayed due to the input filter (nominally 50 ns). Normal master and slave mode operation should not require this feature.

[0] An inactive “one” is currently being sensed on the serial clock.

[1] An active “zero” is currently being sensed on the serial clock. The source of the active driver is not known and can be internal or external.

- **Serial data sense** (SDASEN)

This status bit can be used when direct sensing of the serial data line is required. The register value is delayed due to the input filter (nominally 50 ns). Normal master and slave mode operation should not require this feature.

[0] An inactive “one” is currently being sensed on the serial data line.

[1] An active “zero” is currently being sensed on the serial data line. The source of the active driver is not known and can be internal or external.

- **Buffer write error** (BUFWRERR)

[0] The current master receive has not detected a receive buffer write error.

[1] The current master transfer was aborted due to a receive buffer write error. The receive buffer and receive shift register were both full at the same time. This bit is W1C.

- **Buffer read error** (BUFRDERR)

[0] The current master transmit has not detected a buffer read error.

[1] The current master transfer was aborted due to a transmit buffer read error. At the time data was required by the transmit shift register the buffer was empty. This bit is W1C.

- **Data not acknowledged** (DNAK)

[0] The current master receive has not detected a NAK during data transmission.

[1] The current master transfer was aborted due to the detection of a NAK during data transmission. This bit is W1C.

- **Address not acknowledged** (ANAK)

[0] The current master transmit has not detected NAK during addressing.

[1] The current master transfer was aborted due to the detection of a NAK during the address phase of the transfer. This bit is W1C.

Register Descriptions

- **Lost arbitration** (LOSTARB)

[0] The current transfer has not lost arbitration with another master.

[1] The current transfer was aborted due to the loss of arbitration with another master. This bit is W1C.

- **Master transfer in progress** (MPROG)

[0] Currently no transfer is taking place. This can occur once a transfer is complete or while an enabled master is waiting for an idle bus.

[1] A master transfer is in progress.

TWI FIFO Control Register (TWI_FIFO_CTL)

The TWI_FIFO_CTL register control bits affect only the FIFO and are not tied in any way with master or slave mode operation.

TWI FIFO Control Register (TWI_FIFO_CTL)

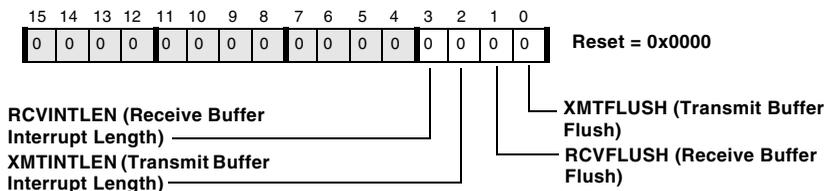


Figure 16-22. TWI FIFO Control Register

Additional information for the `TWI_FIFO_CTL` register bits includes:

- **Receive buffer interrupt length** (`RCVINTLEN`)

This bit determines the rate at which receive buffer interrupts are to be generated. Interrupts may be generated with each byte received or after two bytes are received.

[0] An interrupt (`RCVSERV`) is set when `RCVSTAT` indicates one or two bytes in the FIFO are full (01 or 11).

[1] An interrupt (`RCVSERV`) is set when the `RCVSTAT` field in the `TWI_FIFO_STAT` register indicates two bytes in the FIFO are full (11).

- **Transmit buffer interrupt length** (`XMTINTLEN`)

This bit determines the rate at which transmit buffer interrupts are to be generated. Interrupts may be generated with each byte transmitted or after two bytes are transmitted.

[0] An interrupt (`XMTSERV`) is set when `XMTSTAT` indicates one or two bytes in the FIFO are empty (01 or 00).

[1] An interrupt (`XMTSERV`) is set when the `XMTSTAT` field in the `TWI_FIFO_STAT` register indicates two bytes in the FIFO are empty (00).

- **Receive buffer flush** (`RCVFLUSH`)

[0] Normal operation of the receive buffer and its status bits.

[1] Flush the contents of the receive buffer and update the `RCVSTAT` status bit to indicate the buffer is empty. This state is held until this bit is cleared. During an active receive the receive buffer in this state responds to the receive logic as if it is full.

Register Descriptions

- **Transmit buffer flush** (XMTFLUSH)

[0] Normal operation of the transmit buffer and its status bits.

[1] Flush the contents of the transmit buffer and update the XMTSTAT status bit to indicate the buffer is empty. This state is held until this bit is cleared. During an active transmit the transmit buffer in this state responds as if the transmit buffer is empty.

TWI FIFO Status Register (TWI_FIFO_STAT)

TWI FIFO Status Register (TWI_FIFO_STAT)

All bits are RO.

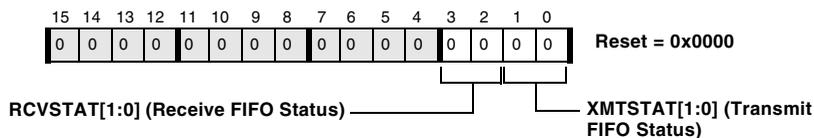


Figure 16-23. TWI FIFO Status Register

TWI FIFO Status

The fields in the TWI_FIFO_STAT register indicate the state of the FIFO buffers' receive and transmit contents. The FIFO buffers do not discriminate between master data and slave data. By using the status and control bits provided, the FIFO can be managed to allow simultaneous master and slave operation.

- **Receive FIFO status** (RCVSTAT[1:0])

The RCVSTAT field is read only. It indicates the number of valid data bytes in the receive FIFO buffer. The status is updated with each FIFO buffer read using the peripheral data bus or write access by the receive shift register. Simultaneous accesses are allowed.

[00] The FIFO is empty.

[01] The FIFO contains one byte of data. A single byte peripheral read of the FIFO is allowed.

[10] Reserved

[11] The FIFO is full and contains two bytes of data. Either a single or double byte peripheral read of the FIFO is allowed.

- **Transmit FIFO status** (`XMTSTAT[1:0]`)

The `XMTSTAT` field is read only. It indicates the number of valid data bytes in the FIFO buffer. The status is updated with each FIFO buffer write using the peripheral data bus or read access by the transmit shift register. Simultaneous accesses are allowed.

[00] The FIFO is empty. Either a single or double byte peripheral write of the FIFO is allowed.

[01] The FIFO contains one byte of data. A single byte peripheral write of the FIFO is allowed.

[10] Reserved

[11] The FIFO is full and contains two bytes of data.

TWI Interrupt Mask Register (`TWI_INT_MASK`)

The `TWI_INT_MASK` register enables interrupt sources to assert the interrupt output. Each mask bit corresponds with one interrupt source bit in the `TWI_INT_STAT` register. Reading and writing the `TWI_INT_MASK` register does not affect the contents of the `TWI_INT_STAT` register.

Register Descriptions

TWI Interrupt Mask Register (TWI_INT_MASK)

For all bits, 0 = Interrupt generation disabled, 1 = Interrupt generation enabled.

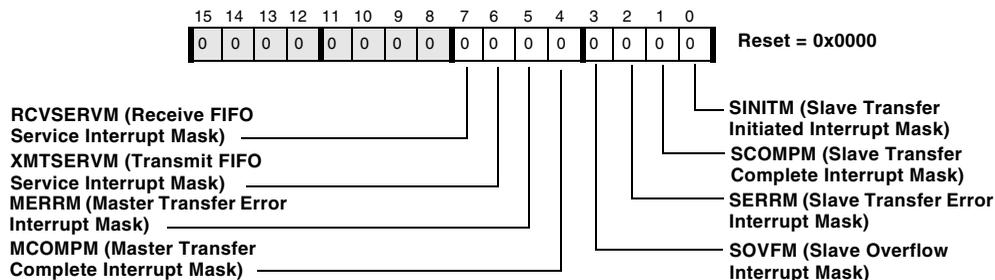


Figure 16-24. TWI Interrupt Mask Register

TWI Interrupt Status Register (TWI_INT_STAT)

TWI Interrupt Status Register (TWI_INT_STAT)

All bits are sticky and W1C.

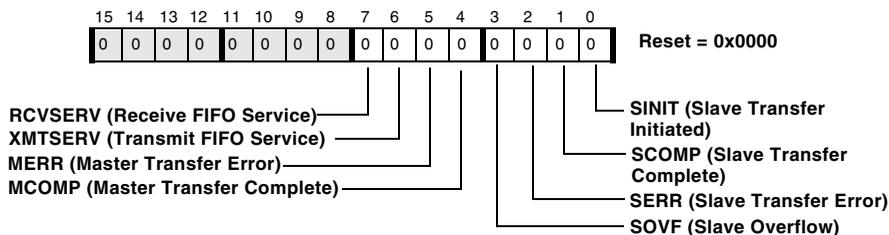


Figure 16-25. TWI Interrupt Status Register

The TWI_INT_STAT register contains information about functional areas requiring servicing. Many of the bits serve as an indicator to further read and service various status registers. After servicing the interrupt source

associated with a bit, the user must clear that interrupt source bit by writing a 1 to it.

- **Receive FIFO service (RCVSERV)**

If `RCVINTLEN` in the `TWI_FIFO_CTL` register is 0, this bit is set each time the `RCVSTAT` field in the `TWI_FIFO_STAT` register is updated to either 01 or 11. If `RCVINTLEN` is 1, this bit is set each time `RCVSTAT` is updated to 01 or 11.

[0] No errors have been detected.

[1] The FIFO does not require servicing or the `RCVSTAT` field has not changed since this bit was last cleared.

- **Transmit FIFO service (XMTSERV)**

If `XMTINTLEN` in the `TWI_FIFO_CTL` register is 0, this bit is set each time the `XMTSTAT` field in the `TWI_FIFO_STAT` register is updated to either 01 or 00. If `XMTINTLEN` is 1, this bit is set each time `XMTSTAT` is updated to 00.

[0] FIFO does not require servicing or `XMTSTAT` field has not changed since this bit was last cleared.

[1] The transmit FIFO buffer has one or two 8-bit locations available to be written.

- **Master transfer error (MERR)**

[0] No errors have been detected.

[1] A master error has occurred. The conditions surrounding the error are indicated by the master status register (`TWI_MASTER_STAT`).

Register Descriptions

- **Master transfer complete** (MCOMP)
 - [0] The completion of a transfer has not been detected.
 - [1] The initiated master transfer has completed. In the absence of a repeat start, the bus has been released.
- **Slave overflow** (SOVF)
 - [0] No overflow has been detected.
 - [1] The slave transfer complete (SCOMP) bit was set at the time a subsequent transfer has acknowledged an address phase. The transfer continues, however, it may be difficult to delineate data of one transfer from another.
- **Slave transfer error** (SERR)
 - [0] No errors have been detected.
 - [1] A slave error has occurred. A restart or stop condition has occurred during the data receive phase of a transfer.
- **Slave transfer complete** (SCOMP)
 - [0] The completion of a transfer has not been detected.
 - [1] The transfer is complete and either a stop, or a restart was detected.
- **Slave transfer initiated** (SINIT)
 - [0] A transfer is not in progress. An address match has not occurred since the last time this bit was cleared.
 - [1] The slave has detected an address match and a transfer has been initiated.

TWI FIFO Transmit Data Single Byte Register (TWI_XMT_DATA8)

The `TWI_XMT_DATA8` register holds an 8-bit data value written into the FIFO buffer. Transmit data is entered into the corresponding transmit buffer in a first-in first-out order. For 16-bit PAB writes, a write access to `TWI_XMT_DATA8` adds only one transmit data byte to the FIFO buffer. With each access, the transmit status (`XMTSTAT`) field in the `TWI_FIFO_STAT` register is updated. If an access is performed while the FIFO buffer is full, the write is ignored and the existing FIFO buffer data and its status remains unchanged.

TWI FIFO Transmit Data Single Byte Register (TWI_XMT_DATA8)

All bits are WO.

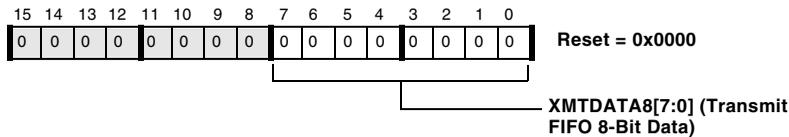


Figure 16-26. TWI FIFO Transmit Data Single Byte Register

TWI FIFO Transmit Data Double Byte Register (TWI_XMT_DATA16)

The `TWI_XMT_DATA16` register holds a 16-bit data value written into the FIFO buffer. To reduce interrupt output rates and peripheral bus access times, a double byte transfer data access can be performed. Two data bytes can be written, effectively filling the transmit FIFO buffer with a single access.

The data is written in little endian byte order as shown in [Figure 16-27](#) where byte 0 is the first byte to be transferred and byte 1 is the second byte to be transferred. With each access, the transmit status (`XMTSTAT`) field in the `TWI_FIFO_STAT` register is updated. If an access is performed while the

Register Descriptions

FIFO buffer is not empty, the write is ignored and the existing FIFO buffer data and its status remains unchanged.

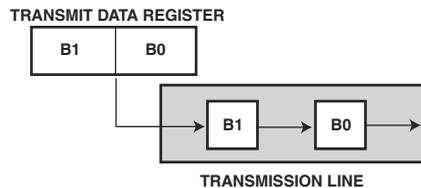


Figure 16-27. Transmit Little Endian Byte Order

TWI FIFO Transmit Data Double Byte Register (TWI_XMT_DATA16)

All bits are WO. This register always reads as 0x0000.

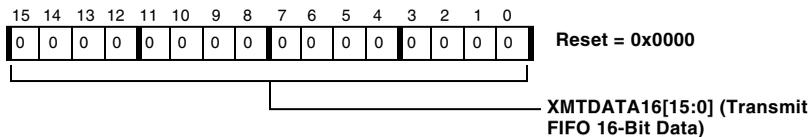


Figure 16-28. TWI FIFO Transmit Data Double Byte Register

TWI FIFO Receive Data Single Byte Register (TWI_RCV_DATA8)

The `TWI_RCV_DATA8` register holds an 8-bit data value read from the FIFO buffer. Receive data is read from the corresponding receive buffer in a first-in first-out order. Although peripheral bus reads are 16 bits, a read access to `TWI_RCV_DATA8` will access only one transmit data byte from the FIFO buffer. With each access, the receive status (`RCVSTAT`) field in the `TWI_FIFO_STAT` register is updated. If an access is performed while the FIFO buffer is empty, the data is unknown and the FIFO buffer status remains indicating it is empty.

TWI FIFO Receive Data Single Byte Register (TWI_RCV_DATA8)

All bits are RO.

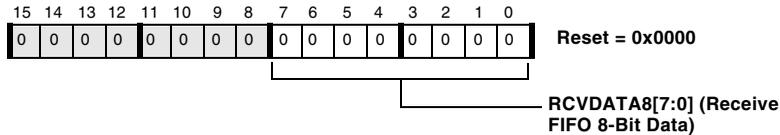


Figure 16-29. TWI FIFO Receive Data Single Byte Register

TWI FIFO Receive Data Double Byte Register (TWI_RCV_DATA16)

The `TWI_RCV_DATA16` register holds a 16-bit data value read from the FIFO buffer. To reduce interrupt output rates and peripheral bus access times, a double byte receive data access can be performed. Two data bytes can be read, effectively emptying the receive FIFO buffer with a single access.

The data is read in little endian byte order as shown in [Figure 16-30](#) where byte 0 is the first byte received and byte 1 is the second byte received. With each access, the receive status (`RCVSTAT`) field in the `TWI_FIFO_STAT` register is updated to indicate it is empty. If an access is performed while the FIFO buffer is not full, the read data is unknown and the existing FIFO buffer data and its status remains unchanged.

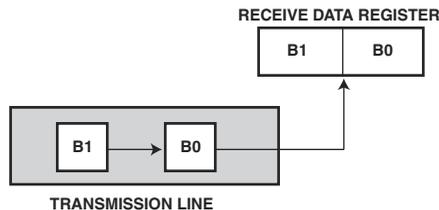


Figure 16-30. Receive Little Endian Byte Order

Programming Examples

TWI FIFO Receive Data Double Byte Register (TWI_RCV_DATA16)

All bits are WO.

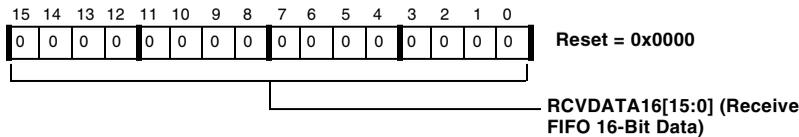


Figure 16-31. TWI FIFO Receive Data Double Byte Register

Programming Examples

The following sections include programming examples for general setup, slave mode, and master mode, as well as guidance for repeated start conditions.

Master Mode Setup

[Listing 16-1](#) shows how to initiate polled receive and transmit transfers in master mode.

Listing 16-1. Master Mode Receive/Transmit Transfer

```
/******  
Macro for the count field of the TWI_MASTER_CTL register  
x can be any value between 0 and 0xFE (254). A value of  
0xFF disables the counter.  
*****/  
#define TWICount(x) (DCNT & ((x) << 6))  
  
.section L1_data_b;  
.byte TX_file[file_size] = "DATA.hex";  
.BYTE RX_CHECK[file_size];
```

```
.byte rcvFirstWord[2];

.SECTION program;
_main:
/*****
TWI Master Initialization subroutine
*****/

TWI_INIT:
/*****
Enable the TWI controller and set the Prescale value
Prescale = 10 (0xA) for an SCLK = 100 MHz (CLKIN = 50MHz)
Prescale = SCLK / 10 MHz

P1 points to the base of the system MMRs
*****/

R1 = TWI_ENA | 0xA (z);
W[P1 + LO(TWI_CONTROL)] = R1;

/*****
Set CLKDIV:
For example, for an SCL of 400 KHz (period = 1/400 KHz = 2500 ns)
and an internal time reference of 10 MHz (period = 100 ns):
CLKDIV = 2500 ns / 100 ns = 25
For an SCL with a 30% duty cycle, then CLKLOW = 17 (0x11) and
CLKHI = 8.
*****/
R5 = CLKHI(0x8) | CLKLOW(0x11) (z);
W[P1 + LO(TWI_CLKDIV)] = R5;

/*****
enable these signals to generate a TWI interrupt: optional
*****/
```

Programming Examples

```
R1 = RCVSERV | XMTSERV | MERR | MCOMP (z);
W[P1 + LO(TWI_INT_MASK)] = R1;

/*****
The address needs to be shifted one place to the right
e.g., 1010 001x becomes 0101 0001 (0x51) the TWI controller
will actually send out 1010 001x where x is either a 0 for
writes or 1 for reads
*****/
R6 = 0xBF;
R6 = R6 >> 1;
TWI_INIT.END: W[P1 + LO(TWI_MASTER_ADDR)] = R6;

/***** END OF TWI INIT *****/

/*****
Starting the Read transfer
Program the Master Control register with:
1. the number of bytes to transfer: TWICount(x)
2. Repeated Start (RESTART): optional
3. speed mode: FAST or SLOW
4. direction of transfer:
    MDIR = 1 for reads, MDIR = 0 for writes
5. Master Enable MEN. This will kick off the master transfer
*****/
R1 = TWICount(0x2) | FAST | MDIR | MEN;
W[P1 + LO(TWI_MASTER_CTL)] = R1;
ssync;

/*****
Poll the FIFO Status register to know when
2 bytes have been shifted into the RX FIFO
*****/
Rx_stat:
```

```

R1 = W[P1 + LO(TWI_FIFO_STAT)](Z);
R0 = 0xC;
R1 = R1 & R0;
CC = R1 == R0;
IF ! cc jump Rx_stat;
R0 = W[P1 + LO(TWI_RCV_DATA16)](Z); /* Read data from the RX fifo
*/
ssync;

/*****
check that master transfer has completed
MCOMP will be set when Count reaches zero
*****/
M_COMP:
    R1 = W[P1 + LO(TWI_INT_STAT)](z);
    CC = BITTST (R1, bitpos(MCOMP));
    if ! CC jump M_COMP;
M_COMP.END: W[P1 + LO(TWI_INT_STAT)] = R1;

/* load the pointer with the address of the transmit buffer */
P2.H = TX_file;
P2.L = TX_file;

/*****
Pre-load the tx FIFO with the first two bytes: this is
necessary to avoid the generation of the Buffer Read Error
(BUFRDERR) which occurs whenever a transmit transfer is
initiated while the transmit buffer is empty
*****/
R3 = W[P2++](Z);
W[P1 + LO(TWI_XMT_DATA16)] = R3;

/*****
Initiating the Write operation
*****/

```

Programming Examples

Program the Master Control register with:

1. the number of bytes to transfer: TWICount(x)
2. Repeated Start (RESTART): optional
3. speed mode: FAST or Standard
4. direction of transfer:
MDIR = 1 for reads, MDIR = 0 for writes
5. Master Enable MEN. Setting this bit will kick off the transfer

```
*****/
```

```
R1 = TWICount(0xFE) | FAST | MEN;
```

```
W[P1 + LO(TWI_MASTER_CTL)] = R1;
```

```
SSYNC;
```

```
/*****
```

```
loop to write data to a TWI slave device P3 times
```

```
*****/
```

```
P3 = length(TX_file);
```

```
LSETUP (Loop_Start1, Loop_End1) LC0 = P3;
```

```
Loop_Start1:
```

```
/******
```

```
check that there's at least one byte location empty in  
the tx fifo
```

```
*****/
```

```
XMTSERV_Status:
```

```
R1 = W[P1 + LO(TWI_INT_STAT)](z);
```

```
CC = BITTST (R1, bitpos(XMTSERV)); /* test XMTSERV bit */
```

```
if ! CC jump XMTSERV_Status;
```

```
W[P1 + LO(TWI_INT_STAT)] = R1; /* clear status */
```

```
SSYNC;
```

```
/*****
```

```
write byte into the transmit FIFO
```

```
*****/
```

```
R3 = B[P2++](Z);
```

```

        W[P1 + LO(TWI_XMT_DATA8)] = R3;
    Loop_End1:  SSYNC;

    /* check that master transfer has completed */
M_COMP1:
    R1 = W[P1 + LO(TWI_INT_STAT)](z);
    CC = BITTST (R1, bitpos(MCOMP1));
    if ! CC jump M_COMP;
M_COMP1.END:W[P1 + LO(TWI_INT_STAT)] = R1;

    idle;
_main.end:

```

Slave Mode Setup

[Listing 16-2](#) shows how to configure the slave for interrupt based transfers. The interrupts are serviced in the subroutine `_TWI_ISR` shown in [Listing 16-3](#).

Listing 16-2. Slave Mode Setup

```

#include <defBF527.h>
/*BF527 is used here as an example—change as appropriate.*/
#include "startup.h"

#define file_size 254
#define SYSMMR_BASE 0xFFC00000
#define COREMMR_BASE 0xFFE00000

.GLOBAL _main;
.EXTERN _TWI_ISR;

.section L1_data_b;
.BYTE TWI_RX[file_size];

```

Programming Examples

```
.BYTE TWI_TX[file_size] = "transmit.dat";

.section L1_code;
_main:

/*****
TWI Slave Initialization subroutine
*****/
TWI_SLAVE_INIT:

/*****
Enable the TWI controller and set the Prescale value
Prescale = 10 (0xA) for an SCLK = 100 MHz (CLKIN = 50MHz)
Prescale = SCLK / 10 MHz
P1 points to the base of the system MMRs
P0 points to the base of the core MMRs
*****/
R1 = TWI_ENA | 0xA (z);
W[P1 + LO(TWI_CONTROL)] = R1;

/*****
Slave address
program the address to which this slave will respond to.
this is an arbitrary 7-bit value
*****/
R1 = 0x5F;
W[P1 + LO(TWI_SLAVE_ADDR)] = R1;

/*****
Pre-load the TX FIFO with the first two bytes to be
transmitted in the event the slave is addressed and a transmit
is required
*****/
R3=0xB537(Z);
```

```
W[P1 + LO(TWI_XMT_DATA16)] = R3;
```

```

/*****
FIFO Control determines whether an interrupt is generated
for every byte transferred or for every two bytes.
A value of zero which is the default, allows for single byte
events to generate interrupts
*****/

```

```
*****/
```

```
R1 = 0;
```

```
W[P1 + LO(TWI_FIFO_CTL)] = R1;
```

```

/*****
enable these signals to generate a TWI interrupt
*****/

```

```
*****/
```

```
R1 = RCVSERV | XMTSERV | SOVF | SERR | SCOMP | SINIT (z);
```

```
W[P1 + LO(TWI_INT_MASK)] = R1;
```

```

/*****
Enable the TWI Slave

```

```
Program the Slave Control register with:
```

1. Slave transmit data valid (STDVAL) set so that the contents of the TX FIFO can be used by this slave when a master requests data from it.

2. Slave Enable SEN to enable Slave functionality

```
*****/
```

```
R1 = STDVAL | SEN;
```

```
W[P1 + LO(TWI_SLAVE_CTL)] = R1;
```

```
TWI_SLAVE_INIT.END:
```

```
P2.H = HI(TWI_RX);
```

```
P2.L = LO(TWI_RX);
```

```
P4.H = HI(TWI_TX);
```

```
P4.L = LO(TWI_TX);
```

Programming Examples

```

/*****
Remap the vector table pointer from the default __I10HANDLER
to the new _TWI_ISR interrupt service routine
*****/
R1.H = HI(_TWI_ISR);
R1.L = LO(_TWI_ISR);
[P0 + LO(EVT10)] = R1; /* note that P0 points to the base of
the core MMR registers */

/*****
ENABLE TWI generate to interrupts at the system level
*****/
R1 = [P1 + LO(SIC_IMASK)];
BITSET(R1,BITPOS(IRQ_TWI));
[P1 + LO(SIC_IMASK)] = R1;

/*****
ENABLE TWI to generate interrupts at the core level
*****/
R1 = [P0 + LO(IMASK)];
BITSET(R1,BITPOS(EVT_IVG10));
[P0 + LO(IMASK)] = R1;

/*****
wait for interrupts
*****/
idle;

_main.END:

```

Listing 16-3. TWI Slave Interrupt Service Routine

```

/*****
Function:  _TWI_ISR
Description: This ISR is executed when the TWI controller
detects a slave initiated transfer. After an interrupt is ser-
viced, its corresponding bit is cleared in the TWI_INT_STAT
register. This done by writing a 1 to the particular bit posi-
tion. All bits are write 1 to clear.
*****/
#include <defBF527.h>
/*BF527 is used here as an example—change as appropriate.*/

.GLOBAL _TWI_ISR;

.section L1_code;
_TWI_ISR:

/*****
read the source of the interrupt
*****/
R1 = W[P1 + LO(TWI_INT_STAT)](z);

/*****
Slave Transfer Initiated
*****/
CC = BITTST(R1, BITPOS(SINIT));
if ! CC JUMP RECEIVE;
R0 = SINIT (Z);
W[P1 + LO(TWI_INT_STAT)] = R0; /* clear interrupt source bit */
ssync;

/*****
Receive service

```

Programming Examples

```
*****/
RECEIVE:
CC = BITTST(R1, BITPOS(RCVSERV));
if ! CC JUMP TRANSMIT;
R0 = W[P1 + LO(TWI_RCV_DATA8)] (Z); /* read data */
B[P2++] = R0; /* store bytes into a buffer pointed to by P2 */
R0 = RCVSERV(Z);
W[P1 + LO(TWI_INT_STAT)] = R0; /*clear interrupt source bit */
ssync;
JUMP _TWI_ISR.END; /* exit */

/*****
Transmit service
*****/
TRANSMIT:
CC = BITTST(R1, BITPOS(XMTSERV));
if ! CC JUMP SlaveError;
R0 = B[P4++](Z);
W[P1 + LO(TWI_XMT_DATA8)] = R0;
R0 = XMTSERV(Z);
W[P1 + LO(TWI_INT_STAT)] = R0; /* clear interrupt source bit */
ssync;
JUMP _TWI_ISR.END; /* exit */

/*****
slave transfer error
*****/
SlaveError:
CC = BITTST(R1, BITPOS(SERR));
if ! CC JUMP SlaveOverflow;
R0 = SERR(Z);
W[P1 + LO(TWI_INT_STAT)] = R0; /* clear interrupt source bit */
ssync;
```

```

JUMP _TWI_ISR.END;    /* exit */

/*****
slave overflow
*****/
SlaveOverflow:
CC = BITTST(R1, BITPOS(SOVF));
if !CC JUMP SlaveTransferComplete;
R0 = SOVF(Z);
W[P1 + LO(TWI_INT_STAT)] = R0;    /* clear interrupt source bit */
ssync;
JUMP _TWI_ISR.END;    /* exit */

/*****
slave transfer complete
*****/
SlaveTransferComplete:
CC = BITTST(R1, BITPOS(SCOMP));
if !CC JUMP _TWI_ISR.END;
R0 = SCOMP(Z);
W[P1 + LO(TWI_INT_STAT)] = R0;    /* clear interrupt source bit */
ssync;
/* Transfer complete read receive FIFO buffer and set/clear sema-
phores etc.... */
R0 = W[P1 + LO(TWI_FIFO_STAT)](z);
CC = BITTST(R0, BITPOS(RCV_HALF));    /* BIT 2 indicates whether
there's a byte in the FIFO or not */
if !CC JUMP _TWI_ISR.END;
R0 = W[P1 + LO(TWI_RCV_DATA8)](Z);    /* read data */
B[P2++] = R0;    /* store bytes into a buffer pointed to by P2 */

_TWI_ISR.END:RTI;

```

Electrical Specifications

All logic complies with the Electrical Specification outlined in the *Philips I²C Bus Specification version 2.1* dated January 2000.

Unique Information for the ADSP-BF51x Processor

None.

17 SPI-COMPATIBLE PORT CONTROLLER

This chapter describes the serial peripheral interface (SPI) port. Following an overview and a list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of SPIs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For SPI DMA channel assignments, refer to [Table 6-7 on page 6-104 in Chapter 6, “Direct Memory Access”](#).

For SPI interrupt vector assignments, refer to [Table 5-3 on page 5-19 in Chapter 5, “System Interrupts”](#).

To determine how each of the SPIs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5 through Table 9-4 on page 9-7 in Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each SPI, refer to [Appendix A, “System MMR Assignments”](#).

SPI behavior for the ADSP-BF51x that differs from the general information in this chapter can be found in the section [“Unique Information for the ADSP-BF51x Processor” on page 17-52](#).

Overview

The SPI port 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 signals, a device select signal, and a clock signal. SPI is a full-duplex synchronous serial interface, supporting master modes, slave modes, and multimaster environments. The SPI-compatible peripheral implementation also supports programmable bit rate and clock phase/polarities. The SPI features the use of open drain drivers to support the multimaster scenario and to avoid data contention.

Features

The SPI includes these features:

- Full duplex, synchronous serial interface
- Supports 8- or 16-bit word sizes
- Programmable baud rate, clock phase, and polarity
- Supports multimaster environments
- Integrated DMA controller
- Double-buffered transmitter and receiver
- One SPI device select input and multiple chip select outputs
- Programmable shift direction of MSB or LSB first
- Interrupt generation on mode fault, overflow, and underflow
- Shadow register to aid debugging

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

Interface Overview

[Figure 17-1](#) 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.

Interface Overview

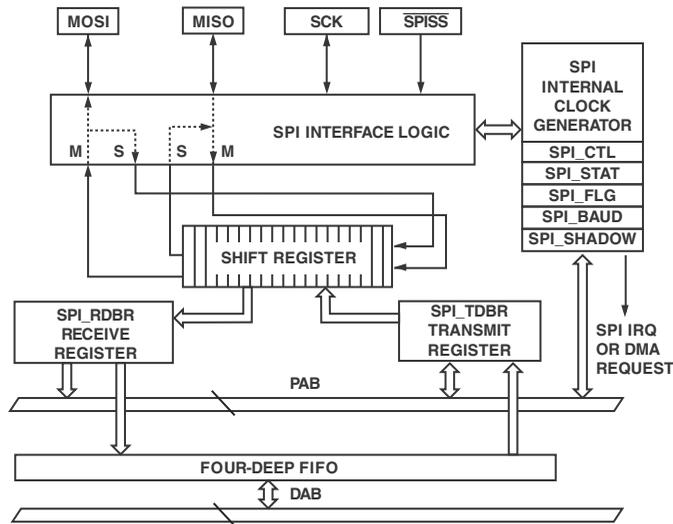


Figure 17-1. SPI Block Diagram

External Interface

The following sections describe the external interface of the SPI.

SPI Clock Signal (SCK)

The SCK signal is the serial 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 bit 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 SPISS input 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_CTL` register and define the transfer format.

Master-Out, Slave-In (MOSI) Signal

The master-out, slave-in (`MOSI`) signal is one of the bidirectional I/O data pins. If the processor is configured as a master, the `MOSI` pin transmits data out. If the processor is configured as a slave, the `MOSI` pin receives data in. 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) Signal

The master-in, slave-out (`MISO`) signal is one of the bidirectional I/O data pins. If the processor is configured as a master, the `MISO` pin receives data in. If the processor is configured as a slave, the `MISO` pin transmits data out. 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 17-2](#) 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 through its SPI interface to allow user application code and data to be downloaded before runtime.

Interface Overview

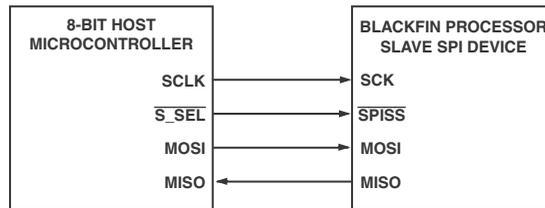


Figure 17-2. Blackfin Processor as Slave SPI Device

SPI Slave Select Input Signal (SPISS)

The SPISS signal is the SPI 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 a multimaster environment. In multimaster mode, if the 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 enable lead time (T_1), the enable lag time (T_2), and the sequential transfer delay time (T_3) each must always be greater than or equal to one-half the SCK period. See [Figure 17-3](#). The minimum time between successive word transfers (T_4) 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 is inactive (high) for at least one-half the SCK period. In this case, T_1 and T_2 will each always be equal to one-half the SCK period.

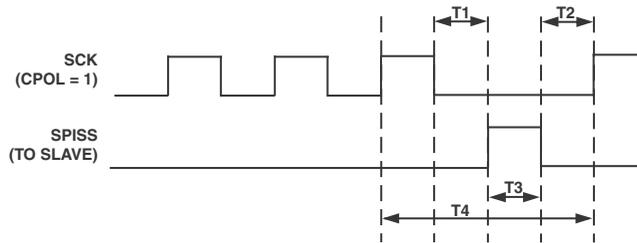


Figure 17-3. SPI Timing

SPI Slave Select Enable Output Signals

When operating in master mode, Blackfin processors may use any GPIO pin to enable individual SPI slave devices by software. In addition, the SPI module provides hardware support to generate up to seven slave select enable signals automatically (depending upon the configuration of the specific processor). See [Figure 17-14 on page 17-37](#) for details.

These signals are always active low in the SPI protocol. Since the respective pins are not driven during reset, it is recommended to pull them up by a resistor.

If enabled as a master, the SPI uses the `SPI_FLG` register to enable general-purpose port pins to be used as individual slave select lines. Before manipulating this register, the port pins that are to be used as SPI slave-select outputs must first be configured as such. To work as SPI output pins, the port pins must be enabled for use by SPI in the appropriate `PORT_MUX` register.

In slave mode, the `SPI_FLG` bits have no effect, and each SPI uses the `SPISS` input as a slave select. Just as in the master mode case, the port pin associated with `SPISS` must first be configured appropriately before use. [Figure 17-14 on page 17-37](#) shows the `SPI_FLG` register diagram.

Interface Overview

Slave Select Inputs

If the SPI is in slave mode, `SPISS` acts as the slave select input. When enabled as a master, `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 `SPISS` input is the master mode error input. Otherwise, `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 master processor equipped with seven slave selects, 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 port pins that can be configured as SPI master mode slave-select output pins can be connected to each of the slave SPI device's `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 through 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 17-4 shows one processor as a master with three processors (or other SPI-compatible devices) as slaves.

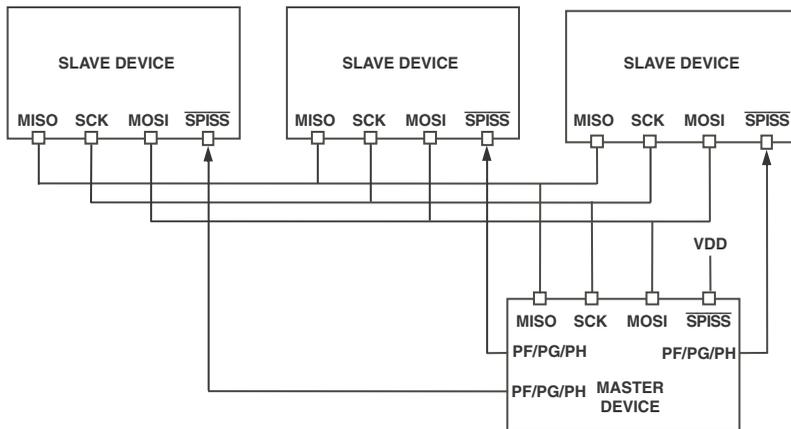


Figure 17-4. Single-Master, Multiple-Slave Configuration

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.

i The `SPIF` bit in the `SPI_STAT` register 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 in the `SPI_STAT` register are initially cleared upon entering DMA mode.

Interface Overview

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 two 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.

Internal Interfaces

The SPI has dedicated connections to the processor's peripheral bus (PAB) and DAB.

The low-latency PAB bus is used to map the SPI resources into the system MMR space. For PAB accesses to SPI MMRs, the primary performance criteria is latency, not throughput. Transfer latencies for both read and write transfers on the peripheral bus are two `SCLK` cycles.

The DAB bus provides a means for DMA SPI transfers to gain access to on-chip and off-chip memory with little or no degradation in core bandwidth to memory. The SPI peripheral, as a DMA master, is capable of sourcing DMA accesses. The arbitration policy for access to the DAB is described in the *Chip Bus Hierarchy* chapter.

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 DAB.

i 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 two 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.

Description of Operation

The following sections describe the operation of the SPI.

SPI Transfer Protocols

The SPI protocol supports four different combinations of serial clock phase and polarity (SPI modes 0, 1, 2, 3). These combinations are selected using the `CPOL` and `CPHA` bits in `SPI_CTL` as shown in [Figure 17-5](#).

[Figure 17-6](#) and [Figure 17-7](#) 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 `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.

Description of Operation

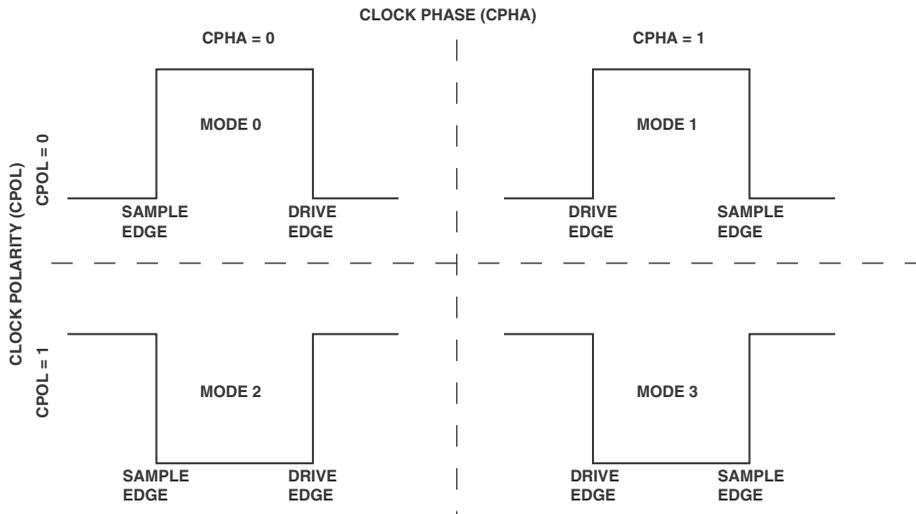


Figure 17-5. SPI Modes of Operation

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, $SPISS$, must be inactive (high) between each serial transfer. This is controlled automatically by the SPI hardware logic. When $CPHA = 1$, $SPISS$ may either remain active (low) between successive transfers or be inactive (high). This must be controlled by the software through manipulation of the SPI_FLG register.

Figure 17-6 shows the SPI transfer protocol for $CPHA = 0$. Note SCK starts toggling in the middle of the data transfer, $SIZE = 0$, and $LSBF = 0$.

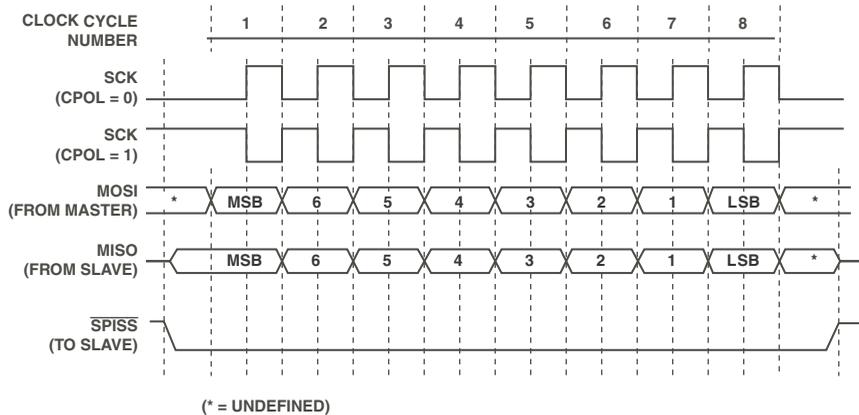


Figure 17-6. SPI Transfer Protocol for $CPHA = 0$

Figure 17-7 shows the SPI transfer protocol for $CPHA = 1$. Note SCK starts toggling at the beginning of the data transfer, $SIZE = 0$, and $LSBF = 0$.

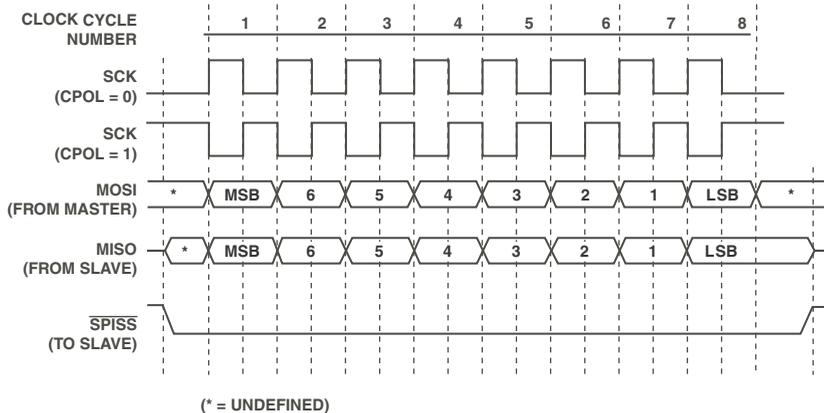


Figure 17-7. SPI Transfer Protocol for $CPHA = 1$

Description of Operation

SPI General Operation

The SPI can be used in single master as well as multimaster environments. 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 in the `SPI_CTL` register 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.

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 (`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, 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 through 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 up to seven programmable flags that are dedicated SPI slave select signals for the SPI slave devices.



At reset, the SPI is disabled and configured as a slave.

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 in the `SPI_BAUD` register. 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 serial 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

Functional Description

clock phase relative to data are programmable in the `SPI_CTL` register and define the transfer format. See [Figure 17-5 on page 17-12](#).

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 `TIMOD` field in the `SPI_CTL` register. In DMA mode (`TIMOD = b#1X`), the data interrupt acts as a DMA request and is generated when the DMA FIFO is ready to be written to (`TIMOD = b#11`) or read from (`TIMOD = b#10`). In non-DMA mode (`TIMOD = 0X`), a data interrupt is generated when the `SPI_TDBR` register is ready to be written to (`TIMOD = b#01`) or when the `SPI_RDBR` register is ready to be read from (`TIMOD = b#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 = b#11`) or an overflow (`RBSY` when `TIMOD = b#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 \(`SPI_CTL`\) Register” on page 17-34](#).

Functional Description

The following sections describe the functional operation of the SPI.

Master Mode Operation (Non-DMA)

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 the appropriate port register(s) to properly configure the SPI interface for master mode operation. The required pins are configured for SPI use as slave-select outputs.
2. 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.
3. 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.
4. If the `CPHA` bit in the `SPI_CTL` register = 1, the core activates the desired slaves by clearing one or more of the SPI flag bits (`FLGx`) of `SPI_FLG`.
5. 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 `SPI_TDBR` register or a data read of the `SPI_RDBR` register.
6. 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 the `SPI_RDBR` register.
7. With each new transfer initiate command, the SPI continues to send and receive words, according to the SPI transfer initiate mode.

See [Figure 17-8 on page 17-29](#) for additional information.

Functional Description

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 zeros 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` register. 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 the `SPI_RDBR` register or a write to the `SPI_TDBR` register. This is summarized in [Table 17-1](#).

 If the SPI port is enabled with `TIMOD = b#01` or `TIMOD = b#11`, the hardware immediately issues a first interrupt or DMA request.

Table 17-1. Transfer Initiation

TIMOD	Function	Transfer Initiated Upon	Action, Interrupt
b#00	Transmit and receive	Initiate new single word transfer upon read of SPI_RDBR and previous transfer completed.	Interrupt is active when the receive buffer is full. Read of SPI_RDBR clears interrupt.
b#01	Transmit and receive	Initiate new single word transfer upon write to SPI_TDBR and previous transfer completed.	Interrupt is active when the transmit buffer is empty. Writing to SPI_TDBR clears interrupt.
b#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.	Request DMA reads as long as the SPI DMA FIFO is not empty.
b#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.	Request DMA writes as long as the SPI DMA FIFO is not full.

Slave Mode Operation (Non-DMA)

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 `SPISS` select signal to the active state (low), or by the first active edge of the clock (`SCK`), depending on the state of the `CPHA` bit in the `SPI_CTL` register.

Functional Description

These steps illustrate SPI operation in the slave mode:

1. The core writes to the appropriate port register(s) to properly configure the SPI for slave mode operation.
2. The core writes to `SPI_CTL` to define the mode of the serial link to be the same as the mode set up in the SPI master.
3. To prepare for the data transfer, the core writes data to be transmitted into `SPI_TDBR`.
4. Once the `SPISS` falling edge is detected, the slave starts shifting data out on `MISO` and in from `MOSI` on `SCK` edges, depending upon the states of `CPHA` and `CPOL`.
5. Reception/transmission continues until `SPISS` is released or until the slave has received the proper number of clock cycles.
6. The slave device continues to receive/transmit with each new falling edge transition on `SPISS` and/or `SCK` clock edge.

See [Figure 17-8 on page 17-29](#) for additional information.

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 zeros 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 the `SPI_RDBR` register. If `GM = 0` and the receive buffer is full, the incoming data is discarded, and the `SPI_RDBR` register is not updated.

Slave Ready for a Transfer

When a device is enabled as a slave, the actions shown in [Table 17-2](#) are necessary to prepare the device for a new transfer.

Table 17-2. Transfer Preparation

TIMOD	Function	Action, Interrupt
b#00	Transmit and receive	Interrupt is active when the receive buffer is full. Read of SPI_RDBR clears interrupt.
b#01	Transmit and receive	Interrupt is active when the transmit buffer is empty. Writing to SPI_TDBR clears interrupt.
b#10	Receive with DMA	Request DMA reads as long as SPI DMA FIFO is not empty.
b#11	Transmit with DMA	Request DMA writes as long as SPI DMA FIFO is not full.

Programming Model

The following sections describe the SPI programming model.

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, which `CPHA` mode is selected, and which 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 `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

Programming Model

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 = b#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 (`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 = b#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 in the `SPI_CTL` register 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 that the SPI port be disabled 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 = b#00` or `TIMOD = b#10` mode. With `TIMOD = b#01` or `TIMOD = b#11`, the interrupt is requested while the transfer is still in progress.

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 core writes to the appropriate port register(s) to properly configure the SPI for master mode operation. The appropriate pins can be configured for SPI use as slave-select outputs.
2. 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 the *Direct Memory Access* chapter.

Programming Model

3. The processor core writes to the `SPI_FLG` register, setting one or more of the SPI flag select bits (`FLSx`).
4. 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 = b#10`) or “transmit with DMA” (`TIMOD = b#11`) mode.
5. 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.

6. 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.
7. 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.

See [Figure 17-9 on page 17-30](#) for additional information.

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 in the `SPI_CTL` register. 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 zeros 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

Programming Model

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.

DMA requests are generated when the DMA FIFO is not empty (when `TIMOD = b#10`), or when the DMA FIFO is not full (when `TIMOD = b#11`).

Error interrupts are generated when there is an `RBSY` overflow error condition (when `TIMOD = b#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 `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 core writes to the appropriate port register(s) to properly configure the SPI for slave mode operation.
2. 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 the *Direct Memory Access* chapter.
3. The processor core writes to the `SPI_CTL` register to define the mode of the serial link to be the same as the mode set up in the SPI master. The `TIMOD` field will be configured to select either “receive with DMA” (`TIMOD = b#10`) or “transmit with DMA” (`TIMOD = b#11`) mode.

4. 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.

5. 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.

See [Figure 17-9 on page 17-30](#) for additional information.

Programming Model

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 in the SPI_CTL register. 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 zeros 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 zeros 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.

DMA requests are generated when the DMA FIFO is not empty (when TIMOD = b#10), or when the DMA FIFO is not full (when TIMOD = b#11).

Error interrupts are generated when there is an RBSY overflow error condition (when TIMOD = b#10), or when there is a TXE underflow error condition (when TIMOD = b#11).

SPI-Compatible Port Controller

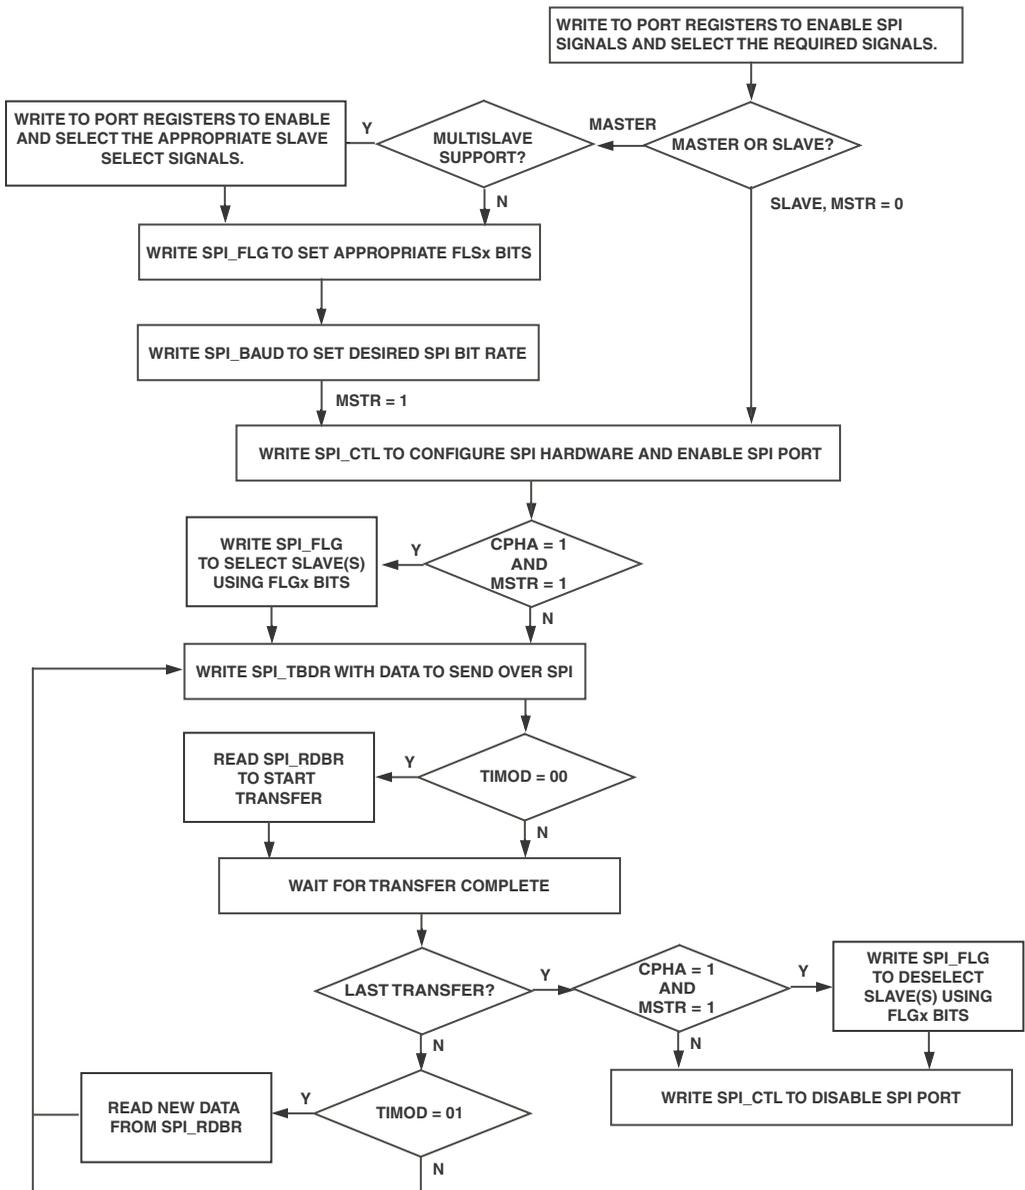


Figure 17-8. Core-Driven SPI Flow Chart

Programming Model

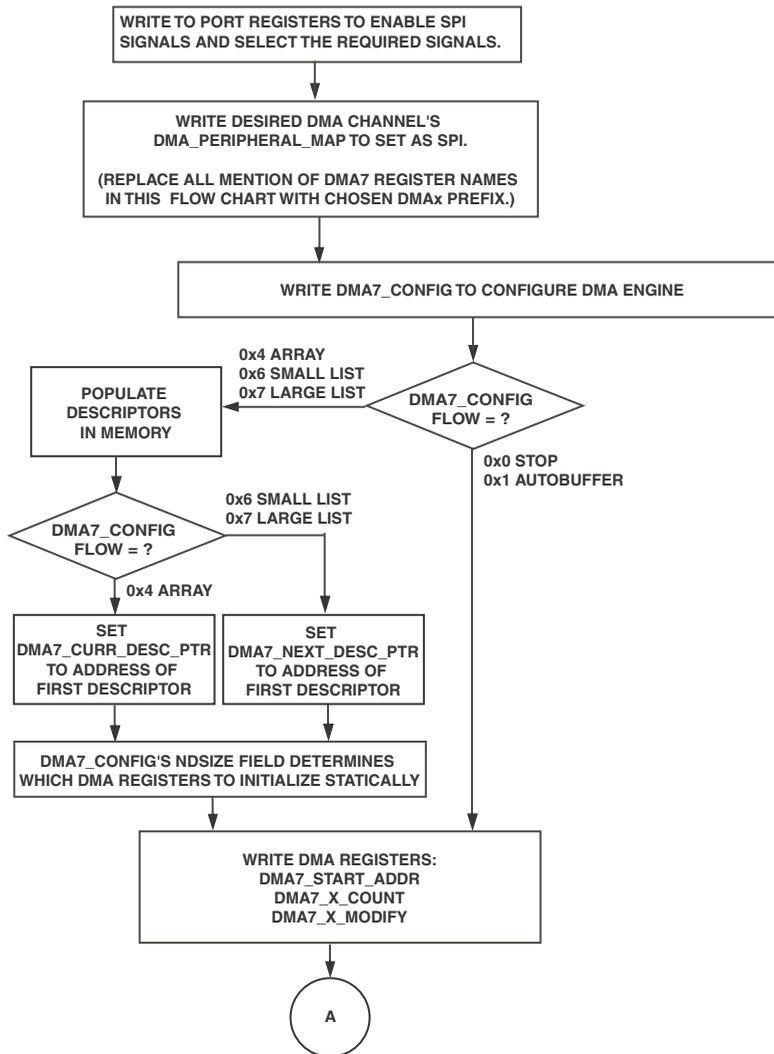


Figure 17-9. SPI DMA Flow Chart (Part 1 of 3)

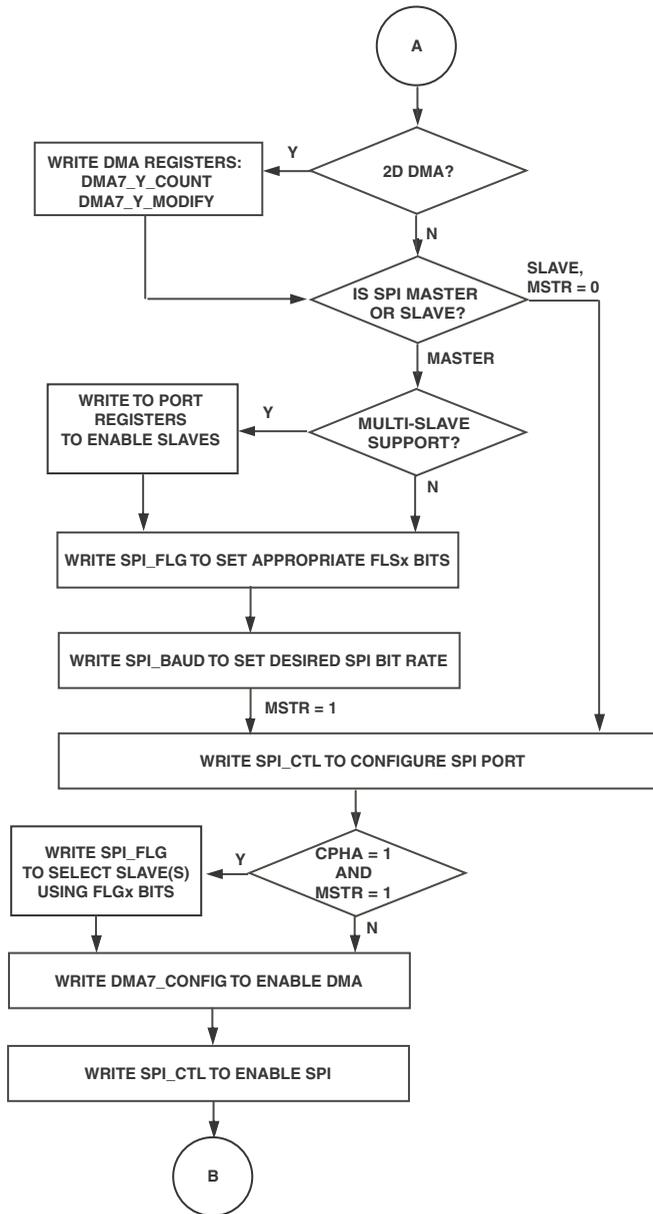


Figure 17-10. SPI DMA Flow Chart (Part 2 of 3)

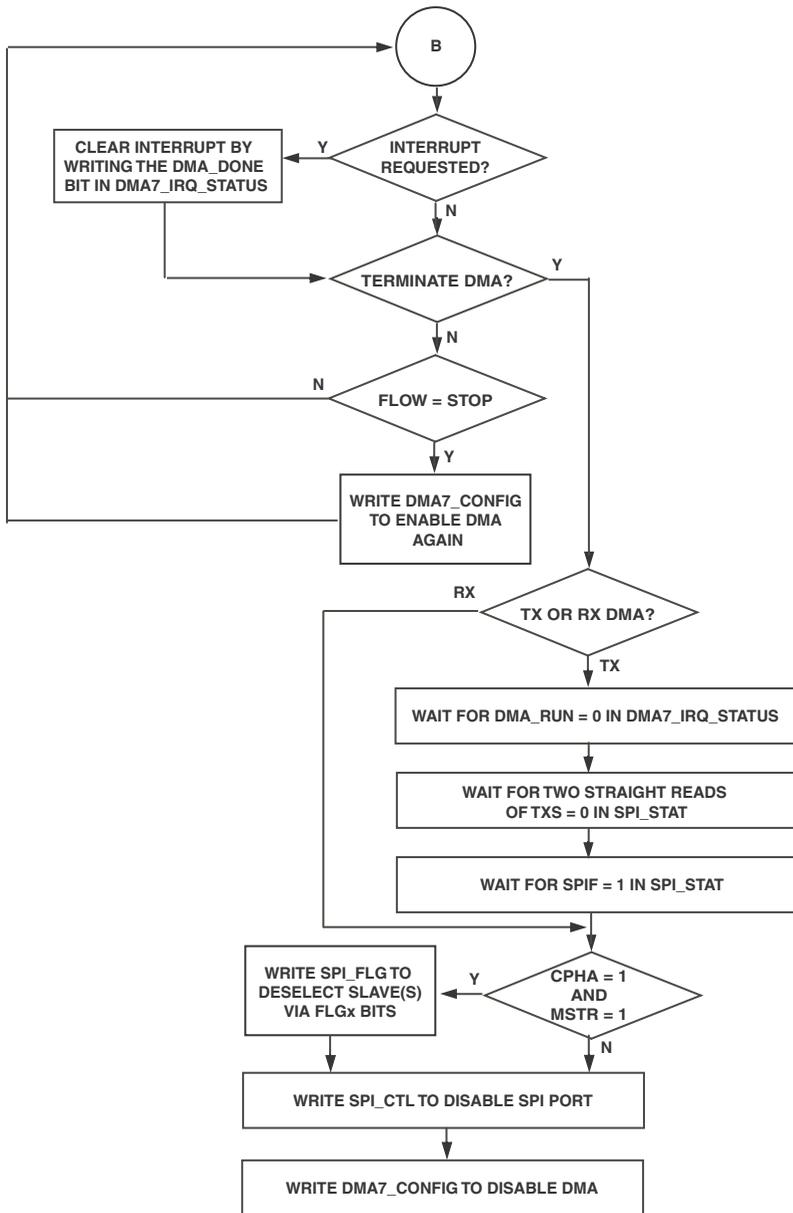


Figure 17-11. SPI DMA Flow Chart (Part 3 of 3)

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. The shift register, SFDR, is internal to the SPI module and is not directly accessible.

Table 17-3 shows the functions of the SPI registers. Figure 17-12 through Figure 17-18 on page 17-43 provide details.

Table 17-3. SPI Register Mapping

Register Name	Function	Notes
SPI_BAUD	SPI port baud control	Value of “0” or “1” disables the serial clock
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 can be triggered
SPI_SHADOW	SPI port data	Register has the same contents as SPI_RDBR, but no action is taken when it is read

SPI Baud Rate (SPI_BAUD) Register

The SPI_BAUD register is used to set the bit transfer rate for a master device. When configured as a slave, the value written to this register is ignored.

SPI Registers

The serial clock frequency is determined by this formula:

$$\text{SCK frequency} = (\text{peripheral clock frequency } \text{SCLK}) / (2 \times \text{SPI_BAUD})$$

Writing a value of “0” or “1” to the register disables the serial clock. Therefore, the maximum serial clock rate is one-fourth the system clock rate.

Table 17-4 lists several possible baud rate values for SPI_BAUD.

Table 17-4. 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 Baud Rate Register (SPI_BAUD)

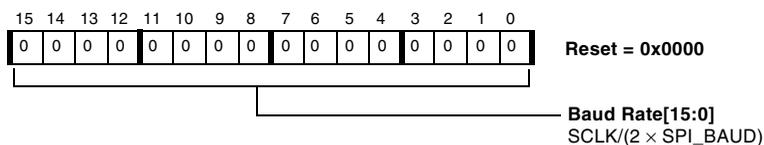


Figure 17-12. SPI Baud Rate Register

SPI Control (SPI_CTL) Register

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 `b#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 `b#01`, the transaction is initiated when the transmit buffer is written. A value of `b#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` register. A value of `11` selects DMA transmit mode and the transaction is initiated by a DMA write of the `SPI_TDBR` register.

The `PSSE` bit is used to enable the `SPISS` input for an external master. When not used, `SPISS` can be disabled, freeing up a pin for an alternate function.

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 `SPI_STAT` register is set. See [“Mode Fault Error \(MODF\)” on page 17-40](#).

SPI Registers

Figure 17-13 provides the bit descriptions for SPI_CTL.

SPI Control Register (SPI_CTL)

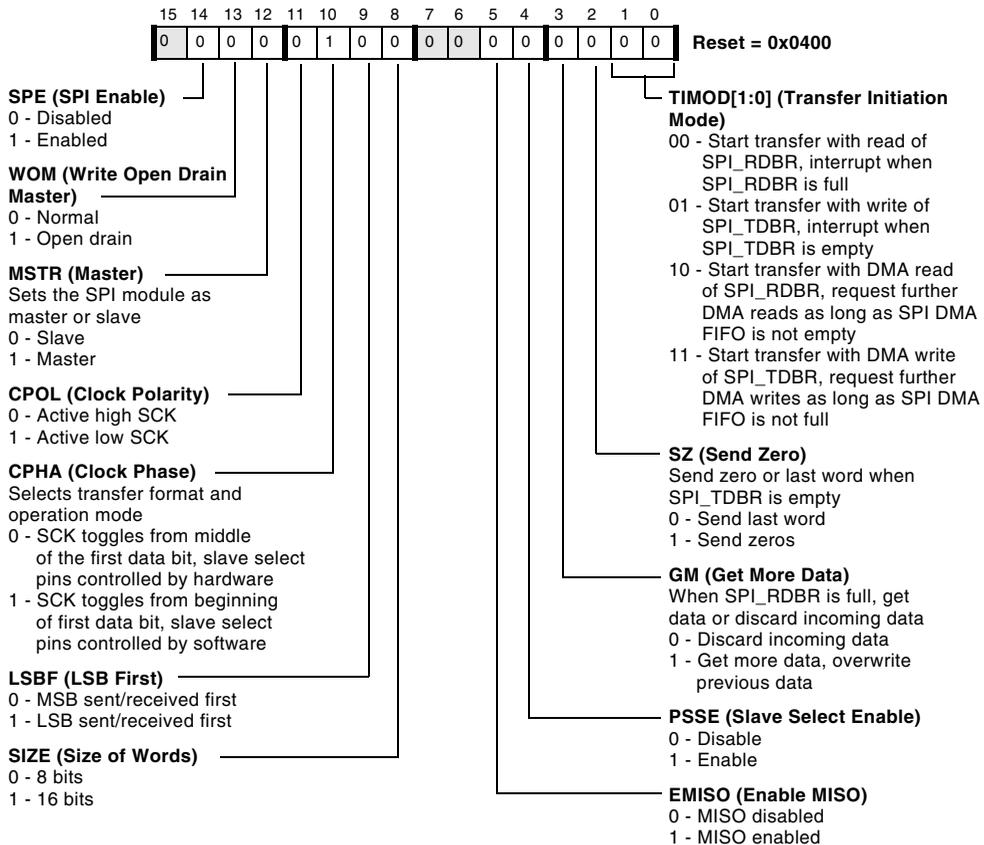


Figure 17-13. SPI Control Register

SPI Flag (SPI_FLG) Register

The SPI_FLG register consists of two sets of bits that function as follows.

SPI Flag Register (SPI_FLG)

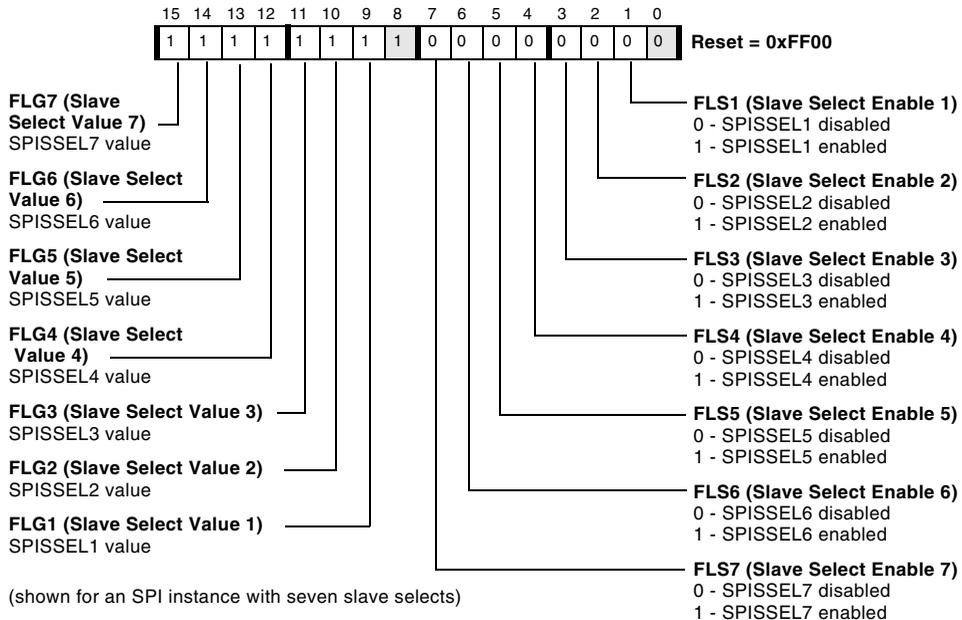


Figure 17-14. SPI Flag Register (example with 7 slave selects)

- Slave select enable (FLS_x) bits

Each FLS_x bit corresponds to a general purpose port pin. When an FLS_x bit is set, the corresponding port pin is driven as a slave select. For example, if FLS1 is set in SPI_FLG, the port pin corresponding to SPISEL1 is driven as a slave select.

SPI Registers

If the `FLSx` bit is not set, the general-purpose port registers configure and control the corresponding port pins.

- Slave select value (`FLGx`) bits

When a port 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, setting `FLS3` in the `SPI_FLG` register drives the `SPISSSEL3` pin as a slave select. Then, clearing `FLG3` in the `SPI_FLG` register drives the pin low, and setting `FLG3` drives it high. The pin can be cycled high and low between transfers by setting and clearing `FLG3`. Otherwise, the pin 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 the slave select function on a port pin to which it is mapped, it is only necessary to set the appropriate `FLS` bit in `SPI_FLG`. It is not necessary to write to an `FLG` bit, because the SPI hardware automatically drives the port pin.

SPI Status (SPI_STAT) Register

The SPI_STAT register is used to detect when an SPI transfer is complete or if transmission/reception errors occur. The SPI_STAT register can be read at any time.

SPI Status Register (SPI_STAT)

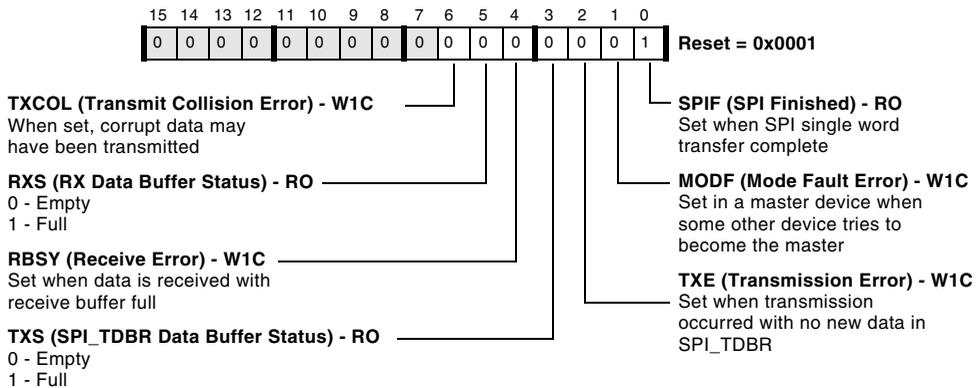


Figure 17-15. SPI Status Register

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.

See [Figure 17-15 on page 17-39](#) for more information.

Mode Fault Error (MODF)

The `MODF` bit is set in `SPI_STAT` when the `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 `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 general-purpose I/O port 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 general-purpose I/O port registers appropriately.

When enabling the `MODF` feature, the program must configure as inputs all of the port pins that will be used as slave selects. Programs can do this by

configuring the direction of the port pins prior to configuring the SPI. This ensures that, once the `MODF` error occurs and the slave selects are automatically reconfigured as port 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 by 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`).

SPI Transmit Data Buffer (SPI_TDBR) Register

The `SPI_TDBR` register is a 16-bit read-write register. 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 the internal shift register `SFDR`. A read of `SPI_TDBR` can occur at any time and does not interfere with or initiate SPI transfers.

SPI Registers

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 `SZ` control bit in the `SPI_CTL` register is set, `SPI_TDBR` may be reset to zero 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.

SPI Transmit Data Buffer Register (SPI_TDBR)

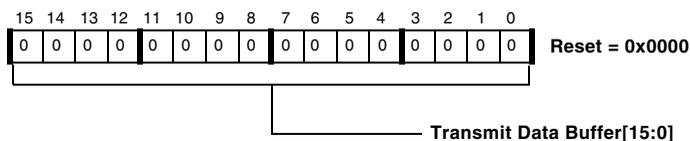


Figure 17-16. SPI Transmit Data Buffer Register

SPI Receive Data Buffer (SPI_RDBR) Register

The `SPI_RDBR` register is a 16-bit read-only register. 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 controller. When `SPI_RDBR` is read by software, the `RXS` bit in the `SPI_STAT` register is cleared and an SPI transfer may be initiated (if `TIMOD = b#00`).

SPI Receive Data Buffer Register (SPI_RDBR)

Read Only

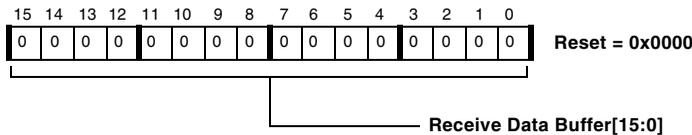


Figure 17-17. SPI Receive Data Buffer Register

SPI RDBR Shadow (SPI_SHADOW) Register

The SPI_SHADOW register is provided for use in debugging software. This register is at a different address than the receive data buffer, 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 = b#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)

Read Only

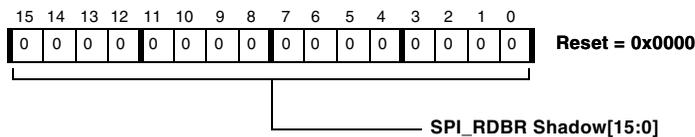


Figure 17-18. SPI RDBR Shadow Register

Programming Examples

This section includes examples ([Listing 17-1](#) through [Listing 17-8 on page 17-51](#)) for both core-generated and DMA-based transfers. Each code example assumes that the appropriate processor header files are included.

Core-Generated Transfer

The following core-driven master-mode SPI example shows how to initialize the hardware, signal the start of a transfer, handle the interrupt and issue the next transfer, and generate a stop condition.

Initialization Sequence

Before the SPI can transfer data, the registers must be configured as follows.

Listing 17-1. SPI Register Initialization

```
SPI_Register_Initialization:
    P0.H = hi(SPI_FLG);
    P0.L = lo(SPI_FLG);
    R0 = W[P0] (Z);
    BITSET (R0,0x7);      /* FLS7 */
    W[P0] = R0;          /* Enable slave-select output pin */

    P0.H = hi(SPI_BAUD);
    P0.L = lo(SPI_BAUD);
    R0.L = 0x208E;       /* Write to SPI Baud rate register */
    W[P0] = R0.L; ssync; /* If SCLK = 133 MHz, SPI clock ~= 8 kHz
*/

/* Setup SPI Control Register */
/*****
* TIMOD [1:0] = 00 : Transfer On RDBR Read.
* SZ [2]      = 0 : Send Last Word When TDBR Is Empty
* GM [3]      = 1 : Overwrite Previous Data If RDBR Is Full
* PSSE [4]    = 0 : Disables Slave-Select As Input (Master)
* EMISO [5]   = 0 : MISO Disabled For Output (Master)
* [7] and [6] = 0 : RESERVED
* SIZE [8]    = 1 : 16 Bit Word Length Select
*****/
```

```

* LSBF [9]      = 0 : Transmit MSB First
* CPHA [10]     = 0 : Hardware Controls Slave-Select Outputs
* CPOL [11]     = 1 : Active Low SCK
* MSTR [12]     = 1 : Device Is Master
* WOM [13]      = 0 : Normal MOSI/MISO Data Output (No Open Drain)
* SPE [14]      = 1 : SPI Module Is Enabled
* [15]         = 0 : RESERVED
*****/
    P0.H = hi(SPI_CTL) ;
    P0.L = lo(SPI_CTL) ;
    R0 = 0x5908;
    W[P0] = R0.L; ssync;    /* Enable SPI as MASTER */

```

Starting a Transfer

After the initialization procedure in the given master mode, a transfer begins following a dummy read of `SPI_RDBR`. Typically, known data which is desired to be transmitted to the slave is preloaded into the `SPI_TDBR`. In the following code, `P1` is assumed to point to the start of the 16-bit transmit data buffer and `P2` is assumed to point to the start of the 16-bit receive data buffer. In addition, the user must ensure appropriate interrupts are enabled for SPI operation.

Listing 17-2. Initiate Transfer

```

Initiate_Transfer:
    P0.H = hi(SPI_FLG);
    P0.L = lo(SPI_FLG);
    R0 = W[P0] (Z);
    BITCLR (R0,0xF);    /* FLG7 */
    W[P0] = R0;        /* Drive 0 on enabled slave-select pin */

    P0.H = hi(SPI_TDBR); /* SPI Transmit Register */
    P0.L = lo(SPI_TDBR);

```

Programming Examples

```
R0 = W[P1++] (z);      /* Get First Data To Be Transmitted
And Increment Pointer */
W[P0] = R0;           /* Write to SPI_TDBR */

P0.H = hi(SPI_RDBR);
P0.L = lo(SPI_RDBR);
R0 = W[P0] (z); /* Dummy read of SPI_RDBR kicks off transfer */
```

Post Transfer and Next Transfer

Following the transfer of data, the SPI generates an interrupt, which is serviced if the interrupt is enabled during initialization. In the interrupt routine, software must write the next value to be transmitted prior to reading the byte received. This is because a read of the SPI_RDBR initiates the next transfer.

Listing 17-3. SPI Interrupt Handler

```
SPI_Interrupt_Handler:
Process_SPI_Sample:
    P0.H = hi(SPI_TDBR);      /* SPI transmit register */
    P0.L = lo(SPI_TDBR);
    R0 = W[P1++] (z);        /* Get next data to be transmitted */
    W[P0] = R0.L;           /* Write that data to SPI_TDBR */

Kick_Off_Next:
    P0.H = hi(SPI_RDBR);     /* SPI receive register */
    P0.L = lo(SPI_RDBR);
    R0 = W[P0] (z); /* Read SPI receive register (also kicks off
next transfer) */
    W[P2++] = R0;           /* Store received data to memory */
    RTI;                   /* Exit interrupt handler */
```

Stopping

In order for a data transfer to end after the user has transferred all data, the following code can be used to stop the SPI. Note that this is typically done in the interrupt handler to ensure the final data has been sent in its entirety.

Listing 17-4. Stopping SPI

```
Stopping_SPI:
    P0.H = hi(SPI_CTL);
    P0.L = lo(SPI_CTL);
    R0 = W[P0];
    BITCLR(R0, 14);          /* Clear SPI enable bit */
    W[P0] = R0.L; ssync;    /* Disable SPI */
```

DMA-Based Transfer

The following DMA-driven master-mode SPI autobuffer example shows how to initialize DMA, initialize SPI, signal the start of a transfer, and generate a stop condition.

DMA Initialization Sequence

The following code initializes the DMA to perform a 16-bit memory read DMA operation in autobuffer mode, and generates an interrupt request when the buffer has been sent. This code assumes that P1 points to the start of the data buffer to be transmitted and that NUM_SAMPLES is a defined macro indicating the number of elements being sent.

Listing 17-5. DMA Initialization

```
Initialize_DMA:    /* Assume DMA7 as the channel for SPI DMA */
    P0.H = hi(DMA7_CONFIG);
    P0.L = lo(DMA7_CONFIG);
```

Programming Examples

```
R0 = 0x1084(z); /* Autobuffer mode, IRQ on complete, linear
16-bit, mem read */
w[P0] = R0;

P0.H = hi(DMA7_START_ADDR);
P0.L = lo(DMA7_START_ADDR);
[p0] = p1; /* Start address of TX buffer */

P0.H = hi(DMA7_X_COUNT);
P0.L = lo(DMA7_X_COUNT);
R0 = NUM_SAMPLES;
w[p0] = R0; /* Number of samples to transfer */

R0 = 2;
P0.H = hi(DMA7_X_MODIFY);
P0.L = lo(DMA7_X_MODIFY);
w[p0] = R0; /* 2 byte stride for 16-bit words */

R0 = 1; /* single dimension DMA means 1 row */
P0.H = hi(DMA7_Y_COUNT);
P0.L = lo(DMA7_Y_COUNT);
w[p0] = R0;
```

SPI Initialization Sequence

Before the SPI can transfer data, the registers must be configured as follows.

Listing 17-6. SPI Initialization

```
SPI_Register_Initialization:
P0.H = hi(SPI_FLG);
P0.L = lo(SPI_FLG);
R0 = W[P0] (Z);
```

SPI-Compatible Port Controller

```
BITSET (R0,0x7);    /* FLS7 */
W[P0] = R0;        /* Enable slave-select output pin */

P1.H = hi(SPI_BAUD);
P1.L = lo(SPI_BAUD);
R0.L = 0x208E;     /* Write to SPI baud rate register */
W[P0] = R0.L; ssync; /* If SCLK = 133MHz, SPI clock ~= 8kHz */

/* Setup SPI Control Register */
/*****
* TIMOD [1:0] = 11 : Transfer on DMA TDBR write
* SZ [2]      = 0 : Send last word when TDBR is empty
* GM [3]      = 1 : Discard incoming data if RDBR is full
* PSSE [4]    = 0 : Disables slave-select as input (master)
* EMISO [5]   = 0 : MISO disabled for output (master)
* [7] and [6] = 0 : RESERVED
* SIZE [8]    = 1 : 16 Bit word length select
* LSBF [9]    = 0 : Transmit MSB first
* CPHA [10]   = 0 : SCK starts toggling at START of first data
bit
* CPOL [11]   = 1 : Active HIGH serial clock
* MSTR [12]   = 1 : Device is master
* WOM [13]    = 0 : Normal MOSI/MISO data output (no open
drain)
* SPE [14]    = 1 : SPI module is enabled
* [15]        = 0 : RESERVED
*****/
/* Configure SPI as MASTER */
R1 = 0x190B(z);    /* Leave disabled until DMA is enabled */
P1.L = lo(SPI_CTL);
W[P1] = R1; ssync;
```

Programming Examples

Starting a Transfer

After the initialization procedure in the given master mode, a transfer begins following enabling of SPI. However, the DMA must be enabled before enabling the SPI.

Listing 17-7. Starting a Transfer

```
Initiate_Transfer:
    P0.H = hi(DMA7_CONFIG);
    P0.L = lo(DMA7_CONFIG);
    R2 = w[P0](z);
    BITSET (R2, 0);    /*Set DMA enable bit */
    w[p0] = R2.L;     /* Enable TX DMA */

    P4.H = hi(SPI_CTL);
    P4.L = lo(SPI_CTL);
    R2=w[p4](z);
    BITSET (R2, 14); /* Set SPI enable bit */
    w[p4] = R2;     /* Enable SPI */
```

Stopping a Transfer

In order for a data transfer to end after the DMA has transferred all required data, the following code is executed in the SPI DMA interrupt handler. The example code below clears the DMA interrupt, then waits for the DMA engine to stop running. When the DMA engine has completed, `SPI_STAT` is polled to determine when the transmit buffer is empty. If there is data in the SPI Transmit FIFO, it is loaded as soon as the `TXS` bit clears. A second consecutive read with the `TXS` bit clear indicates the FIFO is empty and the last word is in the shift register. Finally, polling for the `SPIF` bit determines when the last bit of the last word has been shifted out. At that point, it is safe to shut down the SPI port and the DMA engine.

Listing 17-8. Stopping a Transfer

```

SPI_DMA_INTERRUPT_HANDLER:
    P0.L = lo(DMA7_IRQ_STATUS);
    P0.H = hi(DMA7_IRQ_STATUS);
    R0 = 1 ;
    W[P0] = R0 ; /* Clear DMA interrupt */

    /* Wait for DMA to complete */
    P0.L = lo(DMA7_IRQ_STATUS);
    P0.H = hi(DMA7_IRQ_STATUS);
    R0 = DMA_RUN; /* 0x08 */

CHECK_DMA_COMPLETE: /* Poll for DMA_RUN bit to clear */
    R3 = W[P0] (Z);
    R1 = R3 & R0;
    CC = R1 == 0;
    IF !CC JUMP CHECK_DMA_COMPLETE;

    /* Wait for TXS to clear */
    P0.L = lo(SPI_STAT);
    P0.H = hi(SPI_STAT);
    R1 = TXS; /* 0x08 */

Check_TXS: /* Poll for TXS = 0 */
    R2 = W[P0] (Z);
    R2 = R2 & R1;
    CC = R0 == 0;
    IF !CC JUMP Check_TXS;

    R2 = W[P0] (Z); /* Check if TXS stays clear for 2 reads */
    R2 = R2 & R1;
    CC = R0 == 0;

```

Unique Information for the ADSP-BF51x Processor

```
IF !CC JUMP Check_TXS;

/* Wait for final word to transmit from SPI */
Final_Word:
R0 = W[P0](Z);
R2 = SPIF; /* 0x01 */
R0 = R0 & R2;
CC = R0 == 0;
IF CC JUMP Final_Word;

Disable_SPI:
P0.L = lo(SPI_CTL);
P0.H = hi(SPI_CTL);
R0 = W[P0](Z);
BITCLR (R0,0xe); /* Clear SPI enable bit */
W[P0] = R0; /* Disable SPI */

Disable_DMA:
P0.L = lo(DMA7_CONFIG);
P0.H = hi(DMA7_CONFIG);
R0 = W[P0](Z);
BITCLR (R0,0x0); /* Clear DMA enable bit */
W[P0] = R0; /* Disable DMA */

RTI; /* Exit Handler */
```

Unique Information for the ADSP-BF51x Processor

None.

18 SPI SERIAL FLASH

The ADSP-BF512F/ADSP-BF514F/ADSP-BF516F/ADSP-BF518F processors contain a SPI flash memory (internal functional block diagram, see [Figure 18-1](#)) within the package of the processor and connected to `SPI0` (see [Figure 18-2](#)). See “[SPI0 and Internal Flash Usage](#)” in [Chapter 9, General-Purpose Ports](#) for information on configuring the ports to communicate with the SPI flash.

The SPI flash memory has a 4M bit capacity and 1.8 V (nominal) operating voltage. The program/erase endurance is 100,000 cycles per block, and this memory has greater than 100 years data retention capability. Also included are support for software write protection and support for fast erase and byte-program.

The processors internally connect to the flash memory die with the `SPIOSCK`, `SPI0SEL4` or `PH8`, `SPI0MOSI`, and `SPI0MISO` signals similar to an external SPI flash (for signal descriptions, see [Table 18-1](#)). To further provide a secure processing environment, these internally connected signals are not exposed outside of the package. For this reason, programming the ADSP-BF51xF flash memory is performed by running code on the processor. It cannot be programmed from external signals and data transfers between the SPI flash and the processor cannot be probed externally. The flash memory has the following additional features:

- Serial Interface Architecture—SPI compatible with mode 0 and mode 3
- Superior Reliability—endurance of 100,000 cycles and greater than 100 years data retention

- Flexible Erase Capability—uniform 4K byte sectors and uniform 32 and 64K byte overlay blocks
- Fast Erase and Byte-Program—chip-erase time=125 ms (typical), sector-/block-erase time=62 ms (typical), byte-program time=50 μ S (typical)
- Auto Address Increment (AAI) Programming—decreases total chip programming time over byte-program operations
- End-of-Write Detection—software polling the *BUSY* bit in status register, busy status readout on *S0* pin
- Software Write Protection—write protection through block-protection bits in status register

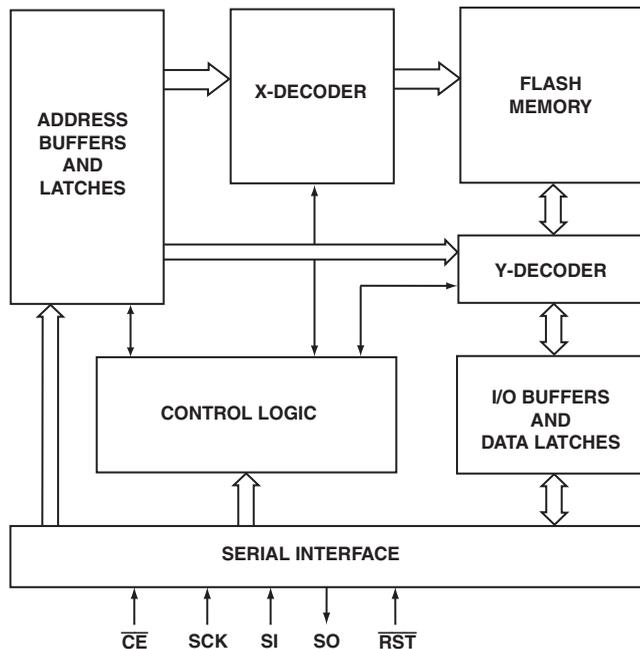


Figure 18-1. SPI Flash Internal Functional Block Diagram

Combinational Logic Truth Table	SEL4 or PH8	MISO_EXT, SPICLK_EXT, MOSI_EXT
	0	Three-state
	1	As programmed

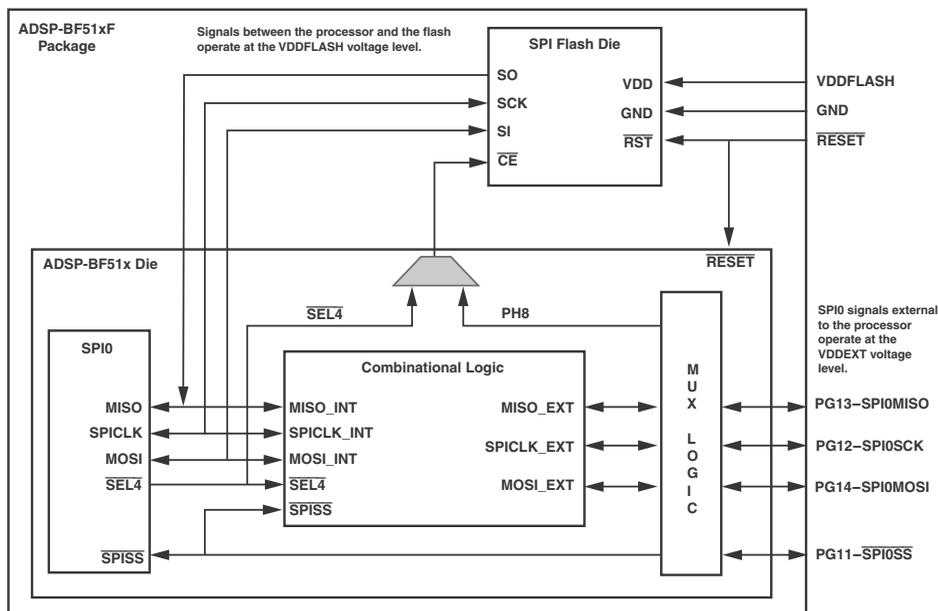


Figure 18-2. Flash Memory Block Diagram

Table 18-1. Internal Flash Memory Signal Descriptions

Symbol	Pin Name	Function
SCK	Serial Clock	To provide the timing of the serial interface. Commands, addresses, or input data are latched on the rising edge of the clock input, while output data is shifted out on the falling edge of the clock input.
SI	Serial Data Input	Transfers commands, addresses, or data serially into the device. Inputs are latched on the rising edge of the serial clock.
SO	Serial Data Output	Transfers data serially out of the device. Data is shifted out on the falling edge of the serial clock. Flash busy status pin in AAI mode if SO is configured as a hardware RY/BY pin.

Memory Organization

Table 18-1. Internal Flash Memory Signal Descriptions (Continued)

Symbol	Pin Name	Function
\overline{CE}	Chip Enable	The device is enabled by a high-to-low transition on \overline{CE} . \overline{CE} must remain low for the duration of any command sequence.
\overline{RST}	Reset	Resets the operation of the device and the internal logic. This signal is tied to the ADSP-BF51x \overline{RESET} signal.

Memory Organization

The memory array is organized in uniform 4K byte sectors with 16K byte, 32K byte, and 64K byte overlay erasable blocks.

Device Operation

The memory is accessed through the SPI (serial peripheral interface) bus compatible protocol. The SPI bus consist of four control lines; chip enable (\overline{CE}) is used to select the device, and data is accessed through the serial data input (SI), serial data output (SO), and serial clock (SCK).

The memory supports both mode 0 (0,0) and mode 3 (1,1) of SPI bus operations. The difference between the two modes, as shown in [Figure 18-3](#), is the state of the SCK signal when the bus master is in standby mode and no data is being transferred. The SCK signal is low for mode 0 and SCK signal is high for mode 3. For both modes, the serial data in (SI) is sampled at the rising edge of the SCK clock signal and the serial data output (SO) is driven after the falling edge of the SCK clock signal.

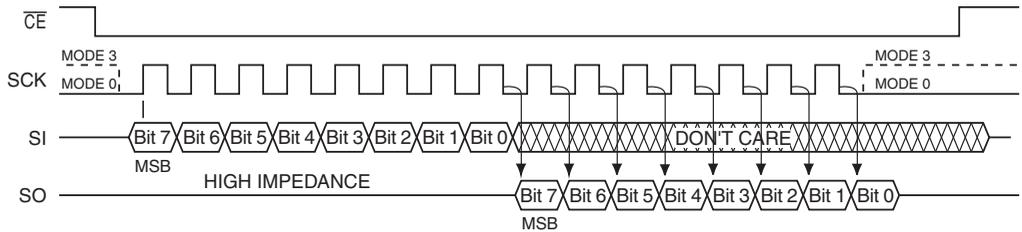


Figure 18-3. SPI Protocol

Reset Mode

Driving the \overline{RESET} pin low resets the flash device. Driving the \overline{RESET} pin high puts the device in normal operating mode. The S0 pin is in high impedance state while the device is in reset. The status register is reset to its power-up state upon a successful reset. See [Table 18-4](#) for default power-up modes. A device reset during an active program or erase operation aborts the operation and data of the targeted address range may be corrupted or lost due to the aborted erase or program operation. The device exits AAI programming mode in progress and places the S0 pin in high impedance state.

Reset Mode

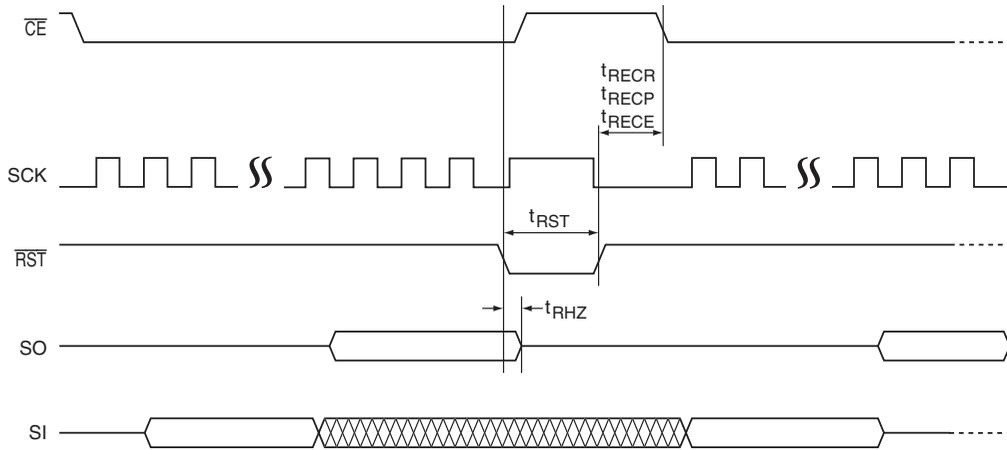


Figure 18-4. Reset Timing Diagram

Table 18-2. Reset Timing Parameters

Symbol	Parameter	Min	Max	Units
t_{RECR}	Reset Recovery from Read		100	ns
t_{RECP}	Reset Recovery from Program		10	μ s
t_{RECE}	Reset Recovery from Erase		1	ms

Table 18-3. Conditions to Execute Write-Status-Register (WRSR) Instruction

BPL	Execute WRSR Instruction
1	Not Allowed
0	Allowed

Status Register

The software status register provides status on whether the flash memory array is available for any read or write operation, whether the device is write enabled, and the state of the memory write protection. During an internal erase or program operation, the status register may be read only to determine the completion of an operation in progress. [Table 18-4](#) describes the function of each bit in the software status register.

Table 18-4. Software Status Register

Bit	Name	Function	Default at Power-up	Read/Write
0	BUSY	1 = Internal write operation is in progress 0 = No internal write operation is in progress	0	R
1	WEL	1 = Device is memory write-enabled 0 = Device is not memory write-enabled	0	R
2	BP0	Indicate current level of block write protection (See Table 18-5 .)	1	R/W
3	BP1	Indicate current level of block write protection (See Table 18-5 .)	1	R/W
4	BP2	Indicate current level of block write protection (See Table 18-5 .)	1	R/W
5	RES	Reserved for future use	0	N/A
6	AAI	Auto address increment programming status 1 = AAI programming mode 0 = Byte-program mode	0	R
7	BPL	1 = BP1 and BP0 are read-only bits 0 = BP1 and BP0 are read/writable	0	R/W

Busy

The `BUSY` bit determines whether there is an internal erase or program operation in progress. A '1' for the `BUSY` bit indicates the device is busy with an operation in progress. A '0' indicates the device is ready for the next valid operation.

Write Enable Latch (WEL)

The write-enable-latch bit indicates the status of the internal write-enable-latch memory. If the `WEL` bit is set to '1', it indicates the device is write-enabled. If the bit is set to '0' (reset), it indicates the device is not write-enabled and does not accept any write (program/erase) commands. The write-enable-latch bit is automatically reset under the following conditions:

- Device reset
- Power-up
- Write-disable (`WRDI`) instruction completion
- Byte-program instruction completion
- Auto address increment (`AAI`) programming is completed or reached its highest unprotected memory address
- Sector-erase instruction completion
- Block-erase instruction completion
- Chip-erase instruction completion
- Write-status-register instructions

Auto Address Increment (AAI)

The auto address increment programming status bit provides status on whether the device is in AAI programming mode or byte-program mode. The default at power-up is byte-program mode.

Block-Protection (BP2, BP1, BP0)

The block-protection (BP1, BP0) bits define the size of the memory area to be software-protected against any memory write (program or erase) operation. The write-status-register (WRSR) instruction is used to program the BP1 and BP0 bits as long as \overline{WP} is high or the block-protect-lock (BPL) bit is '0'. Chip-erase can only be executed if block-protection bits are all '0'. After power-up, BP2, BP1, and BP0 are set to defaults. See [Table 18-4](#) for defaults at power-up.

Block-Protection Lock-Down (BPL)

When BPL is set to '1', it prevents any further alteration of the BPL, BP1, and BP0 bits. After power-up, the BPL bit is reset to '0'.

Table 18-5. Software Status Register Block-Protection

Protection Level	Protected Blocks	Status Register Bit			Protected Memory Address
		BP2	BP1	BP0	4M bit
None	None	0	0	0	None
1 (Upper Eighth Memory)	Blocks 14 through 15	0	0	1	0x70000–0x7FFFF
2 (Upper Quarter Memory)	Blocks 12 through 15	0	1	0	0x60000–0x7FFFF
3 (Upper Half Memory)	Blocks 8 through 15	0	1	1	0x40000–0x7FFFF
4 (Full Memory)	Blocks 0 through 15	1	0	0	0x00000–0x7FFFF
5 (Full Memory)	Blocks 0 through 15	1	0	1	0x00000–0x7FFFF
6 (Full Memory)	Blocks 0 through 15	1	1	0	0x00000–0x7FFFF
7 (Full Memory)	Blocks 0 through 15	1	1	1	0x00000–0x7FFFF

Instructions

Instructions are used to read, write (erase and program), and configure the flash memory. The instruction bus cycles are 8 bits each for commands (op code), data, and addresses. The write-enable (\overline{WREN}) instruction must be executed prior to byte-program, auto-address-increment (AAI) programming, sector-erase, block-erase, write-status-register, or chip-erase instructions. The complete instructions are provided in [Table 18-6](#). All instructions are synchronized off a high-to-low transition of \overline{CE} . Inputs are accepted on the rising edge of SCK starting with the most significant bit. \overline{CE} must be driven low before an instruction is entered and must be driven high after the last bit of the instruction has been shifted in (except for read, read-ID, and read-status-register instructions). Any low-to-high transition on \overline{CE} , before receiving the last bit of an instruction bus cycle, terminates the instruction in progress and returns the device to standby mode. Instruction commands (op code), addresses, and data are all input from the most significant bit (MSB) first.

Table 18-6. Device Operation Instructions

Instruction	Description	Op Code Cycle ¹	Address Cycle(s) ²	Dummy Cycle(s)	Data Cycle(s)	Maximum Frequency
Read	Read memory	b#0000 0011 (0x03)	3	0	1 to ∞	20 MHz
High-Speed Read	Read memory at higher speed	b#0000 1011 (0x0B)	3	1	1 to ∞	25 MHz
4K Byte Sector-Erase ³	Erase 4K byte of memory array	b#0010 0000 (0x20)	3	0	0	
32K Byte Block-Erase ⁴	Erase 32K byte block of memory array	b#0101 0010 (0x52)	3	0	0	

Table 18-6. Device Operation Instructions (Continued)

Instruction	Description	Op Code Cycle ¹	Address Cycle(s) ²	Dummy Cycle(s)	Data Cycle(s)	Maximum Frequency
64K Byte Block-Erase ⁵	Erase 64K byte block of memory array	b#1101 1000 (0xD8)	3	0	0	25 MHz
Chip-Erase	Erase full memory array	b#0110 0000 (0x60) or b#1100 0111 (0xC7)	0	0	0	
Byte-Program	To program one data byte	b#0000 0010 (0x02)	3	0	1	
AAI Word Program ⁶	Auto address increment programming	b#1010 1101 (0xAD)	3	0	2 to ∞	
RDSR ⁷	Read-Status-Register	b#0000 0101 (0x05)	0	0	1 to ∞	
EWSR ⁸	Enable-Write-Status-Register	b#0110 0000 (0x50)	0	0	0	
WRSR	Write-Status-Register	b#0000 0001 (0x01)	0	0	1	
WREN ⁸	Write-Enable	b#0000 0110 (0x06)	0	0	0	
WRDI	Write-Disable	b#0000 0100 (0x04)	0	0	0	
RDID ⁹	Read-ID	b#1001 0000 (0x90) or b#1010 1011 (0xAB)	3	0	1 to ∞	

Instructions

Table 18-6. Device Operation Instructions (Continued)

Instruction	Description	Op Code Cycle ¹	Address Cycle(s) ²	Dummy Cycle(s)	Data Cycle(s)	Maximum Frequency
EBSY	Enable S0 to output $\overline{RY}/\overline{BY}$ status during AAI programming	b#0111 0000 (0x70)	0	0	0	25 MHz
DBSY	Disable S0 to output $\overline{RY}/\overline{BY}$ status during AAI programming	b#1000 0000 (0x80)	0	0	0	
JEDEC-ID	JEDEC ID read	b#1001 1111 (0x9F)	0	0	3 to ∞	

- 1 One bus cycle is eight clock periods.
- 2 Address bits above the most significant bit of each density can be VIL or VIH.
- 3 4K byte sector-erase addresses: use $A_{MS}-A_{12}$, remaining addresses are don't care but must be set either at V_{IL} or V_{IH} .
- 4 32K byte block-erase addresses: use $A_{MS}-A_{15}$, remaining addresses are don't care but must be set either at V_{IL} or V_{IH} .
- 5 64K byte block-erase addresses: use $A_{MS}-A_{16}$, remaining addresses are don't care but must be set either at V_{IL} or V_{IH} .
- 6 To continue programming to the next sequential address location, enter the 8-bit command, ADH, followed by 2 bytes of data to be programmed. Data byte 0 is programmed into the initial address $[A_{23}-A_1]$ with $A_0 = 0$, and data byte 1 is programmed into the initial address $[A_{23}-A_1]$ with $A_0 = 1$.
- 7 The read-status-register is continuous with ongoing clock cycles until terminated by a low-to-high transition on \overline{CE} .
- 8 Either EWSR or WREN followed by WRSR writes to the status register. The EWSR-WRSR sequence provides backward compatibility. The WREN-WRSR sequence is recommended for new designs.
- 9 Manufacturer's ID is read with $A_0 = 0$, and Device ID is read with $A_0 = 1$. All other address bits are 0x00. The manufacturer's ID and device ID output stream is continuous until terminated by a low-to-high transition on \overline{CE} .

Read (20 MHz)

The read instruction, 0x03, supports up to 20 MHz read. The device outputs a data stream starting from the specified address location. The data stream is continuous through all addresses until terminated by a low-to-high transition on \overline{CE} . The internal address pointer automatically increments until the highest memory address is reached. Once the highest memory address is reached, the address pointer automatically increments to the beginning (wraparound) of the address space. For example, for 2M bit density, once the data from the address location 0x3FFFF is read, the next output is from address location 0x000000. The read instruction is initiated by executing an 8-bit command, 0x03, followed by address bits $A_{23}-A_0$. \overline{CE} must remain active low for the duration of the read cycle. See [Figure 18-5](#) for the read sequence.

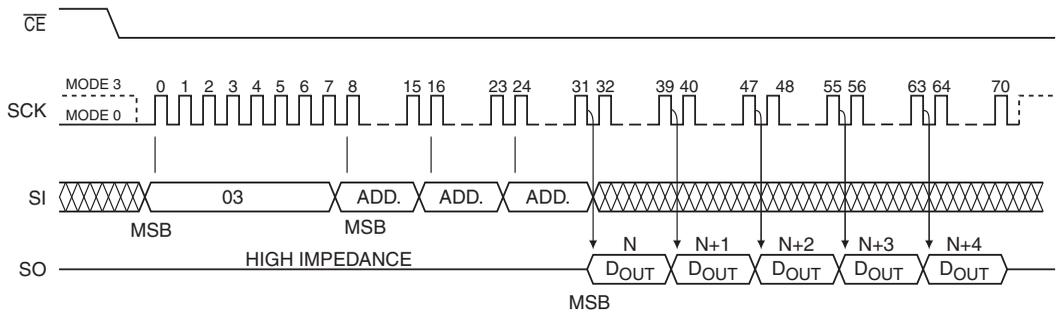


Figure 18-5. Read Sequence

High-Speed-Read (25 MHz)

The high-speed-read instruction supporting up to 25 MHz read is initiated by executing an 8-bit command, 0x0B, followed by address bits [A₂₃-A₀] and a dummy byte. \overline{CE} must remain active low for the duration of the high-speed-read cycle. See [Figure 18-6](#) for the high-speed-read sequence.

Following a dummy cycle, the high-speed-read instruction outputs the data starting from the specified address location. The data output stream is continuous through all addresses until terminated by a low-to-high transition on \overline{CE} . The internal address pointer automatically increments until the highest memory address is reached. Once the highest memory address is reached, the address pointer automatically increments to the beginning (wraparound) of the address space.

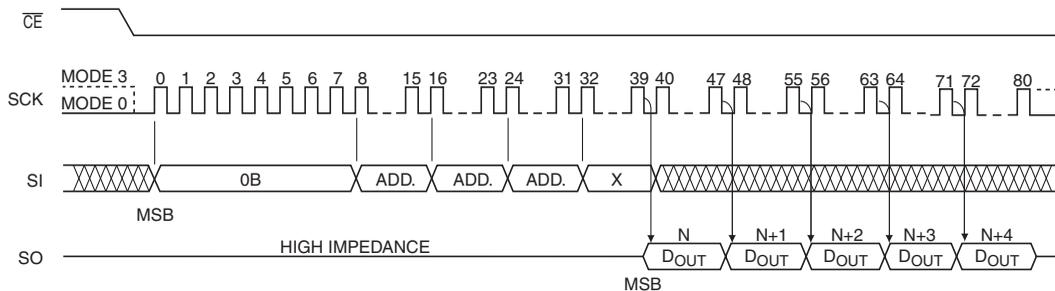


Figure 18-6. High-Speed-Read Sequence

Byte-Program

The byte-program instruction programs the bits in the selected byte to the desired data. The selected byte must be in the erased state (0xFF) when initiating a program operation. A byte-program instruction applied to a protected memory area is ignored.

Prior to any write operation, the write-enable (WREN) instruction must be executed. \overline{CE} must remain active low for the duration of the byte-program instruction. The byte-program instruction is initiated by executing an 8-bit command, 0x02, followed by address bits [A₂₃-A₀]. Following the address, the data is input in order from MSB (bit 7) to LSB (bit 0). \overline{CE} must be driven high before the instruction is executed. The user may poll the *BUSY* bit in the software status register or wait *t_{BP}* for the completion of the internal self-timed byte-program operation. See Figure 18-7 for the byte-program sequence.

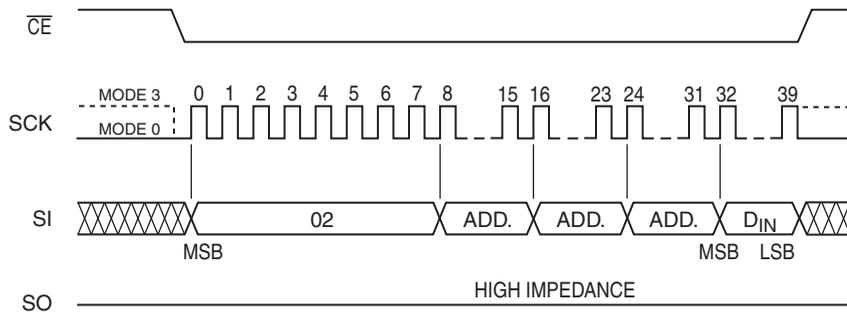


Figure 18-7. Byte-Program Sequence

Auto Address Increment (AAI) Word Program

The AAI program instruction allows multiple bytes of data to be programmed without re-issuing the next sequential address location. This feature decreases total programming time when multiple bytes or the entire memory array is to be programmed. An AAI word program instruction pointing to a protected memory area is ignored. The selected address range must be in the erased state (0xFF) when initiating an AAI word program operation. While within AAI word programming sequence, the only valid instructions are AAI Word (0xAD), RDSR (0x05), or WRDI (0x04). Users have three options to determine the completion of each AAI word program cycle: hardware detection by reading the serial output, software detection by polling the `BUSY` bit in the software status register or wait t_{BP} . Refer to “[End-of-Write Detection](#)” on page 18-17 for details.

Prior to any write operation, the write-enable (`WREN`) instruction must be executed. The AAI word program instruction is initiated by executing an 8-bit command, 0xAD, followed by address bits $[A_{23}-A_0]$. Following the addresses, two bytes of data are input sequentially, each one from MSB (bit 7) to LSB (bit 0). The first byte of data (D0) is programmed into the initial address $[A_{23}-A_1]$ with $A_0=0$, and the second byte of data (D1) is programmed into the initial address $[A_{23}-A_1]$ with $A_0=1$. \overline{CE} must be driven high before the AAI word program instruction is executed. The user must check the busy status before entering the next valid command. Once the device indicates it is no longer busy, data for the next two sequential addresses may be programmed and so on. When the last desired byte had been entered, check the busy status using the hardware method or the RDSR instruction and execute the write-disable (`WRDI`) instruction, 0x04, to terminate AAI. Check the busy status after `WRDI` to determine if the device is ready for any command. See [Figure 18-10 on page 18-19](#) and [Figure 18-11 on page 18-19](#) for AAI word programming sequences.

There is no wrap mode during AAI programming; once the highest unprotected memory address is reached, the device exits AAI operation and resets the write-enable-latch bit (`WEL=0`) and the AAI bit (`AAI=0`).

End-of-Write Detection

There are three methods to determine completion of a program cycle during AAI word programming: hardware detection by reading the serial output, software detection by polling the `BUSY` bit in the software status register or wait t_{BP} .

Hardware End-of-Write Detection

The hardware end-of-write detection method eliminates the overhead of polling the `BUSY` bit in the software status register during an AAI word program operation. The 8-bit command, `0x70`, configures the serial output (`SO`) pin to indicate flash busy status during AAI word programming, as shown in Figure 18-8. The 8-bit command, `0x70`, must be executed prior to executing an AAI word program instruction. Once an internal programming operation begins, asserting \overline{CE} immediately drives the status of the internal flash status on the `SO` pin. A '0' indicates the device is busy and a '1' indicates the device is ready for the next instruction. Deasserting \overline{CE} returns the `SO` pin to tristate.

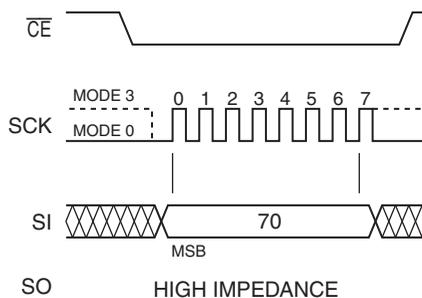


Figure 18-8. Enable `SO` as Hardware `RY/BY` During AAI Programming

Instructions

The 8-bit command, 0x80, disables the serial output (S0) pin to output busy status during AAI word program operation, and reconfigures S0 as an output pin. In this state, the S0 pin functions as a normal serial output pin. At this time, the RDSR command can poll the status of the software status register. This is shown in [Figure 18-9](#).

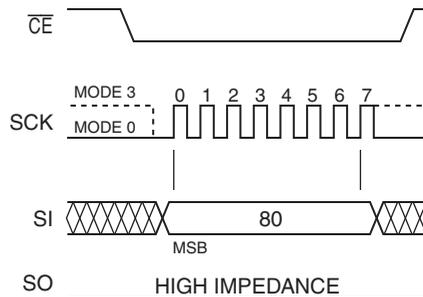
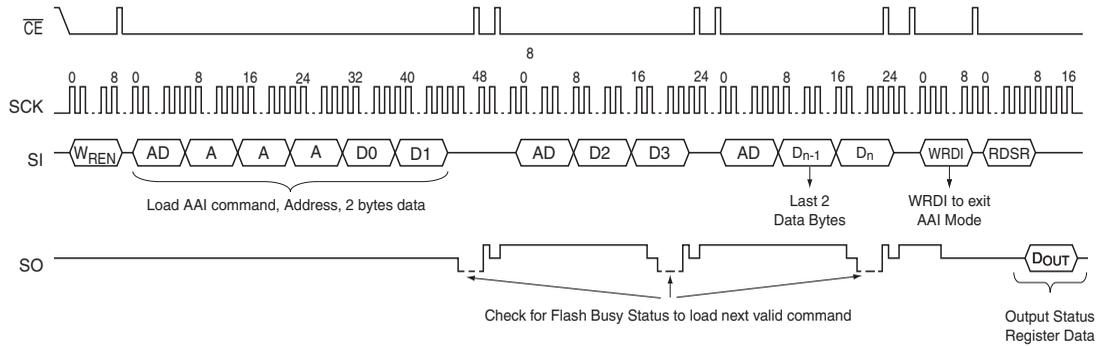
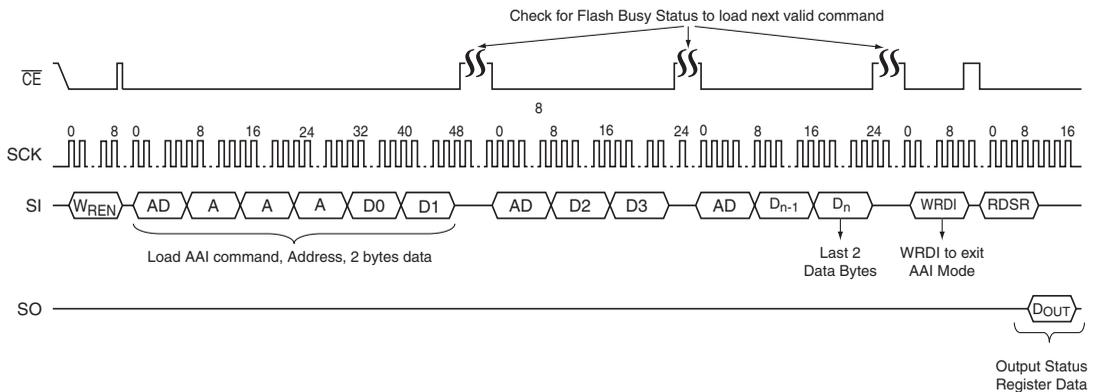


Figure 18-9. Disable S0 as Hardware RY/BY During AAI Programming



- Note: 1. Valid commands during AAI programming: AAI command or \overline{WRDI} command
 2. User must configure the SO pin to output Flash Busy status during AAI programming

Figure 18-10. Auto Address Increment (AAI) Word Program Sequence With Hardware End-of-Write Detection



- Note: Valid commands during AAI programming: AAI command, Read Status Register command, or \overline{WRDI} command

Figure 18-11. Auto Address Increment (AAI) Word Program Sequence With Software End-of-Write Detection

Sector-Erase

The sector-erase instruction clears all bits in the selected 4K byte sector to 0xFF. A sector-erase instruction applied to a protected memory area is ignored. Prior to any write operation, the write-enable (\overline{WREN}) instruction must be executed. \overline{CE} must remain active low for the duration of any command sequence. The sector-erase instruction is initiated by executing an 8-bit command, 0x20, followed by address bits [A₂₃-A₀]. Address bits [A_{MS}-A₁₂] (A_{MS} = most significant address) are used to determine the sector address (SA_X), remaining address bits can be V_{IL} or V_{IH}. \overline{CE} must be driven high before the instruction is executed. The user may poll the \overline{BUSY} bit in the software status register or wait t_{SE} for the completion of the internal self-timed sector-erase cycle. See Figure 18-12 for the sector-erase sequence.



Figure 18-12. Sector-Erase Sequence

32K Byte Block-Erase

The block-erase instruction clears all bits in the selected 32K byte block to 0xFF. A block-erase instruction applied to a protected memory area is ignored. Prior to any write operation, the write-enable (\overline{WREN}) instruction must be executed. \overline{CE} must remain active low for the duration of any command sequence. The block-erase instruction is initiated by executing an 8-bit command, 0x52, followed by address bits [$A_{23}-A_0$]. Address bits [$A_{MS}-A_{15}$] (A_{MS} = most significant address) are used to determine block address (BA_X), remaining address bits can be V_{IL} or V_{IH} . \overline{CE} must be driven high before the instruction is executed. Poll the $BUSY$ bit in the software status register or wait t_{BE} for the completion of the internal self-timed block-erase. See [Figure 18-13](#) for the block-erase sequence.



Figure 18-13. 32K Byte Block-Erase Sequence

64K Byte Block-Erase

The block-erase instruction clears all bits in the selected 64K byte block to 0xFF. A block-erase instruction applied to a protected memory area is ignored. Prior to any write operation, the write-enable (\overline{WREN}) instruction must be executed. \overline{CE} must remain active low for the duration of any command sequence. The block-erase instruction is initialized by executing an 8-bit command, 0xD8, followed by address bits $[A_{23}-A_0]$. Address bits $[A_{MS}-A_{16}]$ (A_{MS} = most significant address) are used to determine block address (BA_X), remaining address bits can be V_{IL} or V_{IH} . \overline{CE} must be driven high before the instruction is executed. Poll the $BUSY$ bit in the software status register or wait t_{BE} for the completion of the internal self-timed block-erase. See [Figure 18-14](#) for the block-erase sequence.

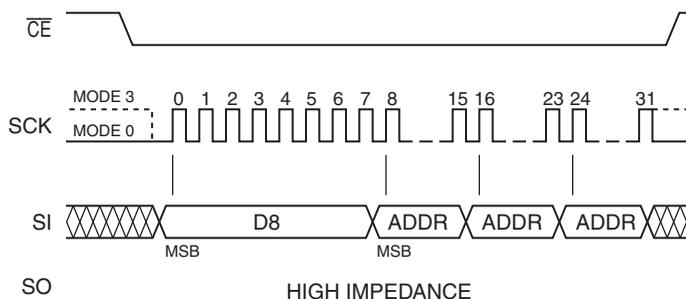


Figure 18-14. 64K Byte Block-Erase Sequence

Chip-Erase

The chip-erase instruction clears all bits in the device to 0xFF. A chip-erase instruction is ignored if any of the memory area is protected. Prior to any write operation, the write-enable (WREN) instruction must be executed. \overline{CE} must remain active low for the duration of the chip-erase instruction sequence. The chip-erase instruction is initiated by executing an 8-bit command, 0x60 or 0xC7. \overline{CE} must be driven high before the instruction is executed. The user may poll the `BUSY` bit in the software status register or wait t_{CE} for the completion of the internal self-timed chip-erase cycle. See [Figure 18-15](#) for the chip-erase sequence.

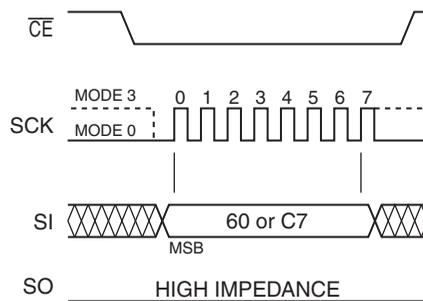


Figure 18-15. Chip-Erase Sequence

Read-Status-Register (RDSR)

The read-status-register (RDSR) instruction, 0x05, allows reading of the status register. The status register may be read at any time, even during a write (program/erase) operation. When a write operation is in progress, the **BUSY** bit may be checked before sending any new commands to assure that the new commands are properly received by the device. \overline{CE} must be driven low before the RDSR instruction is entered and remain low until the status data is read. Read-status-register is continuous with ongoing clock cycles until it is terminated by a low-to-high transition of the \overline{CE} . See [Figure 18-16](#) for the RDSR instruction sequence.

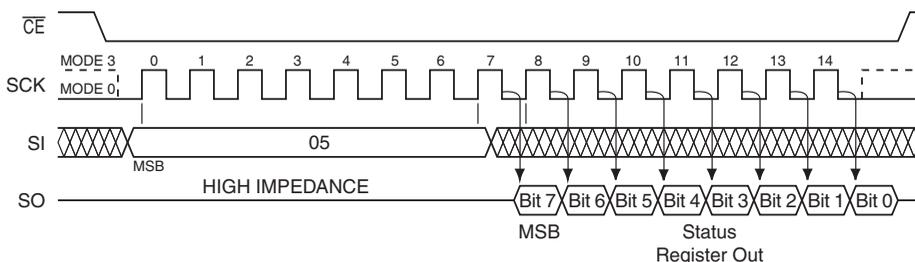


Figure 18-16. Read-Status-Register (RSDR) Sequence

Write-Enable (WREN)

The write-enable (WREN) instruction, 0x06H, sets the write-enable-latch bit in the status register to '1', allowing write operations to occur. The WREN instruction must be executed prior to any write (program/erase) operation. The WREN instruction may also be used to allow execution of the write-status-register (WRSR) instruction; however, the write-enable-latch bit in the status register is cleared upon the rising edge \overline{CE} of the WRSR instruction. \overline{CE} must be driven high before the WREN instruction is executed. See [Figure 18-17](#) for the WREN instruction sequence.

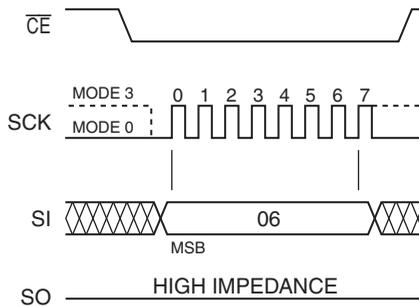


Figure 18-17. Write-Enable (WREN) Sequence

Write-Disable (WRDI)

The write-disable (WRDI) instruction, 0x04, resets the write-enable-latch bit and AAI to '0', disabling any new write operations from occurring. The WRDI instruction does not terminate any programming operation in progress. Any program operation in progress may continue up to t_{BP} after executing the WRDI instruction. \overline{CE} must be driven high before the WRDI instruction is executed. See [Figure 18-18](#) for the WRDI instruction sequence.

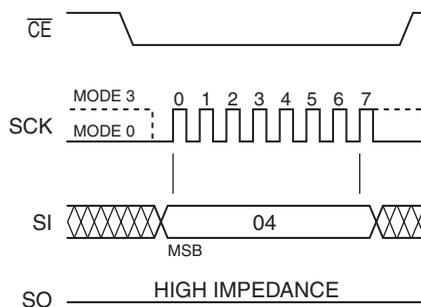


Figure 18-18. Write-Disable (WRDI) Sequence

Enable-Write-Status-Register (EWSR)

The enable-write-status-register (EWSR) instruction arms the write-status-register (WRSR) instruction and opens the status register for alteration. The write-status-register instruction must be executed immediately after the execution of the enable-write-status-register instruction. This two-step instruction sequence of the EWSR instruction followed by the WRSR instruction works like SDP (software data protection) command structure which prevents any accidental alteration of the status register values. \overline{CE} must be driven low before the EWSR instruction is entered and must be driven high before the EWSR instruction is executed. See [Figure 18-19](#) for EWSR instruction followed by WRSR instruction.

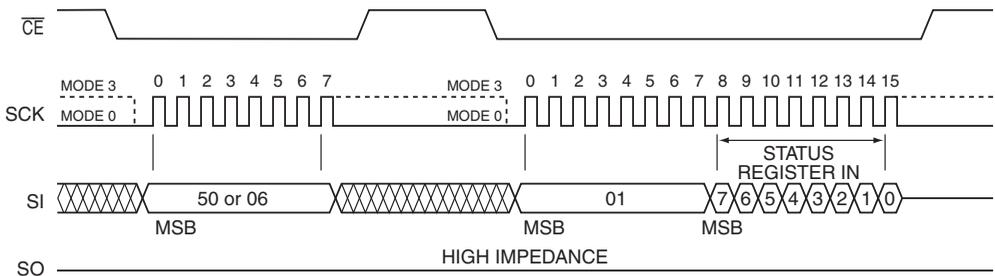


Figure 18-19. Enable-Write-Status-Register (EWSR) or Write-Enable (WREN) and Write-Status-Register (WRSR) Sequence

Write-Status-Register (WRSR)

The write-status-register instruction writes new values to the BP1, BP0, and BPL bits of the status register. \overline{CE} must be driven low before the command sequence of the WRSR instruction is entered and driven high before the WRSR instruction is executed. See [Figure 18-19](#) for EWSR or WREN and WRSR instruction sequences.

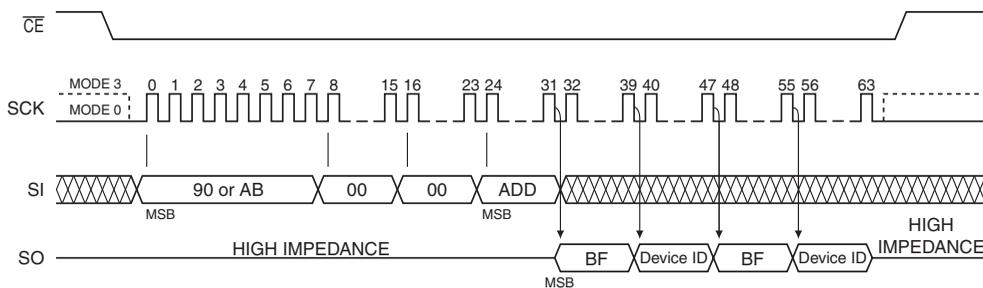
Executing the write-status-register instruction is ignored when the BPL bit is set to '1'. The BPL bit can only be set from '0' to '1' to lock-down the status register, but cannot be reset from '1' to '0'. As long as the BPL bit is set to '0' prior to the low-to-high transition of the \overline{CE} pin at the end of the WRSR instruction, the bits in the status register can all be altered by the WRSR instruction. In this case, a single WRSR instruction can set the BPL bit to '1' to lock down the status register as well as altering the BP0 and BP1 bits at the same time. See [Table 18-3 on page 18-6](#) for a summary description of the BPL functions.

Read-ID

The device information is read by executing an 8-bit command, 0x90 or 0xAB, followed by address bits [A₂₃-A₀]. Following the Read-ID instruction, the manufacturer's ID is located in address 0x000000 and the device ID is located in address 0x000001. Once the device is in read-ID mode, the manufacturer's and device ID output data toggles between address 0x000000 and 0x000001 until terminated by a low-to-high transition on \overline{CE} . See [Table 18-7](#) and [Figure 18-20](#).

Table 18-7. Product Identification

	Address	Data
Manufacturer's ID	0x000000	0xBF
Device ID	0x000001	0x04

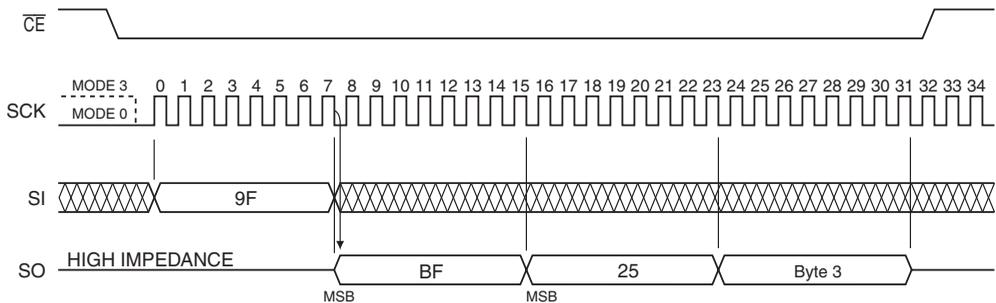


Note: 1. The manufacturer's and device ID output stream is continuous until terminated by a low to high transition on \overline{CE} .
 2. 0x00 will output the manufacturer's ID first and 0x01 will output device ID first before toggling between the two.

Figure 18-20. Read-ID Sequence

JEDEC Read-ID

The device information can be read from executing the 8-bit command, 0x9F. Following the JEDEC read-ID instruction, the 8-bit manufacturer's ID, 0xBF, is output from the device. After that, a 16-bit device ID is shifted out on the S0 pin. The device ID is assigned by the manufacturer and contains the type of memory in the first byte and the memory capacity of the device in the second byte. See Figure 18-21 for the instruction sequence. The JEDEC read-ID instruction is terminated by a low-to-high transition on \overline{CE} at any time during data output.



Note: See Table 1-8 for Device ID Byte 3 value.

Figure 18-21. JEDEC Read-ID Sequence

Table 18-8. JEDEC Read-ID Data-Out

Product	Manufacturer's ID (Byte 1)	Device ID	
		Memory Type (Byte 2)	Memory Capacity (Byte 3)
SPI serial flash	0xBF	0x25	0x4

Instructions

19 SPORT CONTROLLER

This chapter describes the synchronous serial peripheral port (SPORT). Following an overview and a list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of SPORTs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For SPORT DMA channel assignments, refer to [Table 6-7 on page 6-104](#) in [Chapter 6, “Direct Memory Access”](#).

For SPORT interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the SPORTs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each SPORT, refer to [Appendix A, “System MMR Assignments”](#).

SPORT behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the

Overview

section [“Unique Information for the ADSP-BF51x Processor”](#) on page 19-74.

Overview

Unlike the SPI interface which has been designed for SPI-compatible communication only, the SPORT modules support a variety of serial data communication protocols, for example:

- A-law or μ -law companding according to G.711 specification
- Multichannel or time-division-multiplexed (TDM) modes
- Stereo audio I²S mode
- H.100 telephony standard support

In addition to these standard protocols, the SPORT module provides modes to connect to standard peripheral devices, such as ADCs or codecs, without external glue logic. With support for high data rates, independent transmit and receive channels, and dual data paths, the SPORT interface is a perfect choice for direct serial interconnection between two or more processors in a multiprocessor system. Many processors provide compatible interfaces, including processors from Analog Devices and other manufacturers.

Each SPORT has its own set of control registers and data buffers.

Features

A SPORT can operate at up to $\frac{1}{2}$ the system clock (SCLK) rate for an internally generated or external serial clock. The SPORT external clock must always be less than the SCLK frequency. Independent transmit and receive clocks provide greater flexibility for serial communications.

A SPORT 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 I²S 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 signals and buffers to double the total supported datastreams.
- Performs A-law and μ -law hardware companding on transmitted and received words. (See [“Companding” on page 19-27](#) 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.

Interface Overview

- 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).
- Has a multichannel mode for TDM interfaces. A 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.

Interface Overview

A SPORT provides an I/O interface to a wide variety of peripheral serial devices. SPORTs provide synchronous serial data transfer only. Each SPORT has one group of signals (primary data, secondary data, clock, and frame sync) for transmit and a second set of signals for receive. The receive and transmit functions are programmed separately. A SPORT is a full duplex device, capable of simultaneous data transfer in both directions. A SPORT can be programmed for bit rate, frame sync, and number of bits per word by writing to memory-mapped registers.

[Figure 19-1 on page 19-6](#) shows a simplified block diagram of a single SPORT. Data to be transmitted is written from an internal processor register to the `SPORT_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 `DTPRI/DTSEC` pin, MSB first or LSB first, synchronous to the serial clock on the `TSCLK` pin. The receive portion of the SPORT accepts data from the `DRPRI/DRSEC` pin synchronous to the serial clock on the `RSCLK` pin. When an entire word

is received, the data is optionally expanded, then automatically transferred to the `SPORT_RX` register, and then into the RX FIFO where it is available to the processor. [Table 19-1](#) shows the signals for each SPORT.

Table 19-1. SPORT Signals

Pin	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

A SPORT receives serial data on its DRPRI and DRSEC inputs and transmits serial data on its DTPRI and DTSEC outputs. It can receive and transmit simultaneously for full-duplex operation. For transmit, the data bits (DTPRI and DTSEC) are synchronous to the transmit clock (TSCLK). For receive, the data bits (DRPRI and DRSEC) are synchronous to the receive clock (RSCLK). The serial clock is an output if the processor generates it, or an input if the clock is externally generated. Frame synchronization signals RFS and TFS are used to indicate the start of a serial data word or stream of serial words.

The primary and secondary data pins, if enabled by a specific processor port configuration, 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 signals 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 the *Data Address Generators* chapter in the *Blackfin Processor*

Interface Overview

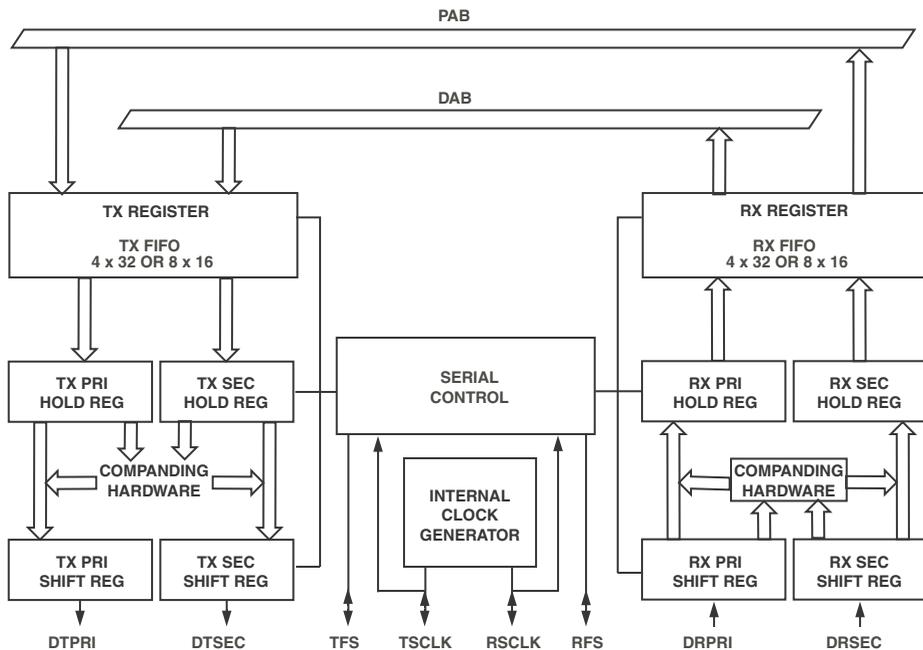


Figure 19-1. SPORT Block Diagram^{1, 2, 3}

- 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.
- 2 TX register is the bottom of the TX fifo, RX register is the top of the RX fifo.
- 3 In multichannel mode, the TFS pin acts as transmit data valid (TDV). For more information, see [“Multichannel Operation”](#) on page 19-14.

Programming Reference. 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.

Figure 19-2 shows a possible port connection for a device with at least two 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 1, 2, ...N.

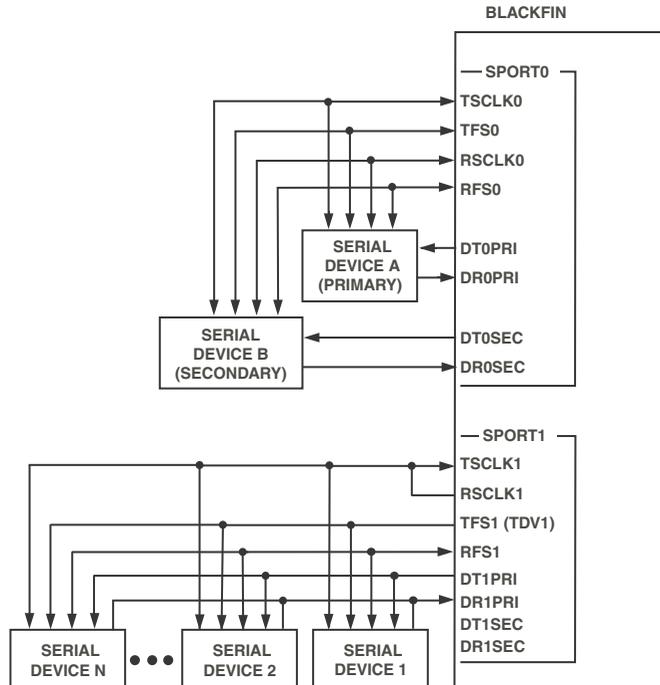


Figure 19-2. Example SPORT Connections
(SPORT0 is Standard Mode, SPORT1 is Multichannel Mode)^{1, 2}

- 1 In multichannel mode, TFS functions as a transmit data valid (TDV) output. See “[Multichannel Operation](#)” on page 19-14.
- 2 Although shown as an external connection, the TSCLK1/RSCLK1 connection is internal in multichannel mode. See “[Multichannel Operation](#)” on page 19-14.

Interface Overview

Figure 19-3 shows an example of a stereo serial device with three transmit and two receive channels connected to a processor with two SPORTs.

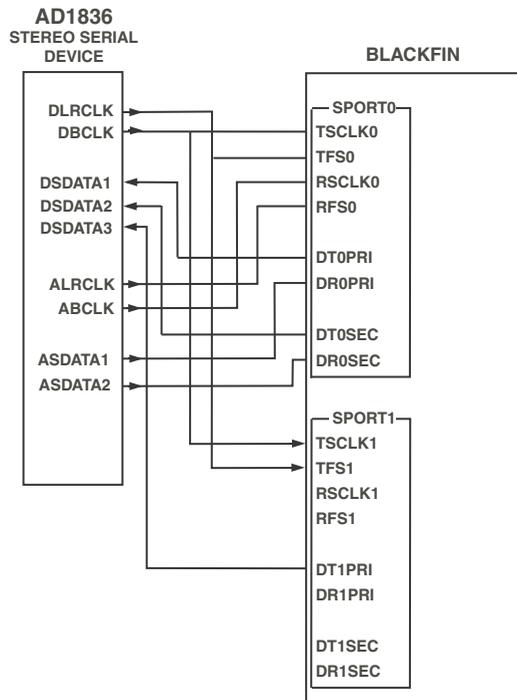


Figure 19-3. Stereo Serial Connection

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.

Description of 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_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 `SPORT_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 `SPORT_TCR1` register. Each bit is shifted out on the driving edge of `TSCLK`. The driving edge of `TSCLK` 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.

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 `SPORT_TCR1` register and `RSPEN` in the `SPORT_RCR1` register, respectively). Each method has a different effect on the SPORT.

A processor reset disables the SPORTs by clearing the `SPORT_TCR1`, `SPORT_TCR2`, `SPORT_RCR1`, and `SPORT_RCR2` registers (including the `TSPEN` and `RSPEN` enable bits) and the `SPORT_TCLKDIV`, `SPORT_RCLKDIV`, `SPORT_TFSDIVx`, and `SPORT_RFSDIVx` clock and frame sync divisor registers. Any ongoing operations are aborted.

Description of Operation

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 `TFS/RFS` and `TSCLK/RSCLK` outputs, if these signals are configured to be generated internally.

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. A SPORT must be configured prior to being enabled. Once the SPORT is enabled, further writes to the SPORT configuration registers are disabled (except for `SPORT_RCLKDIV`, `SPORT_TCLKDIV`, and multichannel mode channel select registers). To change values in all other SPORT configuration registers, disable the SPORT by clearing `TSPEN` in `SPORT_TCR1` and/or `RSPEN` in `SPORT_RCR1`.

Each SPORT has its own set of control registers and data buffers. These registers are described in detail in [“SPORT Registers” on page 19-44](#). All control and status bits in the SPORT registers are active high unless otherwise noted.

Stereo Serial Operation

Several stereo serial modes can be supported by the SPORT, including the popular I²S 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 I²S 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 19-4 on page 19-13](#) and [Figure 19-5 on page 19-14](#) show timing diagrams for stereo serial mode operation.

[Table 19-2](#) shows several modes that can be configured using bits in SPORT_TCR1 and SPORT_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."

 Blackfin SPORTs are designed such that, in I²S master mode, LRCLK is held at the last driven logic level and does not transition, to provide an edge, after the final data word is driven out. Therefore, while transmitting a fixed number of words to an I²S 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 I²S receiver is a Blackfin SPORT.

Table 19-2. Stereo Serial Settings

Bit Field	Stereo Audio Serial Scheme		
	I ² S	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

Description of Operation

Table 19-2. Stereo Serial Settings (Continued)

Bit Field	Stereo Audio Serial Scheme		
	I ² S	Left-Justified	DSP Mode
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 19-2](#) provide glueless interfaces to many popular codecs.

Note $RFSDIV$ or $TFSDIV$ must still be greater than or equal to $SLEN$. For I²S 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 19-24](#)) 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 DRSEC and DTSEC pins are useful extensions of the SPORT which pair well with stereo serial mode. Multiple I²S 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 19-3 on page 19-8](#), which shows multiple stereo serial connections being made between the processor and an AD1836 codec.

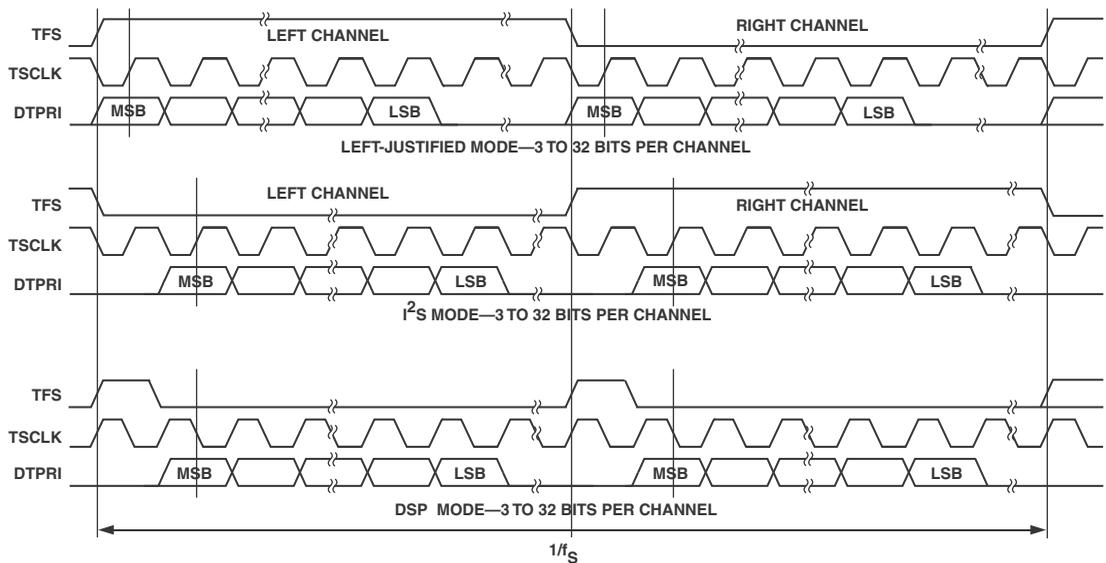


Figure 19-4. SPORT Stereo Serial Modes, Transmit^{1, 2, 3}

- 1 DSP mode does not identify channel.
- 2 TFS normally operates at f_s except for DSP mode which is $2 \times f_s$.
- 3 TSCLK frequency is normally $64 \times f_s$ but may be operated in burst mode.

Description of Operation

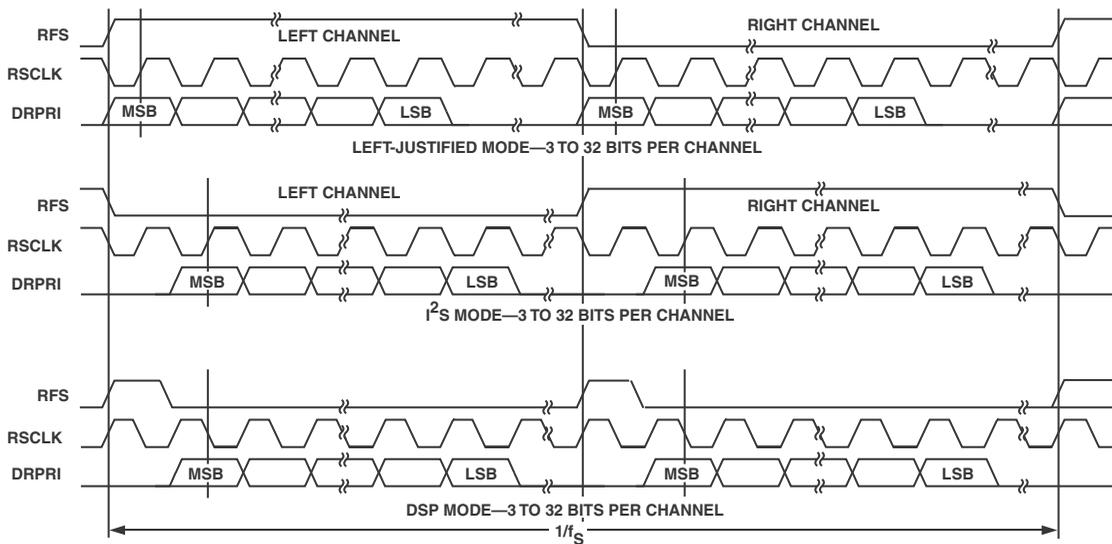


Figure 19-5. SPORT Stereo Serial Modes, Receive^{1, 2, 3}

- 1 DSP mode does not identify channel.
- 2 RFS normally operates at f_s except for DSP mode which is $2 \times f_s$.
- 3 RSCLK frequency is normally $64 \times f_s$ but may be operated in burst mode.

Multichannel Operation

The SPORT offers 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 bitstream 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 `SPORT_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 `SPORT_TCR2` register), unless the SPORT is in multichannel mode and an inactive time slot occurs.



The SPORT multichannel transmit select register and the SPORT multichannel receive select register must be programmed before enabling `SPORT_TX` or `SPORT_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 `SPORT_MCMC2`) must be enabled prior to enabling `SPORT_TX` or `SPORT_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.

Description of Operation

Figure 19-6 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 19-38 for more examples.

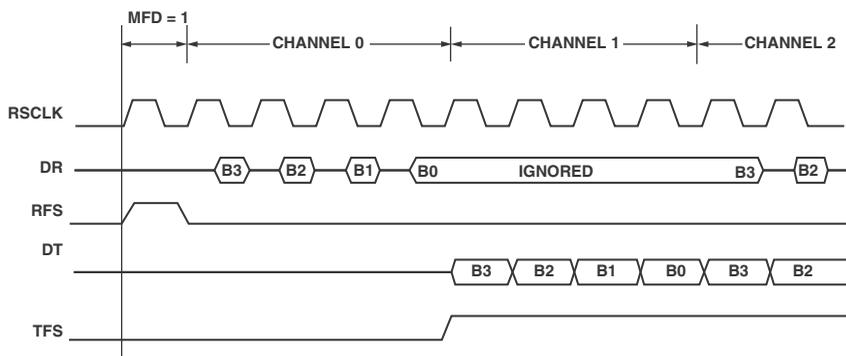


Figure 19-6. Multichannel Operation

Multichannel Enable

Setting the `MCMEN` bit in the `SPORT_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 19-3 shows the dependencies of bits in the SPORT configuration register when the SPORT is in multichannel mode.

Table 19-3. Multichannel Mode Configuration

SPORT_RCR1 or SPORT_RCR2	SPORT_TCR1 or SPORT_TCR2	Notes
RSPEN	TSPEN	Set or clear both
IRCLK	-	Independent
-	ITCLK	Independent
RDTYPE	TDTYPE	Independent
RLSBIT	TLSBIT	Independent
IRFS	-	Independent
-	ITFS	Ignored
RFSR	TFSR	Ignored
-	DITFS	Ignored
LRFS	LTFS	Independent
LARFS	LATFS	Both must be 0
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

Description of Operation

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 `SPORT_TX` and `SPORT_RX` channels of the `SPORT` in multichannel mode configuration, the corresponding bit pairs in `SPORT_RCR1` and `SPORT_TCR1`, and in `SPORT_RCR2` and `SPORT_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 `SPORT_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.

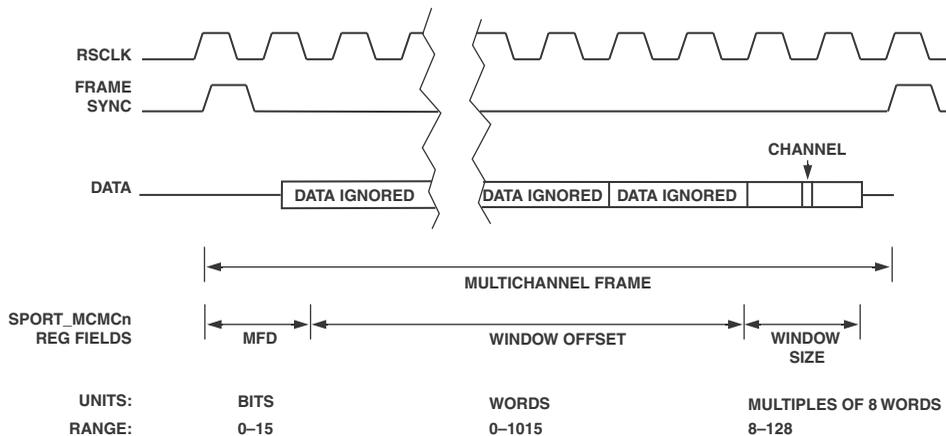
Once the initial `RFS` is received, and a frame transfer has started, all other `RFS` signals are ignored by the `SPORT` until the complete frame has been transferred.

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 19-7](#).



NOTE: FRAME LENGTH IS SET BY FRAME SYNC DIVIDE OR EXTERNAL FRAME SYNC PERIOD.

Figure 19-7. Relationships for Multichannel Parameters

Description of Operation

Multichannel Frame Delay

The 4-bit `MFD` field in `SPORT_MCMC2` specifies a delay between the frame sync pulse and the first data bit in multichannel mode. The value of `MFD` is 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:

$$\text{Number of words in active window} = 8 \times (\text{WSIZE} + 1)$$

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.

Other Multichannel Fields in SPORT_MCMC2

The $FSDR$ bit in the $SPORT_MCMC2$ 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 multibit 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 $SPORT_MRCSn$ and $SPORT_MTCSn$ multichannel select registers are used to enable and disable individual channels; the $SPORT_MRCSn$ registers

Description of Operation

specify the active receive channels, and the `SPORT_MTCsn` registers specify the active transmit channels.

Four registers make up each multichannel select register. 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). See [Figure 19-8](#).

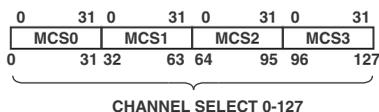


Figure 19-8. Multichannel Select Registers

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 `SPORT_MTCsn` register causes the SPORT to transmit the word in that channel's position of the datastream. Clearing the bit in the `SPORT_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 `SPORT_MRCSn` register causes the SPORT to receive the word in that channel's position of the datastream; the received word is loaded into the `SPORT_RX` buffer. Clearing the bit in the `SPORT_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 `SPORT_TCR1` register and the `RDTYPE` field in the `SPORT_RCR1` register, and applies to all active channels. (See [“Companding” on page 19-27](#) for more information about companding.)

Multichannel DMA Data Packing

Multichannel DMA data packing and unpacking are specified with the `MCDTXPE` and `MCDRXPE` bits in the `SPORT_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 reconfiguration 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 (unless the secondary side is enabled). 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)

Functional Description

- 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)

2× Clock Recovery Control

The SPORT can recover the data rate clock from a provided 2× 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 SPORT_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).

Functional Description

The following sections provide a functional description of the SPORT.

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, SPORT_TCLKDIV and SPORT_RCLKDIV.

$$\text{TSCLK frequency} = (\text{SCLK frequency}) / (2 \times (\text{SPORT_TCLKDIV} + 1))$$

$$\text{RSCLK frequency} = (\text{SCLK frequency}) / (2 \times (\text{SPORT_RCLKDIV} + 1))$$

If the value of SPORT_TCLKDIV or SPORT_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 SPORT_TCR1 register or IRFS = 1 in the SPORT_RCR1 register) and frame syncs are not required, the first frame sync does not update the clock divider if the value in SPORT_TCLKDIV or SPORT_RCLKDIV has changed. The second frame sync will cause the update.

The SPORT_TFSDIV and SPORT_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:

$$\# \text{ of transmit serial clocks between frame sync assertions} = \text{TFSDIV} + 1$$

$$\# \text{ of receive serial clocks between frame sync assertions} = \text{RFSDIV} + 1$$

Use the following equations to determine the correct value of TFSDIV or RFSDIV, given the serial clock frequency and desired frame sync frequency:

$$\text{SPORT TFS frequency} = (\text{TSCLK frequency}) / (\text{SPORT_TFSDIV} + 1)$$

$$\text{SPORT RFS frequency} = (\text{RSCLK frequency}) / (\text{SPORT_RFSDIV} + 1)$$

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

Functional Description

RFSDIV) should not be less than the serial word length minus 1 (the value of the SLEN field in SPORT_TCR2 or SPORT_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-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet* for exact timing specifications.

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 SPORT_TCR2 and SPORT_RCR2 registers of each SPORT determines the word length according to this formula:

$$\text{Serial Word Length} = \text{SLEN} + 1$$

 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 $\text{SLEN} \geq 3$).

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

SPORT_RCR1 and SPORT_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 SPORT_TCR1 register and the RDTYPE field of the SPORT_RCR1 register specify one of four data formats for both single and multichannel operation. See [Table 19-4](#).

Table 19-4. TDTYPE, RDTYPE, and Data Formatting

TDTYPE or RDTYPE	SPORT_TCR1 Data Formatting	SPORT_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 SPORT_RX and SPORT_TX buffers. SPORT_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 SPORT supports 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 SPORT_RX register is the right-justified, expanded value of the eight LSBs received and sign

Functional Description

extended to 16 bits. A write to `SPORT_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 `SPORT_TCR1` and `SPORT_RCR1` registers. Serial clock frequency is configured in the `SPORT_TCLKDIV` and `SPORT_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 `SPORT_TCR1` configuration register and the `IRCLK` bit in the `SPORT_RCR1` configuration register determines the clock source.

When `IRCLK` or `ITCLK` = 1, the clock signal is generated internally by the processor, and the `TSCLK` or `RSCLK` pin is an output. The clock frequency is determined by the value of the serial clock divisor in the `SPORT_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 `SPORT_TCLKDIV`/`SPORT_RCLKDIV` registers are ignored. The externally generated serial clocks do not need to be synchronous with the system clock or

with each other. The system clock must have a higher frequency than RSCLK and TSCLK.

- ⊘ When the SPORT uses external clocks, it must be enabled for a minimal number of stable clock pulses before the first active frame sync is sampled. Failure to allow for these clocks may result in a SPORT malfunction. See the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet* for details.

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.

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 (SPORT_TCR1, SPORT_TCR2, SPORT_RCR1 and SPORT_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 SPORT_TCR1 and SPORT_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 SPORT_TX hold register before the previous word is shifted out and transmitted.

Functional Description

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 19-9 illustrates framed serial transfers, which have these characteristics:

- $TFSR$ and $RFSR$ bits in the $SPORT_TCR1$ and $SPORT_RCR1$ 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 $SPORT_TCR1$ and $SPORT_RCR1$ registers.

See “[Timing Examples](#)” on page 19-38 for more timing examples.

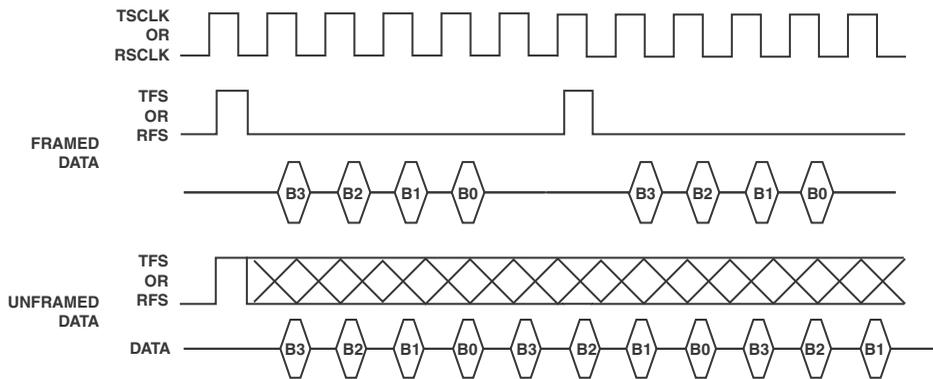


Figure 19-9. Framed Versus Unframed Data

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 `SPORT_TCR1` and `SPORT_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 `SPORT_TFSDIV` or `SPORT_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 `SPORT_TFSDIV`/`SPORT_RFSDIV` registers are ignored.

All of the frame sync options are available whether the signal is generated internally or externally.

Functional Description

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 `SPORT_TCR1` and `SPORT_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 `SPORT_TCR1` and `SPORT_RCR1` registers select the driving and sampling edges of the serial data and frame syncs.

For the SPORT transmitter, setting `TCKFE = 1` in the `SPORT_TCR1` register selects the falling edge of `TSCLK` to drive data and internally generated frame syncs and selects the rising edge of `TSCLK` to sample externally generated frame syncs. Setting `TCKFE = 0` selects the rising edge of `TSCLK` to drive data and internally generated frame syncs and selects the falling edge of `TSCLK` to sample externally generated frame syncs.

For the SPORT receiver, setting `RCKFE = 1` in the `SPORT_RCR1` register selects the falling edge of `RSCLK` to drive internally generated frame syncs and selects the rising edge of `RSCLK` to sample data and externally generated frame syncs. Setting `RCKFE = 0` selects the rising edge of `RSCLK` to drive internally generated frame syncs and selects the falling edge of `RSCLK` 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 `SPORT_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 19-10](#), $TCKFE = RCKFE = 0$ and transmit and receive are connected together to share the same clock and frame syncs.

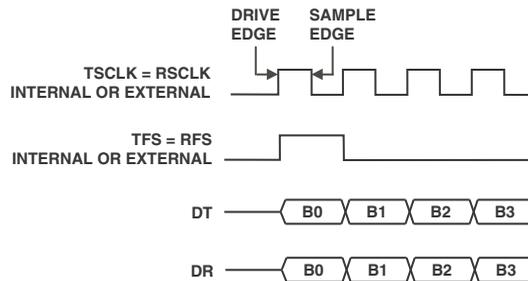


Figure 19-10. Example of $TCKFE = RCKFE = 0$, Transmit and Receive Connected

In [Figure 19-11](#), $TCKFE = RCKFE = 1$ and transmit and receive are connected together to share the same clock and frame syncs.

Functional Description

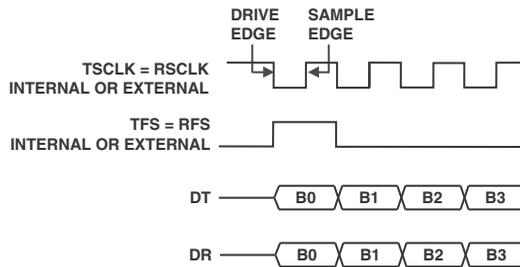


Figure 19-11. Example of TCKFE = RCKFE = 1, Transmit and Receive Connected

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 SPORT_TCR1 and SPORT_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 ($SLEN \geq 3$).

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 19-12 illustrates the two modes of frame signal timing. In summary:

- For the $LATFS$ or $LARFS$ bits of the $SPORT_TCR1$ or $SPORT_RCR1$ 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 ($TLSSBIT = 0$ or $RLSSBIT = 0$) or LSB first ($TLSSBIT = 1$ or $RLSSBIT = 1$).
- Frame sync and clock are generated internally or externally.

See “Timing Examples” on page 19-38 for more examples.

Functional Description

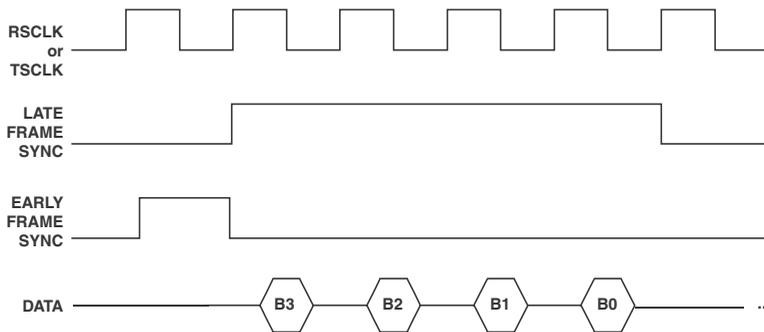


Figure 19-12. Normal Versus Alternate Framing

Data Independent Transmit Frame Sync

Normally the internally generated transmit frame sync signal (TFS) is output only when the `SPORT_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 `SPORT_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 `SPORT_TX` buffer. The next TFS is generated once data is loaded into `SPORT_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 `SPORT_TX` buffer. Whatever data is present in `SPORT_TX` is transmitted again with each assertion of TFS. The `TUVF` (transmit underflow status) bit in the `SPORT_STAT` register is set when this occurs and old data is retransmitted. The `TUVF` status bit is also set if the `SPORT_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 `SPORT_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.

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`.

Functional Description

Peripheral Bus Errors

The SPORT generates a peripheral bus 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

Timing Examples

Several timing examples are included within the text of this chapter (in the sections [“Framed Versus Unframed” on page 19-29](#), [“Early Versus Late Frame Syncs \(Normal Versus Alternate Timing\)” on page 19-34](#), and [“Frame Syncs in Multichannel Mode” on page 19-18](#)). 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-BF512/BF514/BF516/BF518(F) 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 19-13](#) through [Figure 19-18](#) show framing for receiving data.

In [Figure 19-13](#) and [Figure 19-14](#), 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).

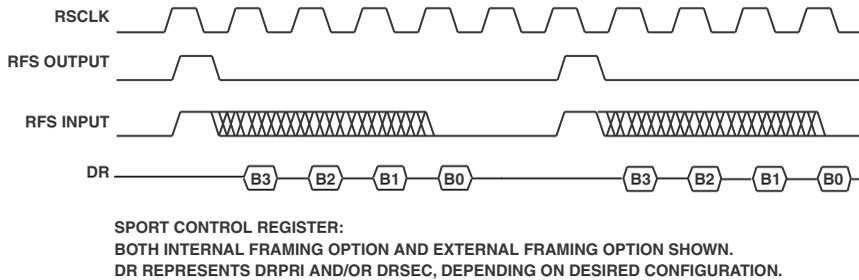


Figure 19-13. SPORT Receive, Normal Framing

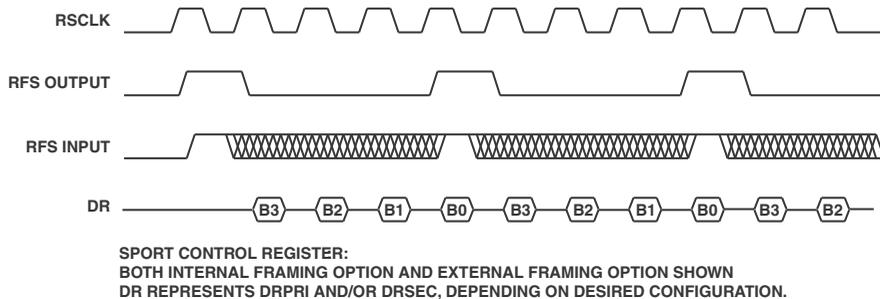


Figure 19-14. SPORT Continuous Receive, Normal Framing

[Figure 19-15](#) and [Figure 19-16](#) 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.

Functional Description

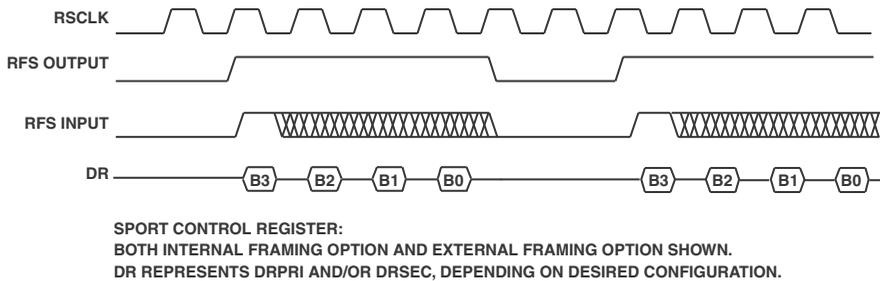


Figure 19-15. SPORT Receive, Alternate Framing

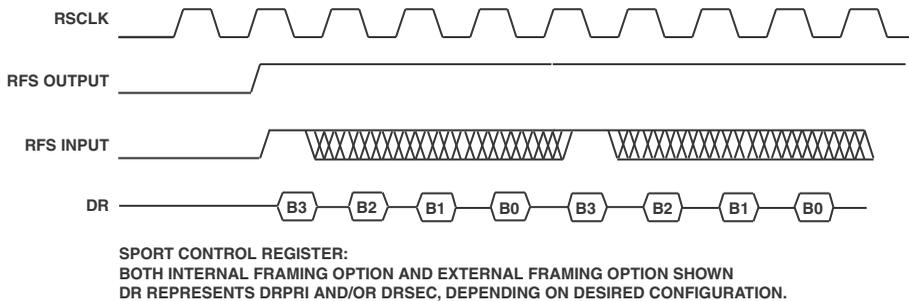


Figure 19-16. SPORT Continuous Receive, Alternate Framing

Figure 19-17 and Figure 19-18 show the receive 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 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 multiword bursts (continuous reception).

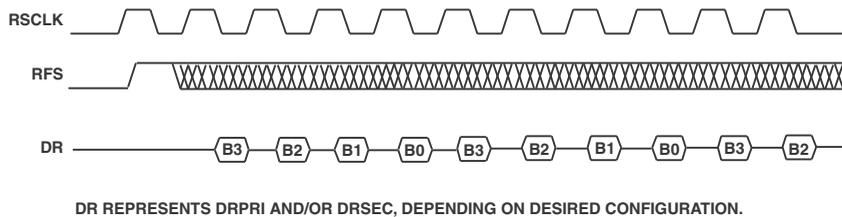


Figure 19-17. SPORT Receive, Unframed Mode, Normal Framing

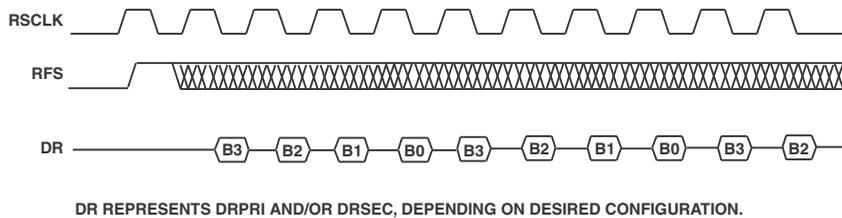


Figure 19-18. SPORT Receive, Unframed Mode, Alternate Framing

Figure 19-19 through Figure 19-24 show framing for transmitting data and are very similar to Figure 19-13 through Figure 19-18.

In Figure 19-19 and Figure 19-20, the normal framing mode is shown for non-continuous data (any number of T_{SCLK} cycles between words) and continuous data (no T_{SCLK} cycles between words). Figure 19-21 and Figure 19-22 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.

Functional Description

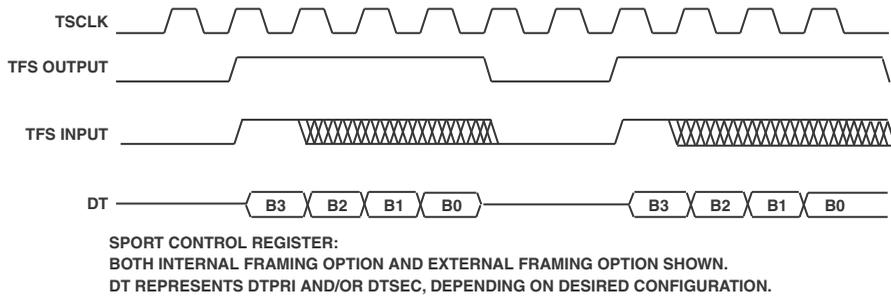


Figure 19-19. SPORT Transmit, Normal Framing

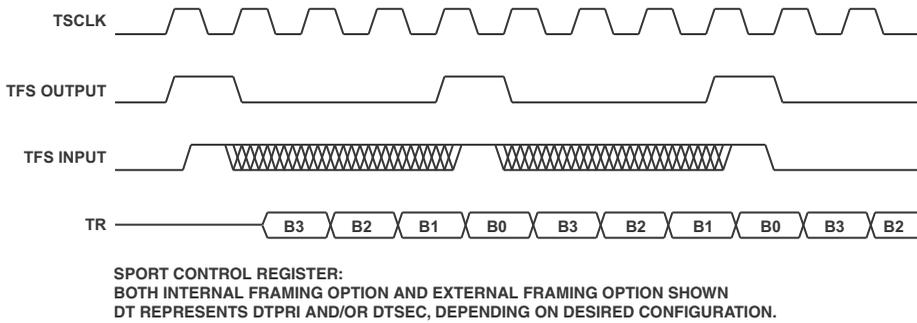


Figure 19-20. SPORT Continuous Transmit, Normal Framing

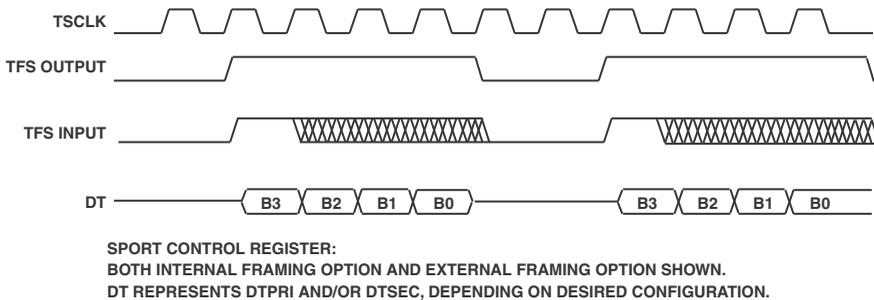


Figure 19-21. SPORT Transmit, Alternate Framing

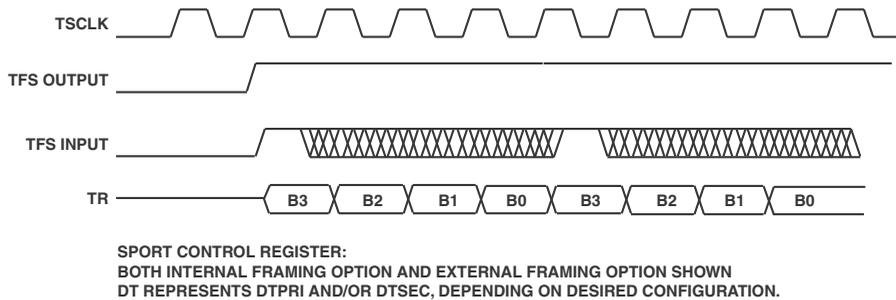


Figure 19-22. SPORT Continuous Transmit, Alternate Framing

Figure 19-23 and Figure 19-24 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).

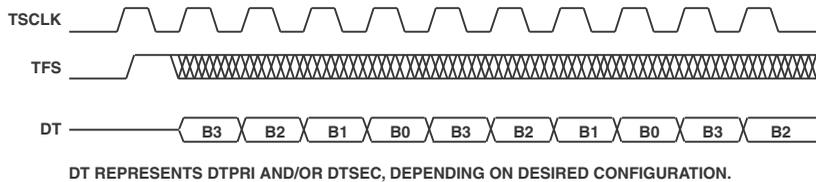


Figure 19-23. SPORT Transmit, Unframed Mode, Normal Framing

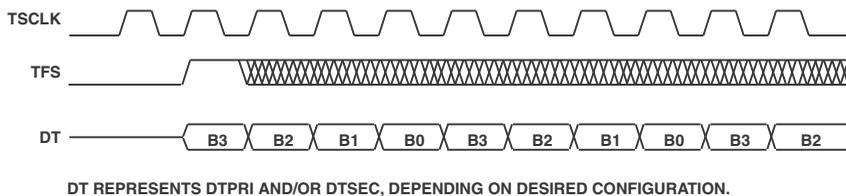


Figure 19-24. SPORT Transmit, Unframed Mode, Alternate Framing

SPORT Registers

The following sections describe the SPORT registers. [Table 19-5](#) provides an overview of the available control registers.

Table 19-5. SPORT Register Mapping

Register Name	Function	Notes
SPORT_TCR1	Primary transmit configuration register	Bits [15:1] can only be written if bit 0 = 0
SPORT_TCR2	Secondary transmit configuration register	
SPORT_TCLKDIV	Transmit clock divider register	Ignored if external SPORT clock mode is selected
SPORT_TFSDIV	Transmit frame sync divider register	Ignored if external frame sync mode is selected
SPORT_TX	Transmit data register	See description of FIFO buffering at “SPORT Transmit Data (SPORT_TX) Register” on page 19-56.
SPORT_RCR1	Primary receive configuration register	Bits [15:1] can only be written if bit 0 = 0
SPORT_RCR2	Secondary receive configuration register	
SPORT_RCLK_DIV	Receive clock divider register	Ignored if external SPORT clock mode is selected
SPORT_RFSDIV	Receive frame sync divider register	Ignored if external frame sync mode is selected
SPORT_RX	Receive data register	See description of FIFO buffering at “SPORT Receive Data (SPORT_RX) Register” on page 19-58.
SPORT_STAT	Receive and transmit status	
SPORT_MCM1	Primary multichannel mode configuration register	Configure this register before enabling the SPORT

Table 19-5. SPORT Register Mapping (Continued)

Register Name	Function	Notes
SPORT_MCM2	Secondary multichannel mode configuration register	Configure this register before enabling the SPORT
SPORT_MRCSn	Receive channel selection registers	Select or deselect channels in a multichannel frame
SPORT_MTCSn	Transmit channel selection registers	Select or deselect channels in a multichannel frame
SPORT_CHNL	Currently serviced channel in a multichannel frame	

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 $SPORT_RCLKDIV$, $SPORT_TCLKDIV$, and multichannel mode channel select registers). The $SPORT_TX$ register writes are always enabled; $SPORT_RX$, $SPORT_CHNL$, and $SPORT_STAT$ 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.

SPORT Transmit Configuration (SPORT_TCR1 and SPORT_TCR2) Registers

The main control registers for the transmit portion of each SPORT are the transmit configuration registers, `SPORT_TCR1` and `SPORT_TCR2`, shown in [Figure 19-25](#) and [Figure 19-26](#).

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 `SPORT_TCLKDIV` and multichannel mode channel select registers. Writes to disallowed registers have no effect. While the SPORT is enabled, `SPORT_TCR1` is not written except for bit 0 (`TSPEN`). For example,

```
write (SPORT_TCR1, 0x0001) ; /* SPORT TX Enabled */
write (SPORT_TCR1, 0xFF01) ; /* ignored, no effect */
write (SPORT_TCR1, 0xFFFF) ; /* SPORT disabled, SPORT_TCR1
                               still equal to 0x0000 */
```

SPORT Transmit Configuration 1 Register (SPORT_TCR1)

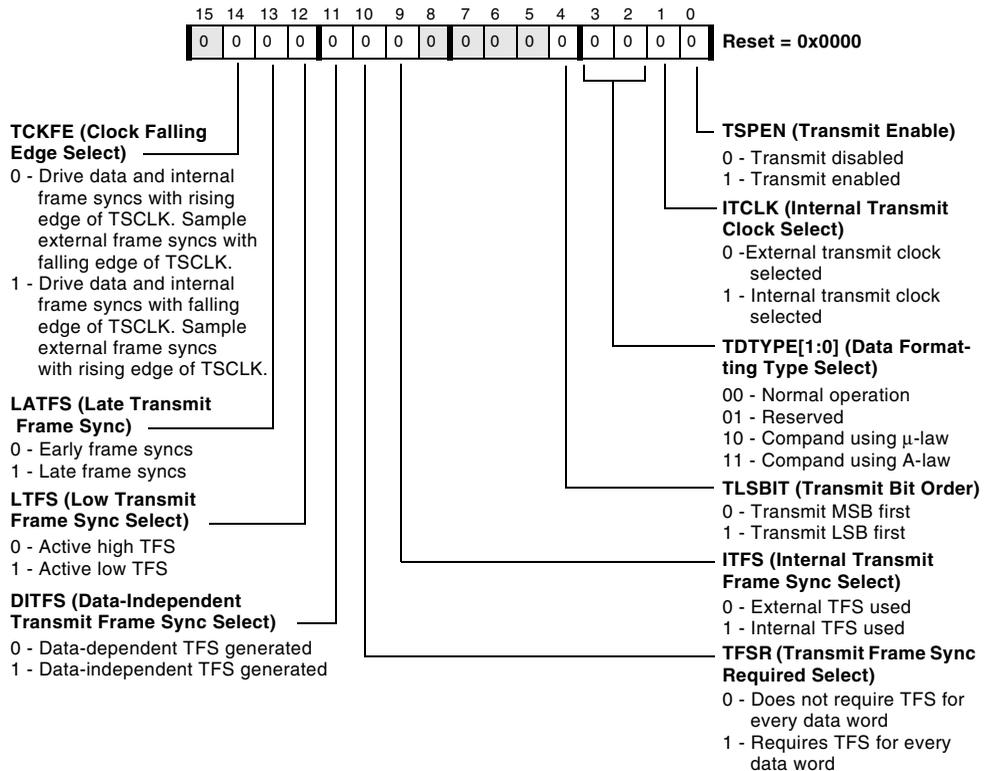


Figure 19-25. SPORT Transmit Configuration 1 Register

SPORT Registers

SPORT Transmit Configuration 2 Register (SPORT_TCR2)

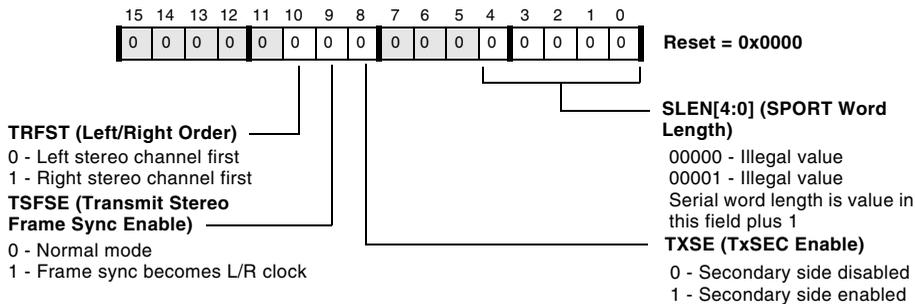


Figure 19-26. SPORT Transmit Configuration 2 Register

Additional information for the `SPORT_TCR1` and `SPORT_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 are 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 `SPORT_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:

$$\text{Serial Word Length} = \text{SLEN} + 1;$$

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 SPORT always transfers the `SLEN+1` lower bits from the transmit buffer.

-  The frame sync signal is controlled by the `SPORT_TFSDIV` and `SPORT_RFSDIV` registers, not by `SLEN`. To produce a frame sync pulse on each byte or word transmitted, the proper frame sync divider

SPORT Registers

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 `SPORT_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 `SPORT_TX` register has been loaded or not; if `DITFS` is cleared, the frame sync pulse is only generated if the `SPORT_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 `SPORT_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 `SPORT_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 SPORT (if set).
- **Stereo serial enable.** (TSFSE). This bit enables the stereo serial operating mode of the SPORT (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.

SPORT Receive Configuration (SPORT_RCR1 and SPORT_RCR2) Registers

The main control registers for the receive portion of each SPORT are the receive configuration registers, `SPORT_RCR1` and `SPORT_RCR2`, shown in [Figure 19-27](#) and [Figure 19-28](#).

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.

SPORT Registers

SPORT Receive Configuration 1 Register (SPORT_RCR1)

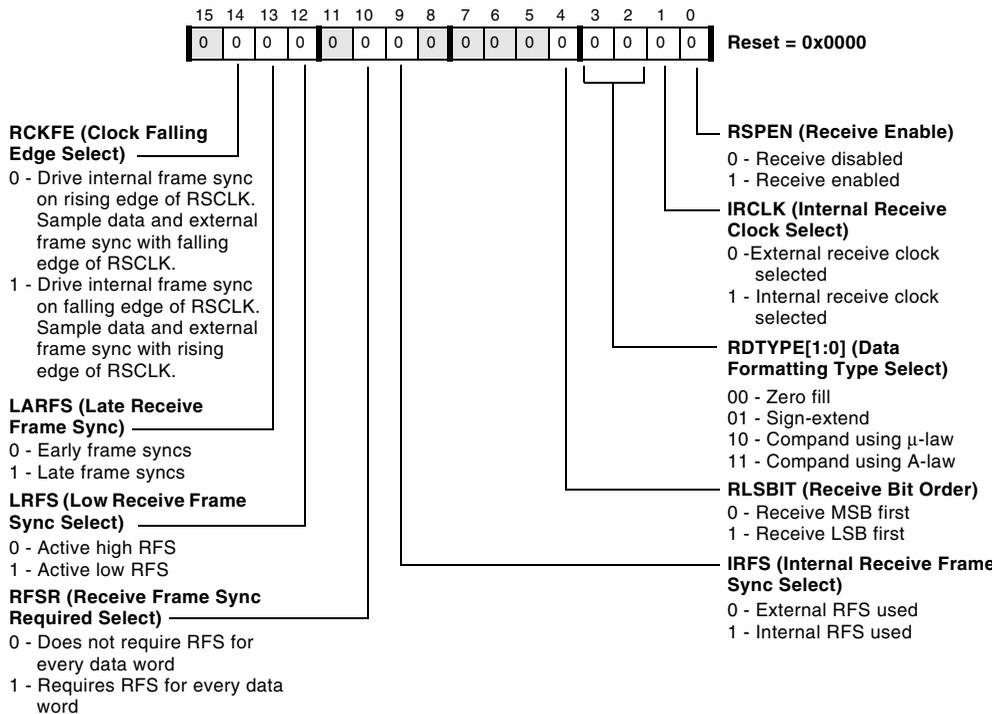


Figure 19-27. SPORT Receive Configuration 1 Register

When the SPORT is enabled to receive (RSPEN set), corresponding SPORT configuration register writes are not allowed except for SPORT_RCLKDIV and multichannel mode channel select registers. Writes to disallowed registers have no effect. While the SPORT is enabled, SPORT_RCR1 is not written except for bit 0 (RSPEN). For example,

```
write (SPORT_RCR1, 0x0001) ; /* SPORT RX Enabled */
write (SPORT_RCR1, 0xFF01) ; /* ignored, no effect */
write (SPORT_RCR1, 0xFFFF0) ; /* SPORT disabled, SPORT_RCR1
                                still equal to 0x0000 */
```

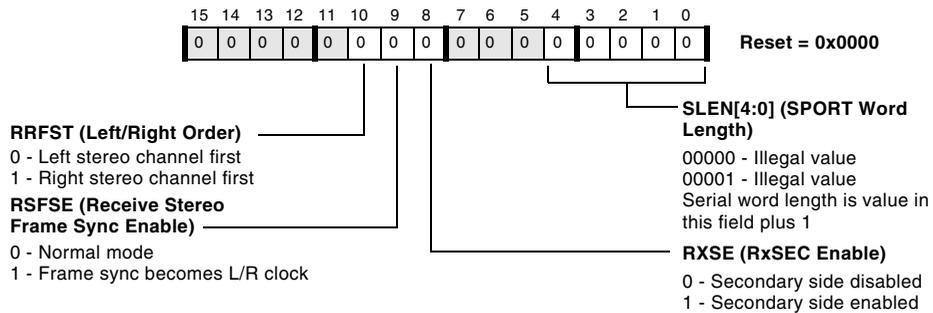
SPORT Receive Configuration 2 Register (SPORT_RCR2)

Figure 19-28. SPORT Receive Configuration 2 Register

Additional information for the `SPORT_RCR1` and `SPORT_RCR2` 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 SPORT receiver, which can generate a SPORT 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.

SPORT Registers

 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 `SPORT_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 `RCLK-DIV` 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 `SPORT_TFSDIV` and `SPORT_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 SPORT (if set).
- **Stereo serial enable.** (RSFSE). This bit enables the stereo serial operating mode of the SPORT (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 SPORT_TCR1, SPORT_TCR2, SPORT_RCR1, and SPORT_RCR2 registers.

SPORT Transmit Data (SPORT_TX) Register

The SPORT_TX register is a write-only register. Reads produce a peripheral bus error. 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 19-29](#). The SPORT_TX register is shown in [Figure 19-30](#).

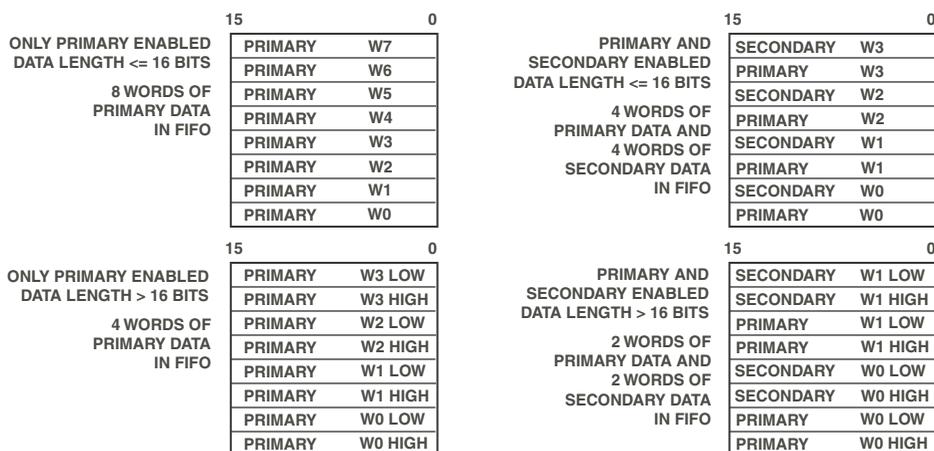


Figure 19-29. SPORT Transmit FIFO Data Ordering

It is important to keep the interleaving of primary and secondary data in the FIFO as shown. This means that peripheral bus/DMA writes to the FIFO must follow an order of primary first, and then secondary, if secondary is enabled. DAB/peripheral bus 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.

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.

The transmit underflow status bit (TUVF) is set in the `SPORT_STAT` 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 SPORT (writing $TXEN = 0$).

If software causes the core processor to attempt a write to a full TX FIFO with a `SPORT_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 `SPORT_TX` register without causing this type of error, read the register's status first. The `TXF` bit in the `SPORT_STAT` register is 0 if space is available for another word in the FIFO.

The `TXF` and `TOVF` status bits in the `SPORT_STAT` register are updated upon writes from the core processor, even when the SPORT is disabled.

SPORT Transmit Data Register (SPORT_TX)

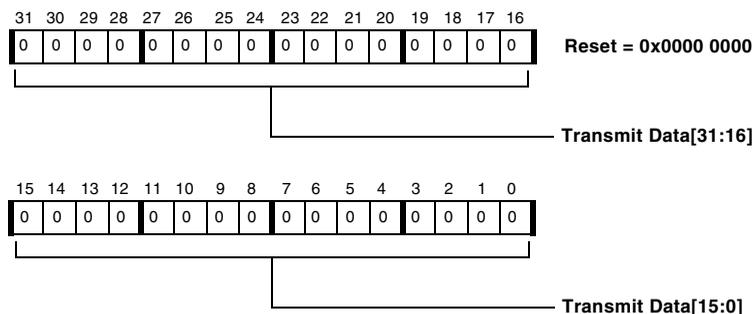


Figure 19-30. SPORT Transmit Data Register

SPORT Receive Data (SPORT_RX) Register

The SPORT_RX register is a read-only register. Writes produce a peripheral bus 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 peripheral bus/DMA reads is important since data is stored in differently depending on the setting of the SLEN and RXSE configuration bits.

Data storage and data ordering in the FIFO are shown in [Figure 19-31](#).

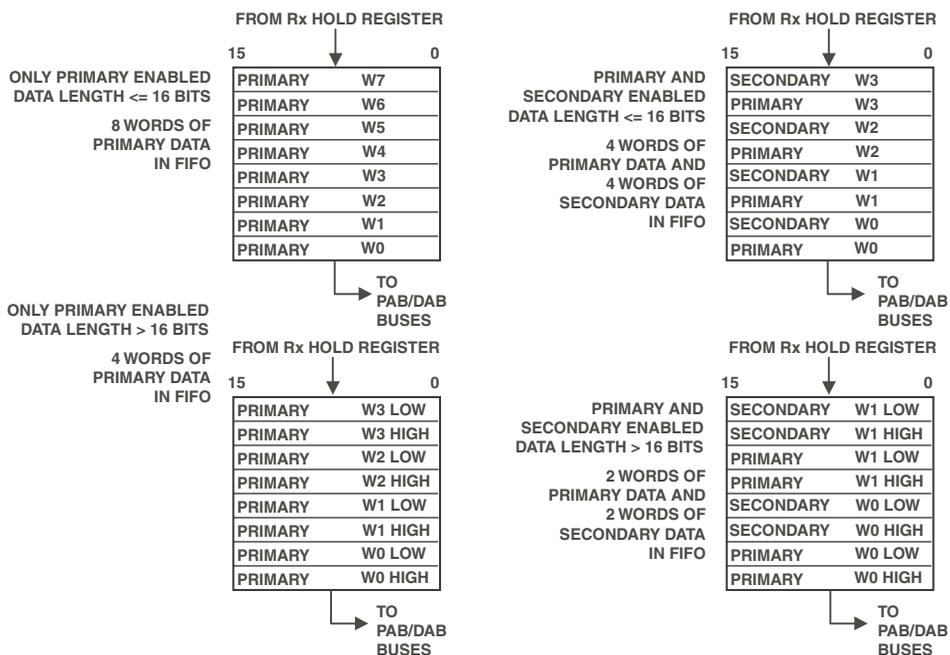


Figure 19-31. SPORT Receive FIFO Data Ordering

The SPORT_RX register is shown in Figure 19-32.

SPORT Receive Data Register (SPORT_RX)

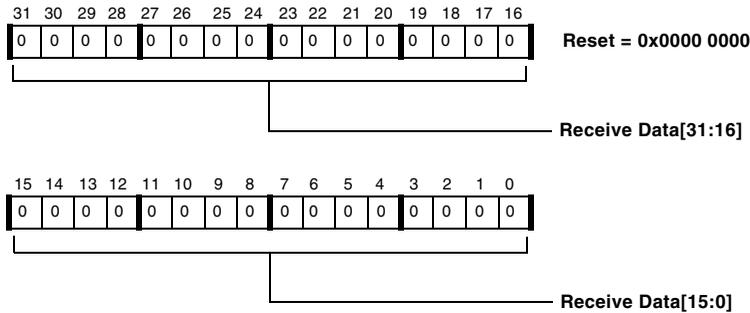


Figure 19-32. SPORT Receive Data Register

When reading from the FIFO for both primary and secondary data, read primary first, followed by secondary. DAB/peripheral bus 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 SPORT_STAT register, and the SPORT error interrupt is asserted. The RUVF bit is a sticky bit and

SPORT Registers

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 SPORT_STAT register). The RUVF status bit is updated even when the SPORT is disabled.

The ROVF status bit is set in the SPORT_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.

SPORT Status (SPORT_STAT) Register

The SPORT_STAT register is used to determine if the access to a SPORT RX or TX FIFO can be made by determining their full or empty status. This register is shown in [Figure 19-33](#).

SPORT Status Register (SPORT_STAT)

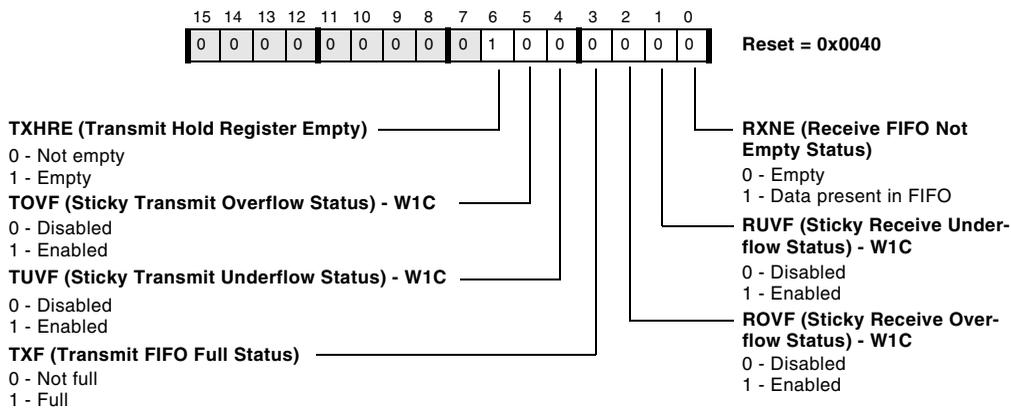


Figure 19-33. SPORT Status Register

The TXF bit in the SPORT_STAT 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 SPORT_TX is empty by clearing the DITFS control bit in the SPORT_TCR1 register. The TUVF status bit is a sticky write-1-to-clear (W1C) bit and is also cleared by disabling the SPORT (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_STAT register. It is a sticky W1C bit and is also cleared by disabling the SPORT (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.

SPORT Transmit and Receive Serial Clock Divider (SPORT_TCLKDIV and SPORT_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 SPORT_TCLKDIV register, shown in [Figure 19-34](#), and the SPORT_RCLKDIV register, shown in [Figure 19-35](#)).

SPORT Registers

SPORT Transmit Serial Clock Divider Register (SPORT_TCLKDIV)

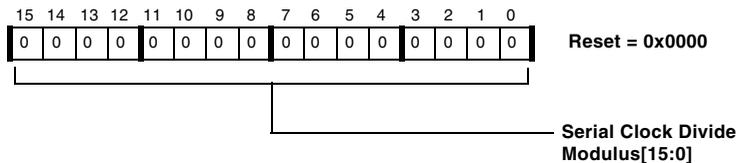


Figure 19-34. SPORT Transmit Serial Clock Divider Register

SPORT Receive Serial Clock Divider Register (SPORT_RCLKDIV)

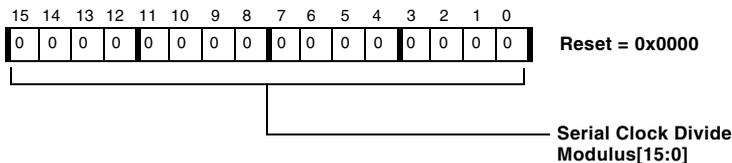


Figure 19-35. SPORT Receive Serial Clock Divider Register

SPORT Transmit and Receive Frame Sync Divider (SPORT_TFSDIV and SPORT_RFSDIV) Registers

The 16-bit `SPORT_TFSDIV` and `SPORT_RFSDIV` registers specify how many transmit or receive clock cycles are counted before generating a TFS or RFS pulse when the frame sync is internally generated. 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. These registers are shown in [Figure 19-36](#) and [Figure 19-37](#).

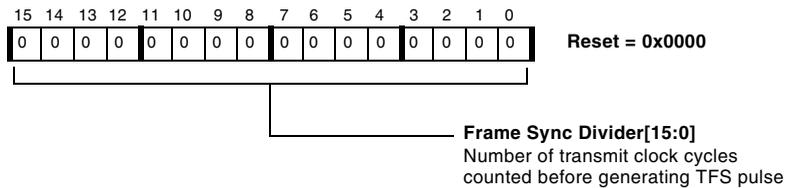
SPORT Transmit Frame Sync Divider Register (SPORT_TFSDIV)

Figure 19-36. SPORT Transmit Frame Sync Divider Register

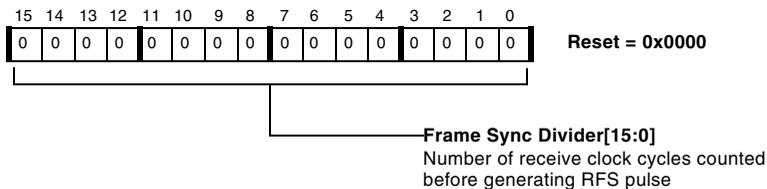
SPORT Receive Frame Sync Divider Register (SPORT_RFSDIV)

Figure 19-37. SPORT Receive Frame Sync Divider Register

SPORT Multichannel Configuration (SPORT_MCMC1 and SPORT_MCMC2) Registers

There are two multichannel configuration registers for each SPORT, shown in [Figure 19-38](#) and [Figure 19-39](#). These registers are used to configure the multichannel operation of the SPORT. The two control registers are shown below.

SPORT Registers

SPORT Multichannel Configuration Register 1 (SPORT_MCMC1)

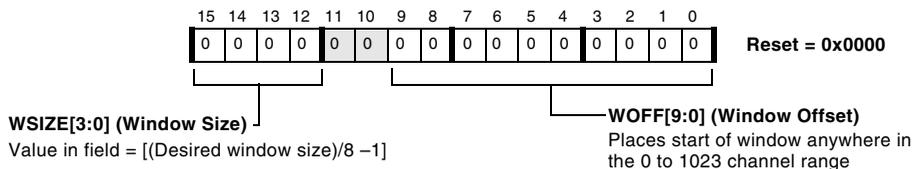


Figure 19-38. SPORT Multichannel Configuration Register 1

SPORT Multichannel Configuration Register 2 (SPORT_MCMC2)

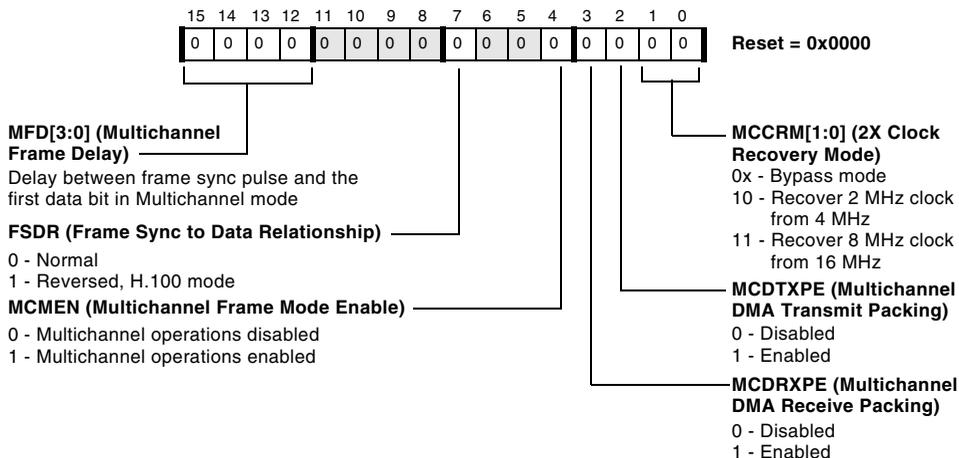


Figure 19-39. SPORT Multichannel Configuration Register 2

SPORT Current Channel (SPORT_CHNL) Register

The 10-bit `CHNL` field in the `SPORT_CHNL` register indicates which channel is currently being serviced during multichannel operation. 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. See [Figure 19-40](#).

SPORT Current Channel Register (`SPORT_CHNL`)

RO

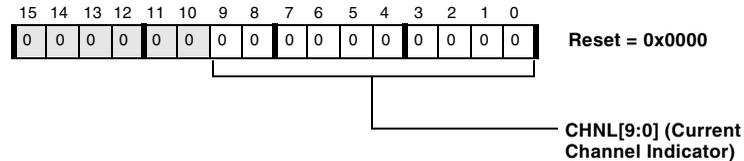


Figure 19-40. SPORT Current Channel Register

SPORT Multichannel Receive Selection (`SPORT_MRCSn`) Registers

The `SPORT_MRCSn` registers (shown in [Figure 19-41](#)) are used to enable and disable individual channels. They specify the active receive channels. There are four registers, each with 32 bits, corresponding to the 128 channels. Setting a bit enables that channel so that the SPORT 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 `SPORT_MRCSn` register causes the SPORT to receive the word in that channel's position of the datastream; the received word is loaded into the RX buffer. When the secondary receive side is enabled by the `RXSE` bit, both inputs are processed on enabled channels. Clearing the bit in the `SPORT_MRCSn` register causes the SPORT to ignore the data on either channel.

SPORT Registers

SPORT Multichannel Receive Select Registers (SPORT_MRCSn)

For all bits, 0 - Channel disabled, 1 - Channel enabled, so SPORT selects that word from multiple word block of data.

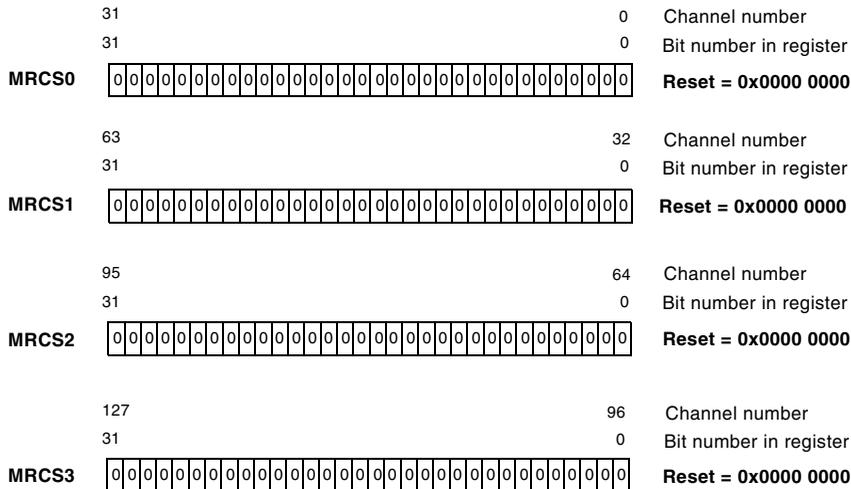


Figure 19-41. SPORT Multichannel Receive Select Registers

SPORT Multichannel Transmit Selection (SPORT_MTCSn) Registers

The `SPORT_MTCSn` registers (shown in [Figure 19-42](#)) are used to enable and disable individual channels. They specify the active transmit channels. There are four registers, each with 32 bits, corresponding to the 128 channels. Setting a bit enables that channel so that the SPORT 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 `SPORT_MTCSn` register causes the SPORT to transmit the word in that channel's position of the datastream. When the secondary transmit side is enabled by the `TXSE` bit, both sides transmit a word on the enabled channel. Clearing the bit in the `SPORT_MTCSn` register

causes a SPORT controllers' data transmit pins to three-state during the time slot of that channel.

SPORT Multichannel Transmit Select Registers (SPORT_MTCSn)

For all bits, 0 - Channel disabled, 1 - Channel enabled, so SPORT selects that word from multiple word block of data.

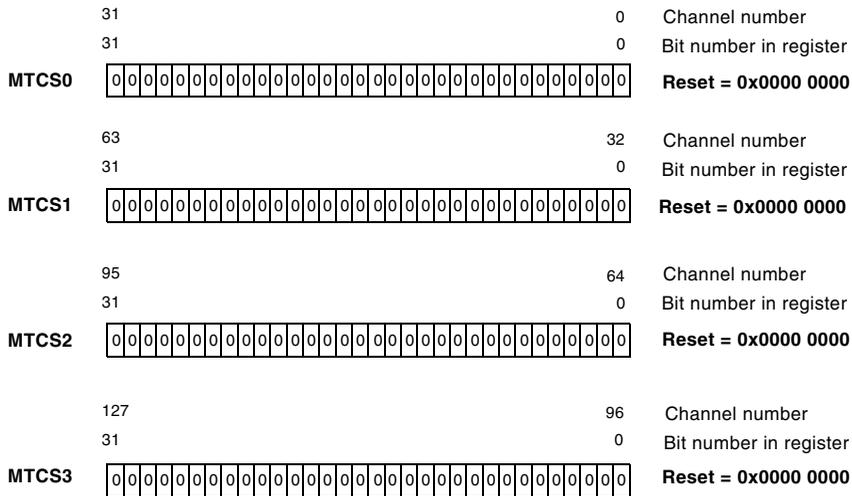


Figure 19-42. SPORT Multichannel Transmit Select Registers

Programming Examples

This section shows an example of typical usage of the SPORT peripheral in conjunction with the DMA controller. See [Listing 19-1](#) through [Listing 19-4](#). These listings assume a processor with at least two SPORTs, SPORT0 and SPORT1.

The SPORT is usually employed for high-speed, continuous serial transfers. The example reflects this, in that the SPORT is set-up for auto-buffered, repeated DMA transfers.

Programming Examples

Because of the many possible configurations, the example uses generic labels for the content of the SPORT's configuration registers (SPORT_RCRn and SPORT_TCRn) and the DMA configuration. An example value is given in the comments, but for the meaning of the individual bits the user is referred to the detailed explanation in this chapter.

The example configures both the receive and the transmit section. Since they are completely independent, the code uses separate labels.

SPORT Initialization Sequence

The SPORT's receiver and transmitter are configured, but they are not enabled yet.

Listing 19-1. SPORT Initialization

```
Program_SPORT_TRANSMITTER_Registers:
    /* Set P0 to SPORT0 Base Address */
    P0.h = hi(SPORT0_TCR1);
    P0.l = lo(SPORT0_TCR1);

    /* Configure Clock speeds */
    R1 = SPORT_TCLK_CONFIG; /* Divider SCLK/TCLK (value 0 to
65535) */
    W[P0 + (SPORT0_TCLKDIV - SPORT0_TCR1)] = R1;
                                     /* TCK divider register */
    /* number of Bitclocks between FrameSynchs -1 (value SPORT_SLEN
to 65535) */
    R1 = SPORT_TFSDIV_CONFIG;
    W[P0 + (SPORT0_TFSDIV - SPORT0_TCR1)] = R1;
                                     /* TFSDIV register */

    /* Transmit configuration */
```

```

    /* Configuration register 2 (for instance 0x000E for 16-bit
wordlength) */
    R1 = SPORT_TRANSMIT_CONF_2;
    W[P0 + (SPORT0_TCR2 - SPORT0_TCR1)] = R1;
    /* Configuration register 1 (for instance 0x4E12 for inter-
nally generated clk and framesync) */
    R1 = SPORT_TRANSMIT_CONF_1;
    W[P0] = R1;
    ssync;
    /* NOTE: SPORT0 TX NOT enabled yet (bit 0 of TCR1 must be zero) */

Program_SPORT_RECEIVER_Registers:
    /* Set P0 to SPORT0 Base Address */
    P0.h = hi(SPORT0_RCR1);
    P0.l = lo(SPORT0_RCR1);

    /* Configure Clock speeds */
    R1 = SPORT_RCLK_CONFIG; /* Divider SCLK/RCLK (value 0 to
65535) */
    W[P0 + (SPORT0_RCLKDIV - SPORT0_RCR1)] = R1; /* RCK divider
register */
    /* number of Bitclock between FrameSyncs -1 (value SPORT_SLEN
to 65535) */
    R1 = SPORT_RFSDIV_CONFIG;
    W[P0 + (SPORT0_RFSDIV - SPORT0_RCR1)] = R1;
                                     /* RFSDIV register */

    /* Receive configuration */
    /* Configuration register 2 (for instance 0x000E for 16-bit
wordlength) */
    R1 = SPORT_RECEIVE_CONF_2;
    W[P0 + (SPORT0_RCR2 - SPORT0_RCR1)] = R1;
    /* Configuration register 1 (for instance 0x4410 for external
clk and framesync) */

```

Programming Examples

```
R1 = SPORT_RECEIVE_CONF_1;
W[P0] = R1;
ssync; /* NOTE: SPORT0 RX NOT enabled yet (bit 0 of RCR1 must
be zero) */
```

DMA Initialization Sequence

Next the DMA channels for receive (channel3 in this example) and for transmit (channel4 in this example) are set up for auto-buffered, one-dimensional, 32-bit transfers. Again, there are other possibilities, so generic labels have been used, with a particular value shown in the comments.

Note that the DMA channels can be enabled at the end of the configuration since the SPORT is not enabled yet. However, if preferred, the user can enable the DMA later, immediately before enabling the SPORT. The only requirement is that the DMA channel be enabled before the associated peripheral is enabled to start the transfer.

Listing 19-2. DMA Initialization

```
Program_DMA_Controller:

/* Receiver (DMA channel 3) */
/* Set P0 to DMA Base Address */
P0.l = lo(DMA3_CONFIG);
P0.h = hi(DMA3_CONFIG);

/* Configuration (for instance 0x108A for Autobuffer, 32-bit
wide transfers) */
R0 = DMA_RECEIVE_CONF(z);
W[P0] = R0; /* configuration register */
```

```

/* rx_buf = Buffer in Data memory (divide count by four because
of 32-bit DMA transfers) */
R1 = (length(rx_buf)/4)(z);
W[P0 + (DMA3_X_COUNT - DMA3_CONFIG)] = R1;
/* X_count register */
R1 = 4(z); /* 4 bytes in a 32-bit transfer */
W[P0 + (DMA3_X_MODIFY - DMA3_CONFIG)] = R1;
/* X_modify register */

/* start_address register points to memory buffer
to be filled */
R1.l = rx_buf;
R1.h = rx_buf;
[P0 + (DMA3_START_ADDR - DMA3_CONFIG)] = R1;

BITSET(R0,0); /* R0 still contains value of CONFIG register -
set bit 0 */
W[P0] = R0; /* enable DMA channel (SPORT not enabled yet) */

/* Transmitter (DMA channel 4) */
/* Set P0 to DMA Base Address */
P0.l = lo(DMA4_CONFIG);
P0.h = hi(DMA4_CONFIG);
/* Configuration (for instance 0x1088 for Autobuffer, 32-bit
wide transfers) */
R0 = DMA_TRANSMIT_CONF(z);
W[P0] = R0; /* configuration register */

/* tx_buf = Buffer in Data memory (divide count by four because
of 32-bit DMA transfers) */
R1 = (length(tx_buf)/4)(z);
W[P0 + (DMA4_X_COUNT - DMA4_CONFIG)] = R1;
/* X_count register */
R1 = 4(z); /* 4 bytes in a 32-bit transfer */

```

Programming Examples

```
W[P0 + (DMA4_X_MODIFY - DMA4_CONFIG)] = R1;
    /* X_modify register */

/* start_address register points to memory buffer to be
   transmitted from */
R1.l = tx_buf;
R1.h = tx_buf;
[P0 + (DMA4_START_ADDR - DMA4_CONFIG)] = R1;

BITSET(R0,0); /* R0 still contains value of CONFIG register -
               set bit 0 */
W[P0] = R0; /* enable DMA channel (SPORT not enabled yet) */
```

Interrupt Servicing

The receive channel and the transmit channel will each generate an interrupt request if so programmed. The following code fragments show the minimum actions that must be taken. Not shown is the programming of the core and system event controllers.

Listing 19-3. Servicing an Interrupt

```
RECEIVE_ISR:
    [--SP] = RETI; /* nesting of interrupts */

/* clear DMA interrupt request */
P0.h = hi(DMA3_IRQ_STATUS);
P0.l = lo(DMA3_IRQ_STATUS);
R1    = 1;
W[P0] = R1.l; /* write one to clear */

RETI = [SP++];
rti;
```

```

TRANSMIT_ISR:
    [--SP] = RETI; /* nesting of interrupts */

    /* clear DMA interrupt request */
    P0.h = hi(DMA4_IRQ_STATUS);
    P0.l = lo(DMA4_IRQ_STATUS);
    R1 = 1;
    W[P0] = R1.l; /* write one to clear */

    RETI = [SP++];
    rti;

```

Starting a Transfer

After the initialization procedure outlined in the previous sections, the receiver and transmitter are enabled. The core may just wait for interrupts.

Listing 19-4. Starting a Transfer

```

/* Enable Sport0 RX and TX */
P0.h = hi(SPORT0_RCR1);
P0.l = lo(SPORT0_RCR1);
R1 = W[P0](Z);
BITSET(R1,0);
W[P0] = R1;
ssync; /* Enable Receiver (set bit 0) */

P0.h = hi(SPORT0_TCR1);
P0.l = lo(SPORT0_TCR1);
R1 = W[P0](Z);
BITSET(R1,0);
W[P0] = R1;
ssync; /* Enable Transmitter (set bit 0) */

```

Unique Information for the ADSP-BF51x Processor

```
/* dummy wait loop (do nothing but waiting for interrupts) */  
wait_forever:  
    jump wait_forever;
```

Unique Information for the ADSP-BF51x Processor

None.

20 PARALLEL PERIPHERAL INTERFACE

This chapter describes the parallel peripheral interface (PPI). Following an overview and a list of key features are a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the number of PPIs for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For PPI DMA channel assignments, refer to [Table 6-7 on page 6-104](#) in [Chapter 6, “Direct Memory Access”](#).

For PPI interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how each of the PPIs is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for each PPI, refer to [Appendix A, “System MMR Assignments”](#).

PPI behavior for the ADSP-BF51x that differs from the general information in this chapter can be found in the section [“Unique Information for the ADSP-BF51x Processor” on page 20-37](#).

Overview

The PPI is a half-duplex, bidirectional port accommodating up to 16 bits of data. It has a dedicated clock pin and three multiplexed frame sync 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).

Features

The PPI includes these features:

- Half-duplex, bidirectional parallel port
- Supports up to 16 bits of data
- Programmable clock and frame sync polarities
- ITU-R 656 support
- Interrupt generation on overflow and underrun

Typical peripheral devices that can be interfaced to the PPI port:

- A/D converters
- D/A converters
- LCD panels
- CMOS sensors
- Video encoders
- Video decoders

Interface Overview

Figure 20-1 shows a block diagram of the PPI.

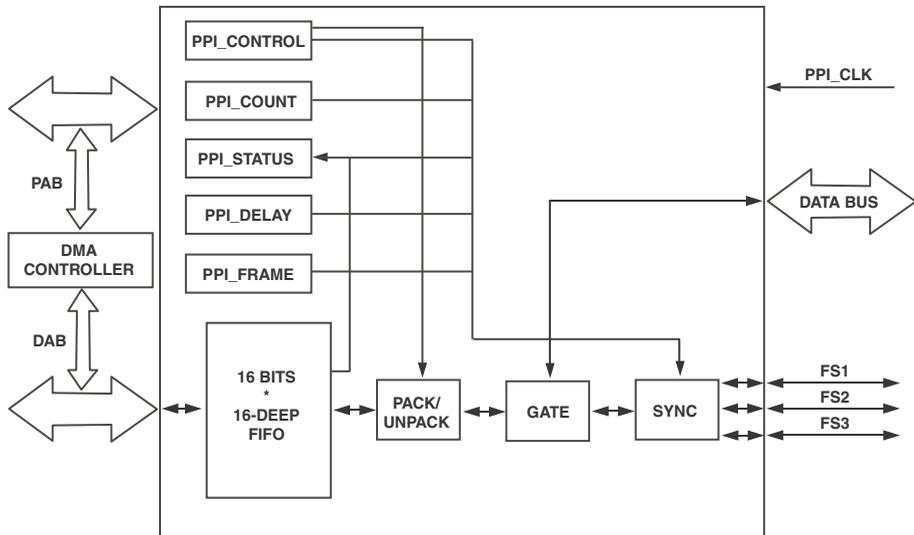


Figure 20-1. PPI Block Diagram

The `PPI_CLK` pin accepts an external clock input. It cannot source a clock internally.

i When the `PPI_CLK` is not free-running, there may be additional latency cycles before data gets received or transmitted. In RX and TX modes, there may be at least 2 cycles latency before valid data is received or transmitted.

The `PPI_CLK` not only supplies the PPI module itself, but it also can clock one or more GP Timers to work synchronously with the PPI. Depending on PPI operation mode, the `PPI_CLK` can either equal or invert the `TMRCLK` input. For more information, see [Chapter 10, “General-Purpose Timers”](#).

Description of Operation

Table 20-1 shows all the possible modes of operation for the PPI.

Table 20-1. PPI Possible Operating Modes

PPI Mode	# 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	0
RX mode, 2 or 3 external frame syncs	3	0	10	11	0 or 1	0 or 1	0
RX mode, 2 or 3 internal frame syncs	3	0	01	11	0 or 1	0 or 1	0
RX mode, ITU-R 656, active field only	embedded	0	00	00	0 or 1	0	0 or 1
RX mode, ITU-R 656, vertical blanking only	embedded	0	00	10	0 or 1	0	0
RX mode, ITU-R 656, entire field	embedded	0	00	01	0 or 1	0	0
TX mode, 0 frame syncs	0	1	00	00	0 or 1	0 or 1	0
TX mode, 1 internal or external frame sync	1	1	00	11	0 or 1	0 or 1	0
TX mode, 2 external frame syncs	2	1	01	11	0 or 1	0 or 1	0

Table 20-1. PPI Possible Operating Modes (Continued)

PPI Mode	# 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	0
TX mode, 2 or 3 internal frame syncs, FS3 sync'd to FS2 assertion	3	1	11	11	0 or 1	0 or 1	0

Functional Description

The following sections describe the function of the PPI.

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 20-2](#), and [Figure 20-3](#) 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

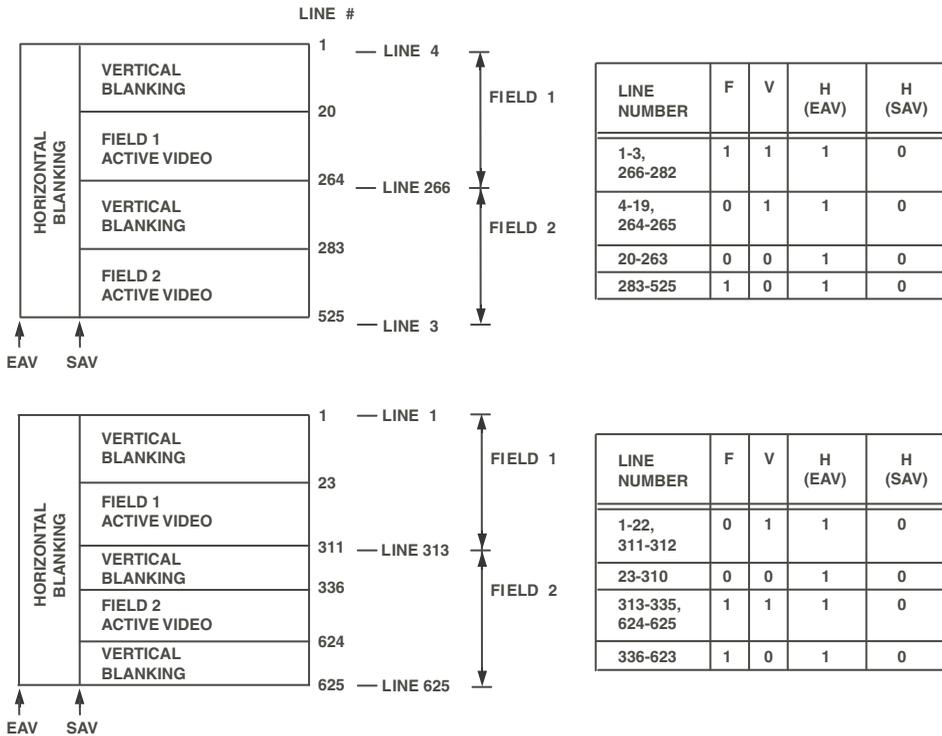


Figure 20-3. Typical Video Frame Partitioning for NTSC/PAL Systems for ITU-R BT.656-4

The SAV and EAV codes are shown in more detail in [Table 20-2](#). 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

Functional Description

- $V = 1$ during vertical blanking
- $V = 0$ when not in vertical blanking
- $H = 0$ at SAV
- $H = 1$ at EAV
- $P3 = V \text{ XOR } H$
- $P2 = F \text{ XOR } H$
- $P1 = F \text{ XOR } V$
- $P0 = F \text{ XOR } V \text{ XOR } H$

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 20-2. Control Byte Sequences for 8-bit 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

Functional Description

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 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 `PPI_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 `PPI_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 does not 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 signaling is embedded in the datastream in ITU-R 656 modes, the `PPI_COUNT` register is not necessary. However, the `PPI_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

Functional Description

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 `PPI_FRAME`. If there is a mismatch, the `FT_ERR` bit in the `PPI_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 `PPI_STATUS` register. This way, an interrupt service routine (ISR) can discern which field was just read in.

General-Purpose PPI Modes

The general-purpose PPI modes are intended to suit a wide variety of data capture and transmission applications. [Table 20-3](#) 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 functions.

Table 20-3. General-Purpose PPI Modes

GP PPI Mode	PPI_FS1 Direction	PPI_FS2 Direction	PPI_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 (if used)	Input
RX mode, 2 or 3 internal frame syncs	Output	Output	Output (if used)	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

Table 20-3. General-Purpose PPI Modes (Continued)

GP PPI Mode	PPI_FS1 Direction	PPI_FS2 Direction	PPI_FS3 Direction	Data Direction
TX mode, 2 external frame syncs	Input	Input	Not used	Output
TX mode, 1 internal frame sync	Output	Not used	Not used	Output
TX mode, 2 or 3 internal frame syncs	Output	Output	Output (if used)	Output

Figure 20-6 illustrates the general flow of the general purpose PPI modes. The top of the diagram shows an example of RX mode with one external frame sync. After the PPI receives the hardware frame sync pulse (PPI_FS1), it delays for the duration of the PPI_CLK cycles programmed into PPI_DELAY. The DMA controller then transfers in the number of samples specified by PPI_COUNT. Every sample that arrives after this, but before the next PPI_FS1 frame sync arrives, is ignored and not transferred onto the DMA bus.

 If the next PPI_FS1 frame sync arrives before the specified PPI_COUNT samples have been read in, the sample counter reinitializes to 0 and starts to count up to PPI_COUNT again. This situation can cause the DMA channel configuration to lose synchronization with the PPI transfer process.

The bottom of Figure 20-6 shows an example of TX mode, one internal frame sync. After PPI_FS1 is asserted, there is a latency of one PPI_CLK cycle, and then there is a delay for the number of PPI_CLK cycles programmed into PPI_DELAY. Next, the DMA controller transfers out the number of samples specified by PPI_COUNT. No further DMA takes place until the next PPI_FS1 sync and programmed delay occur.

Functional Description

 If the next PPI_FS1 frame sync arrives before the specified PPI_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.

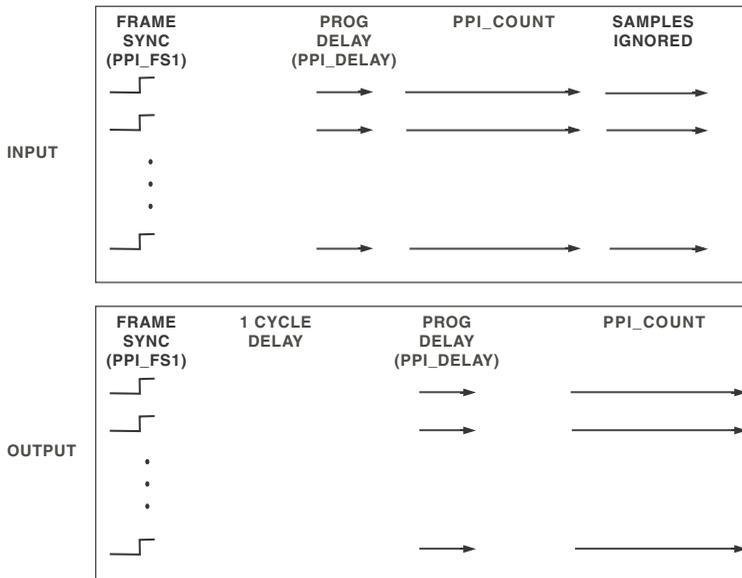


Figure 20-6. General Flow for GP Modes (Assumes Positive Assertion of PPI_FS1)

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 20-1 on page 20-4](#) 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 `PPI_CONTROL` register.

- **External trigger:** An external source sends a single frame sync (tied to `PPI_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 byte 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 frame syncs are level-sensitive signals. The 1-sync mode is intended for analog-to-digital converter (ADC) applications. The top part of [Figure 20-7](#) shows a typical illustration of the system setup for this mode.

The 3-sync mode shown at the bottom of [Figure 20-7](#) supports video applications that use hardware signaling (`HSYNC`, `VSYNC`, `FIELD`) in accordance with the ITU-R 601 recommendation. The mapping for the frame syncs in this mode is `PPI_FS1 = HSYNC`, `PPI_FS2 = VSYNC`, `PPI_FS3 = FIELD`. Please refer to [“Frame Synchronization in GP Modes” on page 20-19](#) for more information about frame syncs in this mode.

Functional Description

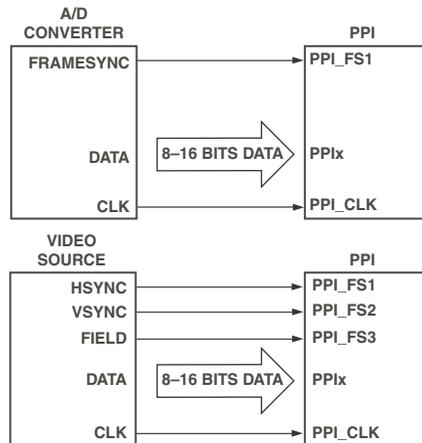


Figure 20-7. RX Mode, External Frame Syncs

A 2-sync mode is supported by not enabling the PPI_FS3 pin. See the *Product Specific Implementation* section for information on how this is achieved on this processor.

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 PPI_FS1 and PPI_FS2, and then reading data into the PPI. The PPI_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 20-8](#) shows a sample application for this mode.

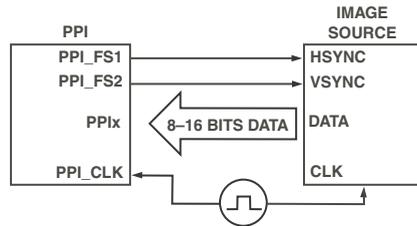


Figure 20-8. RX Mode, Internal Frame Syncs

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 20-1 on page 20-4](#) 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 PPI_CLK. See [Figure 20-9](#) for an illustration of this mode.

- 

In this mode, there is a delay of up to 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).

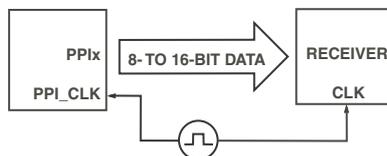


Figure 20-9. TX Mode, 0 Frame Syncs

Functional Description

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 20-10](#) shows the 1-sync case, while the bottom diagram illustrates the 2-sync mode.

⚡ There is a mandatory delay of 1.5 PPI_CLK cycles, plus the value programmed in `PPI_DELAY`, between assertion of the external frame sync(s) and the transfer of valid data out through the PPI.

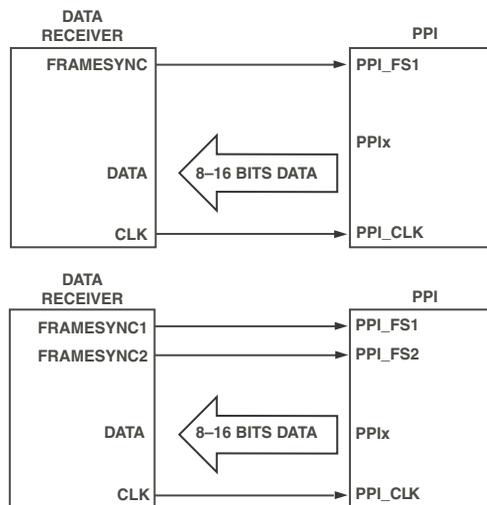


Figure 20-10. TX Mode, 1 or 2 External Frame Syncs

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 20-11](#) 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 20-11](#). A 2-sync mode is implicitly supported by leaving `PPI_FS3` unconnected in this case.

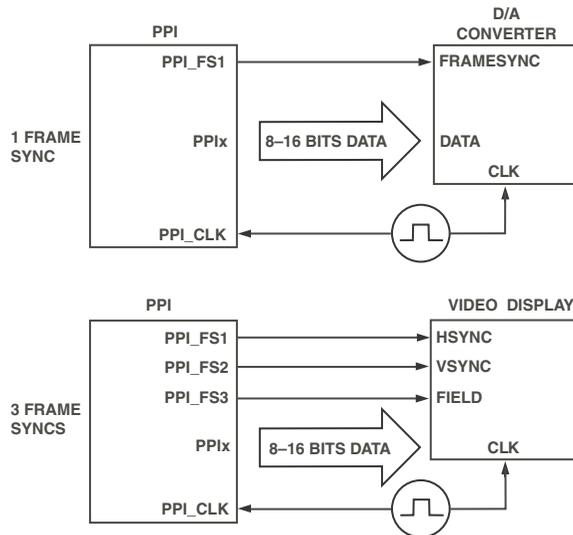


Figure 20-11. PPI GP Output

Frame Synchronization in GP Modes

Frame synchronization in general purpose 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, PPI_FS1 and PPI_FS2 link directly to the pulsewidth modulation (PWM) circuits of general purpose timers. See the [Chapter 10, “General-Purpose Timers”](#) for information on how this is achieved on this processor. 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 PPI_CLK, not by SCLK (as during conventional timer PWM operation). If PPI_FS2 is not used in the configured PPI mode, its corresponding timer operates as it normally would, unrestricted in functionality. The state of PPI_FS3 depends completely on the

Functional Description

state of PPI_FS1 and/or PPI_FS2, so PPI_FS3 has no inherent programmability.



To program PPI_FS1 and/or PPI_FS2 for operation in an internal frame sync mode:

1. Configure and enable DMA for the PPI. See [“DMA Operation” on page 20-22](#).
2. Configure the width and period for each frame sync signal via the appropriate `TIMER_WIDTH` and `TIMER_PERIOD` registers.
3. Set up the appropriate `TIMER_CONFIG` register(s) for `PWM_OUT` mode. This includes setting `CLK_SEL` to 1 and `TIN_SEL` to 1 for each timer involved.
4. Write to `PPI_CONTROL` to configure and enable the PPI.
5. Write to `TIMER_ENABLE` to enable the appropriate timer(s).



It is important to guarantee proper frame sync polarity between the PPI and timer peripherals. To do this, make sure that if `PPI_CONTROL[15:14] = b#10` or `b#11`, the `PULSE_HI` bit is cleared in the appropriate `TIMER_CONFIG` register(s). Likewise, if `PPI_CONTROL[15:14] = b#00` or `b#01`, the `PULSE_HI` bit should be set in the appropriate `TIMER_CONFIG` register(s).

To switch to another PPI mode not involving internal frame syncs:

1. Disable the PPI (using `PPI_CONTROL`).
2. Disable the appropriate timer(s) (using `TIMER_DISABLE`).

Modes With External Frame Syncs

In RX modes with external frame syncs, the `PPI_FS1` and `PPI_FS2` pins become edge-sensitive inputs. In such modes the timers associated with the `PPI_FS1` and `PPI_FS2` pins can still be used for a purpose not involving the actual pin. However, timer access to a `TMRx` pin is disabled when the

PPI is using that pin for a PPI_FSx frame sync input function. For modes that do not require PPI_FS2, the associated timer is not restricted in functionality and can be operated as if the PPI were not being used (that is, the TMR1 pin becomes available for timer use as well). For more information on configuring and using the timers, please refer to the *General-Purpose Timers* chapter.

i In RX mode with 3 external frame syncs, the start of frame detection occurs where a PPI_FS2 assertion is followed by an assertion of PPI_FS1 while PPI_FS3 is low. This happens at the start of field 1. Note that PPI_FS3 only needs to be low when PPI_FS1 is asserted, not when PPI_FS2 asserts. Also, PPI_FS3 is only used to synchronize to the start of the very first frame after the PPI is enabled. It is subsequently ignored.

In TX modes with external frame syncs, the PPI_FS1 and PPI_FS2 pins are treated as edge-sensitive inputs. In this mode, it is not necessary to configure the timer(s) associated with the frame sync(s) as input(s), or to enable them via the TIMER_ENABLE register. Additionally, the actual timers themselves are available for use, even though the timer pin(s) are taken over by the PPI. In this case, there is no requirement that the timebase (configured by TIN_SEL in TIMERx_CONFIG) be PPI_CLK.

However, if using a timer whose pin is connected to an external frame sync, be sure to disable the pin via the OUT_DIS bit in TIMER_CONFIG. Then the timer itself can be configured and enabled for non-PPI use without affecting PPI operation in this mode. For more information, see the *General-Purpose Timers* chapter.

Programming Model

The following sections describe the PPI programming model.

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 the *Direct Memory Access* chapter.

The PPI DMA channel can be configured for either transmit or receive operation, and it has a maximum throughput of $(\text{PPI_CLK}) \times (16 \text{ bits/transfer})$. In modes where data lengths are greater than eight bits, only one element can be clocked in per PPI_CLK cycle, and this results in reduced bandwidth (since no packing is possible). The highest throughput is achieved with 8-bit data and `PACK_EN = 1` (packing mode enabled). Note for 16-bit packing mode, there must be an even number of data elements.

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 `DMA_XCOUNT` and `DMA_YCOUNT` MMRs allows for flexible data interrupt points. For example, assume the DMA registers `XMODIFY = YMODIFY = 1`. Then, if a data frame contains 320 x 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 `DMA_CONFIG`) interrupts on every row transferred, for the entire frame.
- Setting `XCOUNT = 320`, `YCOUNT = 240`, and `DI_SEL = 0` interrupts only on the completion of the frame (when 240 rows of 320 bytes have been transferred).

- Setting `XCOUNT = 38,400 (320 x 120)`, `YCOUNT = 2`, and `DI_SEL = 1` causes an interrupt when half of the frame has been transferred, and again when the whole frame has been transferred.

The general procedure for setting up DMA operation with the PPI follows.

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 `PPI_CONTROL`.
5. If internally generated frame syncs are used, write to the `TIMER_ENABLE` register to enable the timers linked to the PPI frame syncs.

[Figure 20-12](#) shows a flow diagram detailing the steps on how to configure the PPI for the various modes of operation.

Programming Model

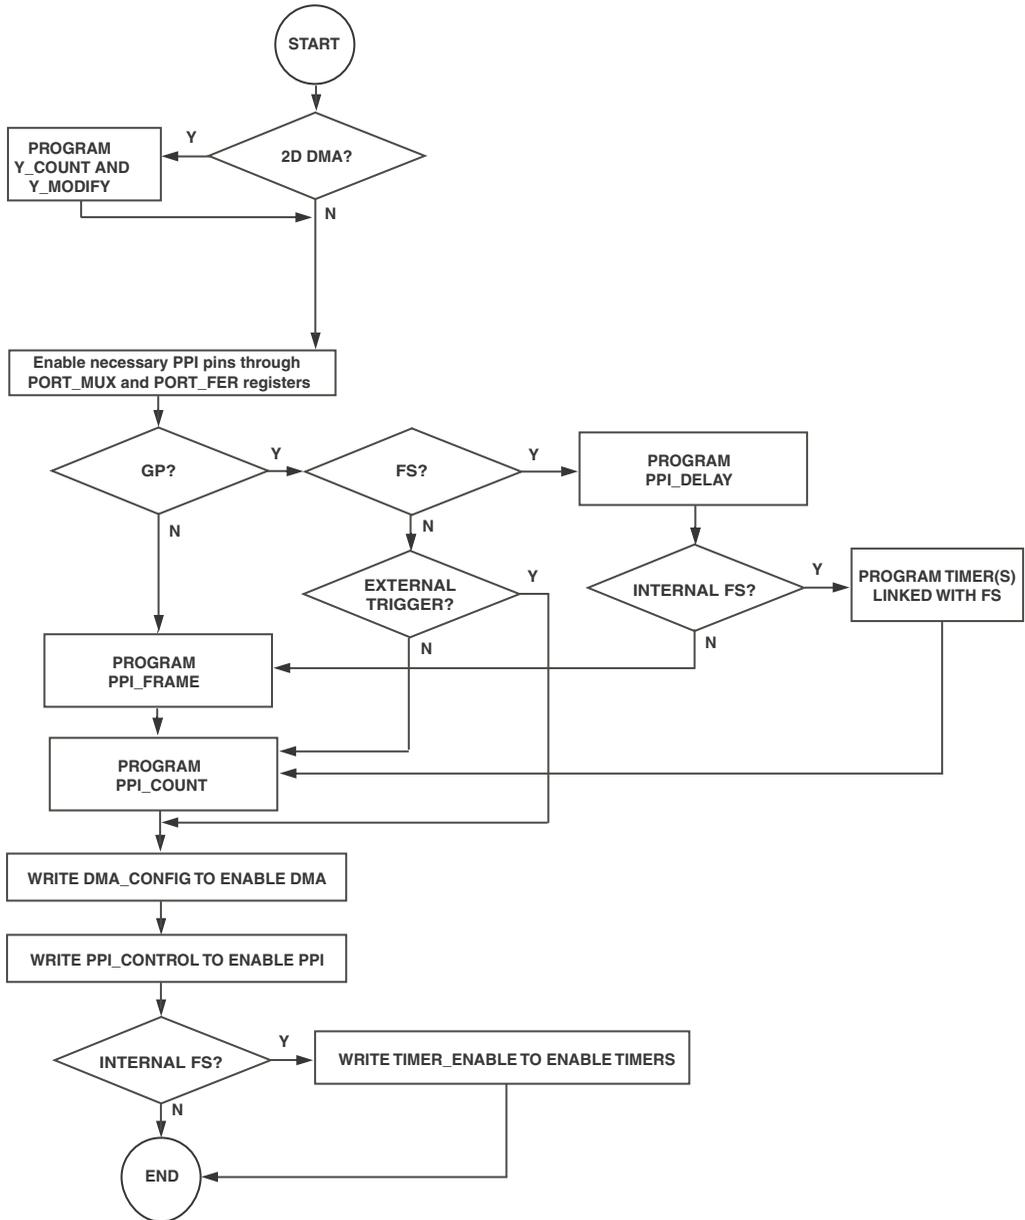


Figure 20-12. PPI Flow Diagram

PPI Registers

The PPI has five memory-mapped registers (MMRs) that regulate its operation. These registers are the PPI control register (`PPI_CONTROL`), the PPI status register (`PPI_STATUS`), the delay count register (`PPI_DELAY`), the transfer count register (`PPI_COUNT`), and the lines per frame register (`PPI_FRAME`).

Descriptions and bit diagrams for each of these MMRs are provided in the following sections.

PPI Control Register (`PPI_CONTROL`)

The `PPI_CONTROL` register configures the PPI for operating mode, control signal polarities, and data width of the port. See [Figure 20-13](#) for a bit diagram of this MMR.

The `POLC` and `POLS` bits allow for selective signal inversion of the `PPI_CLK` and `PPI_FS1/PPI_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 pins unused by the PPI as a result of the `DLEN` setting are free for use in their other functions.



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.

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.

PPI Registers

PPI Control Register (PPI_CONTROL)

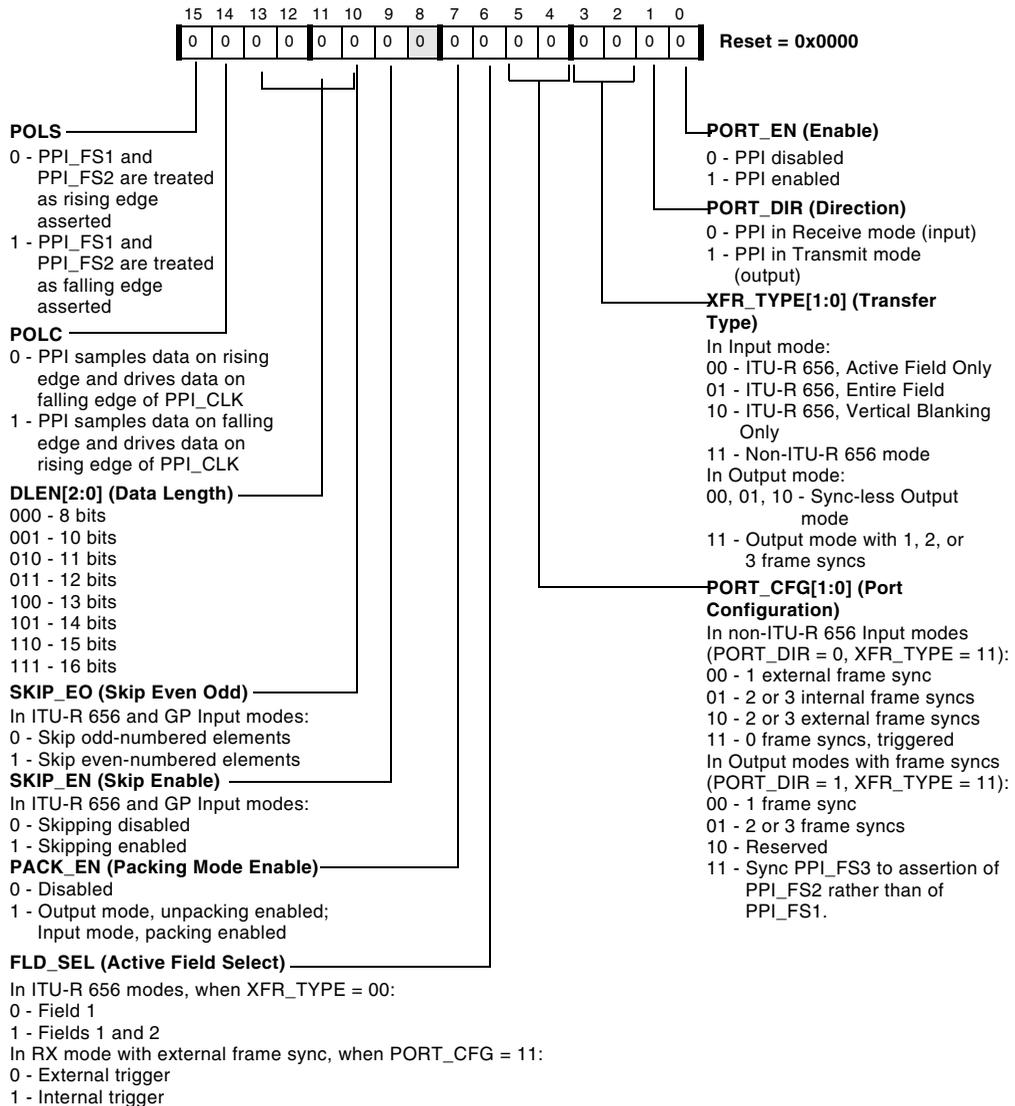


Figure 20-13. PPI Control Register

When the `SKIP_EN` bit is set, the `SKIP_E0` 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 processors to the same incoming video stream. One processor could handle luma processing and the other (whose `SKIP_E0` bit is set differently from the first processor's) 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 only has meaning when the PPI port width (selected by `DLEN[2:0]`) is eight bits. Every `PPI_CLK`-initiated event on the DMA bus (that is, an input or output operation) handles 16-bit entities. In other words, an input port width of ten bits still results in a 16-bit input word for every `PPI_CLK`; the upper 6 bits are 0s. Likewise, a port width of eight bits also results in a 16-bit input word, with the upper eight bits all 0s. In the case of 8-bit data, it is usually more efficient to pack this information so that there are two bytes of data for every 16-bit word. This is the function of the `PACK_EN` bit. When set, it enables packing for all RX modes.

Consider this data transported into the PPI via DMA:

0xCE, 0xFA, 0xFE, 0xCA....

- With `PACK_EN` set:

This is read into the PPI, configured for an 8-bit port width:

0xCE, 0xFA, 0xFE, 0xCA...

- This is transferred onto the DMA bus:

0xFACE, 0xCAFE,...

- With `PACK_EN` cleared:

This is read into the PPI:

0xCE, 0xFA, 0xFE, 0xCA,...

PPI Registers

- This is transferred onto the DMA bus:

0x00CE, 0x00FA, 0x00FE, 0x00CA,...

For TX modes, setting `PACK_EN` enables unpacking of bytes. Consider this data in memory, to be transported out through the PPI via DMA:

0xFACE CAFE....

(0xFA and 0xCA are the two most significant bits (MSBs) of their respective 16-bit words)

- With `PACK_EN` set:

This is DMAed to the PPI:

0xFACE, 0xCAFE,...

- This is transferred out through the PPI, configured for an 8-bit port width (note LSBs are transferred first):

0xCE, 0xFA, 0xFE, 0xCA,...

- With `PACK_EN` cleared:

This is DMAed to the PPI:

0xFACE, 0xCAFE,...

- This is transferred out through the PPI, configured for an 8-bit port width:

0xCE, 0xFE,...

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 20-1 on page 20-4](#) for the possible operating modes for the PPI.

The `XFR_TYPE[1:0]` field configures the PPI for various modes of operation. Refer to [Table 20-1 on page 20-4](#) to see how `XFR_TYPE[1:0]` interacts with other bits in `PPI_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 20-19](#) for more information.

PPI Status Register (PPI_STATUS)

The `PPI_STATUS` register, shown in [Figure 20-14](#), contains bits that provide information about the current operating state of the PPI.

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` and `ERR_DET = 1`, all preamble errors that have occurred have been 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_IMASK` register.

PPI Registers

PPI Status Register (PPI_STATUS)

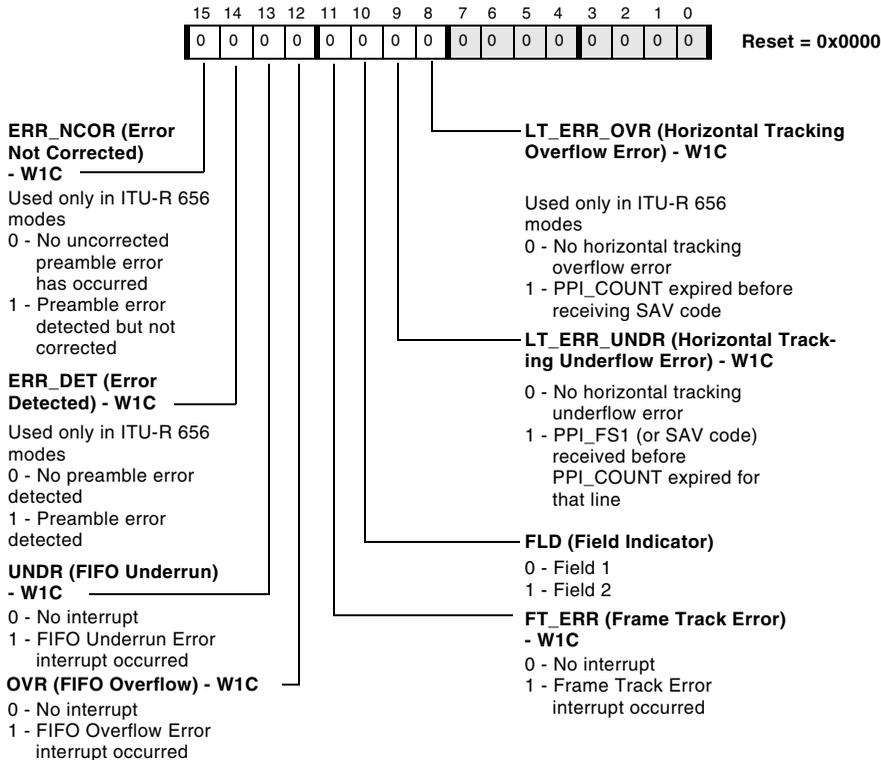


Figure 20-14. PPI Status Register

The FT_ERR bit is sticky and indicates, when set, that a frame track error has occurred. In this condition, the programmed number of lines per frame in PPI_FRAME does not match up with the “frame start detect” condition (see the information note on page 20-34). A frame track error generates a PPI error interrupt, unless this condition is masked off in the SIC_IMASK 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 PPI_FS3 (in other RX modes). It is valid for input modes only. The state of FLD reflects the current state of the F or PPI_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_IMASK` register.

 The PPI FIFO is 16 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_IMASK` register.

The `LT_ERR_OVR` and `LT_ERR_UNDR` bits are sticky and indicate, when set, that a line track error has occurred. These bits are valid for RX modes with recurring frame syncs only. If one of these bits is set, the programmed number of samples in `PPI_COUNT` did not match up with the actual number of samples counted between assertions of `PPI_FS1` (for general-purpose modes) or start of active video (SAV) codes (for ITU-R 656 modes). If the PPI error interrupt is enabled in the `SIC_IMASK` register, an interrupt request is generated when one of these bits is set.

The `LT_ERR_OVR` flag signifies that a horizontal tracking overflow has occurred, where the value in `PPI_COUNT` was reached before a new SAV code was received. This flag does not apply for non ITU-R 656 modes; in this case, once the value in `PPI_COUNT` is reached, the PPI simply stops counting until receiving the next `PPI_FS1` frame sync.

The `LT_ERR_UNDR` flag signifies that a horizontal tracking underflow has occurred, where a new SAV code or `PPI_FS1` assertion occurred before the value in `PPI_COUNT` was reached.

PPI Delay Count Register (PPI_DELAY)

The `PPI_DELAY` register, shown in [Figure 20-15](#), 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 `PPI_CLK` cycles to delay after assertion of `PPI_FS1` before starting to read in or write out data.

 Note in TX modes using at least one frame sync, there is a one-cycle delay beyond what is specified in the `PPI_DELAY` register.

PPI Delay Count Register (PPI_DELAY)

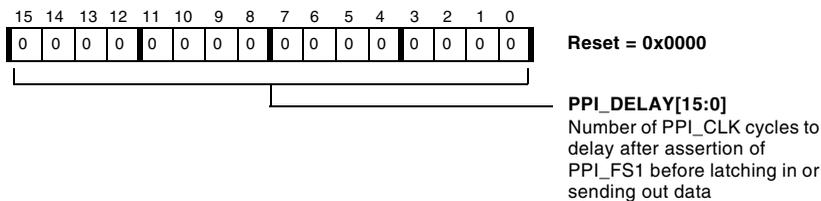


Figure 20-15. PPI Delay Count Register

PPI Transfer Count Register (PPI_COUNT)

The `PPI_COUNT` register, shown in [Figure 20-16](#), is used in all modes except “RX mode with 0 frame syncs, external trigger” and “TX mode with 0 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 `PPI_COUNT` to 99.

 Take care to ensure that the number of samples programmed into `PPI_COUNT` is in keeping with the number of samples expected during the “horizontal” interval specified by `PPI_FS1`.

PPI Transfer Count Register (PPI_COUNT)

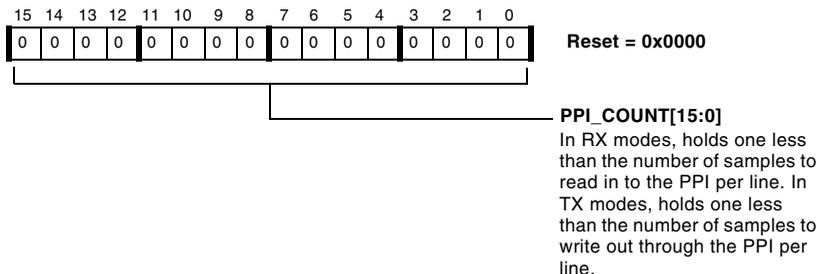


Figure 20-16. PPI Transfer Count Register

PPI Lines Per Frame Register (PPI_FRAME)

The `PPI_FRAME` register, shown in [Figure 20-17](#), is used in all TX and RX modes with two or three 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 `PPI_FS2` assertions, regardless of the state of `PPI_FS3`. A line is defined as a complete `PPI_FS1` cycle. In these modes, `PPI_FS3` is used only to determine the original “frame start” each time the PPI 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 `PPI_FRAME` have been transferred, a frame track error results, and the `FT_ERR` bit in `PPI_STATUS` is set. However, the PPI still automatically reinitializes to count to the value programmed in `PPI_FRAME`, and data transfer continues.

Programming Examples

i 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 three external frame syncs, a frame start detect refers to a condition where a PPI_FS2 assertion is followed by an assertion of PPI_FS1 while PPI_FS3 is low. This occurs at the start of field 1. Note that PPI_FS3 only needs to be low when PPI_FS1 is asserted, not when PPI_FS2 asserts. Also, PPI_FS3 is only used to synchronize to the start of the very first frame after the PPI is enabled. It is subsequently ignored.

When using RX mode with three external frame syncs, and only two syncs are needed, configure the PPI for 3-frame-sync operation and provide an external pull-down to GND for the PPI_FS3 pin.

PPI Lines Per Frame Register (PPI_FRAME)

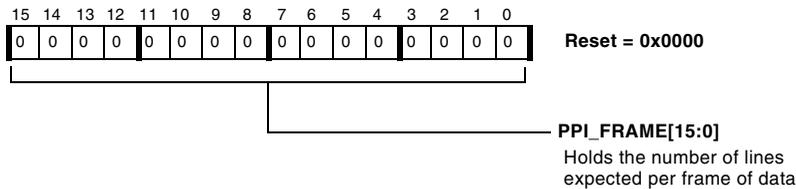


Figure 20-17. PPI Lines Per Frame Register

Programming Examples

The PPI can be configured to receive data from a video source in several RX modes. The following programming examples ([Listing 20-1](#) through [Listing 20-5](#)) describe the ITU-R 656 entire field input mode.

Listing 20-1. Configure DMA Registers

```
config_dma:
/*Assumes PPI is mapped to DMA channel 0.*/
  /* DMA0_START_ADDR */
  R0.L = rx_buffer;
  R0.H = rx_buffer;
  P0.L = lo(DMA0_START_ADDR);
  P0.H = hi(DMA0_START_ADDR);
  [P0] = R0;

  /* DMA0_CONFIG */
  R0.L = DI_EN | WNR;
  P0.L = lo(DMA0_CONFIG);
  P0.H = hi(DMA0_CONFIG);
  W[P0] = R0.L;

  /* DMA0_X_COUNT */
  R0.L = 256;
  P0.L = lo(DMA0_X_COUNT);
  P0.H = hi(DMA0_X_COUNT);
  W[P0] = R0.L;

  /* DMA0_X_MODIFY */
  R0.L = 0x0001;
  P0.L = lo(DMA0_X_MODIFY);
  P0.H = hi(DMA0_X_MODIFY);
  W[P0] = R0.L;
  ssync;
config_dma.END:  RTS;
```

Programming Examples

Listing 20-2. Configure PPI Registers

```
config_ppi:  
  
    /* PPI_CONTROL */  
    PO.L = lo(PPI_CONTROL);  
    PO.H = hi(PPI_CONTROL);  
    RO.L = 0x0004;  
    W[PO] = RO.L;  
    ssync;  
  
config_ppi.END:    RTS;
```

Listing 20-3. Enable DMA

```
    /* DMA0_CONFIG */  
    PO.L = lo(DMA0_CONFIG);  
    PO.H = hi(DMA0_CONFIG);  
    RO.L = W[PO];  
    bitset(RO,0);  
    W[PO] = RO.L;  
    ssync;
```

Listing 20-4. Enable PPI

```
    /* PPI_CONTROL */  
    PO.L = lo(PPI_CONTROL);  
    PO.H = hi(PPI_CONTROL);  
    RO.L = W[PO];  
    bitset(RO,0);  
    W[PO] = RO.L;  
    ssync;
```

Listing 20-5. Clear DMA Completion Interrupt

```
/* DMA0_IRQ_STATUS */  
P2.L = 1o(DMA0_IRQ_STATUS);  
P2.H = hi(DMA0_IRQ_STATUS);  
R2.L = W[P2];  
BITSET(R2,0);  
W[P2] = R2.L;  
ssync;
```

Unique Information for the ADSP-BF51x Processor

None.

Unique Information for the ADSP-BF51x Processor

21 REMOVABLE STORAGE INTERFACE

This chapter describes the ADSP-BF51x Blackfin processor removable storage interface (RSI) and includes the following sections:

- [“Overview”](#)
- [“Interface Overview” on page 21-2](#)
- [“Description of Operation” on page 21-6](#)
- [“Functional Description” on page 21-8](#)
- [“Programming Model” on page 21-31](#)
- [“RSI Registers” on page 21-52](#)

Overview

ADSP-BF51x Blackfin processors provide an RSI interface for multimedia cards (MMC), secure digital memory cards (SD), secure digital input/output cards (SDIO) and consumer electronic ATA devices (CE-ATA). All of these storage solutions use similar interface protocols. The main difference between MMC and SD support is the initialization sequence. The main difference between SD and SDIO support is the use of interrupt and read wait signals for SDIO. CE-ATA devices require handling of larger block sizes of 4K bytes and implement a device interrupt scheme known as the command completion signal (CCS).

Interface Overview

Features of the RSI interface include:

- Support for a single SD or SDIO card
- Support for a single MMC device (removable or embedded)
- Support for 1- and 4-bit SD modes (SPI mode is not supported)
- Support for 1-, 4-, and 8-bit MMC modes (SPI mode is not supported)
- Support for 4- and 8-bit CE-ATA devices
- Programmable clock frequency generated from `SCLK`
- Card detection capabilities
- SDIO interrupt and read wait features
- Command completion signal recognition and disable for CE-ATA device support
- High-capacity card support such as SDHC implemented within software
- 512-bit transmit/receive FIFO
- DMA channel with 32-bit DMA access bus

Interface Overview

The RSI interface handles the multimedia and secure digital card functions. This includes clock generation, power management, command transfer, and data transfer. The bus interface converts 16-bit PAB accesses to 32-bit register accesses to the memory-mapped registers, and generates interrupt requests to the processor core and system. The RSI has two interrupt signals (`IRQ0` and `IRQ1`), that are fed to the system interrupt controller (SIC) `IRQ51` and `IRQ52` respectively.

The RSI block has 22 individual status bits in the `RSI_STATUS` register that can be configured to generate an interrupt. The status bits can be mapped to either of the two interrupts fed to the system interrupt controller. This allows for greater flexibility in system configuration.

To generate an interrupt on `IRQ0`, the interrupt should be enabled by setting the corresponding bit in the `RSI_MASK0` register. Interrupts that are required to be generated on `IRQ1` are enabled by setting the corresponding bit in the `RSI_MASK1` register.

In addition to the status flags in the `RSI_STATUS` register generating interrupts, each of the flags in the `RSI_ESTAT` register are also capable of generating an interrupt. Interrupts for the `RSI_ESTAT` flags are enabled by setting the corresponding bit in the `RSI_EMASK` register. The interrupts are sent to the SIC through `IRQ51`.

The 32-bit DAB bus allows for efficient transfer of data, both to and from internal memory, through DMA channel 4 that is shared with the `SPORT0 TX`. The peripheral used by this DMA channel is the peripheral that is enabled for pin multiplexing.

The RSI ([Figure 21-1](#)) is a 10-pin interface consisting of:

- `RSI_CLK`
The clock signal applied to the card from the RSI. All command and data signal transfers are synchronous to this clock. The frequency can vary between zero and the maximum clock frequency. Refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet* for the maximum supported clock frequencies.
- `RSI_CMD`
A bi-directional command signal used for command transfer and card initialization. The RSI uses this signal to send commands to the cards, and the card uses the signal to send responses back to the RSI. This signal can be configured for both push-pull mode and open-drain mode, but only MMC cards support the open-drain

Interface Overview

mode. The open-drain mode allows for multiple MMC cards to share data and command signals on the RSI interface and allows for the initialization sequence to take place on all cards.

- RSI_DATA7-0
These are the configurable bi-directional data channels used for all data transfers both to and from the card. The data bus width can be configured as 1-, 4-, or 8-bit.

i Although earlier revisions of the MMC specifications allowed for multiple MMC cards to be used to a single RSI interface, it is strongly recommended that only a single device be interfaced to any given RSI interface. Multiple devices sharing a single command and data bus is now actively discouraged by the device specifications.

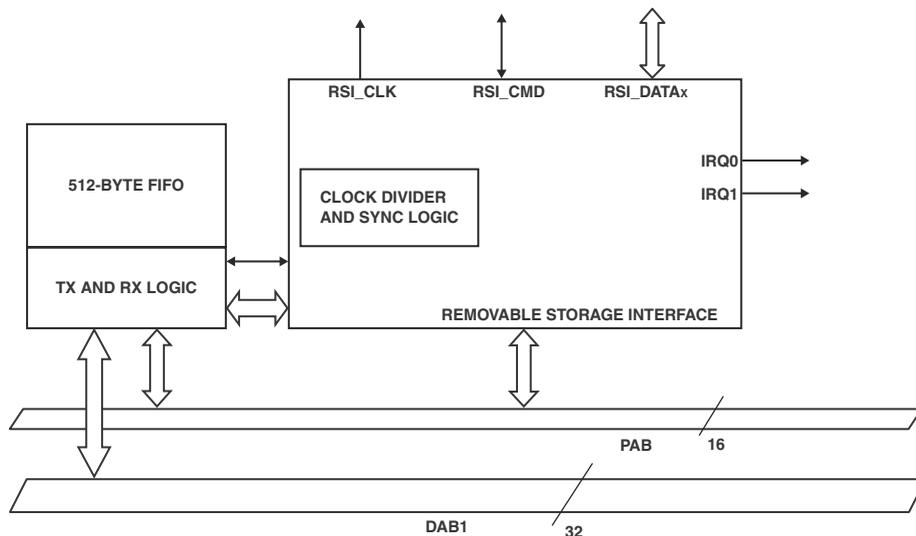


Figure 21-1. RSI Block Diagram

Removable Storage Interface

Table 21-1 and Table 21-2 list the RSI interface pins functional operations for all supported protocol modes.

Table 21-1. RSI Protocol Interface (MMC and CE-ATA)

Signal Name	MMC (1-bit)	MMC (4-bit)	MMC (8-bit)	CE-ATA (4-bit)	CE-ATA (8-bit)	Direction
RSI_DATA7	Not Used	Not Used	Dat7	Not Used	Dat7	Bi-directional
RSI_DATA6	Not Used	Not Used	Dat6	Not Used	Dat6	Bi-directional
RSI_DATA7	Not Used	Not Used	Dat7	Not Used	Dat7	Bi-directional
RSI_DATA6	Not Used	Not Used	Dat6	Not Used	Dat6	Bi-directional
RSI_DATA5	Not Used	Not Used	Dat5	Not Used	Dat5	Bi-directional
RSI_DATA4	Not Used	Not Used	Dat4	Not Used	Dat4	Bi-directional
RSI_DATA3	Not Used/ Card Detect	Dat3/ Card Detect	Dat3/ Card Detect	Dat3	Dat3	Bi-directional
RSI_DATA2	Not Used	Dat2	Dat2	Dat2	Dat2	Bi-directional
RSI_DATA1	Not Used	Dat1	Dat1	Dat1	Dat1	Bi-directional
RSI_DATA0	Dat0	Dat0	Dat0	Dat0	Dat0	Bi-directional
RSI_CMD	Command/ Response	Command/ Response	Command/ Response	Command/ Response/ CCS/CCSD	Command/ Response/ CCS/CCSD	Bi-directional
RSI_CLK	CLK	CLK	CLK	CLK	CLK	Output

Table 21-2. RSI Protocol Interface (SD and SDIO)

Signal Name	SD (1-bit)	SD (4-bit)	SDIO (1-bit)	SDIO (4-bit)	Direction
RSI_DATA7	Not Used	Not Used	Not Used	Not Used	Bi-directional
RSI_DATA6	Not Used	Not Used	Not Used	Not Used	Bi-directional
RSI_DATA5	Not Used	Not Used	Not Used	Not Used	Bi-directional
RSI_DATA4	Not Used	Not Used	Not Used	Not Used	Bi-directional
RSI_DATA3	Not Used/ Card Detect	Dat3/ Card Detect	Not Used/ Card Detect	Dat3/ Card Detect	Bi-directional

Description of Operation

Table 21-2. RSI Protocol Interface (SD and SDIO) (Continued)

Signal Name	SD (1-bit)	SD (4-bit)	SDIO (1-bit)	SDIO (4-bit)	Direction
RSI_DATA2	Not Used	Dat2	Read Wait	Dat2/ Read Wait	Bi-directional
RSI_DATA1	Not Used	Dat1	Interrupt	Dat1/ Interrupt	Bi-directional
RSI_DATA0	Dat0	Dat0	Dat0	Dat0	Bi-directional
RSI_CMD	Command/ Response	Command/ Response	Command	Command	Bi-directional
RSI_CLK	CLK	CLK	CLK	CLK	Output

Description of Operation

The RSI controller is a fast, synchronous peripheral that uses various protocols to communicate with MMC, SD, and SDIO cards as well as CE-ATA hard drives. The RSI is compatible with the following protocols.

- MMC (Multimedia Card) bus protocol
- SD (Secure Digital) bus protocol
- SDIO (Secure Digital Input Output) bus protocol
- CE-ATA (Consumer Electronic ATA)



The RSI does not support the SPI bus protocol.

Communication is through a master and slave configuration, where the RSI is the master and the card is the slave device. The RSI communicates with the device using a message-based bus protocol in which the host sends commands serially using the `RSI_CMD` signal. Some commands require that the card provide a response back to the host. This response is also sent serially by the `RSI_CMD` signal.

Data transfers, both to and from the card, occur using the `RSI_DATAx` signals. The number of data lines used for the data transfer can be configured to 1, 4, or 8 using `RSI_DATA0`, `RSI_DATA3-0`, or `RSI_DATA7-0`. All `RSI_CMD` and `RSI_DATAx` transfers are synchronous with `RSI_CLK`.

Cyclic redundancy codes (CRC) are used to protect commands, responses and data transfers from transmission errors. A CRC7 code is generated for every command sent by the host and for almost every response returned by the card on the `RSI_CMD` signal. A CRC16 code is used on the `RSI_DATAx` signals to protect block data transfers. In 4- and 8-bit bus configurations, CRC16 is calculated for each individual data signal.

When a device connected to the RSI is first powered and detected by the host or has been reset—the device must first be identified and initialized by the host. The software determines whether the device is compatible with the RSI controller and the implemented software drivers. This phase in the procedure is known commonly as the card identification mode.

When a device is in card identification mode, the host may be required to perform the following actions.

- Reset the device
- Validate the device operating voltage range
- Identify the device type
- Assign/request a relative card address (RCA)

All communications between the host and card during the card identification phase occur using the `RSI_CMD` signal. The maximum clock frequencies during this identification phase is typically far lower than the maximum data transfer frequency for the card.

The card will only transition to a stand-by state, known as the data transfer mode, when it has been assigned an RCA. Data transfers can only take place when the device has entered the data transfer mode.

Functional Description

Once the device is in data transfer mode, communication takes place through the `RSI_CMD` and the `RSI_DATAx` signals. The card is further interrogated to identify bus widths, maximum clock frequency, and the device capacity. At this point the bus width can be altered and the clock frequency can be increased.

Data may be written to the device or read from the device using the following two methods.

- Stream reads and writes
- Block reads and writes

Stream transfers produce a continuous stream of data until the RSI stops the transfer by setting a specific command.

For stream read and write operations, additional maximum operating frequency limitations may be imposed by the device. Stream write operations may also have restrictions that are dependent upon writable block boundaries.

Block based transfers result in a block of a pre-configured size being transferred. The block size depends on the device and is obtained by reading registers contained on the device during the device detection procedure.

Functional Description

The following sections describe the functions and features of the RSI controller as well as the MMC, SD, SDIO, and CE-ATA protocols. For detailed information on timing parameters and protocol requirements, refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet* and the following standards and specifications.

- MMCA System Specification
- JESD84 series of JEDEC standards

- SD Specifications Part 1 Physical Layer Specification
- SD Specifications Part 1 Physical Layer Simplified Specification
- SD Specifications Part E1 SDIO Specification

RSI Clock Configuration

The RSI is a fast, synchronous peripheral with a programmable clock frequency that is supplied by the `RSI_CLK` signal. The interface between the RSI and the PAB/DAB busses operates at `SCLK` frequency. Communication between the clock domain that is supplied externally from the RSI on the `RSI_CLK` signal and the internal RSI access to the PAB and DAB busses is accomplished using synchronizers in the RSI module. The `RSI_CLK` frequency is configured by the 8-bit `CLKDIV` field and the `CLKDIV_BYPASS` bit of the `RSI_CLK_CONTROL` register (see “[RSI Clock Control Register \(RSI_CLK_CONTROL\)](#)” on page 21-54).

If `CLKDIV_BYPASS` is set, the clock frequency driven on the `RSI_CLK` signal is derived directly from `SCLK`.

If `CLKDIV_BYPASS` is cleared, the clock divider logic provides an `RSI_CLK` frequency as shown below, where `CLKDIV` is an 8-bit value ranging between 0 and 255.

$$RSI_CLK = \frac{SCLK}{2 \times (CLKDIV + 1)}$$

The `RSI_CLK` output is enabled or disabled by the `CLK_EN` bit in the `RSI_CLK_CONTROL` register and a power save feature is implemented by setting `PWR_SV_EN`, which disables the `RSI_CLK` output when there are no transfers taking place on the RSI interface.

RSI Interface Configuration

The RSI supports multiple card types under various protocols. Different card types may require slightly different interface configurations.

Functional Description

The command signal on MMC cards operates in two different modes depending upon the operating mode of the card. During the card identification mode, the command signal operates in open-drain configuration; but when the card enters data transfer mode, the signal is configured to push-pull mode.

 The internal pull-up resistor of the `RSI_CMD` signal is only intended to keep the signal from floating. The internal pull-up resistor is not sufficient during the card identification phase when the MMC card `RSI_CMD` signal is operating in open-drain mode. If support for MMC devices is required, an external pull-up resistor should be added to the `SD_CMD` signal as detailed in the JEDEC standard.

The bus width used for the data transfers is configurable to 1-bit, 4-bits, or 8-bits using the `BUS_MODE` field in the `RSI_CLK_CONTROL` register (see [“RSI Clock Control Register \(RSI_CLK_CONTROL\)”](#) on page 21-54).

To stop signals from floating when no card is inserted or during times when all card drivers are in a high-impedance mode, various pull-up and pull-down resistor configurations can be enabled on the `RSI_DATAx` signals. The `RSI_CONFIG` register (see [“RSI Configuration Register \(RSI_CONFIG\)”](#) on page 21-77) provides the following options.

- Enable or disable a pull-up resistor on the `RSI_DATA3` signal
- Enable or disable pull-up resistors on the `RSI_DATA7` through `RSI_DATA4` and `RSI_DATA2` through `RSI_DATA0` signals

Card Detection

The RSI allows for software to detect when a card is inserted into its slot. There are a number of ways that this card detection can be performed. See [Figure 21-2](#).

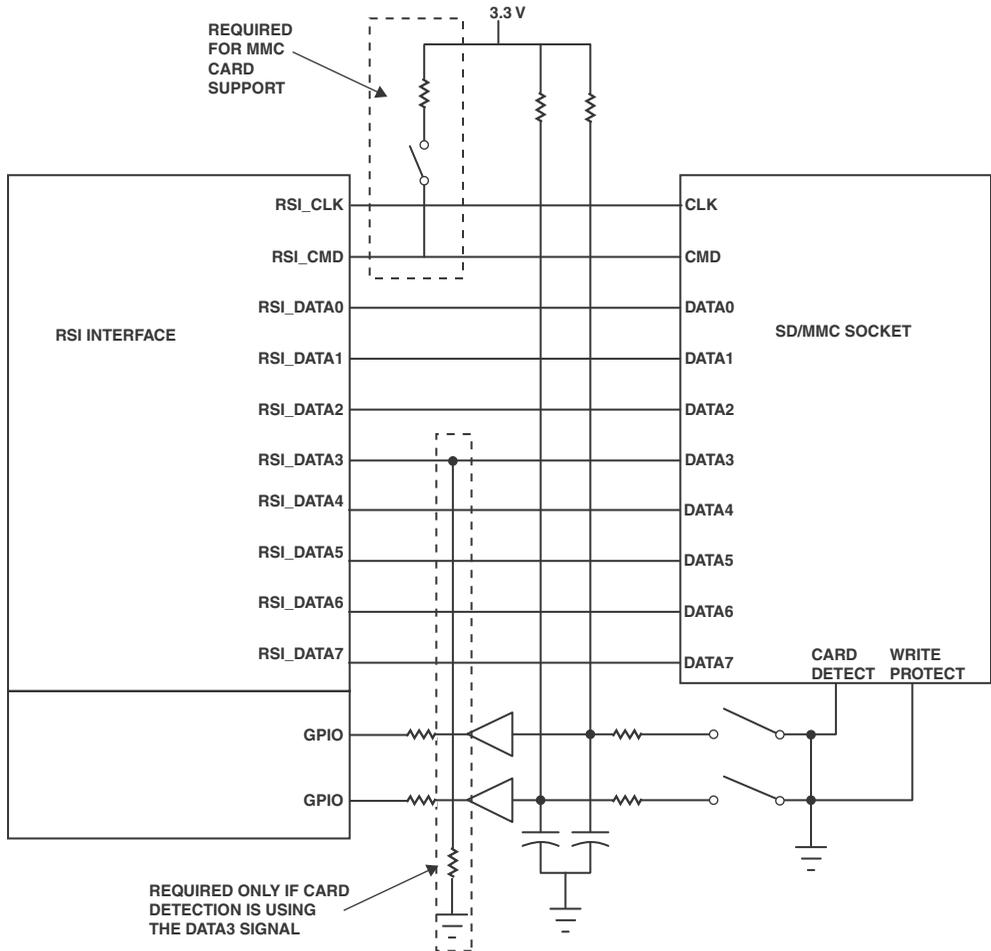


Figure 21-2. RSI Socket Interface

Functional Description

SD and SDIO cards use an internal pull-up resistor on the `RSI_DATA3` line as a card detect signal to indicate to the host that a card is present. In order to use the `RSI_DATA3` signal for card detection, an external pull down resistor should be added to the pin to pull the pin low during the time a card is not inserted. When a card is inserted into the slot, a rising edge is detected on `RSI_DATA3` by the RSI and `SD_CARD_DET` in the `RSI_ESTAT` register is set. Once the card has been correctly identified, the `SD_CARD_DET` interrupt should be cleared and disabled. The pull-up resistor in the SD card should be disabled by issuing the `SET_CLR_CARD_DETECT` command.

 When using the `RSI_DATA3` signal for card detection with an external pull-down resistor, do not enable the internal pull-up resistor by setting `PU_DAT3`.

The recommended method of detecting the insertion of a card is to use the card detect feature that is made available through most sockets. Sockets supporting this feature can have the card detect pin de-bounced and connected to a GPIO pin in order to allow not only interrupt driven card detection but also interrupt driven card removal. This is the most reliable and efficient method of detecting the insertion and removal of a card as some MMC devices may not implement the card detect pull-up resistor on the `RSI_DATA3` signal. Once a card is detected, the GPIO pin can have the interrupt level inverted to then generate an interrupt on card removal.

The final approach to detecting the insertion and removal of a card is to simply use software polling. Software would poll the slot periodically using the card identification commands for the supported card types. Once a card is inserted, this will result in valid responses being sent back to the host; when the card is removed, command and data timeout errors will be observed.

RSI Power Saving Configuration

The RSI requires two internal clock signals that are derived directly from SCLK. One of these clock signals is routed to the clock divider and generates the RSI_CLK clock.

These clocks must be enabled by setting RSI_CLK_EN in the RSI_CONFIG register for the RSI to function. Clearing RSI_CLK_EN disables the RSI regardless of other RSI clock configurations. The RSI_CLK signal can be enabled or disabled using CLK_EN in the RSI_CLK_CONTROL register.

Additional power saving options can be implemented by setting PWR_SV_EN which disables the RSI_CLK output when there are no transfers taking place on the RSI interface. These configurations are shown in [Table 21-3](#).

Table 21-3. RSI Power Saving Configurations

CLKS_EN	CLK_EN	PWR_SV_E	RSI State	RSI_CLK Output
0	0	0	Disabled	No clock
0	0	1	Disabled	No clock
0	1	0	Disabled	No clock
0	1	1	Disabled	No clock
1	0	0	Enabled	No clock
1	0	1	Enabled	No clock
1	1	0	Enabled	Continuous clock ¹
1	1	1	Enabled	Clock only driven during transfers ¹

¹ The PWR_ON field of the RSI_PWR_CTL register must be set to 0x3. If PWR_ON is 0x0—the clock will not output.

RSI Commands and Responses

The RSI sends commands to and receives responses from the card using the RSI_CMD signal. The command to be sent to the card is issued by writing to the RSI_COMMAND register (see “[RSI Command Register \(RSI_COMMAND\)](#)” on page 21-57). This register contains a 6-bit CMD_IDX field that contains the command index to be sent to the card. The command index provides support for a total of 64 commands—0 (CMD0) to 63 (CMD63). Some commands require an argument to be sent along with the command, such as an address for a read transaction. An argument is always sent with the command and it is the responsibility of the card to either ignore or use the argument field based on the command that is received. The argument sent with the command is provided using the RSI_ARGUMENT register (see “[RSI Argument Register \(RSI_ARGUMENT\)](#)” on page 21-56).

All command transfers are protected by a 7-bit cyclic redundancy check (CRC) code, more commonly referred to as a CRC7 checksum. This allows for transmission errors to be detected and for the command to be re-issued to the card in the event of an error. All commands sent to the card are composed of 48-bits as shown in [Table 21-4](#).

Table 21-4. RSI Command Format

Bit Position	Width	Value	Description
47	1	0	Start bit
46	1	1	Transmitter bit
[45:40]	6	-	Command index
[39:8]	32	-	Argument
[7:1]	7	-	CRC7 checksum
0	1	1	End bit

The `RSI_COMMAND` register also provides configuration information about whether a response is to be expected back from the card and the type of response.

The RSI can be configured using the `CMD_RESP` and `CMD_L_RESP` fields of the `RSI_COMMAND` register to expect the following response types.

- No response
- Short response (see [Table 21-5](#))
- Long response (see [Table 21-6](#))

Table 21-5. RSI Short Response Format

Bit Position	Width	Value	Description
47	1	0	Start bit
46	1	0	Transmitter bit
[45:40]	6		Command index or check bits ¹
[39:8]	32		Card status, register contents or argument field
[7:1]	7		CRC7 checksum or check bits ²
0	1	1	End bit

1 Responses that do not contain the command index have b#111111 in the check bits field.

2 Responses that do not contain a CRC7 check sum have b#111111 in the check bits field.

Functional Description

Table 21-6. RSI Long Response Format

Bit Position	Width	Value	Description
135	1	0	Start bit
134	1	0	Transmitter bit
133:128	6	111111	Check bits ¹
127:1	127	-	Register contents including internal CRC7 ²
0	1	1	End bit

1 Responses that do not contain the command index have b#111111 in the check bits field.

2 Responses that do not contain a CRC7 check sum have b#111111 in the check bits field.

Like the commands, all responses are sent on the `RSI_CMD` signal. A response always has a “0” start bit followed by a “0” transmission bit to indicate the transfer is from card to RSI. Unlike the commands issued by the RSI, not all responses are protected by a CRC7 checksum. Refer to the appropriate specification for full details on the response formats for a specific device and whether they are protected by a CRC7 checksum.

When a short response is received, the response is broken down by the RSI. The 32-bit field containing bits[39:8] is stored to `RSI_RESPONSE0`, where bit 39 of the response corresponds to bit 31 of `RSI_RESPONSE0` and bit 8 of the response to bit 0 of `RSI_RESPONSE0`. Bits[45:40] of the response are stored to the `RESP_CMD` field of the `RSI_RESP_CMD` register.

For a long response, bits[127:1] of the response are stored in `RSI_RESPONSE3-0`. Bit 31 of `RSI_RESPONSE0` contains the most significant bit (bit 127) of the response and bit 0 of `RSI_RESPONSE3` contains bit 1 of the response. Bit 31 of `RSI_RESPONSE3` is always zero.

Figure 21-3 shows the command path state machine. For the state machine to be active, the RSI must be enabled through `RSI_CLK_EN`. Disabling the clocks to the RSI results in the state machine returning to the IDLE state.

The command path state machine is responsible for setting and clearing a number of status flags in the `RSI_STATUS` register (see “[RSI Status Register \(RSI_STATUS\)](#)” on page 21-64). [Table 21-7](#) lists the status flags and exception flags that are affected by the command path state machine.

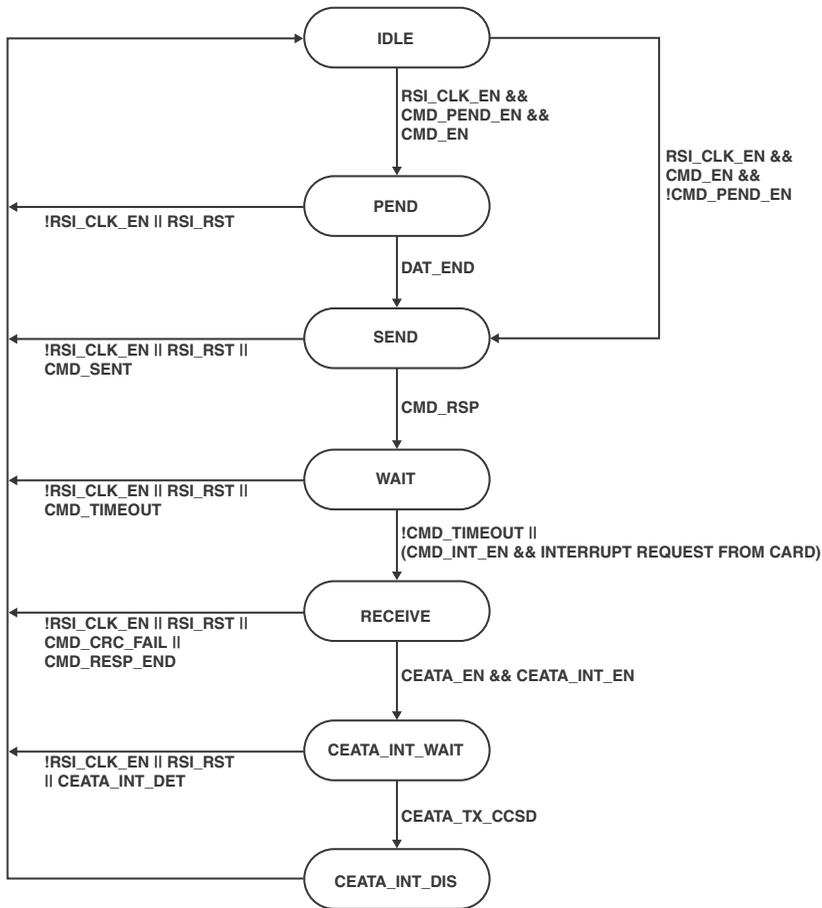


Figure 21-3. RSI Command Path State Machine

Functional Description

Table 21-7. RSI Command Path Status Flags

RSI_STATUS Flag	Description	State Flag Set in
CMD_ACT	Command transfer is in progress	WAIT_S
CMD_SENT	Command without response sent successfully	SEND
CMD_TIMEOUT	Response timeout occurred (64 RSI_CLK cycles)	WAIT_S
CMD_CRC_FAIL	Response CRC failure	RECEIVE
CMD_RSP_END	Response CRC successful	RECEIVE
CEATA_INT_DET	CE-ATA command completion signal detected	CEATA_INT_WAIT

The command path operates in a half-duplex mode, so that commands and responses can either be sent or received. If the state machine is not in the SEND state, the RSI_CMD output is in high impedance state.

Figure 21-4 describes a typical command and response transfer, the RSI_CMD signal is sampled by the card and the host on the rising edge of RSI_CLK.

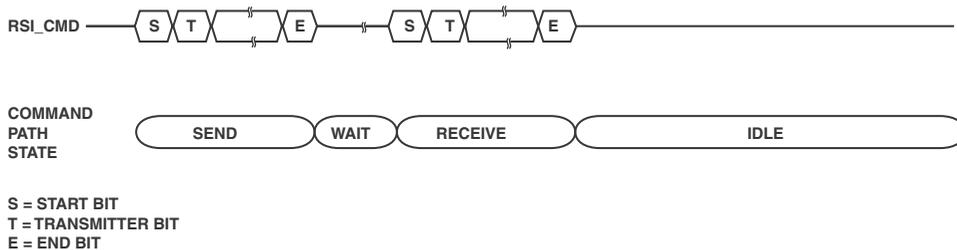


Figure 21-4. RSI Command Transfer

The following sections describes the RSI command path states.

IDLE State

The command path state machine remains in the IDLE state when it is not active. The command path state machine is enabled and leaves the IDLE state when the CMD_EN bit in the RSI_COMMAND register is set. The

state goes to the PENDING state if the `CMD_PENDING_EN` bit in the `RSI_COMMAND` is set; otherwise it enters the SEND state.

When the command path state machine returns to the IDLE state from another state and the return is not due to RSI being disabled or reset—the state machine remains in the IDLE state for at least eight `RSI_CLK` cycles. During this time, the RSI will continue to drive the `RSI_CLK` signal even if the `PWR_SV_EN` feature is enabled. This allows the card to complete the current operation. If enabled again, the state machine leaves the IDLE state only after the eight `RSI_CLK` cycles have passed.

PENDING State

The RSI enters the PENDING state if the `CMD_PENDING_EN` bit in the `RSI_COMMAND` register is set. The state machine remains in the PENDING state until it is notified by the data path sub block that a data transfer has completed. This is indicated by the `DAT_END` flag being set when `RSI_DATA_CNT` decrements to zero. This mode allows for the automatic transmission of a `STOP_TRANSMISSION` command after reading or writing the required amount of data for stream-based transactions.



The `CMD_PENDING_EN` feature is not functional for block-based transfers and cannot be used to automatically issue the `STOP_TRANSMISSION` command for `MULTIPLE_BLOCK_READ` or `MULTIPLE_BLOCK_WRITE` operations.

SEND State

During the SEND state, the RSI sets the `CMD_ACT` flag in the `CMD_STATUS` register to indicate a transfer is in progress. The behavior of the state machine after the command is sent, depends upon whether the command expects a response back from the card.

If no response is expected, the RSI clears the `CMD_ACT` flag and sets the `CMD_SENT` flag to indicate that a command operation without a response has been completed. The state then goes to IDLE.

Functional Description

If a response is expected, the RSI enters the WAIT state.

WAIT State

In the WAIT state, the RSI waits for a response to be received on the RSI_CMD signal. Upon entering this state, an internal timer starts. If the response is not received within 64 RSI_CLK cycles, the CMD_TIMEOUT flag is set and the CMD_ACT flag is cleared. The state machine then enters the IDLE state, awaiting the next action.

A response, sent back from the card and indicated by the "0" start bit on the RSI_CMD signal, transitions the RSI to the RECEIVE state to receive a 48- or 136-bit response.

The WAIT state can also detect card interrupts. This is an optional feature that applies only to MMC cards. This feature is enabled by setting the CMD_INT_EN bit within RSI_COMMAND. When CMD_INT_EN is set, the timeout timer that is normally started upon entry to the WAIT state is disabled. The RSI remains in this state until a card interrupt is detected. Cards that implement this feature may have functions with a delayed response that is triggered by an internal event in the card. Once the event is triggered the card sends the response. The RSI then detects this start bit of the response and proceeds to the RECEIVE state.

RECEIVE State

In the RECEIVE state the RSI reads the response on the RSI_CMD signal from the card. Upon receiving either the short or long response, if the response passes the CRC check, the CMD_ACT flag is cleared and the CMD_RESP_END flag is set. If the CRC check fails, the CMD_CRC_FAIL flag is set. In either case, the state machine then goes to the IDLE state.

Some CE-ATA commands require additional functionality in this state. A command completion signal must be sent back to the host upon completion of a specific task. For commands that require this functionality, the CEATA_EN and CEATA_CCS_EN bits of the RSI_DATA_CONTROL register should

be set prior to enabling the command path state machine. After receiving the response, the state machine then enters the `CEATA_INT_WAIT` state.

CEATA_INT_WAIT State

In this state, the RSI waits for the CE-ATA device to issue the command completion signal. This is indicated by the device sending a “0” on the `RSI_CMD` signal. Upon detection of the command completion signal, the `CMD_ACT` flag is cleared and the `CEATA_INT_DET` flag of the `RSI_ESTAT` register is set. Alternatively, the command completion signal of the CE-ATA device can be disabled by the RSI. This is performed by setting the `CEATA_TX_CCSD` bit of the `RSI_CEATA_CONTROL` register, at which point the state machine enters the `CEATA_INT_DIS` state. The `CEATA_TX_CCSD` bit can be set prior to enabling the command path state machine. The `CCSD` sequence is issued after the response is received rather than having to wait for the response before setting `CEATA_TX_CCSD`.

CEATA_INT_DIS State

In this state, the RSI issues the command completion signal disable sequence on the `RSI_CMD` signal before transitioning to the `IDLE` state and clearing the `CMD_ACT` flag. The command completion signal disable sequence issued is the binary sequence `b#00001`.

RSI Command Path CRC

The command CRC generator of the RSI calculates the 7-bit CRC checksum for all 40 bits preceding the CRC code for both 48-bit commands and 48-bit responses. This includes the start bit, transmitter bit, command index, and command argument (or card status). The 7-bit CRC checksum is calculated for the first 120 bits of the register contents field for the long response format. Note that the start bit, transmitter bit, and the six check bits are not used in the CRC calculation for the long

Functional Description

response. The command and response CRC checksum is a 7-bit value that is calculated as follows:

$$CRC[6:0] = Remainder \quad \frac{x_7 \times M(x)}{G(x)}$$

with:

$$G(x) = x_7 + x_3 + 1$$

and for a short response:

$$M(x) = x_{39} \cdot (\textit{start bit}) + \dots + x_0 \cdot (\textit{last bit before CRC})$$

or for a long response:

$$M(x) = x_{19} \cdot (\textit{start bit}) + \dots + x_0 \cdot (\textit{last bit before CRC})$$

RSI Data

Data transfers both to and from the RSI take place over the RSI data bus signals `RSI_DATA7–0`. The RSI data bus width is configured by the `BUS_MODE` field of the `RSI_CLK_CONTROL` register (see “[RSI Clock Control Register \(RSI_CLK_CONTROL\)](#)” on page 21-54). The default is 1-bit bus mode, where the data is transferred over the `RSI_DATA0` signal.

Or 4-bit mode or 8-bit mode can be enabled after configuring the card for 4-bit or 8-bit mode of operation, respectively.

The RSI data path state machine operates at `RSI_CLK` frequency. The state machine leaves the IDLE state when the `DATA_EN` field of the `RSI_DATA_CONTROL` register is set, enabling the data transfer. The state entered upon leaving the IDLE state is determined by `DATA_DIR`. The data path state machine is shown in [Figure 21-5](#).

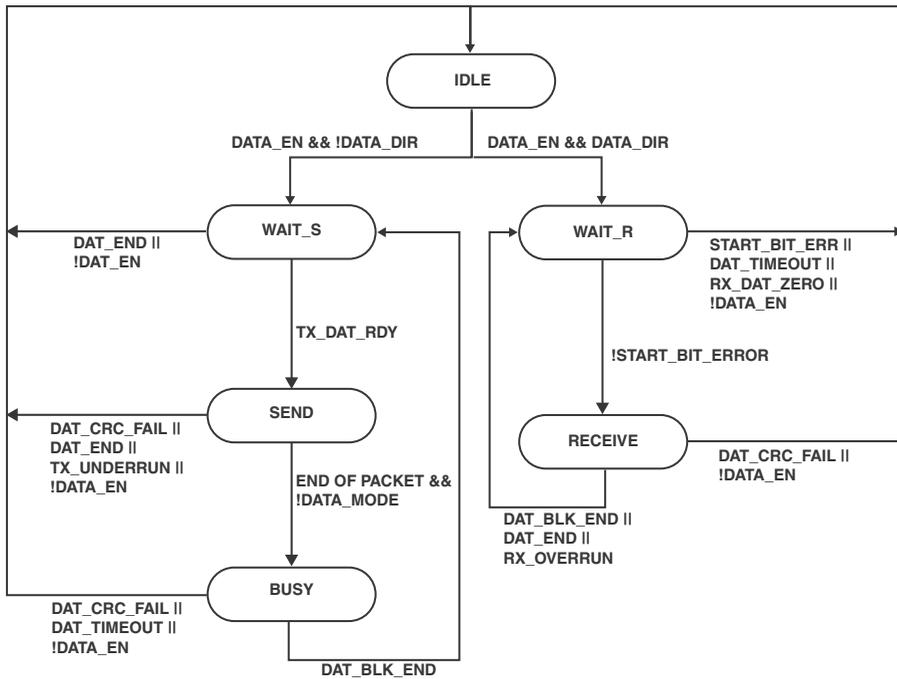


Figure 21-5. RSI Data Path State Machine

The data path status flags are shown in [Table 21-8](#).

Table 21-8. RSI_STATUS Flags

RSI_STATUS Flag	Description	States Flag Set in
TX_ACT	Data transmit in progress	WAIT_S
RX_ACT	Data receive in progress	WAIT_R
DAT_BLK_END	Data block sent successfully and CRC pass token received	BUSY (block transfer mode only)
	Data block received correctly and CRC passed	RECEIVE (block transfer only)

Functional Description

Table 21-8. RSI_STATUS Flags (Continued)

RSI_STATUS Flag	Description	States Flag Set in
DAT_CRC_FAIL	Data block CRC failed on transmit	SEND if transmitted data is not a multiple of DATA_BLK_LGTH. BUSY if CRC token indicates failure.
	Data block CRC failed on receive	RECEIVE
DAT_TIMEOUT	Transmit timeout occurred before card de-asserted busy signal on RSI_DATA0	BUSY
	Receive timeout error occurred before start bit of data detected	WAIT_R
DAT_END	All data sent	SEND
	All data received	RECEIVE
START_BIT_ERR	Start bit not detected on all RSI_DATAx signals	WAIT_R
TX_FIFO_STAT	Transmit FIFO is half empty	SEND
TX_FIFO_FULL	Transmit FIFO is full	SEND
TX_FIFO_EMPTY	Transmit FIFO is empty	SEND
TX_UNDERRUN	Transmit FIFO under run error	SEND
TX_DAT_RDY	Valid data available in the transmit FIFO	SEND
RX_FIFO_STAT	Receive FIFO is half empty	RECEIVE
RX_FIFO_FULL	Receive FIFO is full	RECEIVE
RX_FIFO_EMPTY	Receive FIFO is empty	RECEIVE
RX_OVERRUN	Receive FIFO over run error	RECEIVE
RX_FIFO_RDY	Valid data is available in the receive FIFO	RECEIVE

RSI Data Transmit Path

The transmit path consists of the WAIT_S, SEND, and BUSY states. Both RSI_DATA_LGTH and RSI_DATA_TIMER must be configured before enabling the data path state machine with RSI_DATA_CONTROL. Upon leaving the IDLE state and entering the WAIT_S state, the RSI sets the TX_ACTIVE flag and copies RSI_DATA_LGTH into RSI_DATA_CNT.

The behavior of the SEND state depends on which transfer mode is configured.

- Stream Transfer Mode

If the SDH is configured for stream transfer mode, the SDH sends data to the card until RSI_DATA_CNT expires, at which point the DATA_END flag is set and the state machine returns to the IDLE state.

Additionally the transition of RSI_DATA_CNT to zero activates the command path state machine if it is currently in the PEND state.

If at any point during the stream transfer the transmit FIFO becomes empty and data is not available in the FIFO by the time the next transfer is due to take place, the TX_UNDERRUN flag is set before returning to the IDLE state.

- Block Transfer Mode

In block transfer mode, DATA_BLK_LGTH bytes, as specified during the write to RSI_DATA_CONTROL, are transmitted. Each byte transferred also decrements RSI_DATA_CNT.

Upon completion of the block transfer, the RSI appends an internally generated 16-bit CRC code and an end bit to the data transferred over the RSI_DATAx signals. The SDH then waits for the card token response on the RSI_DATA0 line to indicate whether the data was received correctly by the card or not.

Functional Description

If the CRC response token sent by the card indicates the data was received correctly, the `DAT_BLK_END` flag is set before moving to the BUSY state. If the data was not received correctly, the `DAT_CRC_FAIL` flag is set before returning to the IDLE state.

The decrementing of `RSI_DATA_CNT` to zero results in the `DAT_END` flag being set. If the total number of bytes transmitted for the current block results in the `RSI_DATA_CNT` decrementing to zero and the number of bytes transferred is not equal to `DATA_BLK_LGTH`, the transmission stops and the `DAT_CRC_FAIL` flag is set. The data path returns to the IDLE state.

If at any point during the block transfer the transmit FIFO becomes empty and data is not available in the FIFO by the time the next transfer is due to take place, the `TX_UNDERRUN` flag is set before returning to the IDLE state.

During the BUSY state, the RSI continuously samples `RSI_DATA0` which at this point is driven low by the card to indicate that the card is busy. When a logic high state is detected, indicating that the card is no longer busy, the state machine returns to the `WAIT_S` state. It then either returns to IDLE if all data has been sent or moves back to the `SEND` state to start another block transfer.

Upon entry to the BUSY state, the RSI started decrementing the timeout value specified by `RSI_DATA_TIMER`. If the RSI timeout counter expires before the `RSI_DATA0` signal is detected high, the RSI sets the `DAT_TIMEOUT` flag and returns to the IDLE state.

RSI Data Receive Path

The receive path consists of the `WAIT_R` and the `RECEIVE` states. Both `RSI_DATA_LGTH` and `RSI_DATA_TIMER` must be configured, before enabling the data path state machine with `RSI_DATA_CONTROL`. Upon leaving the IDLE state and entering the `WAIT_R` state, the RSI sets the `RX_ACTIVE`

flag and copies `RSI_DATA_LGTH` into `RSI_DATA_CNT`. The behavior of the `RECEIVE` state is influenced by the transfer mode.

Once the receive path has entered the `WAIT_R` state after being enabled for a receive transaction, the RSI starts decrementing the timeout value supplied by the `RSI_DATA_TIMER`.

If the RSI is configured for a 1-bit data bus, the `DAT_TIMEOUT` flag is set if a start bit is not detected on the `RSI_DATA0` signal before the timeout counter reaches zero. The state machine then returns to the `IDLE` state.

If the RSI is configured for 4-bit bus mode and the start bit is not detected on all four `RSI_DATAx` signals before the timeout counter expires—the `DAT_TIMEOUT` flag is set. The state machine returns to the `IDLE` state. If a start bit is detected on some, but not all, of the `RSI_DATAx` signals on the same sampled clock cycle, then the `START_BIT_ERR` flag is set and the state machine returns to the `IDLE` state. Upon correct detection of the start bit, the state machine moves into the `RECEIVE` state.

The behavior of the `RECEIVE` state differs for stream and block transfers.

- For stream transfers, received data is packed into bytes and written to the data FIFO. Data is continuously received and written to the data FIFO until `RSI_DATA_CNT` decrements to zero.

When the counter reaches zero, the remaining data in the shift register is written into the FIFO, the `DAT_END` flag is set and the state machine transitions to the `WAIT_R` state.

When the receive FIFO is detected empty, the `RX_DAT_ZERO` flag is set and the state goes to the `IDLE` state.

Functional Description

If the data FIFO becomes full and data has not been read from the FIFO prior to the next byte being written to the FIFO, then the `RX_OVERRUN` flag is set. The state goes to the `WAIT_R` state then into the `IDLE` state.

- In block transfer mode, the received data is packed into bytes and written to the data FIFO. When `DATA_BLK_LGTH` bytes have been received, the RSI reads the 16-bit CRC check bits. If the received CRC matches the internally calculated CRC, the `DAT_BLK_END` flag is set and the state transitions to the `WAIT_R` state.

If the `RSI_DATA_CNT` counter expires in alignment with the end of a `DATA_BLK_LGTH` block, the `DAT_END` and `DAT_BLK_END` flags are set, and the state goes to the `WAIT_R` state. When the receive FIFO is detected empty, the `RX_DAT_ZERO` flag is set and the state goes to the `IDLE` state.

If `RSI_DATA_CNT` expires before the end of a `DATA_BLK_LGTH` block, the `DAT_CRC_FAIL` flag is set. The state goes to the `IDLE` state.

RSI Data Path CRC

The data CRC generator of the RSI calculates the 16-bit CRC checksum for all bits sent or received for a given block transaction. The data path CRC generator is not enabled for stream based data transfers. For a 1-bit bus configuration, the 16-bit CRC is calculated for all data sent on the `RSI_DATA0` signal. For a 4-bit-wide data bus, the 16-bit CRC is calculated separately for each `RSI_DATAx` signal. The data path CRC checksum is a 16-bit value calculated as follows.

$$CRC[15:0] = \text{Remainder} \frac{x_{16} \cdot M(x)}{G(x)}$$

with:

$$G(x) = x_{16} + x_{12} + x_5 + 1$$

where:

$$M(x) = x_{((8 \times DTX_BLK_LGTH) - 1)} \cdot (\text{first data bit}) + \dots + x_0 \cdot (\text{last data bit})$$

RSI Data FIFO

The data FIFO is a 32-bit wide, 16-word deep data buffer with transmit and receive logic. The FIFO is configuration depends on the state of the TX_ACT and RX_ACT flags. If TX_ACT is set, the FIFO operates as a transmit FIFO supplying data to the RSI for transfer to the card. If RX_ACT flag is set, the FIFO operates as a receive FIFO, where the RSI writes data received from the card. If neither TX_ACT nor RX_ACT flags are set, then the FIFO is disabled.

When the transmit FIFO is disabled, all the transmit status flags are de-asserted and the transmit read and write pointers are reset. The RSI asserts the TX_ACT flag upon starting a data transfer. During the data transfer the transmit logic maintains the transmit FIFO status flags shown in [Table 21-9](#).

Table 21-9. RSI Transmit FIFO Status Flags

RSI_STATUS Flag	Description
TX_FIFO_STAT	Transmit FIFO is half empty
TX_FIFO_FULL	Transmit FIFO is full
TX_FIFO_EMPTY	Transmit FIFO is empty
TX_UNDERRUN	Transmit FIFO under run error
TX_DAT_RDY	Valid data available in the transmit FIFO

When the receive FIFO is disabled, all receive status flags are de-asserted and the receive read and write pointers are reset. The RSI asserts the RX_ACT flag upon starting a data read transaction. During the data transfer, the receive logic maintains the receive FIFO status flags shown in [Table 21-10](#).

Functional Description

Table 21-10. RSI Receive FIFO Status Flags

RSI_STATUS Flag	Description
RX_FIFO_STAT	Receive FIFO is half empty
RX_FIFO_FULL	Receive FIFO is full
RX_FIFO_EMPTY	Receive FIFO is empty
RX_OVERRUN	Receive FIFO under run error
RX_DAT_RDY	Valid data available in the receive FIFO

SDIO Interrupt and Read Wait Support

Two additional RSI features implement SDIO functionality.

- Hardware interrupt support over the `RSI_DATA1` pin
- Read wait request over the `RSI_DATA2` pin

SDIO devices may have multiple interrupt sources that are mapped to a single interrupt line. The interrupt is level-sensitive, allowing multiple functions to generate an interrupt simultaneously. Thus the interrupt request will continually be asserted until all sources generating an interrupt are determined and cleared by the RSI.

The sources of the interrupts are found by interrogating the SDIO device. The interrupts are cleared through operations unique to each function.

The SDIO device sends an interrupt request to the RSI by asserting the `RSI_DATA1` signal low. The interrupt status is indicated by the `SDIO_INT_DET` bit of the `RSI_ESTAT` register. The status can be configured to interrupt the processor through the `SDIO_INT_DET_MASK` bit of the `RSI_EMASK` register.

When the RSI is configured for 1-bit bus width, the interrupt is generated by the SDIO with no timing constraints since the `RSI_DATA1` signal acts as a dedicated IRQ signal. The RSI should be configured using `RSI_CONFIG` such that pull-up are enabled on all `RSI_DATAx` signals. When the RSI

sampling `RSI_DATA1` low, the RSI asserts the `SDIO_INT_DET` flag. This flag is asserted until the `RSI_DATA1` signal is sampled high again.

When the RSI is configured for 4-bit bus widths, the `RSI_DATA1` signal is shared between the IRQ signal and the `RSI_DATA1` signal. In this configuration, the interrupt is only be recognized by the RSI within a specific interrupt period.

Programming Model

This section contains the following procedures:

- [“Card Identification”](#)
- [“Single Block Write Operations” on page 21-34](#)
- [“Single Block Read Operation” on page 21-38](#)
- [“Multiple Block Write Operation” on page 21-43](#)
- [“Multiple Block Read Operation” on page 21-47](#)

Card Identification

Before data transfers can take place between the RSI and the SD/MMC/SDIO device, the device type must be identified. During this phase, the `RSI_CLK` frequency must be no greater than 400 kHz.

SD Card Identification Procedure

The SD card identification procedure is shown below.

1. Issue the IDLE command to the card using the `RSI_COMMAND` register
2. Issue the `SEND_IF` command through the `RSI_COMMAND` register, supplying the host supply voltage and a check pattern via the `RSI_ARGUMENT` register. The command expects an R7 response type. If a valid response with a compatible voltage range and matching check pattern is received, the card is compliant with SD v2.00 or later. If a response is received with an incompatible voltage range the card cannot be used. If no response is received at all, as indicated by the `CMD_RESP_TIMEOUT` flag on the `RSI_STATUS` register, go to step 5.
3. Issue the `RSI_SEND_OP_COND` command through the `RSI_COMMAND` register, supplying the voltage window supported and whether the host supports high capacity cards using the `RSI_ARGUMENT` register. The RSI expects an R3 response to this command. The RSI can reject the card at this point if the voltage range is not compatible.

If the card returns a response indicating that it is busy, resend the `RSI_SEND_OP_COND` until the card indicates it is ready. If the host does not support the high capacity mode as indicated by setting the HCS bit of the argument to 0—a high capacity card never clears the busy status bit. The card should be identified within one second. If in that time the card is still busy, or no valid responses have been received, the card is rejected.

4. If the host supports high-capacity cards, verify whether the response in the `RSI_RESPONSE0` register indicates that the card capacity status (CCS) bit is set. If CCS is set, an SD v2.00 or later *high capacity* SD memory card is

present—proceed to step 6. If the CCS bit is cleared, then the card is an SD VeSDH on 2.00 or later *standard capacity* memory card—proceed to step 6.

5. Issue the `RSI_SEND_OP_COND` command via the `RSI_COMMAND` register, supplying the voltage window supported and with the high-capacity support (HCS) bit set to 0 via the `RSI_ARGUMENT` register. The RSI expects an R3 response to this command, at which time the card can be rejected if the voltage range is not compatible. If the card returns a response indicating that it is busy, resend the `RSI_SEND_OP_CMD` until the card indicates that it is ready. The card should be identified within 1 second. If in that timeframe the card is still busy or no valid responses have been received, the card should be rejected. Once the response indicates that the card is ready, the card type has been identified as an SD Version 1.x standard-capacity memory card.
6. Issue the `ALL_SEND_CID` command to which an R2 response type is expected. This will result in the card sending the 128-bit card identification (CID) register and transitioning from ready to identification mode.
7. Issue the `SEND_RELATIVE_ADDR` command to which an R6 response type is expected. This will result in the card issuing a new relative address which must be used in order to select the card in the future for data transfers. The card will then move into standby mode completing the identification procedure.

MMC Identification Procedure

The MMC identification procedure is as follows:

1. Issue the `IDLE` command to the card via the `RSI_COMMAND` register.
2. Issue the `SEND_OP_COND` command to the card via the `RSI_COMMAND` register, supplying the operating voltage window that the host is compatible with and the access mode that the host supports (byte or sector) via the `RSI_ARGUMENT` register. The RSI expects an `R3` type response. This allows the host to reject the card if it is not compatible with the supply voltage or if the access mode is not supported by the host software. If the card returns an indication that it is busy, repeat this step until the card is either rejected or not busy.
3. Issue the `ALL_SEND_CID` command via the `RSI_COMMAND` register. The RSI expects an `R2` response to this command. This will result in the card sending the 128-bit card identification (CID) register and transitioning from ready to identification mode.
4. Issue the `SET_RELATIVE_ADDR` command, providing a 16-bit relative card address (RCA) via the `RSI_ARGUMENT` register that will get assigned to the card. An `R1` response type is expected for this command. This will result in the card being assigned with the provided RCA, which must be used in order to select the card in the future for data transfers. The card will then move into standby mode, completing the identification procedure.

Single Block Write Operations

Block write operations typically consist of 512 bytes of data per block. If the card is found to support other block lengths or the default block length as specified in the CID register is not 512, the block length of the RSI must be configured accordingly. The block length of the card and the block length of the RSI must be configured for the same block size at all

times. The block length of the RSI is configured via the `DATA_BLK_LGTH` field of the `RSI_DATA_CONTROL` register.

 It is important to pay attention as to when the data path state machine is enabled and when data is written to the FIFO for transfer to the card. Write transactions require that data be written after the response has completed for the `WRITE_BLOCK` command. If the data path state machine is enabled prior to sending the `WRITE_BLOCK` command, data must not be written to the transmit FIFO via the DMA or core until after the response has been received as indicated by the `CMD_RESP_END` flag. Failure to adhere to this procedure can result in data being written to the card in violation to the block write timing parameters, resulting in a data CRC failure.

Using Core

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register.
4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 512 bytes for a single block.

Programming Model

6. Write the appropriate timeout value for a write operation to the `RSI_DATA_TIMER` register.
7. Write the destination start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512 byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or sector-addressable MMC card.
8. Write the `WRITE_BLOCK` command to the `RSI_COMMAND` register, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
9. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register.
10. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` should also be set to enable the data path state machine. All other fields of the `RSI_DATA_CONTROL` register should be zero.
11. Write data to the `RSI_FIFO` register until the FIFO becomes full as indicated by the `TX_FIFO_FULL` flag of the `RSI_STATUS` register. Continue to write data to the FIFO as long as the FIFO is not full or write data in blocks of eight 32-bit words if polling on the `TX_FIFO_STAT` bit indicating the transmit FIFO is half empty. Continue until all 128 32-bit words (512 bytes) have been transferred.
12. Wait for the card to respond with the CRC token by waiting for the `DAT_BLK_END` flag to be set. `DAT_END` will also be set at this time if the `RSI_DATA_LGTH` register was set to 512 bytes in step 5.
13. Clear the `DAT_BLK_END` and `DAT_END` flags via the `RSI_STATUSCL` register.

Using DMA

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Configure the DMA channel assigned to the RSI controller. Write `DMAx_START_ADDR` with the address of the first byte of data to be written to the card. The `DMAx_X_COUNT` register should be set to 128, and the `DMAx_X_MODIFY` register to 4. The `DMAx_CONFIG` register should be set for DMA enable (a word size of 32-bits).
6. Once the DMA channel has been configured and enabled, write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 512 bytes for a single block.
7. Write the appropriate timeout value for a write operation to the `RSI_DATA_TIMER` register.
8. Write the destination start address to the `RSI_ARGUMENT` register. The address supplied must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or sector-addressable MMC card.

Programming Model

9. Write the `WRITE_BLOCK` command to `RSI_COMMAND`, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
10. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
11. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` and `DATA_DMA_EN` should also be set to enable the data path state machine and to allow the DMA controller to access the transmit FIFO. All other fields of the `RSI_DATA_CONTROL` register should be zero.
12. Wait for the card to respond with the CRC token by waiting for the `DAT_BLK_END` flag to be set. `DAT_END` will also be set at this point if the `RSI_DATA_LGTH` register was set to 512 bytes in step 5.
13. Clear the `DAT_BLK_END` and `DAT_END` flags via the `RSI_STATUSCL` register. Also clear the `DMA_DONE` bit of the `DMAx_IRQ_STATUS` register, if applicable.

Single Block Read Operation

Block read operations typically consist of 512 bytes of data per block. If the card is found to support other block lengths or the default block length as specified in the CID register is not 512, the block length of the RSI must be configured accordingly. The block length of the card and the block length of the RSI must be configured for the same block size at all times. The block length of the RSI is configured via the `DATA_BLK_LGTH` field of the `RSI_DATA_CONTROL` register.



It is important to pay attention as to when the data path state machine is enabled and when data is read from the receive FIFO for data transfers from the card to the RSI. Read transactions can

occur on the `RSI_DATAx` signals prior to the response of the command being received. It is therefore advisable to enable the data path state machine, and DMA controller if being used, either:

- Prior to issuing a command that involves a data read packet
- Immediately after the command has been issued but prior to pending on the `CMD_RESP_END` flag



If the core is being used to read the receive FIFO, it is advised not to pend on the `CMD_RESP_END` flag. It is possible for data to be driven on the `RSI_DATAx` signals two `RSI_CLK` cycles after the end bit of the command. At minimum, an additional 48 `RSI_CLK` cycles will pass before the response is received, during which time the receive buffer may potentially have received 24 bytes of data on a 4-bit bus and will be approaching the half full state. Software should ensure that the receive buffer does not become full prior to data being read from the receive FIFO.

Using Core

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register.

Programming Model

4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 512 bytes for a single block.
6. Write the appropriate timeout value for a read operation to the `RSI_DATA_TIMER` register.
7. Write the destination start address to the `RSI_ARGUMENT` register. The address supplied must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or sector-addressable MMC card.
8. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` and `DATA_DIR` should also be set to enable the data path state machine and indicate the transfer direction is from card to controller. All other fields of the `RSI_DATA_CONTROL` register should be zero.
9. Write the `READ_SINGLE_BLOCK` command to the `RSI_COMMAND` register, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
10. In order to meet some timing restrictions related to block read operations, it is advisable to not wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register but instead move immediately on the next step. This is due to the card being able to send data before a response can be completed on the `RSI_CMD` signal, moving immediately onto step 11 will ensure a receive FIFO overflow does not occur.

11. Poll the `RX_FIFO_RDY` bit or the `RX_DAT_ZERO` bit of `RSI_STATUS` indicating the receive FIFO has data available, or the receive FIFO is empty. As long as the receive FIFO is not empty, read data from the `RSI_FIFO` register until all 512 bytes have been read
12. Once all bytes have been read, wait for the `DAT_BLK_END` flag to indicate that the data was received correctly and passed the CRC check. The `DAT_END` flag may also be set, depending on the value written to `RSI_DATA_LGTH`.
13. Clear the `DAT_BLK_END` and `DAT_END` flags via the `RSI_STATUSCL` register.

Using DMA

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Configure the DMA channel assigned to the RSI controller. Write `DMAx_START_ADDR` with the address of the first byte of where the received data is to be stored. The `DMAx_X_COUNT` register should be

Programming Model

set to 128 and the `DMAX_X_MODIFY` register to 4. The `DMAX_CONFIG` register should be set for DMA enable (a word size of 32-bits and direction set to memory write).

6. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 512 bytes for a single block.
7. Write the appropriate timeout value for a read operation to the `RSI_DATA_TIMER` register.
8. Write the source start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or sector-addressable MMC card.
9. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN`, `DATA_DIR`, and `DATA_DMA_EN` should also be set to enable the data path state machine, set the transfer direction from card to controller and allow the DMA controller access to the receive FIFO. All other fields of the `RSI_DATA_CONTROL` register should be zero.
10. Write the `READ_SINGLE_BLOCK` command to the `RSI_COMMAND` register, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is `R1`.
11. Unlike core accesses, it is safe to poll on `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register. The DMA controller enabled in step 5 will ensure any data sent to the receive FIFO prior to the `CMD_RESP_END` flag being set is received correctly.

12. Wait for the `DAT_BLK_END` flag to indicate that the data was received correctly and passed the CRC check. The `DAT_END` flag may also be set, depending on the value written to `RSI_DATA_LGTH`.
13. Clear the `DAT_BLK_END` and `DAT_END` flags via the `RSI_STATUSCL` register. Also clear the `DMA_DONE` bit of the `DMAx_IRQ_STATUS` register, if applicable.

Multiple Block Write Operation

Block write operations typically consist of 512 bytes of data per block. If the card is found to support other block lengths or the default block length as specified in the CID register is not 512, the block length of the RSI must be configured accordingly. The block length of the card and the block length of the RSI must be configured for the same block size at all times. The block length of the RSI is configured via the `DATA_BLK_LGTH` field of the `RSI_DATA_CTL` register.

Using Core

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register.

Programming Model

4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. For example, write 4096 to write eight blocks of 512 bytes.
6. Write the appropriate timeout value for a write operation to the `RSI_DATA_TIMER` register.
7. Write the destination start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or a sector-addressable MMC card.
8. Write the `WRITE_MULTIPLE_BLOCK` command to `RSI_COMMAND`, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
9. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
10. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` should also be set to enable the data path state machine. All other fields of the `RSI_DATA_CONTROL` register should be zero.
11. Write data to the `RSI_FIFO` register until the FIFO becomes full as indicated by the `TX_FIFO_FULL` flag of the `RSI_STATUS` register. Continue to write data to the FIFO as long as the FIFO is not full or write data in blocks of eight 32-bit words if polling on the `TX_FIFO_STAT` bit indicating the transmit FIFO is half empty. Continue until all 128 32-bit words (512 bytes) have been transferred.

12. Wait for the card to respond with the CRC token by waiting for the `DAT_BLK_END` flag to be set.
13. Clear the `DAT_BLK_END` flag.
14. Repeat steps 11 to 13 for the number of blocks to be transferred or until `DAT_END` flag is set.
15. Write the `RSI_COMMAND` register with the `STOP_TRANSMISSION` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is `R1`.
16. Clear the `DAT_END` and `CMD_RESP_END` flags via the `RSI_STATUSCL` register.

Using DMA

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is `R1b`.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
4. Ensure that the device is not busy and no errors occurred by verifying the response contained in `RSI_RESPONSE0`.

Programming Model

5. Configure the DMA channel assigned to the RSI controller. Write `DMAx_START_ADDR` with the address of the first byte of data to be written to the card. The `DMAx_X_COUNT` register should be set to the overall number of 32-bit words to be written; for example, write 1024 to transfer 4096 bytes. The `DMAx_X_MODIFY` register should be set to 4. The `DMAx_CONFIG` register should be set for DMA enable and a word size of 32-bits.
6. Once the DMA channel has been configured and enabled, write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. For example, write 4096 to write eight blocks of 512 bytes.
7. Write the appropriate timeout value for a write operation to the `RSI_DATA_TIMER` register.
8. Write the destination start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or sector addressable MMC card.
9. Write the `WRITE_MULTIPLE_BLOCK` command to the `RSI_COMMAND`, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
10. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCL` register.
11. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` and `DATA_DMA_EN` should also be set to enable the data path state machine and to allow the DMA controller to access the transmit FIFO. All other fields of the `RSI_DATA_CONTROL` register should be zero.

12. Poll for the `DAT_END` flag or alternatively poll for each instance of the `DAT_BLK_END` flag that will be set on successful completion of each block transfer. For a 4096 byte transfer, `DAT_BLK_END` will be set eight times and should be cleared after it is detected via the `RSI_STATUSCL` register.
13. Write the `RSI_COMMAND` register with the `STOP_TRANSMISSION` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
14. Clear the `DAT_END` and `CMD_RESP_END` flag via the `RSI_STATUSCL` register. Also clear the `DMA_DONE` bit of the `DMAx_IRQ_STATUS` register, if applicable.

Multiple Block Read Operation

Block read operations typically consist of 512 bytes of data per block. If the card is found to support other block lengths or the default block length as specified in the CID register is not 512, the block length of the RSI must be configured accordingly. The block length of the card and the block length of the RSI must be configured for the same block size at all times. The block length of the RSI is configured via the `DATA_BLK_LGTH` field of the `RSI_DATA_CONTROL` register.

Using Core

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCLR` register.
4. Ensure that the device is not busy and there are no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 4096 for a transfer of eight 512 byte blocks.
6. Write the appropriate timeout value for a read operation to the `RSI_DATA_TIMER` register.
7. Write the destination start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or a sector-addressable MMC card.
8. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN` and `DATA_DIR` should also be set to enable the data path state machine and indicate the transfer direction is from card to controller. All other fields of the `RSI_DATA_CONTROL` register should be zero.

9. Write the `READ_MULTIPLE_BLOCK` command to the `RSI_COMMAND` register, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is `R1`.
10. In order to meet some timing restrictions related to block read operations, it is advisable to not wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register but instead move immediately on the next step. This is due to the card being able to send data before a response can be completed on the `RSI_CMD` signal, moving immediately onto step 11 will ensure a receive FIFO overflow does not occur.
11. Poll the `RX_FIFO_RDY` bit or the `RX_DAT_ZERO` bit of `RSI_STATUS` indicating the receive FIFO has data available, or the receive FIFO is empty. As long as the receive FIFO is not empty, read data from the `RSI_FIFO` register until 512 bytes have been read.
12. Once the block has been read, wait for the `DAT_BLK_END` flag to indicate that the data was received correctly and passed the CRC check.
13. Clear the `DAT_BLK_END` flag via `RSI_STATUSCL`.
14. Repeat steps 11 to 13 until the required number of blocks have been read or until the `DAT_END` flag has been set.
15. Write the `RSI_COMMAND` register with the `STOP_TRANSMISSION` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is `R1`.
16. Clear the `DAT_END` and `CMD_RESP_END` flags via the `RSI_STATUSCL` register.

Using DMA

The procedure is as follows:

1. Write the `RSI_ARGUMENT` register with the cards RCA. The 16-bit RCA should be written to the upper 16-bits of the `RSI_ARGUMENT` register.
2. Write the `RSI_COMMAND` register with the `SELECT/DESELECT_CARD` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1b.
3. Wait for the `CMD_RESP_END` indication within the `RSI_STATUS` register, and clear the status bit once detected via the `RSI_STATUSCLR` register.
4. Ensure that the device is not busy and that no errors occurred by verifying the response contained in `RSI_RESPONSE0`.
5. Configure the DMA channel assigned to the RSI controller. Write `DMAx_START_ADDR` with the address of the first byte of where the received data is to be stored. The `DMAx_X_COUNT` register should be set to the number of 32-bit words to be read, which would be 1024 for a 4096 byte read transfer. The `DMAx_X_MODIFY` register should be set to 4. The `DMAx_CONFIG` register should be set for DMA enable (a word size of 32-bits and direction set to memory write).
6. Write the number of bytes to be transferred to the `RSI_DATA_LGTH` register. This will be 4096 for eight blocks of 512 bytes.
7. Write the appropriate timeout value for a read operation to the `RSI_DATA_TIMER` register.
8. Write the source start address to the `RSI_ARGUMENT` register. The supplied address must be aligned to a 512-byte boundary if misaligned accesses are not enabled and the card is not a high-capacity SD card or a sector-addressable MMC card.

9. Enable the data path state machine by writing to the `RSI_DATA_CONTROL` register with `DATA_BLK_LGTH` set to 9 for a 512-byte block. `DATA_EN`, `DATA_DIR`, and `DATA_DMA_EN` should also be set to enable the data path state machine. Set the transfer direction from card to controller and allow the DMA controller access to the receive FIFO. All other fields of the `RSI_DATA_CONTROL` register should be zero.
10. Write the `READ_MULTIPLE_BLOCK` command to the `RSI_COMMAND` register, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
11. Unlike core accesses, it is safe to poll on `CMD_RESP_END` indication within the `RSI_STATUS` register and clear the status bit once detected via the `RSI_STATUSCL` register. The DMA controller, enabled in step 5 will ensure any data sent to the receive FIFO prior to the `CMD_RESP_END` flag being set is received correctly.
12. Poll for the `DAT_END` flag or alternatively poll for each instance of the `DAT_BLK_END` flag that will be set on successful completion of each block transfer. For a 4096-byte transfer, `DAT_BLK_END` will be set eight times and should be cleared after it is detected via the `RSI_STATUSCL` register.
13. Write the `RSI_COMMAND` register with the `STOP_TRANSMISSION` command, configuring the command path state machine to expect a short response by setting `CMD_RESP` and clearing `CMD_L_RESP`. The response type is R1.
14. Clear the `DAT_END` and `CMD_RESP_END` flags via the `RSI_STATUSCL` register. Also clear the `DMA_DONE` bit of the `DMAx_IRQ_STATUS` register, if applicable.

RSI Registers

Table 21-11 summarizes the RSI registers together with their function, memory-mapped address, type, and access.

Table 21-11. RSI Module Registers

Register Name	Function	Address	Type	Access
RSI_PWR_CONTROL	RSI power control register on page 21-53	0xFFC03800	R/W	16-bit
RSI_CLK_CONTROL	RSI clock control register on page 21-54	0xFFC03804	R/W	16-bit
RSI_ARGUMENT	RSI argument register on page 21-56	0xFFC03808	R/W	32-bit
RSI_COMMAND	RSI command register on page 21-57	0xFFC0380C	R/W	16-bit
RSI_RESP_CMD	RSI response command register on page 21-58	0xFFC03810	R	16-bit
RSI_RESPONSE0 RSI_RESPONSE1 RSI_RESPONSE2 RSI_RESPONSE3	RSI response registers on page 21-59	0xFFC03814 0xFFC03818 0xFFC0381C 0xFFC03820	R	32-bit
RSI_DATA_TIMER	RSI data timer register on page 21-60	0xFFC03824	R/W	32-bit
RSI_DATA_LGTH	RSI data length register on page 21-61	0xFFC03828	R/W	16-bit
RSI_DATA_CONTROL	RSI data control register on page 21-62	0xFFC0382C	R/W	16-bit
RSI_DATA_CNT	RSI data counter register on page 21-63	0xFFC03830	R	16-bit
RSI_STATUS	RSI status register on page 21-64	0xFFC03834	R	32-bit
RSI_STATUSCL	RSI status clear register on page 21-68	0xFFC03838	W1A	16-bit

Table 21-11. RSI Module Registers (Continued)

Register Name	Function	Address	Type	Access
RSI_MASK0 RSI_MASK1	RSI IRQ0 mask registers on page 21-70	0xFFC0383C 0xFFC03840	R/W	32-bit
RSI_FIFO_CNT	RSI FIFO counter register on page 21-73	0xFFC03848	R	16-bit
RSI_CEATA_CONTROL	RSI CE-ATA control register on page 21-73	0xFFC0384C	R/W1A/W	16-bit
RSI_FIFO	RSI data FIFO register on page 21-74	0xFFC03880	R/W	32-bit
RSI_ESTAT	RSI exception status register on page 21-75	0xFFC038C0	R/W1C	16-bit
RSI_EMASK	RSI exception mask register on page 21-76	0xFFC038C4	R/W	16-bit
RSI_CONFIG	RSI configuration register on page 21-77	0xFFC038C8	R/W	16-bit
RSI_RD_WAIT_EN	RSI read wait enable register on page 21-79	0xFFC038CC	R/W1A/W	16-bit
RSI_PID0 RSI_PID1 RSI_PID2 RSI_PID3	RSI peripheral identifica- tion registers on page 21-80	0xFFC038D0 0xFFC038D4 0xFFC038D8 0xFFC038DC	R	16-bit

RSI Power Control Register (RSI_PWR_CONTROL)

The `RSI_PWR_CONTROL` register contains bits that control the power to the RSI module as well as the open-drain configuration for the `RSI_CMD` signal. The `PWR_ON` field must be set to “11” in order for the RSI to be enabled. The `RSI_CMD_OD` bit, when set, results in the RSI driving the `RSI_CMD` signal in open-drain mode. The default mode of operation is push-pull. After a data write, data cannot be written to this register for a five `SCLK` cycles.

RSI Registers



The `RSI_CLK_EN` bit in the `RSI_CONFIG` register must be set before writing to other registers; except for the `RSI_EMASK`, `RSI_STAT`, `RSI_CONFIG` and `RSI_RD_WAIT_EN` registers.

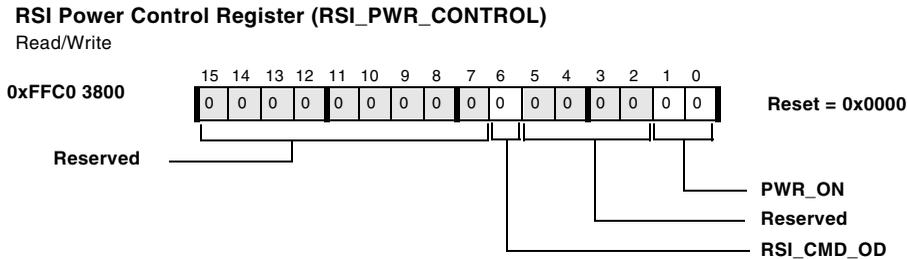


Figure 21-6. RSI Power Control Register

Table 21-12. RSI_PWR_CONTROL Register

Bit	Name	Function	Type	Default
1:0	PWR_ON	Power on 00 = RSI disabled 01 = Reserved 10 = Reserved 11 = RSI enabled	RO	0
5:2	Reserved	Reserved	RO	0
6	RSI_CMD_OD	RSI_CMD open drain 0 = Disabled (push-pull) 1 = Enabled	RO	0
15:7	Reserved	Reserved	RO	0

RSI Clock Control Register (RSI_CLK_CONTROL)

The `RSI_CLK_CONTROL` register provides control functionality for the RSI clock. `RSI_CLK` can be derived directly from the `SCLK` signal by enabling `CLKDIV_BYPASS`; otherwise, `RSI_CLK` frequency is determined from the current `SCLK` frequency and the `CLKDIV` field as shown in [Equation 21-1](#).

Equation 21-1.

$$RSI_CLK = \frac{SCLK}{2 \times (CLKDIV + 1)}$$

In order to conserve power, the RSI clock can be disabled without disabling the entire RSI interface via the `CLK_EN` bit; additionally the `PWR_SV_EN` bit, when set, results in the `RSI_CLK` signal only been driven when the RSI is performing a transfer either to or from the card. In addition to clock control functionality, the data bus width of the RSI interface is also controlled from this register.

RSI Clock Control Register (RSI_CLK_CONTROL)

Read/Write

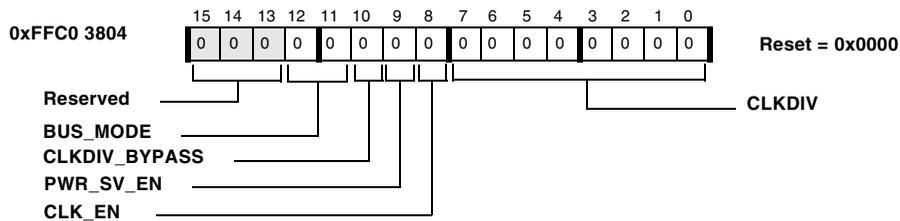


Figure 21-7. RSI Clock Control Register

Table 21-13. RSI_CLK_CONTROL Register

Bit	Name	Function	Type	Default
7:0	CLKDIV	Clock divisor 0x0 to 0xFF (see Equation 21-1)	R/W	0
8	CLK_EN	RSI_CLOCK enable 0 = Disable RSI_CLK 1 = Enable RSI_CLK	R/W	0
9	PWR_SV_EN	Power save enable 0 = Disabled (RSI_CLK always driven) 1 = Enabled (RSI_CLK only enabled when bus is active)	R/W	0

RSI Registers

Table 21-13. RSI_CLK_CONTROL Register (Continued)

Bit	Name	Function	Type	Default
10	CLKDIV_BYPASS	Bypass clock divisor 0 = Disabled (do not bypass clock divisor) 1 = Enabled (RSI_CLK derived directly from SCLK)	R/W	0
12:11	BUS_MODE	Data bus width 00 = 1-bit data bus 01 = 4-bit data bus 10 = 8-bit data bus 11 = Reserved	R/W	0
15:13	Reserved	Reserved	RO	0

RSI Argument Register (RSI_ARGUMENT)

The `RSI_ARGUMENT` register contains the 32-bit argument that is sent on the `RSI_CMD` signal as part of a command message. If a command requires an argument, the argument must first be loaded into the `RSI_ARGUMENT` register prior to writing and enabling the command in the `RSI_COMMAND` register.

RSI Argument Register (RSI_ARGUMENT)

Read/Write

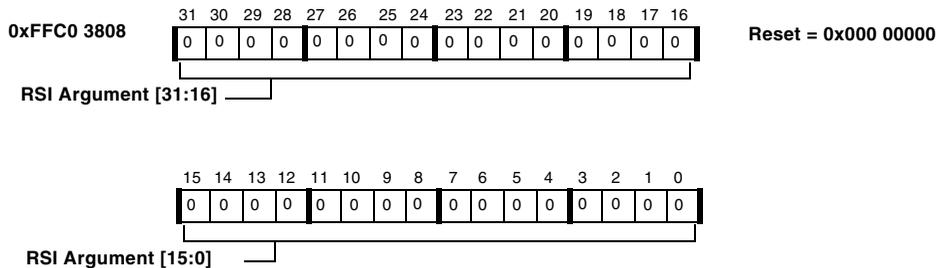


Figure 21-8. RSI Argument Register

RSI Command Register (RSI_COMMAND)

The `RSI_COMMAND` register is responsible for controlling the command path state machine. The `CMD_IDX` field contains the index of the command to be issued via the RSI as part of the command message. If the command requires a response, this is indicated via `CMD_RSP_EN`.

The length of the response (short or long) is controlled with `CMD_LRSP_EN`. The command path state machine becomes active once the `CMD_EN` bit is set and is disabled if this bit is cleared.

i It is not required to manually clear the `CMD_EN` bit after the command sequence has completed. The command path state machine will automatically terminate and become IDLE once the operation has completed.

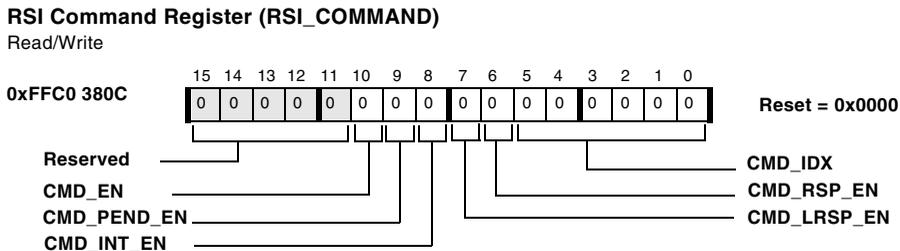


Figure 21-9. RSI Command Register

Table 21-14. RSI_COMMAND Register

Bit	Name	Function	Type	Default
5:0	CMD_IDX	Command index 0x3F - 0x00 (Command number to be issued)	R/W	0
6	CMD_RSP_EN	Wait for response 0 = Disabled 1 = Enabled	R/W	0

RSI Registers

Table 21-14. RSI_COMMAND Register (Continued)

Bit	Name	Function	Type	Default
7	CMD_LRSP_EN	Long response enable 0 = Disabled (short response expected) 1 = Enabled (long response expected)	R/W	0
8	CMD_INT_EN	Command interrupt enable 0 = Disabled (timeout after 64 RSI_CLK cycles) 1 = Enabled (disable timeout counter and wait for interrupt)	R/W	0
9	CMD_PEND_EN	Pend enable 0 = Disabled (send command immediately) 1 = Enabled (wait for DAT_END before sending command)	R/W	0
10	CMD_EN	Command enable 0 = Disable command path state machine 1 = Enable command path state machine	R/W	0
15:11	Reserved	Reserved	RO	0

RSI Response Command Register (RSI_RESP_CMD)

The RSI_RESP_CMD register contains the command index field of the last response received. If the command response does not contain a command index field (as is the case with a long response), the RESP_CMD field would typically be ignored. In this situation, it will likely contain “0x3F”, which is the value of the reserved field of the response.

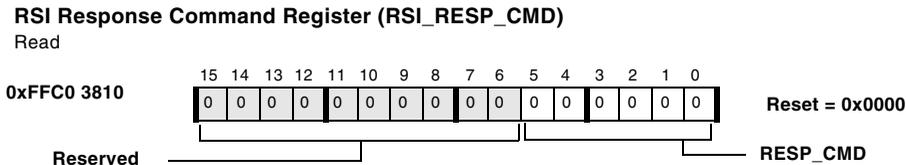


Figure 21-10. RSI Response Command Register

Table 21-15. RSI_RESP_CMD Register

Bit	Name	Function	Type	Default
5:0	RESP_CMD	Command index of last received response 0x3F - 0x00 (command index)	RO	0
15:6	Reserved	Reserved	RO	0

RSI Response Registers (RSI_RESPONSEx)

The four `RSI_RESPONSEx` registers (`RSI_RESPONSE0`, `RSI_RESPONSE1`, `RSI_RESPONSE2`, and `RSI_RESPONSE3`), contain the response information received back from a card for a given command message. The received response may be 32 or 127 bits in length, depending on whether the response type is short or long. The most significant bit of the response is received first and is located in bit 31 of the `RSI_RESPONSE0` register. Bit 0 of `RSI_RESPONSE3` is always zero. [Table 21-16](#) shows the RSI response registers contents for the two types of responses.

RSI Registers

RSI Response Registers (RSI_RESPONSEx)

Read

RSI_RESPONSE0 = 0xFFC0 3814
 RSI_RESPONSE1 = 0xFFC0 3818
 RSI_RESPONSE2 = 0xFFC0 381C
 RSI_RESPONSE3 = 0xFFC0 3820

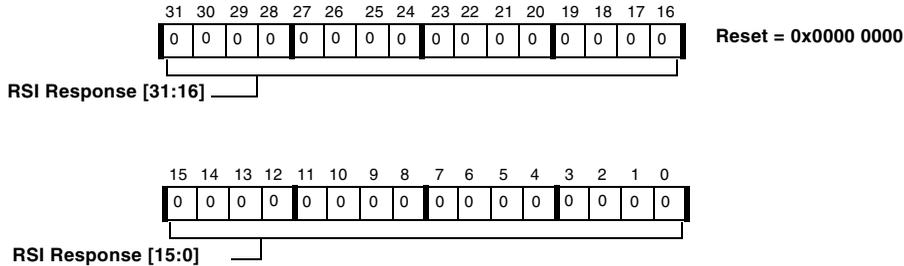


Figure 21-11. RSI Response Registers

Table 21-16. RSI Response Registers Content

Response Register	Short Response	Long Response
RSI_RESPONSE0	Response bits [31:0]	Response bits [127:96]
RSI_RESPONSE1	Not used	Response bits [95:64]
RSI_RESPONSE2	Not used	Response bits [63:32]
RSI_RESPONSE3	Not used	Response bits [31:1] ¹

1 Bits 31:1 of the long response are stored in bits 30:0 of the RSI_RESPONSE3 register. Bit 31 of the RSI_RESPONSE3 register is not used and is always zero.

RSI Data Timer Register (RSI_DATA_TIMER)

The RSI_DATA_TIMER register contains a 32-bit value for the data timeout period (RSI_CLK cycles). An internal counter loads the value from this register, and starts to decrement when the data path state machine enters the WAIT_R or the BUSY states. If the timer decrements to zero while the data path state machine is still in either of these two states, the DAT_TIMEOUT flag of the RSI_STATUS register is set. The RSI_DATA_TIMER

and the `RSI_DATA_LGTH` registers must both be written to prior to starting a data transfer via the `RSI_DATA_CONTROL` register.

RSI Data Timer Register (`RSI_DATA_TIMER`)

Read

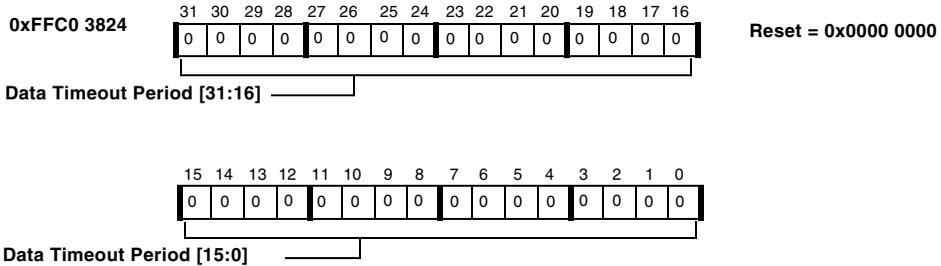


Figure 21-12. RSI Data Timer Register

RSI Data Length Register (`RSI_DATA_LGTH`)

The `RSI_DATA_LGTH` register contains a 16-bit value for the number of data bytes to be transferred before setting the `DAT_END` flag of the `RSI_STATUS` register. The value loaded to this register is copied into the `RSI_DATA_CNT` register when the data path state machine is enabled and starts the transfer.

RSI Data Length Register (`RSI_DATA_LGTH`)

Read

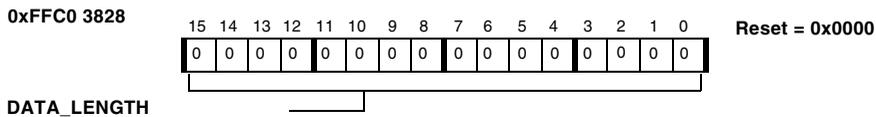


Figure 21-13. RSI Data Length Register

RSI Registers

Table 21-17. RSI_DATA_LGTH Register

Bit	Name	Function	Type	Default
15:0	DATA_LENGTH	Number of bytes to be transferred	R/W	0

RSI Data Control Register (RSI_DATA_CONTROL)

The `RSI_DATA_CONTROL` register largely controls the data path state machine. The state machine becomes enabled once the `DATA_EN` bit is set. The direction of the transfer is determined by `DATA_DIR`. If the DMA channel is to be used for the data transfer, the `DATA_DMA_EN` bit must be set; otherwise, the RSI FIFO is only accessible via the core. For block transfers, the block length must be specified via `DATA_BLK_LGTH`, where the block length is $2^{\text{DATA_BLK_LGTH}}$. Two bits (`CEATA_CCS_EN` and `CEATA_EN`) in this register configure the behavior of the command path state machine for communication with CE-ATA devices. After a data write, data cannot be written to this register for five `SCLK` cycles.

RSI Data Control Register (RSI_DATA_CONTROL)

Read/Write

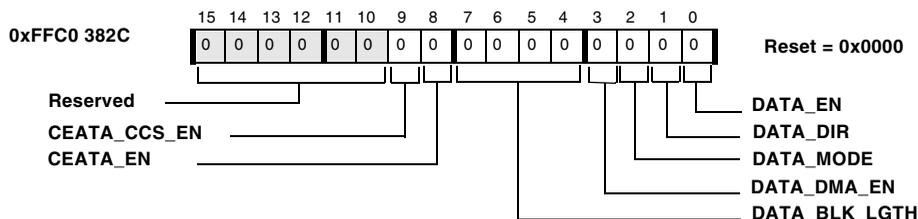


Figure 21-14. RSI Data Control Register

Table 21-18. RSI_DATA_CONTROL Register

Bit	Name	Function	Type	Default
0	DATA_EN	Data enable 0 = Disabled (disables data path state machine) 1 = Enabled (enables data path state machine)	R/W	0
1	DATA_DIR	Data transfer direction 0 = From RSI to card 1 = From card to RSI	R/W	0
2	DATA_MODE	Data transfer mode 0 = Block transfer 1 = Stream transfer	R/W	0
3	DATA_DMA_EN	Data DMA enable 0 = Disabled (use core to read/write RSI_FIFO) 1 = Enabled (use DMA controller to read/write RSI_FIFO)	R/W	0
7:4	DATA_BLK_LGTH	Data block length 0x0 - 0xC data block length (2^0 to 2^{12})	R/W	0
8	CEATA_EN	CE-ATA mode enable 0 = Disabled 1 = Enabled	R/W	0
9	CEATA_CCS_EN	Command completion signal enable 0 = Disabled 1 = Enabled (wait for command completion signal)	R/W	0
15:10	Reserved	Reserved	R/W	0

RSI Data Counter Register (RSI_DATA_CNT)

The RSI_DATA_CNT register is loaded from the RSI_DATA_LGTH register when the data path state machine becomes enabled and moves from the IDLE state to the WAIT_S or WAIT_R states. As the data is transferred,

RSI Registers

the counter decrements; upon decrementing to zero, the state machine then moves back to the IDLE state and the DAT_END flag of the RSI_STATUS register is set.

RSI Data Counter Register (RSI_DATA_CNT)

Read

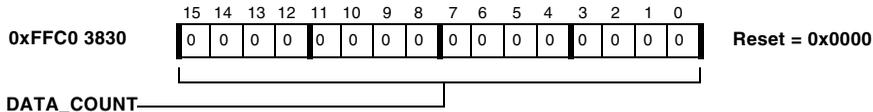


Figure 21-15. RSI Data Counter Register

Table 21-19. RSI_DATA_CNT Register

Bit	Name	Function	Type	Default
15:0	DATA_COUNT	Number of bytes still to be transferred	RO	0

RSI Status Register (RSI_STATUS)

The RSI_STATUS register contains both static and dynamic flags that indicate the status of the RSI. The static flags (bits [10:0]) remain asserted and are required to be cleared by writing to the RSI_STATUSCL register. The dynamic flags (bits [21:11]) change state, depending on the state of the underlying logic. The transmit and receive FIFO logic controls bits [21:12], which will vary depending on the state of the FIFO and whether the FIFO is currently enabled for a transmit or receive operation.

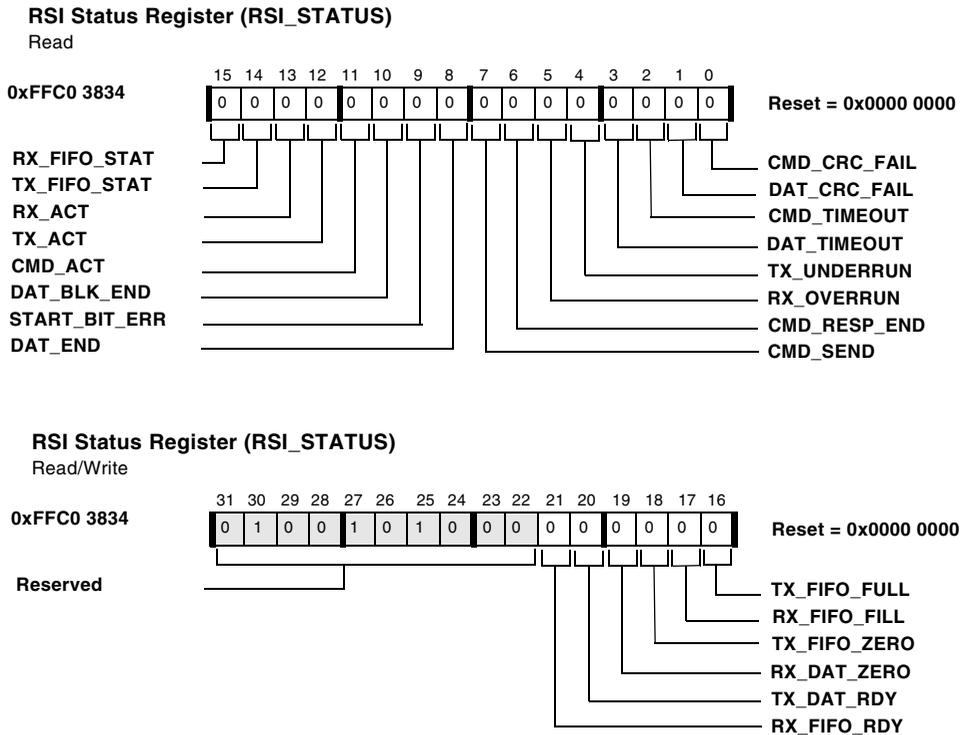


Figure 21-16. RSI Status Register

Table 21-20. RSI_STATUS Register

Bit	Name	Function	Type	Default
0	CMD_CRC_FAIL	Command response CRC fail 0 = No CRC received 1 = CRC failed on command response	RO	0
1	DAT_CRC_FAIL	Data CRC failure 0 = No CRC received on data block 1 = CRC failed on data block	RO	0

RSI Registers

Table 21-20. RSI_STATUS Register (Continued)

Bit	Name	Function	Type	Default
2	CMD_TIMEOUT	Command timeout 0 = Command response not timed out 1 = Command response timed out	RO	0
3	DAT_TIMEOUT	Data timeout 0 = Data not timed out 1 = Data timed out	RO	0
4	TX_UNDERRUN	Transmit FIFO underrun error 0 = No error 1 = Underrun error	RO	0
5	RX_OVERRUN	Receive FIFO overrun error 0 = No error 1 = Overrun error	RO	0
6	CMD_RESP_END	Command response received 0 = No response received 1 = Response received and CRC passed	RO	0
7	CMD_SENT	Command sent 0 = No command sent 1 = Command sent (no response required)	RO	0
8	DAT_END	End of data 0 = Not end of data 1 = End of data	RO	0
9	START_BIT_ERR	Start bit error 0 = No start bit error 1 = Start bit error (start bit not detected on all enabled data signals)	RO	0
10	DAT_BLK_END	Data block end 0 = No data block end 1 = End of data block and CRC passed	RO	0

Table 21-20. RSI_STATUS Register (Continued)

Bit	Name	Function	Type	Default
11	CMD_ACT	Command active 0 = No command active 1 = Command transfer in progress	RO	0
12	TX_ACT	Data transmit active 0 = No data transmit in progress 1 = Data transmit in progress	RO	0
13	RX_ACT	Data receive active 0 = No data receive in progress 1 = Data receive in progress	RO	0
14	TX_FIFO_STAT	Transmit FIFO watermark 0 = No FIFO watermark detected 1 = Transmit FIFO half empty	RO	0
15	RX_FIFO_STAT	Receive FIFO watermark 0 = No FIFO watermark detected 1 = Receive FIFO half full	RO	0
16	TX_FIFO_FULL	Transmit FIFO full 0 = Not full 1 = Transmit FIFO full	RO	0
17	RX_FIFO_FULL	Receive FIFO full 0 = Not full 1 = Receive FIFO full	RO	0
18	TX_FIFO_ZERO	Transmit FIFO empty 0 = Not empty 1 = Transmit FIFO empty	RO	0
19	RX_FIFO_ZERO	Receive FIFO empty 0 = Not empty 1 = Receive FIFO empty	RO	0
20	TX_DAT_RDY	Transmit data available 0 = No data 1 = Data available in transmit FIFO	RO	0

RSI Registers

Table 21-20. RSI_STATUS Register (Continued)

Bit	Name	Function	Type	Default
21	RX_DAT_RDY	Receive data available 0 = No data 1 = Data available in receive FIFO	RO	0
31:22	Reserved	Reserved	RO	0

RSI Status Clear Register (RSI_STATUSCL)

The RSI_STATUSCL register is used to clear the static flags of the RSI_STATUS register. Write a “1” to any of the bits to clear the corresponding flag in the RSI_STATUS register.

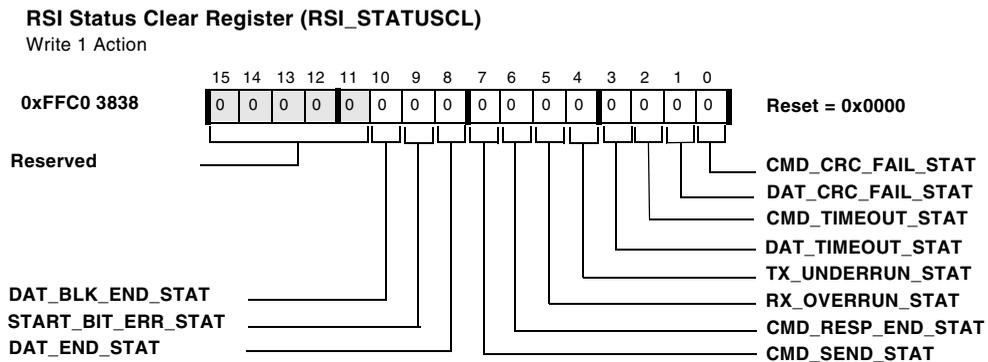


Figure 21-17. RSI Status Clear Register

Table 21-21. RSI_STATUSCL Register

Bit	Name	Function	Type	Default
0	CMD_CRC_FAIL_STAT	Clear command response CRC fail 0 = No effect 1 = Clear CMD_CRC_FAIL	W1A	0
1	DAT_CRC_FAIL_STAT	Clear data CRC failure 0 = No effect 1 = Clear DAT_CRC_FAIL	W1A	0

Table 21-21. RSI_STATUSCL Register (Continued)

Bit	Name	Function	Type	Default
2	CMD_TIMEOUT_STAT	Clear command timeout 0 = No effect 1 = Clear CMD_TIMEOUT	W1A	0
3	DAT_TIMEOUT_STAT	Clear data timeout 0 = No effect 1 = Clear DAT_TIMEOUT	W1A	0
4	TX_UNDERRUN_STAT	Clear transmit FIFO underrun error 0 = No effect 1 = Clear TX_UNDERRUN	W1A	0
5	RX_OVERRUN_STAT	Clear receive FIFO overrun error 0 = No effect 1 = Clear RX_OVERRUN	W1A	0
6	CMD_RESP_END_STAT	Clear command response received 0 = No effect 1 = Clear CMD_RSEP_END	W1A	0
7	CMD_SENT_STAT	Clear command sent 0 = No effect 1 = Clear CMD_SENT	W1A	0
8	DAT_END_STAT	Clear end of data 0 = No effect 1 = Clear DAT_END	W1A	0
9	START_BIT_ERR_STAT	Clear start bit error 0 = No effect 1 = Clear START_BIT_ERR	W1A	0
10	DAT_BLK_END_STAT	Clear data block end 0 = No effect 1 = Clear DAT_BLK_END	W1A	0
15:11	Reserved	Reserved	W1A	0

RSI Interrupt Mask Registers (RSI_MASKx)

The RSI_MASKx registers (RSI_MASK0 and RSI_MASK1) determine which of the static and dynamic flags of the RSI_STATUS register generate an interrupt request to the SIC via one of the two available RSI interrupts.

An interrupt is enabled by setting the corresponding bit in the RSI_MASKx register to 1. Interrupts enabled in the RSI_MASK0 register will result in an IRQ being sent via the IRQ0 signal of the RSI, and interrupts enabled in the RSI_MASK1 register generate an IRQ on the IRQ0 signal of the RSI.

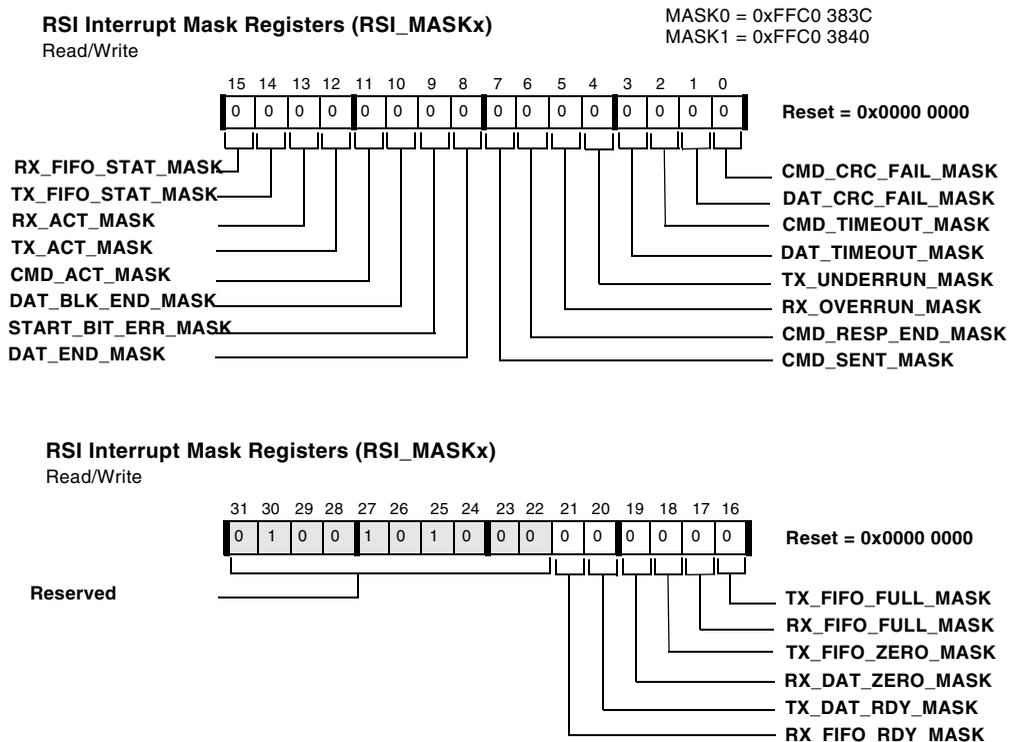


Figure 21-18. RSI Interrupt Mask Registers

Table 21-22. RSI_MASKx Registers

Bit	Name	Function	Type	Default
0	CMD_CRC_FAIL_MASK	Command response CRC fail 0 = Disable interrupt 1 = Enable interrupt	R/W	0
1	DAT_CRC_FAIL_MASK	Data CRC failure 0 = Disable interrupt 1 = Enable interrupt	R/W	0
2	CMD_TIMEOUT_MASK	Command timeout 0 = Disable interrupt 1 = Enable interrupt	R/W	0
3	DAT_TIMEOUT_MASK	Data time out 0 = Disable interrupt 1 = Enable interrupt	R/W	0
4	TX_UNDERRUN_MASK	Transmit FIFO underrun error 0 = Disable interrupt 1 = Enable interrupt	R/W	0
5	RX_OVERRUN_MASK	Receive FIFO overrun error 0 = Disable interrupt 1 = Enable interrupt	R/W	0
6	CMD_RESP_END_MASK	Command response received 0 = Disable interrupt 1 = Enable interrupt	R/W	0
7	CMD_SENT_MASK	Command sent 0 = Disable interrupt 1 = Enable interrupt	R/W	0
8	DAT_END_MASK	End of data 0 = Disable interrupt 1 = Enable interrupt	R/W	0
9	START_BIT_ERR_MASK	Start bit error 0 = Disable interrupt 1 = Enable interrupt	R/W	0
10	DAT_BLK_END_MASK	Data block end 0 = Disable interrupt 1 = Enable interrupt	R/W	0

RSI Registers

Table 21-22. RSI_MASKx Registers (Continued)

Bit	Name	Function	Type	Default
11	CMD_ACT_MASK	Command active 0 = Disable interrupt 1 = Enable interrupt	R/W	0
12	TX_ACT_MASK	Data transmit active 0 = Disable interrupt 1 = Enable interrupt	R/W	0
13	RX_ACT_MASK	Data receive active 0 = Disable interrupt 1 = Enable interrupt	R/W	0
14	TX_FIFO_STAT_MASK	Transmit FIFO watermark 0 = Disable interrupt 1 = Enable interrupt	R/W	0
15	RX_FIFO_STAT_MASK	Receive FIFO watermark 0 = Disable interrupt 1 = Enable interrupt	R/W	0
16	TX_FIFO_FULL_MASK	Transmit FIFO full 0 = Disable interrupt 1 = Enable interrupt	R/W	0
17	RX_FIFO_FULL_MASK	Receive FIFO full 0 = Disable interrupt 1 = Enable interrupt	R/W	0
18	TX_FIFO_ZER/W_MASK	Transmit FIFO empty 0 = Disable interrupt 1 = Enable interrupt	R/W	0
19	RX_FIFO_ZER/W_MASK	Receive FIFO empty 0 = Disable interrupt 1 = Enable interrupt	R/W	0
20	TX_DAT_RDY_MASK	Transmit data available 0 = Disable interrupt 1 = Enable interrupt	R/W	0
21	RX_DATA_RDY_MASK	Receive data available 0 = Disable interrupt 1 = Enable interrupt	R/W	0
31:22	Reserved	Reserved	R/W	0

RSI FIFO Counter Register (RSI_FIFO_CNT)

The `RSI_FIFO_CNT` register contains a value indicating the number of 32-bit words still to be read from or written to the FIFO. `RSI_FIFO_CNT` is loaded from the `RSI_DATA_LGTH` register when the `DATA_EN` bit of the `RSI_DATA_CONTROL` register is set. If the data length is not word-aligned (multiple of 4), the remaining 1 to 3 bytes are regarded as a word.

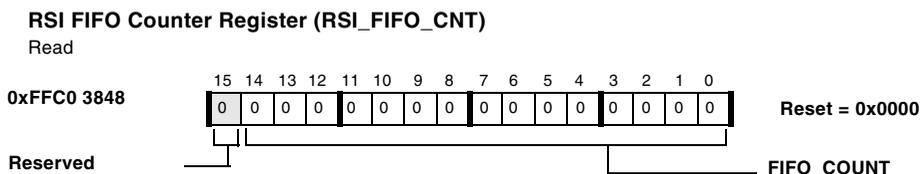


Figure 21-19. RSI FIFO Counter Register

Table 21-23. RSI_FIFO Register

Bit	Name	Function	Type	Default
14:0	FIFO_COUNT	Number of 32-bit words remaining	RO	0
15	Reserved	Reserved	RO	0

RSI CE-ATA Control Register (RSI_CEATA_CONTROL)

The `RSI_CEATA_CONTROL` register contains bits applicable to CE-ATA mode of operation. `CEATA_TX_CCSD`, when set, results in the RSI sending the command completion signal disable sequence to the CE-ATA device to notify the device not to send back the command completion signal. The `CEATA_TX_CCSD` bit is a write-1-action bit and remains set until actively cleared. If the bit is set prior to enabling the command path state machine, the CCSD signal will automatically be sent after the response is received from the CE-ATA device and the command path state machine will return to the IDLE state.

RSI Registers

RSI CE_ATA Control Register (RSI_CEATA_CONTROL)

Read/Write 1 Action/Write

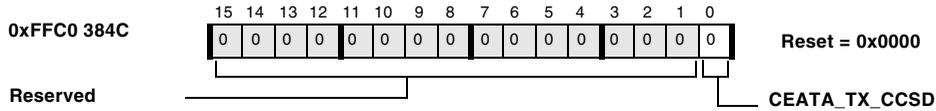


Figure 21-20. RSI CE_ATA Control Register

Table 21-24. RSI_CEATA_CONTROL Register

Bit	Name	Function	Type	Default
0	CEATA_TX_CCSD	Transmit command completion signal disable 0 = No action 1 = Send command completion signal disable sequence	R/W1A/W	0
15	Reserved	Reserved	-	0

RSI Data FIFO Register (RSI_FIFO)

The RSI_FIFO register provides access to the 16-entry transmit and receive FIFO. The register is accessed as a 32-bit word.

RSI Data FIFO Register (RSI_FIFO)

Read/Write

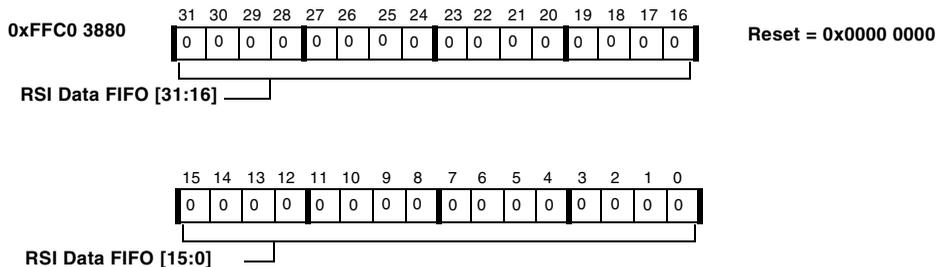


Figure 21-21. RSI Data FIFO Register

RSI Exception Status Register (RSI_ESTAT)

The `RSI_ESTAT` register contains exception status bits for SDIO cards, CE-ATA devices, and the card detection logic. These status bits can be used to generate an interrupt request via the `IRQ0` signal by enabling the interrupt in the `RSI_EMASK` register. All bits in this register are write-1-to-clear bits. The SDIO interrupt is an interrupt generated by SDIO cards on the `RSI_DATA1` signal. The `SD_CARD_DET` bit is set when a rising edge is detected on the `RSI_DATA3` signal and is intended for use with MMC devices that support card detection using this signal.

`CEATA_INT_DET` indicates whether the command completion response has been received from the attached CE-ATA device, indicating that the ATA operation has completed successfully or that ATA command termination has occurred as the result of an error condition.

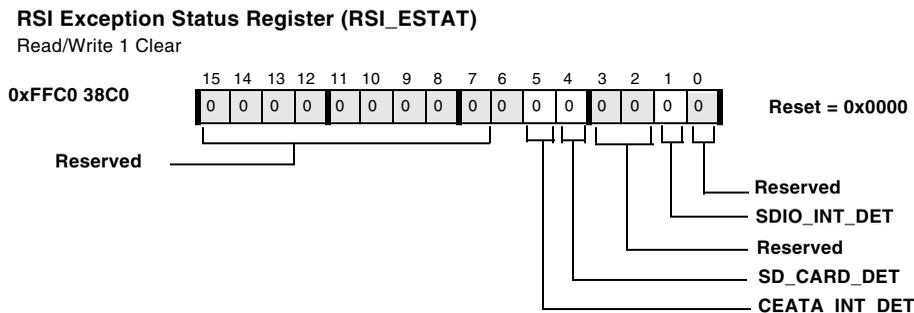


Figure 21-22. RSI Exception Status Register

Table 21-25. RSI_ESTAT Register

Bit	Name	Function	Type	Default
0	Reserved	Reserved	RO	0
1	SDIO_INT_DET	SDIO interrupt detect 0 = No interrupt detected 1 = Interrupt detected	R/W1C	0
3:2	Reserved	Reserved	RO	0

RSI Registers

Table 21-25. RSI_ESTAT Register (Continued)

Bit	Name	Function	Type	Default
4	SD_CARD_DET	Card detect interrupt 0 = No interrupt detected 1 = Interrupt detected	R/W1C	0
5	CEATA_INT_DET	Command completion signal detect 0 = No CCS detected 1 = CCS detected	R/W1C	0
15:6	Reserved	Reserved	RO	0

RSI Exception Mask Register (RSI_EMASK)

The RSI_EMASK register contains mask bits for the RSI_ESTAT status bits. Writing a “1” to the RSI_EMASK bit enables the interrupt for the corresponding bit in the RSI_ESTAT register.

RSI Exception Mask Register (RSI_EMASK)

Read/Write

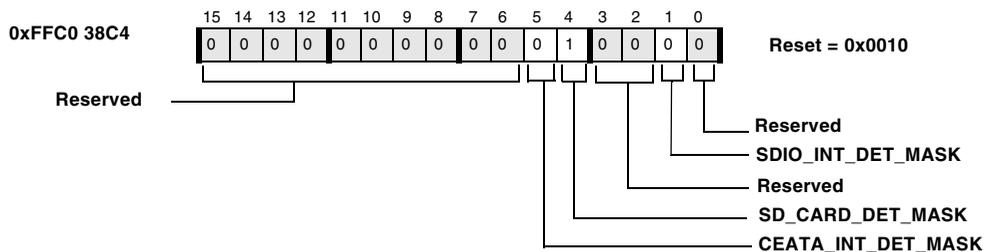


Figure 21-23. RSI Exception Mask Register

Table 21-26. RSI_EMASK Register

Bit	Name	Function	Type	Default
0	Reserved	Reserved	R/W	0
1	SDIO_INT_DET_MASK	SDIO interrupt enable 0 = Interrupt disabled 1 = Interrupt enabled	R/W	0

Table 21-26. RSI_EMASK Register (Continued)

Bit	Name	Function	Type	Default
3:2	Reserved	Reserved	R/W	0
4	SD_CARD_DET_MASK	Card detect interrupt enable 0 = Interrupt disabled 1 = Interrupt enabled	R/W	0
5	CEATA_INT_DET_MASK	Command completion signal detect enable 0 = Interrupt disabled 1 = Interrupt enabled	R/W	0
15:6	Reserved	Reserved	RO	0

RSI Configuration Register (RSI_CONFIG)

The `RSI_CONFIG` register controls bits that enable and disable portions of the RSI module. The `RSI_CLK_EN` bit must be set in order to enable the RSI for operation. After reset, `PD_DAT3` is set. This bit is mutually exclusive with `PU_DAT3`, therefore if the pull-up resistor is to be enabled on `RSI_DATA3` then `PD_DAT3` should be cleared. If an external pull-down resistor is used for implementing card detection on the `RSI_DATA3` signal then `PU_DAT3` should not be set. The pull-up and pull-down resistors on the `RSI_DATAx` signals become active only when the corresponding GPIO pins are configured for RSI functionality via the pin multiplexing. For example, if only the 4-bit data bus is enabled in the pin multiplexing, setting `PU_DAT` will only enable the pull-up resistors on the signals that are configured for RSI use. The `RSI_CONFIG` register also provides additional functionality for SDIO support. To enable SDIO 4-bit mode, in addition to setting the bus width to 4-bit via the `BUS_MODE` field of the `RSI_CLK_CONTROL` register, `SDIO4_EN` should be set. The `MW_EN` bit, when set, allows for SDIO interrupts to be detected outside the specified one-cycle window and is set when interrupt support is required during multiple block read transactions from SDIO. The RSI can also be reset with the `RSI_RST` bit. Writing this bit resets the RSI module and returns all registers to their default values.

RSI Registers

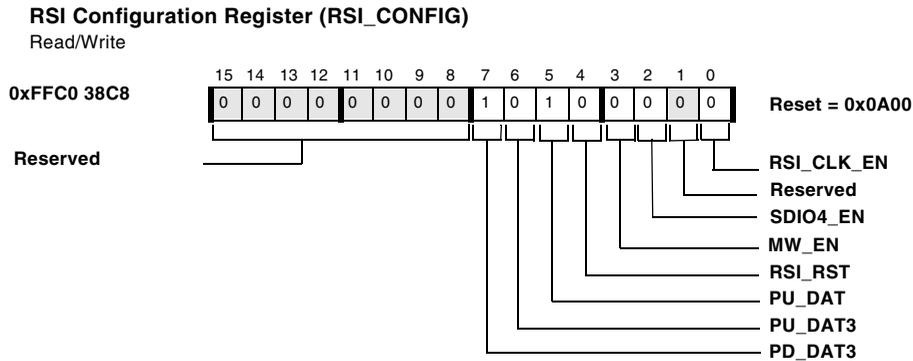


Figure 21-24. RSI Configuration Register

Table 21-27. RSI_CONFIG Register

Bit	Name	Function	Type	Default
0	RSI_CLK_EN	RSI clocks enable 0 = Disable internal RSI clocks 1 = Enable internal RSI clocks	R/W	0
1	Reserved	Reserved	R/W	0
2	SDIO4_EN	SDIO 4-bit enable 0 = Disable SDIO 4-bit mode 1 = Enable SDIO 4-bit mode	R/W	0
3	MW_EN	SDIO interrupt moving window enable 0 = Disabled 1 = Enabled (required when using SDIO multiple block read operations)	R/W	0
4	RSI_RST	RSI reset 0 = No action 1 = Reset the RSI	R/W	0
5	PU_DAT	Pull-up enable 0 = Disable pull-up resistor on RSI_DATA7-4 and RSI_DATA2-0 1 = Enable pull-up resistor on RSI_DATA7-4 and RSI_DATA2-0	R/W	0

Table 21-27. RSI_CONFIG Register (Continued)

Bit	Name	Function	Type	Default
6	PU_DAT3	RSI_DATA3 pull-up enable 0 = Disable pull-up resistor on RSI_DATA3 1 = Enable pull-up resistor on RSI_DATA3	R/W	0
7	PD_DAT3	RSI_DATA3 pull-down enable 0 = Disable pull-down resistor on RSI_DATA3 1 = Enable pull-down resistor on RSI_DATA3 For added system flexibility, no internal pull-down resistor is present. An external pull-down resistor is required for card detection capability on the RSI_DATA3 signal.	R/W	0
15:8	Reserved	Reserved	RO	0

RSI Read Wait Enable Register (RSI_RD_WAIT_EN)

The RSI_RD_WAIT_EN register contains the SDIO_RWR bit that, when set, issues a read wait request to an SDIO card. Once software is ready to resume the data transfer, this bit must be cleared. The functionality applies to both 1-bit and 4-bit SDIO modes.

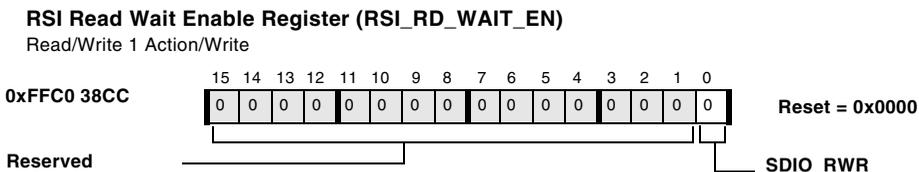


Figure 21-25. RSI Read Wait Enable Register

Table 21-30. Peripheral IDs

RSI Peripheral ID Register	RSI_PID Value
RSI_PID0	0x80
RSI_PID1	0x11
RSI_PID2	0x04
RSI_PID3	0x00

RSI Registers

22 ETHERNET MAC

This chapter describes the Ethernet Media Access Controller (EMAC) peripheral. Following an overview and list of key features is a description of operation and functional modes of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For details regarding the EMAC for the ADSP-BF51x product, please refer to the *ADSP-BF512/BF514/BF516/BF518(F) Embedded Processor Data Sheet*.

For EMAC DMA channel assignments, refer to [Table 6-7 on page 6-104](#) in [Chapter 6, “Direct Memory Access”](#).

For EMAC interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

To determine how the EMAC is multiplexed with other functional pins, refer to [Table 9-2 on page 9-5](#) through [Table 9-4 on page 9-7](#) in [Chapter 9, “General-Purpose Ports”](#).

For a list of MMR addresses for the EMAC, refer to [Appendix A, “System MMR Assignments”](#).

EMAC behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the

Overview

section [“Unique Information for the ADSP-BF51x Processor”](#) on page 22-127.

Overview

The Ethernet MAC provides a 10/100M bit/s Ethernet interface, compliant to IEEE Std. 802.3-2002, between an MII (Media Independent Interface) and the Blackfin peripheral subsystem.

Features

The Ethernet MAC includes these features:

- Independent DMA-driven RX and TX channels
- MII/RMII interface
- 10M bit/s and 100M bit/s operation (full or half duplex)
- VLAN support (full or half duplex)
- Automatic network monitoring statistics
- Flexible address filtering
- Flexible event detection for interrupt handling
- Validation of IP and TCP (payload) checksum
- Remote-wakeup Ethernet frames
- Network-aware system power management

The MAC is fully compliant to IEEE Std. 802.3-2002.

Interface Overview

Figure 22-1 illustrates the overall architecture of the Ethernet controller. The central MAC block implements the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol for both half-duplex and full-duplex modes.

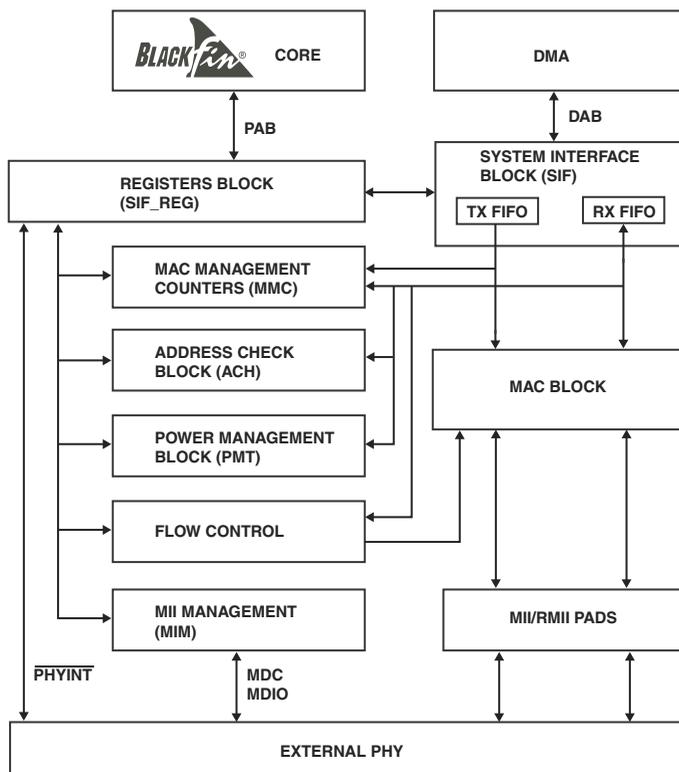


Figure 22-1. Ethernet MAC Block Diagram

The System Interface (SIF) block contains FIFOs for RX and TX data and handles the synchronization of data between the MAC RX and TX data streams and the Blackfin DMA controller.

Interface Overview

The System Interface Registers (SIF_REG) block is an interface from the Blackfin peripheral access bus to the internal registers in the MAC. This block also generates the Ethernet event interrupt, and supports the PHYINT pin by which the PHY can notify the Blackfin processor when the PHY detects changes to the link status, such as auto-negotiation or duplex mode change.

The MAC Management Counters (MMC) block is an extended set of registers that collect various statistics compliant with IEEE 802.3 definitions regarding the operation of the interface. They are updated for each new transmitted or received frame.

The Power Management (PMT) block adds support for wakeup frames and magic packet technology that allows waking up the processor from low power operating modes. Further details regarding these low-power operating modes and voltage regulator wakeup functionality can be found in [Chapter 8, “Dynamic Power Management”](#).

The Address Check (ACH) block checks the destination address field of all incoming packets. Based on the type of address filtering selected, this indicates the result of the address checking to the MAC block.

The MII Management (MIM) block handles all transactions to the control and status registers on the external PHY.

External Interface

The following sections describe the external interface.

Clocking

The Ethernet MAC is clocked internally from SCLK on the processor. A buffered version of CLKIN may be used to drive the external PHY via the CLKBUF pin. See [Figure 22-2](#).

The CLKBUF signal is not generated by a PLL and supports jitter and stability functions comparable to XTAL. The CLKBUF pin is enabled by the CLKBUFOE bit in the VR_CTL register. See [Chapter 8, “Dynamic Power Management”](#) for more information.

A 25 MHz clock (whether driven with the CLKBUF pin or an external crystal) should be used with an MII PHY. A 50 MHz clock source is required to drive an RMII PHY.

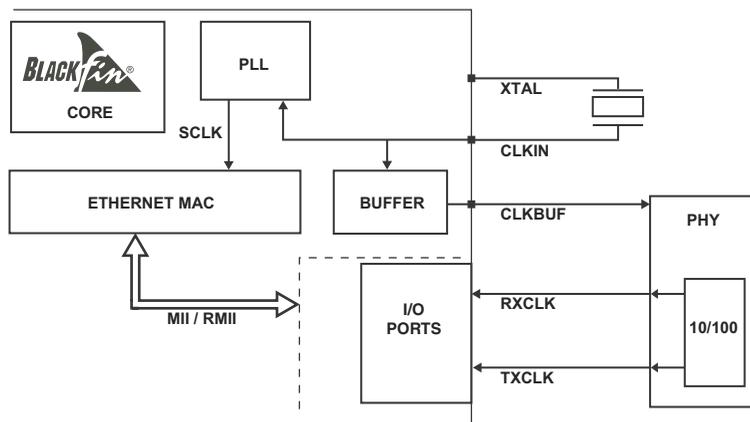


Figure 22-2. Clock Function Diagram

Pins

MII and RMII peripherals are multiplexed into the general-purpose ports. To use MII and RMII operations, set the appropriate I/O Port registers accordingly. See [Chapter 9, “General-Purpose Ports”](#) for more information.

- i IEEE802.3-2002, section two, clause 22.2.1.6, characterizes the MII TX_ER pin as an option useful only for certain applications (for example, repeater applications). Therefore, the TX_ER pin is not present in this design.

Interface Overview

Table 22-1 shows the pins for the MAC.

Table 22-1. Ethernet MAC Pins

MII Signal Name	MII Input/Output	RMII Multiplexed Name	RMII Input/Output	Description
MII CRS	I	RMII CRS_DV	I	Ethernet MII carrier sense/RMII carrier sense and receive data valid
MII RXER	I	RMII RXER	I	Ethernet MII or RMII receive error
MDIO	I/O	MDIO	I/O	Ethernet management channel serial data
MII TXEN	O	RMII TXEN	O	Ethernet MII or RMII transmit enable
MII TXCLK	I	RMII REFCLK	I	Ethernet MII transmit clock/RMII reference clock
MII TXD0	O	RMII TXD0	O	Ethernet MII or RMII transmit D0
MII RXD0	I	RMII RXD0	I	Ethernet MII or RMII receive D0
MII TXD1	O	RMII TXD1	O	Ethernet MII or RMII transmit D1
MII RXD1	I	RMII RXD1	I	Ethernet MII or RMII receive D1
MII TXD2	O			Ethernet MII transmit D2
MII RXD2	I			Ethernet MII receive D2
MII TXD3	O			Ethernet MII transmit D3
MII RXD3	I			Ethernet MII receive D3
MII RXCLK	I			Ethernet MII receive clock
MII RXDV	I			Ethernet MII receive data valid
MII COL	I			Ethernet collision
MDC	O	MDC	O	Ethernet management channel clock
MII PHYINT	I/O	RMII MDINT	I	Ethernet MII PHY interrupt/RMII management data interrupt

Internal Interface

Communication between the MAC and the Blackfin processor peripheral subsystem takes place over the peripheral bus and the DMA Access Bus (DAB). The peripheral bus is used by the Blackfin processor core to configure and monitor the peripheral's control and status registers. All data transfers to and from the peripheral are handled by the Blackfin DMA controller and take place via the DAB.

Power Management

The processor provides power management states which allow programming the MAC to wake the processor upon reception of specific Ethernet frames and/or upon selected events detected by the PHY. The MAC itself requires no additional power management intervention; its internal clocks power down automatically when not required. The MAC clocks run in any of these conditions (provided the processor is in the sleep, active, or full on state):

1. Either the receiver or transmitter is enabled (RE or $TE = 1$)
2. During an MII Management transfer (on MDC/MDIO)
3. During a core access to an MAC control/status register
4. While PHY interrupts are enabled in the MAC ($PHYIE$ in the `EMAC_SYSCTL` register is set)

Description of Operation

The following sections describe the operation of the MAC.

Description of Operation

Protocol

The Ethernet MAC complies with IEEE Std. 802.3-2002. The MII management interface is described below.

MII Management Interface

The IEEE 802.3 MII management interface, also known as the MDIO station management interface, allows the Blackfin processor to monitor and control one or more external Ethernet physical-layer transceivers (PHYs). The MII management interface physically consists of a 2-wire serial connection composed of the `MDC` (management data clock) output signal and the `MDIO` (management data input/output) bidirectional data signal. See [Figure 22-3](#) and [Figure 22-4](#).

The MII management logical interface specifies:

- A set of 16-bit device control/status registers within PHYs, including both required registers with standardized bit definitions as well as optional vendor-specified registers
- A 5-bit device addressing scheme which allows the MAC to select one of up to 32 externally-connected PHY devices
- A 5-bit register addressing scheme for selecting the target register within the addressed device
- A transfer frame protocol for 16-bit read and write accesses to PHY registers via the `MDC` and `MDIO` signals under control of the MAC (PHY devices may not directly initiate `MDIO` transfers.)

Standard PHY control and status registers provide device capability status bits (for example, auto-negotiation, duplex modes, 10/100 speeds and protocols), device status bits (for example, auto-negotiation complete, link status, remote fault), and device control bits (for example, reset, speed selection, loopback, and auto-negotiation start).

The transfer frame protocol defines a MDC clock at a nominal period of 400ns, and an MDIO frame up to 64 bits in length. The MDIO frame consists of an optional 32-bit preamble driven by the MAC, 14 control bits driven by the MAC including the opcode and addresses, a 2-bit turn-around sequence, and a 16-bit data transfer driven either by the MAC or the PHY. Note that various PHYs support optional features such as reduced preamble or increased clock rate.

The features supported by the PHY may be determined at powerup by a MDIO read access (at default rates) of device capabilities in PHY status registers.

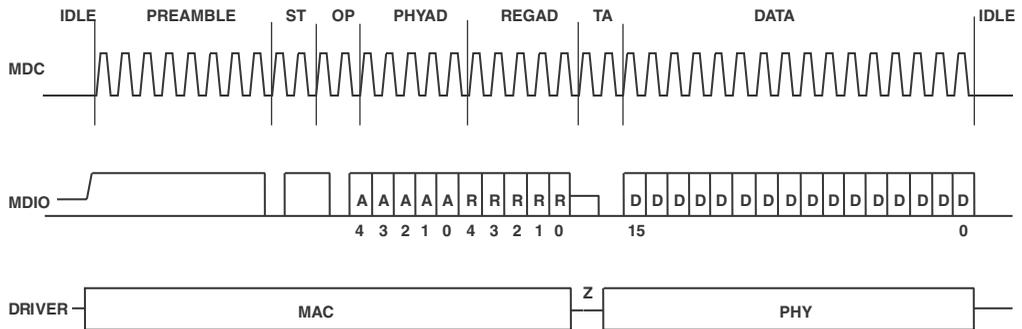


Figure 22-3. Station Management Read

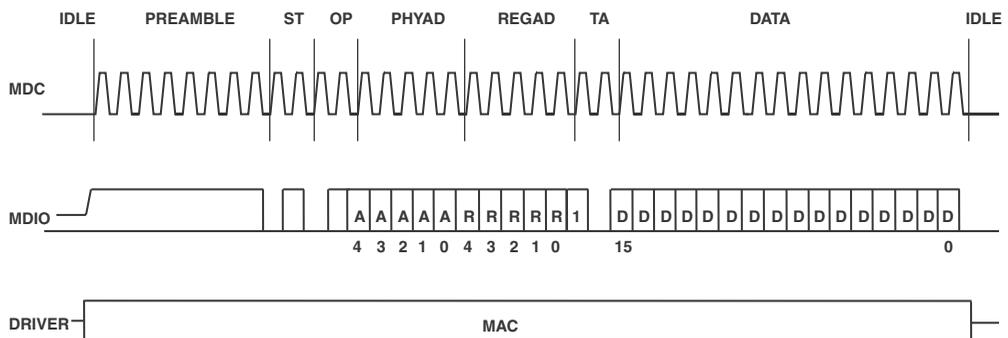


Figure 22-4. Station Management Write

Description of Operation

Operation

The following sections describe the detailed operation of the Ethernet MAC peripheral.

MII Management Interface Operation

The MAC peripheral performs MDIO-protocol transfers in response to register read/write commands issued by the Blackfin processor. Three registers are provided to support MII management transfers:

- The `EMAC_SYSCTL` register contains the `MDCDIV` field which specifies the frequency of the MDC clock output in a ratio to the `SCLK` frequency, and must be initialized before any transfers.
- The `EMAC_STADAT` register holds the 16-bit data for read or write transfers.
- The `EMAC_STAADD` register supports several functions.
 - It commands the access—writes to it may initiate station management transfers, provided the `STABUSY` bit is set and provided that the interface is not already busy.
 - It selects the addressed device, register, and direction of the access.
 - It provides mode controls for MDIO preamble generation and station management transfer done interrupt.
 - It provides the `STABUSY` status bit indicating whether the interface is still busy performing a prior transfer.

As these serial accesses may require significant time (25.6 μ s, or several thousand processor clock cycles at default rates), the Blackfin MAC provides an end-of-transfer interrupt to allow the processor to perform other functions while station management transfers are in progress.

Alternatively, the processor may determine the status of the transfer in progress by reading the `STABUSY` bit in the `EMAC_STAADD` register.

Receive DMA Operation

Data flow between the MAC and the Blackfin peripheral subsystem takes place via bidirectional descriptor based DMA. The element size for any DMA transfer to and from the Ethernet MAC is restricted to 32 bits. In the receive case, a queue or ring of DMA descriptor pairs are normally used, as illustrated in [Figure 22-5](#). In the figure, data descriptors are labeled with an “A” and status descriptors are labeled with a “B.”

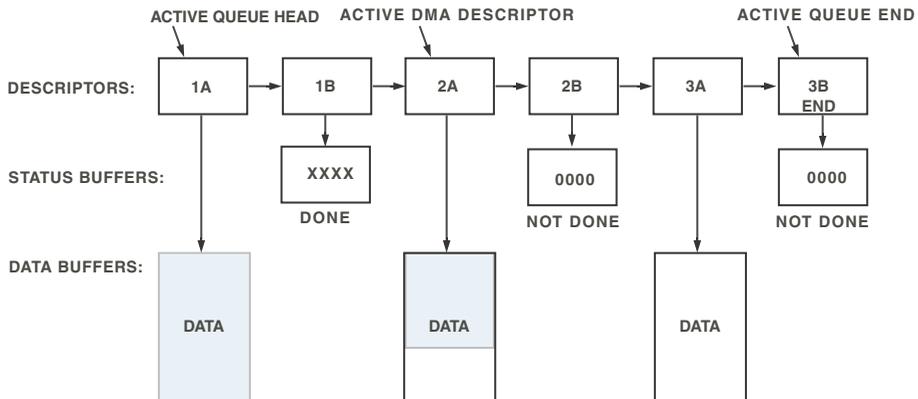


Figure 22-5. Ethernet MAC Receive DMA Operation

Description of Operation

Receive DMA works with a queue or ring of DMA descriptor pairs structured as data and status.

- **Data** – The first descriptor in each pair points to a data buffer that is at least 1556 (0x614) bytes long and is 32-bit aligned. The descriptor `XCOUNT` field should be set to 0, because the MAC controls the actual buffer length.
- **Status** – The second descriptor points to a status buffer of either 4 or 8 bytes. The descriptor `XCOUNT` field should be set to 0, because the MAC controls the actual buffer length. After receiving and accepting any RX frame, the MAC writes a status word and optionally two IP checksum words to this status buffer. The `RXCKS` bit in the `EMAC_SYSCTL` register controls the generation of the two checksum words.

Status words written by the MAC after frame reception have the same format as the `EMAC_RX_STAT` register, and always have the receive complete bit set to 1. If the driver software initializes the length/status words to 0, it can reliably interrogate (poll) an RX frame's length/status word to determine if the DMA transfer of the data buffer is complete. Alternatively, status descriptors may be individually enabled to signal an interrupt when frame reception is complete.

The MAC and DMA operate on the active queue in this manner:

- **Start** – The queue is activated by initializing the DMA next descriptor pointer and then writing the `DMA_CONFIG` register. Meanwhile, the MAC listens to the MII, looking for a frame that passes its address filter.
- **Data** – When a matching frame is seen, the MAC transfers the frame data into the data buffer. The MAC does not initiate the DMA transfer until either the destination address filtering is complete, or the frame ends (if a runt frame).

- **End of frame** – At the end of the frame, the MAC issues a finish command to the DMA controller, causing it to advance to the next (status) descriptor.
- **Status** – The MAC then transfers the frame status into the status buffer. The frame status structure contains the length of the frame data. The MAC then issues another finish command to complete the status DMA buffer.
- **Interrupt** – Upon completion, the DMA may issue an interrupt, if the descriptor was programmed to do so. The DMA then advances to the next (data) descriptor, if any.

Frame Reception and Filtering

Frame data written to memory normally includes the Ethernet header (destination MAC address, source MAC address, and length/type field), the Ethernet payload, and the Frame Check Sequence (FCS) checksum, but not the preamble. If the `RXDWA` bit in `EMAC_SYSCCTL` is 1, then the first 16-bit word is all-zero to pad the frame. The data written includes all complete bytes for which the received data valid (`ERxDV`) pin on the MII interface was asserted after but not including the start of frame delimiter (SFD) nibble (1011). The preamble and any other nibbles prior to the SFD are also not included.

The MAC applies two filtering mechanisms to received frames: the address filter and the frame filter. The address filter considers only the destination MAC address and provides control over the reception of unicast, multicast, and broadcast addresses. The frame filter considers the entire frame and provides control over reception of frames with errors and of MAC control frames.

The address filter is evaluated in the following sequence. Note that this sequence is in the same order as the related bits in the `EMAC_OPMODE`

Description of Operation

register, from LSB to MSB: HU, HM, PAM, PR, IFE, and DBF. The first few filter decisions are additive, while the last two are subtractive.

1. Initially, the address filter is true if the frame's MAC destination address (DA) is either the broadcast address (all 1s) or exactly matches the 48-bit station MAC address in the EMAC_ADDRHI and EMAC_ADDRLO registers.
2. **HU (hash unicast)** – If the HU bit is 1 and the DA is a unicast address which matches the hash table, the address filter is set to true.
3. **HM (hash multicast)** – If the HM bit is 1 and if the DA is a multicast address which matches the hash table, the address filter is set to true.
4. **PAM (pass all multicast)** – If the PAM bit is 1 and the DA is any multicast address, the address filter is set to true.
5. **PR (promiscuous)** – If the PR bit is 1, the address filter is set to true regardless of the frame DA.
6. **FLCE (flow control enable)** – If the FLCE bit in the EMAC_FLC register is 1, and if the DA is an exact match to either the global multicast pause address or to the station MAC address, the address filter is set to true.
7. **IFE (inverse filter)** – If the IFE bit is 1 and the DA exactly matches the 48-bit station MAC address, the address filter is set to false.
8. **DBF (disable broadcast frames)** – If the DBF bit is 1 and the DA is the broadcast address, the address filter is set to false.

The hash table address filtering is configured with the EMAC_HASHLO and EMAC_HASHHI registers described [on page 22-69](#).

The frame filter is evaluated in the following sequence. Note that the frame filter is updated as each byte of data is received. The frame filter can

change from true to false during a frame, for example, upon DMA overrun, but can never change from false back to true.

1. Initially, the frame filter is set to true if the address filter is true, otherwise the frame filter is set to false.
2. **PCF (pass control frames)** – If the PCF bit is 0 and the frame is any valid supported MAC control frame (destination address is either the MAC address or the global multicast pause address; and the length/type field = 88-08, opcode = 0001, length = 64 bytes, and receiveOK = 1), then the frame filter is set to false.
3. **PBF (pass bad frames)** – If the PBF bit is 0 and the frame has any type of error except a frame fragment error, the frame filter is set to false. This rejects any frame for which any of these status bits are set: frame too long, alignment error, frame-CRC error, length error, or unsupported control frame. The frame filter does not reject frames on the basis of the out of range length field status bit. Note that this step may reject MAC control frames passed by PCF.
4. **PSF (pass short frames)** – If the PSF bit is 0 and the frame has a frame fragment error (frame contains less than 64 bytes), the frame filter is set to false. This step may reject frames which were passed by PCF or PBF.
5. **DMA RX overrun** – If the RX DMA FIFO overflows, the frame filter is set to false. If the FIFO overflows at a point where it contains parts of two frames, that is, the last data and status of frame A and the beginning data of frame B, then frame B is rejected by the frame filter and the MAC continues to try to deliver frame A's data and status.

Discarded Frames

Frames that fail the address filter are discarded immediately after the destination address is received, and neither their data nor their status values are written to memory via DMA. Frames that pass the address filter but fail

Description of Operation

the frame filter before 32 bytes are received are also discarded immediately. Once at least 32 bytes of a frame have been received, and if the address and frame filters both pass, the MAC begins to write the frame to memory via DMA RX.

Aborted Frames

Frames that fail the frame filter after 32 bytes have been received are aborted. The MAC issues a restart DMA control command, causing the current RX data DMA descriptor to be reinitialized with its starting address and counts. The aborted frame's status is not written to memory. Instead, the current DMA data and status buffers are recycled for the next RX frame. For all frames that pass both the address and frame filters, both data and status are written to memory via DMA.

Control Frames

If the `FLCE` (flow control enable) bit is set, MAC control frames (with the control type 88-08) whose DAs match either the station MAC address (with inverse filtering disabled) or the global pause multicast address will pass the address filter, and thus may also have status of receiveOK. If the frame also is a supported pause control frame (with length = 64 bytes, and opcode = pause = 00-01, and in full-duplex mode), then the frame filter condition is determined by the `PCF` (pass control frames) bit. If the frame is not also a supported pause control frame, then it is in error, and its frame filter condition depends on the `PBF` (pass bad frames) bit.

Examples

- To perform standard IEEE-802.3 filtering, clear the `EMAC_OPMODE` register bits `HU`, `PR`, `IFE`, `DBF`, `PBF`, and `PSF`. With these selections, the Ethernet MAC accepts error-free broadcast frames and only those error-free unicast frames that exactly match the station MAC address. Set `PAM` to accept all multicast addresses, or set `HM` and program the `EMAC_HASHHI` and `EMAC_HASHLO` registers to accept only a subset of multicast addresses.

- To accept all addresses, set `PR` and clear `IFE` and `DBF` in the `EMAC_OPMODE` register.
- To accept a set of several unicast addresses, set the `HU` bit and set the multicast hash table register bits which correspond to the desired addresses. Note that there is one set of hash table registers that apply to both unicast and multicast addresses, as selected by the `HU` or `HM` bits.
- To reject all addresses, set `IFE` and `DBF`, and clear `HU`, `HM`, `PAM`, and `PR` in the `EMAC_OPMODE` register.

RX Automatic Pad Stripping

If the `ASTP` bit in the `EMAC_OPMODE` register is set, the pad bytes and FCS are stripped from any IEEE-type frame which was lengthened (padded) to reach the minimum Ethernet frame length of 64 bytes. This applies to frames where the Ethernet length/type field is less than 46 bytes, since the Ethernet header and FCS add 18 bytes. When pad stripping occurs, only the first Length/Type + 14 bytes are written to memory via DMA, and the frame length reported in the `EMAC_RX_STAT` register and in the RX status DMA buffer will be Length/Type + 14 rather than the actual number of received bytes.

Pad bytes are never stripped from typed Ethernet frames. Typed Ethernet frames are frames with a length/type field that takes the type interpretation because it is greater than or equal to 0x600 (1536).

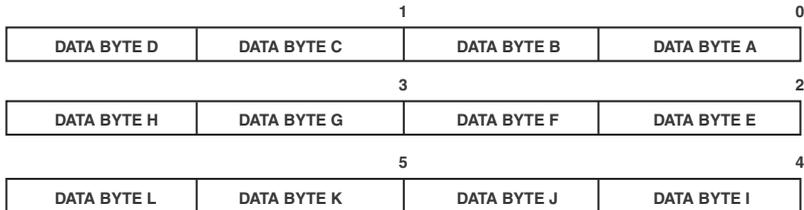
RX DMA Data Alignment

If the `RXDWA` bit in the `EMAC_SYSCCTL` register is clear, the MAC delivers the frame data via DMA to a 32-bit-aligned buffer in memory, including the Ethernet header and FCS. Because the Ethernet header is an odd number of 16-bit words long, this results in the frame payload being odd-aligned, which may be inconvenient for later processing.

Description of Operation

If the `RXDWA` bit is set, however, the MAC prefixes one 16-bit pad word to the frame data with value `0x0000`, resulting in a frame payload aligned on an even 16-bit boundary. See [Figure 22-6](#).

EVEN WORD ALIGNMENT, `RXDWA = 0`



ODD WORD ALIGNMENT, `RXDWA = 1`

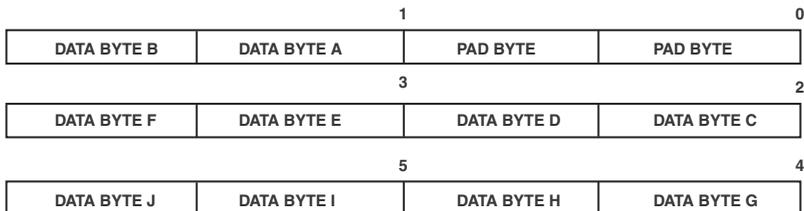


Figure 22-6. RX DMA Data Alignment

RX DMA Buffer Structure

The length of each RX DMA buffer must be at least 1556 (`0x614`) bytes. This is the maximum number of bytes that the MAC can deliver by DMA on any receive frame. Frames longer than the 1556-byte hardware limit are truncated by the MAC. The 1556-byte hardware limit accommodates the longest legal Ethernet frames (1518 bytes for untagged frames, or 1522 bytes for tagged 802.1Q frames) plus a small margin to accommodate future standards extensions.

The MAC does not support RX DMA data buffers composed of more than one descriptor.

RX Frame Status Buffer

The RX frame status buffer is always an integer multiple of 32-bit words in length (either 1 or 2) and must always be aligned on a 32-bit boundary. The RX frame status buffer always contains a frame status word, and may also contain two 16-bit IP checksum words if the `RXCKS` bit in the `EMAC_SYSCCTL` register is set.

To synchronize RX DMA and software, the `RX_COMP` semaphore bit may be used in the RX frame status word. This word is always the last word written via DMA in both status buffer formats, so a transition from 0 to 1 as seen by the processor always means that both the RX data and the status buffers are entirely valid.

[Table 22-2](#) and [Table 22-3](#) describe each of the status buffer formats.

Table 22-2. Receive Status DMA Buffer Format (Without IP Checksum)

Offset	Size	Description
0	32	RX frame status (Same format as the <code>EMAC_RX_STAT</code> register)

Table 22-3. Receive Status DMA Buffer Format (With IP Checksum)

Offset	Size	Description
0	16	IP header checksum
2	16	IP payload checksum
4	32	RX frame status (Same format as the <code>EMAC_RX_STAT</code> register)

RX Frame Status Classification

The RX frame status buffer and the `EMAC_RX_STAT` register provide a convenient classification of each received frame, representing the IEEE-802.3 “receive status” code. The bit layout in the RX frame status buffer is identical to that in the `EMAC_RX_STAT` register, and is arranged so that exactly one status bit is asserted for each of the possible receive status codes

Description of Operation

defined in IEEE-802.3 section 4.3.2. Note in the case of a frame that does not pass the frame filter, neither the frame data nor the status are delivered by DMA into the RX frame status buffer.

The priority order for determination of the receive status code is shown in [Table 22-4](#).

Table 22-4. RX Receive Status Priority

Priority	Bit	Bit Name	IEEE receive status	Condition
1	20	DMA overrun	Undefined	The frame was not completely delivered by DMA
2	18	Frame fragment	Not received	The frame was less than the minimum 64 bytes and was discarded without reporting any other error
3	19	Address filter failed	Not received	The frame did not pass the address filter
4	14	Frame too long	Frame too long	The frame size was more than the maximum allowable frame size (1518, 1522, or 1538 bytes for normal, VLAN1, or VLAN2 frames)
5	15	Alignment error	Alignment error	The frame did not contain an integer number of bytes, and also failed the CRC check
6	16	Frame CRC error	Frame check error	The frame failed CRC validation, and/or RX_ER was asserted during reception of the frame
7	17	Length error	Length error	The frame's length/type field was < 0x600 but did not match the actual length of the data received
8	13	Receive OK	receiveOK	The frame had none of the above conditions

RX IP Frame Checksum Calculation

The MAC calculates TCP/IP-style “raw” checksums of two useful segments of the frame data. Checksum calculation is enabled when the `RXCKS` bit is set to 1 in the `EMAC_SYSCTL` register.

The two checksum segments correspond to the typical position of the IP header and of the IP payload (see [Table 22-5](#)). The checksums are computed as a 16-bit one’s-complement sum of the selected big-endian data words. In each summand, the most significant byte is stored in `byte[1]` and the least significant byte is stored in `byte[2]`, counting bytes starting at 1. If an odd number of data bytes is to be summed, the final value is stored in the most significant byte and zero is stored in the least significant byte. One’s complement addition can be done in ordinary unsigned integer arithmetic by adding the two numbers, followed by adding the carry-out bit value in at the least significant bit. This gives one’s-complement addition the property of being endian invariant, which makes it possible for software running on Blackfin’s little-endian architecture to adjust the sums without explicit byte swapping. See also *RFC 1624* and its references.

The checksum calculation hardware provides an enormous boost to TCP/IP throughput and bandwidth, but requires checksum corrections in software to properly adapt to the details of each packet protocol. For example, TCP packets require the payload checksum to include a TCP pseudo-header made up of certain fields of the IP header. These fields should be added to the “raw” hardware-generated checksum. Similarly, the Ethernet FCS at the end of the frame should be deducted. These adjustments must be made before the IP checksum can be validated.

Description of Operation

Table 22-5. IP Checksum Byte Ranges

Byte Number	Description	Included in IP Header Checksum?	Included in IP Payload Checksum?
1–14	Standard Ethernet header: dest address, src address, length/type	No	No
15–34	Typical IP header, without IP header options	Yes	No
35–N	IP payload, including Ethernet FCS	No	Yes

RX DMA Direction Errors

The RX DMA channel halts immediately after any transfer that sets the `RXDMAERR` bit in the `EMAC_SYSTAT` register. This bit is set if an RX data or RX status DMA request is granted by the RX DMA channel, but the DMA channel is programmed to transfer in the wrong (memory-read) direction. This could indicate a software problem in managing the RX DMA descriptor queue.

In order to facilitate software debugging, the RX DMA channel guarantees that the last transfer to occur is the one with the direction error. On an error, usually the current frame is corrupted. All later frames are ignored until the error is cleared. Since the MAC may have lost synchronization with the DMA descriptor queue, the RX channel must be disabled in order to clear the error condition.

To clear the error and resume operation, perform these steps:

1. Disable the MAC RX channel (clear the `RE` bit in the `EMAC_OPCODE` register).
2. Disable the DMA channel.
3. Clear the `RXDMAERR` bit in the `EMAC_SYSTAT` register by writing 1 to it.

4. Reconfigure the MAC and the DMA engine as if starting from scratch.
5. Re-enable the DMA channel.
6. Re-enable the MAC RX channel.

Transmit DMA Operation

Figure 22-7 shows the transmit DMA operation.

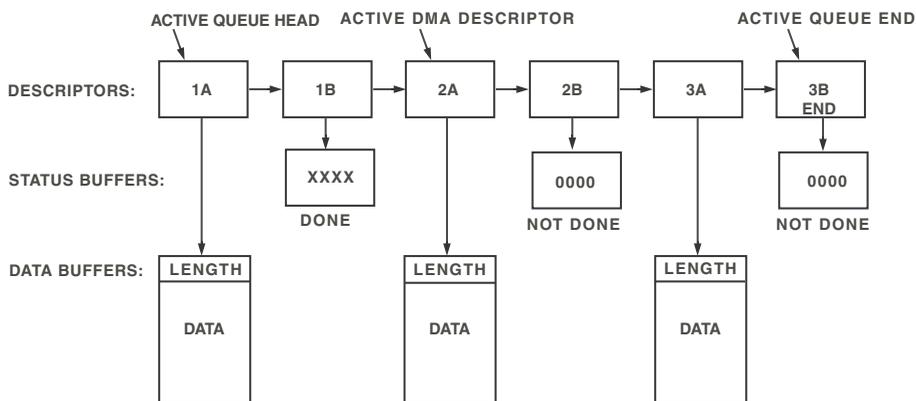


Figure 22-7. Ethernet MAC Transmit DMA Operation

Transmit DMA normally works with a queue or ring of DMA descriptor pairs.

- **Data** – The first descriptor in each pair points to a memory-read data buffer aligned on a 32-bit boundary. The first 16-bit word contains the length in bytes of the frame data, not including the length word or FCS. The descriptor `XCOUNT` field should be set to 0.

Description of Operation

- **Status** – The second descriptor points to a 4-byte status buffer which is written via DMA at the end of the frame. The descriptor `XCOUNT` field should be set to 0, because the MAC controls the termination of the status buffer DMA. The driver software should initialize the status words to zero in advance.

Status words written by the MAC after frame reception have the same format as the `EMAC_TX_STAT` register and always have the transmit complete bit set to 1. Software can therefore interrogate (poll) a TX frame's status word to determine if the transmission of its frame data is complete. Alternatively, status descriptors can be individually enabled to signal an interrupt when frame transmission is complete.

The MAC and DMA operate on the active queue in this manner:

- **Start** – The queue is activated by initializing the DMA `NEXT_DESC_PTR` register and then writing the `DMA_CONFIG` register.
- **Data** – The MAC transfers the frame length word and the first bytes of frame data into its TX data FIFO via DMA. When 32 bytes of data are present in the FIFO, and if the medium is unoccupied, the MAC begins transmission on the MII.
- **Collisions** – The MAC transfers data from memory via DMA into its FIFO, and then from the FIFO over the MII to the PHY. Collisions (in half-duplex mode) can occur at any time in the first 64 bytes of MII transmission, however, the MAC does not discard any of the data in its 96-byte TX FIFO until the first 64 bytes have been successfully transmitted. If a collision occurs during this collision window, and if retry is enabled (`DRTY = 0`), the MAC rewinds its FIFO pointer back to the start of the frame data and begins transmission again. No redundant DMA transfers are performed in such collisions. The MAC makes up to 16 attempts to transmit the frame in response to collisions (if not disabled by `DRTY`), each time backing off and waiting. After the 16th attempt, the frame is

aborted—the MAC terminates data transmission by sending a finish command to the DMA controller, then sending frame status, and then proceeding to the next frame data.

- **Late collisions** – After the collision window is passed, the MAC allows DMA into the FIFO to resume and to overwrite older data. If a collision occurs after the 96th byte has been transferred into the FIFO by DMA (that is, after the FIFO has “wrapped around”), then the MAC issues a restart command to the DMA controller to repeat the DMA of the current descriptor’s data buffer (if enabled by the `LCTRE` bit).
- **End of frame** – At the end of the frame, the MAC issues a finish command to the DMA controller, causing it to advance to the next (status) descriptor. If the TX frame exceeds the maximum length limit (1560 bytes, or `0x618`), the frame’s DMA transfer is truncated. Only 1543 (`0x607`) are transmitted on the MII.
- **Status** – The MAC transfers the frame status into the status buffer.
- **Interrupt** – Upon completion, the DMA may issue an interrupt, if the descriptor was programmed to do so. The DMA then advances to the next (data) descriptor, if any.

Figure 22-8 shows an alternative descriptor structure. The frame length value and Ethernet MAC header are separated from the data payload in each frame.

Description of Operation

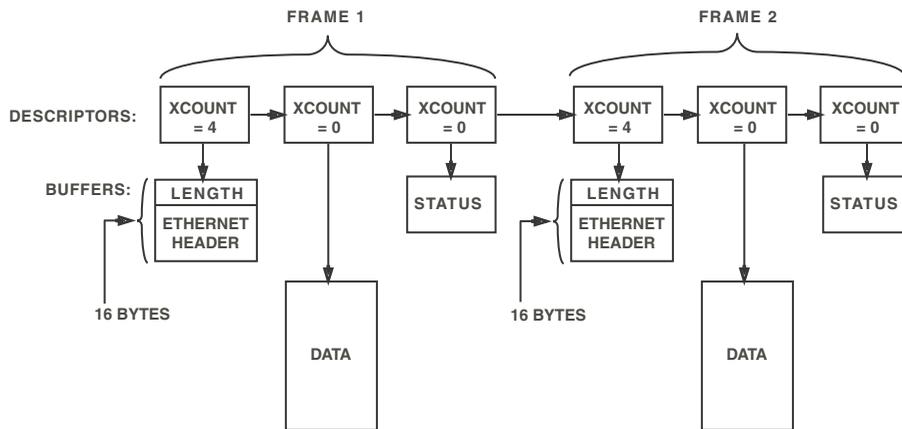


Figure 22-8. Alternative Descriptor Structure

Flexible Descriptor Structure

The Blackfin processor's DMA structure allows flexibility in the arrangement of TX frame data in memory. The frame data can be partitioned into segments, each with a separate DMA descriptor, which allows any of the first 88 bytes of DMA data (86 bytes of frame data) to reside in a separate data segment from the remainder of the frame. This permits the frame length word, the Ethernet MAC header, and even some higher level stack headers to be in one area of memory, while the payload data might be in another. The header and payload may even be in different memory spaces (some internal, some external). Each data buffer segment must be 32-bit aligned. In each frame, the `XCOUNT` field of all but the last data descriptor should be set to the actual length of the data buffers that they reference. As usual, the `XCOUNT` field of the last data descriptor should be set to 0 and the `XCOUNT` field of the status descriptor should be set to 0. The data after the first 88 bytes must all be contained in the data buffer of the last descriptor in the packet.

Multi-descriptor data formatting is not supported if retry is enabled upon late collisions (`LCRTE = 1` in the `EMAC_OPMODE` register). The `LCRTE` bit must be 0 in order to use multiple DMA descriptors for transmit.

TX DMA Data Alignment

The MAC receives TX frame data via DMA from a 32-bit-aligned buffer in memory. If the `TXDWA` bit in the `EMAC_SYSCTL` register is clear, the first word of the MAC frame destination address should immediately follow the TX DMA length word. The MAC frame header starts at an odd word address and the MAC frame payload starts at an even word address.

If the `TXDWA` bit is set, the 16-bit TX DMA length word should be followed by a 16-bit pad word that the MAC ignores. The pad word is transferred over DMA but is not transmitted by the MAC to the PHY. The first word of the MAC frame destination address should immediately follow the pad word. The MAC frame header starts at an even word address and the MAC frame payload starts at an odd word address.

In all cases, the TX DMA length word specifies the number of bytes to be transferred via DMA, excluding the TX DMA length word itself. Specifically, when `TXDWA` is set, the TX DMA length word includes the length of the two pad bytes. See [Figure 22-9](#).

Late Collisions

If a frame's transmission is interrupted (for example, by a late collision) after the transmission of the first 64 bytes, the MAC can be programmed to either automatically retry the frame or to discard the frame. If the `LCRTE` bit in the `EMAC_OPMODE` register is set, the MAC issues a restart command to the TX DMA channel and resets the DMA current address pointer to the start of the current DMA descriptor. This requires the frame data to be entirely contained in a single DMA descriptor.

If the `LCRTE` bit is clear and a late collision is detected, the MAC issues a finish command to the TX DMA controller, advancing the DMA channel to the status descriptor. The MAC then transfers the TX frame status to memory and advances to the next frame descriptor for data.

Description of Operation

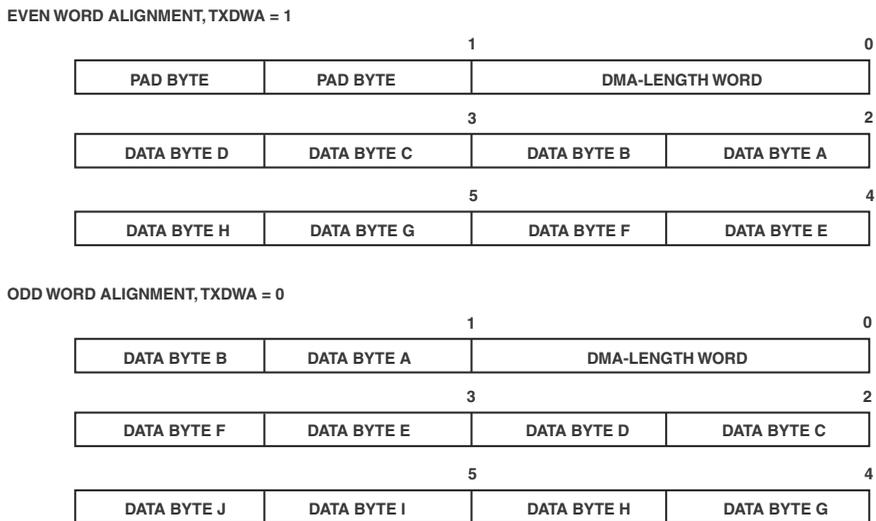


Figure 22-9. TX DMA Data Alignment

TX Frame Status Classification

The TX frame status buffer and the `EMAC_TX_STAT` register provide a convenient classification of each received frame, representing the IEEE-802.3 “transmit status” code. The bit layout in the TX frame status buffer is identical to that in the `EMAC_TX_STAT` register, and is arranged so that exactly one status bit is asserted for each of the possible transmit status codes defined in IEEE-802.3 section 4.3.2.

The priority order for determination of the transmit status code is shown in [Table 22-6](#).

Table 22-6. TX Transmit Status Priority

Priority	Bit	Bit Name	IEEE transmit status	Condition
1	4	DMA underrun	Undefined	The frame was not completely delivered by DMA
2	2	Excessive collision	Excessive collision error	The frame was aborted because of too many (16) collisions, or because of excessive deferral
3	3	Late collision error	Late collision error status	The frame was aborted because of a late collision
4	14, 13	Loss of carrier, no carrier		Carrier sense was deasserted during some or all of the frame transmission (half-duplex only, MII mode only).
5	1	Transmit OK	Transmit OK	The frame had none of the above conditions

TX DMA Direction Errors

The TX DMA channel halts immediately after any transfer that sets the `TXDMAERR` bit in the `EMAC_SYSTAT` register. This bit is set if a TX data or status DMA request is granted by the DMA channel, but the DMA channel is programmed to transfer in the wrong direction. Data DMA should be memory-read; status DMA should be memory-write. TX DMA errors could indicate a software problem in managing the TX DMA descriptor queue.

In order to facilitate software debugging, the TX DMA channel guarantees that the last transfer to occur is the one with the direction error. On an error, usually the current frame is corrupted. Any later frames in the descriptor queue are not sent until the error is cleared. Since the MAC may have lost synchronization with the DMA descriptor queue, the TX channel must be disabled in order to clear the error condition.

Description of Operation

To clear the error and resume operation, perform these steps:

1. Disable the MAC TX channel (clear the `TE` bit in the `EMAC_OPCODE` register).
2. Disable the DMA channel.
3. Clear the `TXDMAERR` bit in the `EMAC_SYSSTAT` register by writing 1 to it.
4. Reconfigure the MAC and the DMA engine as if starting from scratch.
5. Re-enable the DMA channel.
6. Re-enable the MAC TX channel.

Power Management

The Blackfin MAC can be programmed to trigger the following two types of power state transitions:

1. Wake from hibernate

When the processor is in hibernate state (V_{DDINT} powered off) or any higher state, a low level on the `PHYINT` pin can wake the processor to the full on state (via `RESET`). This transition is enabled by setting the `PHYWE` bit to 1 in the `VR_CTL` register prior to powerdown (See “[Dynamic Supply Voltage Control](#)” on page 8-15.)

This pin may be connected to an `INT` output of the external PHY, if applicable. Many PHY devices provide such a pin (sometimes called `MDINT` or `INTR`). PHYs with interrupt capability may be programmed in advance via the MII management interface (MDC/MDIO) to assert the `INT` pin asynchronously upon detecting various conditions. Examples of `INT` conditions include link up, remote fault, link status change, auto-negotiation complete, and duplex and speed status change.

Note that the `PHYINT` pin is general-purpose, and may be driven by any external device or left unused (pulled up to V_{DDEXT}). It is not limited to use with external PHYs.

When the processor is in either the hibernate or deep sleep state, the MAC is powered down. It is not possible to receive or transmit Ethernet frames in these states.

2. Wake from sleep

When the processor is in the sleep state (or any higher state), the Ethernet MAC can remain powered up and can wake the processor to the active or full on states upon signalling an Ethernet event interrupt. The Ethernet event interrupts most useful for power management include:

- Remote wakeup frame received, matching one of four programmable frame filters (see [“Remote Wake-up Filters” on page 22-34](#)).
- Magic Packet™ detected (see [“Magic Packet Detection” on page 22-33](#)).
- Any of the RX or TX frame status interrupts. Examples of these interrupts include: frame received (any frame), Broadcast frame received, VLAN1 frame received, and good frame received (which includes passing the address filters).

For example, the MAC could be programmed to wake the system upon receiving a frame with a particular group destination address, by setting the multicast frame received interrupt enable bit in the `EMAC_RX_IRQE` register and by selecting the appropriate address hash bins in the `EMAC_HASHLO/HI` registers.

Description of Operation

Ethernet Operation in the Sleep State

When the processor is in the sleep state, the Ethernet MAC supports several levels of operation.

- The MAC may be powered down, by clearing `RE` and `TE` in the `EMAC_OPMODE` register. In this lowest-power state, the MAC's internal clocks do not run, and the MAC neither transmits nor responds to received frames. Note that the MAC will not receive a PAUSE control frame in this state.
- The MAC receiver may be partially powered up in a “wake-detect-only” state, but without enabling either the MAC transmitter or MAC DMA. This state is selected by:
 1. Setting `RE` and clearing `TE` in the `EMAC_OPMODE` register.
 2. Setting either the `MPKE` (magic packet wake enable) or `RWKE` (remote wakeup frame enable) bits in the `EMAC_WKUP_CTL` register.
 3. Clearing the capture wakeup frame (`CAPWKFRM`) bit in `EMAC_WKUP_CTL`.

When in the wake-detect-only state, the MAC receiver disables its DMA interface, and does not request any DMA transfers (whether data or status). Instead, the MAC receiver processes good incoming frames through its remote wake-up and/or Magic Packet filters. When a match is detected, the MAC signals a `WAKEDET` interrupt (setting the `WAKEDET` status bit in the `EMAC_SYSSTAT` register). DMA transfers do not resume until the `CAPWKFRM` bit is cleared.

- The MAC receiver may be fully powered up to both receive and/or transmit frames, provided that only external memory (for example, SDRAM) is used. Both the DMA data buffers and descriptor structures must be in external memory, since internal SRAM is unavailable when core clocks are stopped.

This state is intended to be used with very restricted receive-frame filters, so that only certain specific frames are stored via DMA—perhaps only the frame(s) which caused the wakeup event itself. The transmit functionality permits the processor to enqueue a list of final frame transmissions before going to sleep.

The MAC can only transmit frames contained in DMA buffers set up by the processor prior to entering the sleep state. Once the last transmit frame has been sent, the transmitter and DMA channel pauses. Note that if the last TX DMA descriptor was programmed to signal an interrupt, the processor wakes from sleep at the conclusion of that transmission.

Similarly, the MAC can only receive as many frames as can be contained in the DMA buffers and descriptors allocated by the processor prior to entering the sleep state. Once the last receive frame has been filled, the DMA channel pauses, and if any further frames are received (beyond the capacity of the MAC RX FIFO), a DMA overrun occurs. Note that if the last RX DMA descriptor was programmed to signal an interrupt, the processor wakes from sleep after that frame was received.

Magic Packet Detection

The MAC can be programmed to detect a Magic Packet as a wakeup event. This is enabled by setting the `MPKE` bit (Magic Packet enable) bit in the `EMAC_WKUP_CTL` register. When the MAC receives the Magic Packet, it sets the `MPKS` (Magic Packet status) bit in the `EMAC_WKUP_CTL` register, which causes the Ethernet event interrupt to be asserted. The associated ISR should clear the interrupt by writing a 1 to the `MPKS` bit; writing a 0 has no effect.

A Magic Packet is any valid Ethernet frame which contains a specific 102-byte pattern derived from the MAC's 48-bit MAC address anywhere within the frame after the 12th byte (after the destination and source

Description of Operation

address fields). This byte pattern consists of 6 consecutive bytes of 0xFFs followed by sixteen consecutive repeats of the MAC address of the MAC which is targeted for wakeup. See [Figure 22-10](#).

Good Magic Packet frames exclude frame-too-short error, frame-too-long error, FCS error, Alignment error, and PHY error conditions.

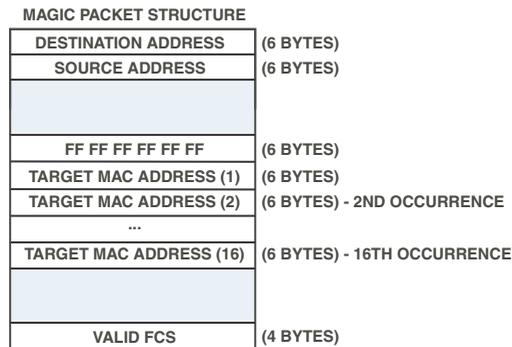


Figure 22-10. Magic Packet Structure

Remote Wake-up Filters

The Blackfin Ethernet MAC provides four independent remote wakeup frame filters for use while in powerdown. See [Figure 22-11](#). These filters are enabled by setting the `RWKE` (remote wakeup enable) bit in the `EMAC_WKUP_CTL` register. Each filter works in parallel, simultaneously examining each incoming frame for a specific byte pattern. Each pattern is described by a byte offset to the start of the pattern within the frame, a 32-bit byte mask selecting bytes at that offset to include in the pattern, and a CRC-16 hash value of the selected bytes which identifies the pattern.

Each of the four filters sets a separate status bit (`RWKS0–RWKS3`) in the `EMAC_WKUP_CTL` register upon detection of their programmed frame pattern. The Ethernet event interrupt is asserted when any of these four status bits is set to 1; the `WAKEDET` bit in the `EMAC_SYSSTAT` register indicates the logical OR of all four of these bits and the `MPKS` (Magic Packet status) bit.

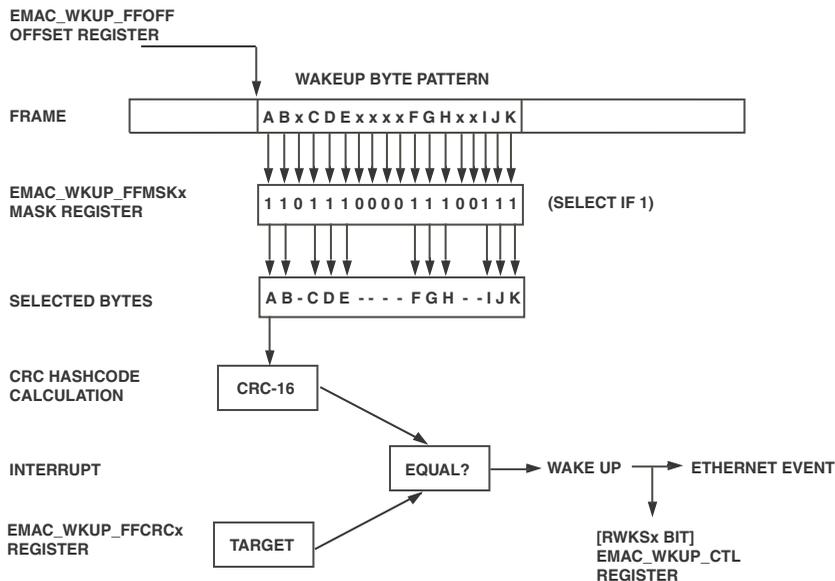


Figure 22-11. Remote Wakeup Filters

The remote wakeup interrupt is cleared by writing a 1 to the appropriate RWKS0–RWKS3 status bit(s). The WAKEDET bit is read-only and does not need to be explicitly cleared.

To program each remote wakeup filter:

1. The RWKE bit in the EMAC_WKUP_CTL register must be set to 1 (enables all four filters.).
2. The enable wakeup filter N bit in the EMAC_WKUP_FF CMD register must be set to 1 to enable filter N.
3. The wakeup filter N address type bit in the EMAC_WKUP_FF CMD register selects whether the target frame is unicast (if 0) or multi-cast (if 1).

Description of Operation

4. The 8-bit pattern offset N field in the `EMAC_WKUP_FF0FF` register selects the starting byte offset for the target data pattern, counting from 0 for the first byte of the MAC frame. The preamble and SFD bytes are not included.
5. The 32-bit `EMAC_WKUP_FFMSKn` register selects which of the 32 bytes starting at the selected offset into the frame will be considered in the pattern match. If the `EMAC_WKUP_FF0FF` register field contains the value K, then bit J of the `EMAC_WKUP_FFMSKn` register set controls whether byte (J + K) of the frame will be compared, counting from 0. A value of 1 in the mask bit enables comparison.
6. The 16-bit wakeup filter N pattern CRC field in the `EMAC_WKUP_FFRCRC0/1` register specifies the 16-bit CRC hash value expected for the wake-up pattern.

Each filter has a separate 16-bit CRC state register which is independently updated as the frame is received. The CRC state for filter N is only updated when an enabled byte is received; the CRC state remains unchanged if the current byte is not enabled by the filter's byte offset and mask registers.

Good frames whose CRC-16 value matches the specified value at the end of the selected pattern window will cause a wake-up event at the end of the frame. Good wake-up frames exclude frame-too-short error, frame-too-long error, alignment error, FCS error, PHY error, and length error conditions.

The CRC-16 hash value for a sequence of bytes may be calculated serially, with each byte processed LSB-first. The initial value of the CRC state is 0xFFFF (all 1s). For each input bit, the LFSR is shifted left one position, and the bit shifted out is XORed with the new input bit. The resulting feedback bit is then XORed into the LFSR at bit positions 15, 2, and 1. Thus the generator polynomial for this CRC is:

$$G(x) = x^{16} + x^{15} + x^2 + 1$$

For example, if the wakeup pattern specified the single byte 0x12, or 0100_1000 (LSB first), the calculation of the wakeup CRC_16 is performed as shown in [Table 22-7](#):

G polynomial = 1000 0000 0000 0101

Table 22-7. CRC-16 Hash Value Calculation

Bit In	XOR	MSB Bit	Feedback Bit	CRC State
				1111 1111 1111 1111, Initial = 0xFFFF
0		1	1	0111 1111 1111 1011
1		0	1	0111 1111 1111 0011
0		0	0	1111 1111 1110 0110
0		1	1	0111 1111 1100 1001
1		0	1	0111 1111 1001 0111
0		0	0	1111 1111 0010 1110
0		1	1	0111 1110 0101 1001
0		0	0	1111 1100 1011 0010, Final = 0xFCB2

Description of Operation

Ethernet Event Interrupts

The Ethernet event interrupt is signalled to indicate that any or all of the conditions listed below are pending. [Figure 22-12](#) shows the Ethernet event interrupts.

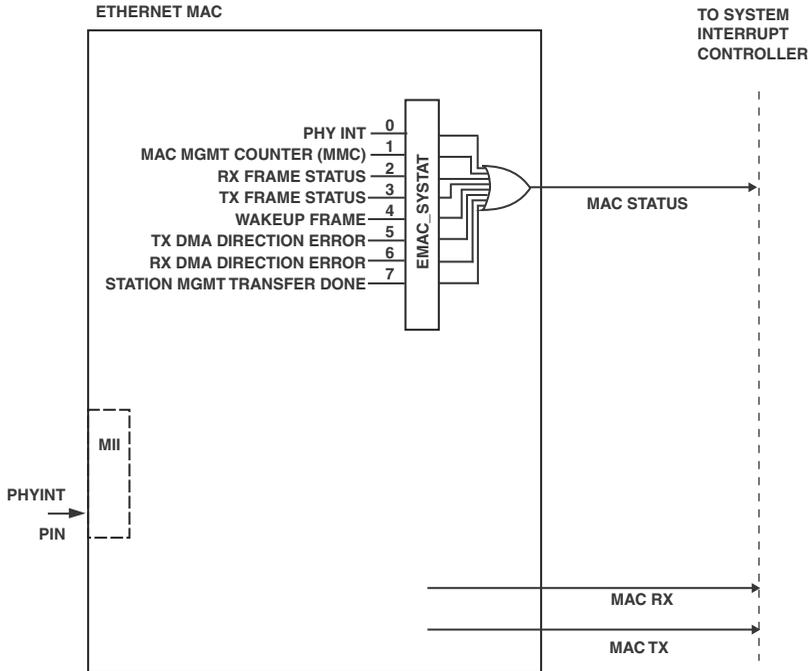


Figure 22-12. Ethernet MAC Event Interrupt

The handler for the peripheral interrupt ID corresponding to the MAC STATUS assignment should interrogate each of the peripherals assigned to that peripheral interrupt ID to determine which peripheral or peripherals are asserting an interrupt. To interrogate the Ethernet MAC, the handler should read the `EMAC_SYSTAT` register, as all of the MAC Ethernet event interrupt condition types are represented in that register.

These conditions result in an Ethernet event interrupt:

- **PHYINT interrupt** – Whenever the asynchronous PHYINT pin is asserted low, the PHYINT sticky bit in the EMAC_SYSTAT register is set to 1. The PHYINT interrupt condition is asserted whenever the logical AND of the PHYINT bit and the PHYIE enable bit in the EMAC_SYSCTL register is 1. This condition is cleared by writing a 1 to the PHYINT bit.
- **MAC management counter (MMC) interrupt** – When any MMC counter reaches half of its maximum value (that is, transitions from 0x7FFF FFFF to 0x8000 0000), the corresponding bit in the MMC RX or TX interrupt status register is set. An MMC interrupt is asserted whenever either:
 - the logical AND of the EMAC_MMC_RIRQS register and the EMAC_MMC_RIRQE register is nonzero, or
 - the logical AND of the EMAC_MMC_TIRQS register and the EMAC_MMC_TIRQE register is nonzero.

The MMC interrupt condition is cleared by writing 1s to all of the MMC RX and/or TX interrupt status register bits which are enabled in the MMC RX/TX interrupt enable register.

- **RX frame status interrupt** – The RX frame status interrupt condition is signalled whenever the logical AND of the EMAC_RX_STKY register and the EMAC_RX_IRQE register is nonzero. This condition is cleared by writing 1s to all of the EMAC_RX_STKY register bits that are enabled in the EMAC_RX_IRQE register.
- **TX frame status interrupt** – The TX frame status interrupt condition is signalled whenever the logical AND of the EMAC_TX_STKY register and the EMAC_TX_IRQE register is nonzero. This condition is cleared by writing 1s to all of the EMAC_TX_STKY register bits that are enabled in the EMAC_TX_IRQE register.

Description of Operation

- **Wakeup frame detected** – This bit is set when a wakeup event is detected by the MAC core (either a magic packet or a remote wakeup packet is accepted by the wakeup filters). This condition is cleared by writing a 1 to the `MPKS` and/or `RWKS` status bits in the `EMAC_WKUP_CTL` register.
- **RX DMA direction error detected** – This bit is set if an RX data or status DMA request is granted by the DMA channel, but the DMA is programmed to transfer in the wrong (memory-read) direction. This could indicate a software problem in managing the RX DMA descriptor queue. This interrupt is non-maskable in the MAC and must always be handled. This condition is cleared by writing a 1 to the `RXDMAERR` bit in the `EMAC_SYSTAT` register.
- **TX DMA direction error detected** – This bit is set if a TX data or status DMA request is granted by the DMA channel, but the DMA is programmed to transfer in the wrong direction. Data DMA should be memory-read, status DMA should be memory-write. This could indicate a software problem in managing the TX DMA descriptor queue. This interrupt is non-maskable in the MAC and must always be handled. This condition is cleared by writing a 1 to the `TXDMAERR` bit in the `EMAC_SYSTAT` register.
- **Station management transfer done** – This bit is set when a station management transfer (on MDC/MDIO) has completed, provided the `STAIE` interrupt enable control bit is set in the `EMAC_STAADD` register.



When the MAC DMA engine is disabled, all the MAC peripheral requests are routed directly into the interrupt controller. This can manifest itself at startup as spurious interrupts. The solution is to configure the system in such a way that the DMA controller is always enabled before the MAC peripheral.

RX/TX Frame Status Interrupt Operation

The contents of the `EMAC_RX_STAT` register indicate the result of the most recent frame receive operation. The register contents are updated just after the end of the frame is received on the MII and synchronized into the system clock domain.

The contents of the `EMAC_RX_STKY` register are updated at the same time. Each applicable bit in the `EMAC_RX_STKY` register is set if the corresponding bit in the `EMAC_RX_STAT` register is set, otherwise the bit in the `EMAC_RX_STKY` register keeps its prior value.

The `EMAC_RX_IRQE` register is continuously bitwise ANDed with the contents of the `EMAC_RX_STKY` register, and then all of the resulting bits are OR'ed together to produce the RX frame status interrupt condition. The state of the RX frame status interrupt condition is readable in the `RXFSINT` bit of the `EMAC_SYSTAT` register. This interrupt condition is cleared by writing 1s to all the bits in the `EMAC_RX_STKY` register for which corresponding bits are set in the `EMAC_RX_IRQE` register. Do not attempt to clear this interrupt condition by writing a 1 to the read only `RXFSINT` bit; such a write has no effect.

The three `EMAC_TX_STAT` registers (`EMAC_TX_STAT`, `EMAC_TX_STKY`, and `EMAC_TX_IRQE`) operate in a similar manner.

RX Frame Status Register Operation at Startup and Shutdown

After the `RE` bit in the `EMAC_OPMODE` register is cleared, the `EMAC_RX_STAT` register, the `EMAC_RX_STKY` register, and the `EMAC_RX_IRQE` register hold their last state. Of course, the two writable registers can still be written.

In order to not confuse status from old and new frames, the `EMAC_RX_STAT` register and the `EMAC_RX_STKY` register are automatically cleared at a 0-to-1 transition of the `RE` bit. The `EMAC_RX_IRQE` register is not cleared when the `RE` bit transitions from 0 to 1. It changes state only when written. All three of these registers are cleared at system reset.

Description of Operation

TX Frame Status Register Operation at Startup and Shutdown

After the `TE` bit in the `EMAC_OPMODE` register is cleared, the `EMAC_TX_STAT` register, the `EMAC_TX_STKY` register, and the `EMAC_TX_IRQE` register hold their last state. Of course, the two writable registers can still be written.

In order to not confuse status from old and new frames, the `EMAC_TX_STAT` register and the `EMAC_TX_STKY` register are automatically cleared at a 0-to-1 transition of the `TE` bit. The `EMAC_TX_IRQE` register is not cleared when the `TE` bit transitions from 0 to 1. It changes state only when written.

All three of these registers are cleared at system reset.

MAC Management Counters

The Blackfin Ethernet MAC provides a comprehensive set of 32-bit read-only MAC management counters, 24 for receive and 23 for transmit, in accordance with the “Layer Management for DTEs” specification in IEEE 802.3 Sec. 30.3. When enabled by setting the `MMCE` bit in the `EMAC_MMC_CTL` register, the counters are updated automatically at the conclusion of each frame. The counters may be read at any time, but may not be written. The counters can be reset to zero all at once by writing the `RSTC` bit to 1.

The counters can be configured to be cleared individually after each read access if the `CCOR` bit is set to 1. This mode guarantees that no counts are dropped between the value returned by the read and the value remaining in the register.



Although this read operation has a side effect, the speculative read operation of the Blackfin core pipeline is properly handled by the MAC. During the time between the speculative read stage and the commit stage of the read instruction, the MMC block freezes the addressed counter so that intervening updates are deferred until the MMR read instruction is resolved.

For best results, to minimize the amount of time that any given MMC counter is frozen, it is suggested not to intentionally place MMC counter read instructions in positions that result in frequent speculative reads which are not ultimately executed. For example, MMC counter reads should not be placed in the shadow of frequently mispredicted flow-of-control operations.



Continuous polling of any MMC register is not recommended. The MMC update process requires at least one `SCLK` cycle between successive reads to the same register, which may not occur if the register read is placed in a tight code loop. If the polling operation excludes the MMC update process, loss of information results.

The overflow behavior of the counters is configurable using the `CROLL` bit. The counters may be configured either to saturate at maximum value (`CROLL = 0`) or to roll over to zero and continue counting (`CROLL = 1`).

The range of the counters can be extended into software-managed counters (for example, 64-bit counters) by use of selectable MMC interrupts. The `EMAC_MMC_RIRQE` and `EMAC_MMC_TIRQE` registers allow the programmer to select which counters should signal an MMC interrupt on the Ethernet event interrupt line when they pass half of the maximum counter value. Even if interrupt latency is large, this mechanism makes it unlikely that any counter data is lost to overrun.

A recommended structure for the ISR for the MMC interrupt would be as follows. In this example, the `CCOR` (clear counter on read) bit is set to 1, and the `CROLL` (counter rollover) bit may also be set to 1.

1. In the ISR, read the SIC to determine which peripheral ID caused the interrupt.
2. If an Ethernet MAC event interrupt is pending, then read the `EMAC_SYSTAT` register. If any of the interrupt bits are set, then an Ethernet event interrupt is pending.

Description of Operation

3. If the MMCINT bit is set, then read the EMAC_MMC_RIRQS and EMAC_MMC_TIRQS registers. Then, for each bit that is set, read the corresponding MMC counter using CCOR (clear counter on read) mode, and add the result to the software-maintained counter.

As an option, if the CROLL bit is set to 1, the ISR can check the count value to see if it is less than 0x8000 0000. This would indicate that the counter has somehow incremented beyond the maximum value (0xFFFF FFFF) and wrapped around to zero while the interrupt awaited servicing. In this case, the software could add an additional 2^{32} to its extended counter to repair the count deficit.

4. Write the interrupt-status values previously read from EMAC_MMC_RIRQS and EMAC_MMC_TIRQS back to those same registers, so that the bits which were 1 cause the corresponding interrupt status bits to be cleared in a write-1-to-clear operation. This guarantees that all the counter interrupts that are cleared are those that correspond to counters that have been read by the interrupt handler. If other counter(s) cross the half-maximum interrupt threshold after the “snapshot” of the EMAC_MMC_RIRQS and EMAC_MMC_TIRQS was taken, then those interrupts are still correctly pending at the RTI; the interrupt handler is then re-entered and the remaining counter interrupts are handled in a second pass.

Programming Model

The following sections describe the Ethernet MAC programming model for a typical system. The initialization sequence can be summarized as follows.

1. Configure MAC MII pins.
 - Multiplexing scheme
 - CLKBUF
2. Configure interrupts.
3. Configure MAC registers.
 - MAC address
 - MII station management
4. Configure PHY.
5. Receive and transmit data through the DMA engine.

Configure MAC Pins

The first step is to configure the hardware interface between the MAC and the external PHY device.

Multiplexing Scheme

The MII interface pins are multiplexed with GPIO pins on the I/O ports. To configure a pin on these ports for Ethernet MAC functionality, see [Chapter 9, “General-Purpose Ports”](#).

CLKBUF

The external PHY chip can be clocked with the buffered clock (CLKBUF) output from the Blackfin processor. In order to enable this clock output, the CLKBUFOE bit in the VR_CTL register must be set. Note that writes to VR_CTL take effect only after the execution of a PLL programming sequence.

Configure Interrupts

Next, the MAC interrupts and MAC DMA interrupts need to be configured properly. Interrupt service routines should be installed to handle all applicable events. Refer to [Figure 22-12 on page 22-38](#) for a graphical representation of how event signals are propagated through the interrupt controller. The status of the MAC interrupts can be sensed with the EMAC_SYSTAT register. However, the process of enabling these interrupts is achieved through a number of different registers.

- The PHYINT interrupt is enabled by setting the PHYIE bit in the EMAC_SYSCTL register.
- The MAC management counter (MMC) interrupt can be enabled through the EMAC_MMC_RIRQE and EMAC_MMC_TIRQE registers.
- The RX frame status and TX frame status interrupts can be enabled through the EMAC_RX_IRQE and EMAC_TX_IRQE registers, respectively.
- The wakeup frame events are controlled through the EMAC_WKUP_CTL register.
- The TX DMA direction error detected and RX DMA direction error detected interrupts are non-maskable. Therefore, an interrupt service routine to handle them should always be installed.
- The station management transfer done interrupt is enabled through the STAIE bit of the EMAC_STAADD register.

The interrupts for the DMA channels corresponding to the Ethernet MAC transfers should be unmasked and a corresponding ISR should be installed if a polling technique is not used.

Configure MAC Registers

After the interrupts are set up correctly, the MAC address registers and the MII protocol must be initialized.

MAC Address

Set the MAC address by writing to the `EMAC_ADDRHI` and `EMAC_ADDRLO` registers. Since the MAC address is a unique number, it is usually stored in a non-volatile memory like a flash device. In this way, every system using the Blackfin MAC peripheral can be easily programmed with a different MAC address during mass production.

MII Station Management

The following procedure should be used to set up the MII communications protocol with the external PHY device.

To perform a station management write transfer:

1. Initialize `MDCDIV` in the `EMAC_SYSCTL` register. The frequency of the MDC clock is $SCLK/[2 * (MDCDIV + 1)]$. Thus $MDCDIV = (SCLK_Freq / MDC_Freq)/2 - 1$.

For example, for a typical 400ns (2.5 MHz) MDC rate at $SCLK = 125$ MHz, set `MDCDIV` to $(125 \text{ MHz}/2.5 \text{ MHz})/2 - 1 = 50/2 - 1 = 24$.

2. Write the data into `EMAC_STADAT`.

Programming Model

3. Write `EMAC_STAADD` with the PHY address, register address, `STAOP = 1`, `STABUSY = 1`, and desired selections for preamble enable and interrupt enable.
4. Do not initiate another read or write access until `STABUSY` reads 0 or until the station management done interrupt (if enabled) has been received. Accesses attempted while `STABUSY = 1` are discarded.

To perform a station management read transfer:

1. Initialize `MDCDIV`.
2. Write `EMAC_STAADD` with the PHY address, register address, `STAOP = 0`, `STABUSY = 1`, and desired selections for preamble enable and interrupt enable.
3. Wait either while polling `STABUSY` or until the station management done interrupt (if enabled) has been received. Note that subsequent accesses attempted while `STABUSY = 1` are discarded. Proceed when `STABUSY` reads 0.
4. Read the data from `EMAC_STADAT`.

Configure PHY

After the MII interface is configured, the PHY can be programmed with the `EMAC_STAADD` and `EMAC_STADAT` registers. Before configuration, the PHY is usually issued a soft reset. Depending on the capabilities of the specific PHY device, the configurable options might include auto-negotiation, link speed, and whether the transfers are full-duplex or half-duplex. The PHY device may also be set up to assert an interrupt on certain conditions, such as a change of the link status.

Receive and Transmit Data

Data transferred over the MAC DMA must be handled with a descriptor-based DMA queue. Refer to [Figure 22-5 on page 22-11](#) and [Figure 22-7 on page 22-23](#) for a graphical representation of a receive queue and transmit queue, respectively.

An Ethernet frame header is placed in front of the payload of each data buffer. The data buffer structure is described in [Table 22-8](#).

Table 22-8. Frame Header

Field	Size in Bytes
Frame size (Tx only)	2
Destination MAC address	6
Source MAC address	6
Length/type	2
Data Payload	Determined by the length/type field

Receiving Data

In order to receive data, memory buffers must be allocated to construct a queue of DMA data and status descriptors. If the `RXDWA` bit in `EMAC_SYSTCL` is 0, then the first item in the receive frame header is the destination MAC address. If the `RXDWA` bit in `EMAC_SYSTCL` is 1, then the first 16-bit word is all-zero to pad the frame, and the second item is the destination MAC address. The DMA engine is then configured through the `DMA_CONFIG` register. After the DMA is set up, the MAC receive functionality is enabled by setting the `RE` bit in `EMAC_OPMODE`. Completion can be signaled by interrupts or by polling the DMA status registers.

Ethernet MAC Register Definitions

Transmitting Data

To transmit data, memory buffers must be allocated to construct a queue of DMA data and status descriptors. The first 16-bit word of the data buffers is written to signify the number of bytes in the frame. The DMA engine is then configured through the `DMA_CONFIG` register. After the DMA is set up, the MAC transmit functionality is enabled by setting the `TE` bit in `EMAC_OPMODE`. Completion can be signaled by interrupts or by polling the DMA status registers.

Ethernet MAC Register Definitions

The MAC register set is broken up into three groups corresponding to the peripheral's major system blocks:

- Control-status register group (MAC block)
- System interface register group (SIF block)
- MAC management counter register group (MMC block)

Most registers require 32-bit accesses, but certain registers have only 16 or fewer functional bits and can be accessed with either 16-bit or 32-bit MMR accesses.

[Table 22-9](#) shows the functions of the MAC registers. MMC counter registers are found in [Table 22-10 on page 22-53](#).

Table 22-9. MAC Register Mapping

Register Name	Function	Notes
Control-Status Register Group		
EMAC_OPMODE	MAC operating mode	Enables the Ethernet MAC transmitter.
EMAC_ADDRLO	MAC address low	Used with EMAC_ADDRHI to set the MAC address.
EMAC_ADDRHI	MAC address high	Used with EMAC_ADDRLO to set the MAC address.
EMAC_HASHLO	MAC multicast hash table low	Used with EMAC_HASHHI to hold the multicast hash table.
EMAC_HASHHI	MAC multicast hash table high	Used with EMAC_HASHLO to hold the multicast hash table.
EMAC_STAADD	MAC station management address	
EMAC_STADAT	MAC station management data	
EMAC_FLC	MAC flow control	
EMAC_VLAN1	MAC VLAN1 tag	
EMAC_VLAN2	MAC VLAN2 tag	
EMAC_WKUP_CTL	MAC wakeup frame control and status	
EMAC_WKUP_FFMSK0	MAC wakeup frame 0 byte mask	
EMAC_WKUP_FFMSK1	MAC wakeup frame 1 byte mask	
EMAC_WKUP_FFMSK2	MAC wakeup frame 2 byte mask	
EMAC_WKUP_FFMSK3	MAC wakeup frame 3 byte mask	
EMAC_WKUP_FFCMD	MAC wakeup frame filter commands	
EMAC_WKUP_FFOFF	MAC wakeup frame filter offsets	
EMAC_WKUP_FFCRC0	MAC wakeup frame filter CRC0/1	
EMAC_WKUP_FFCRC1	MAC wakeup frame filter CRC2/3	

Ethernet MAC Register Definitions

Table 22-9. MAC Register Mapping (Continued)

Register Name	Function	Notes
System Interface Register Group		
EMAC_SYSCTL	MAC system control	
EMAC_SYSTAT	MAC system status	
EMAC_RX_STAT	Ethernet MAC RX current frame status	
EMAC_RX_STKY	Ethernet MAC RX sticky frame status	
EMAC_RX_IRQE	Ethernet MAC RX frame status interrupt enable	
EMAC_TX_STAT	Ethernet MAC TX current frame status	
EMAC_TX_STKY	Ethernet MAC TX sticky frame status	
EMAC_TX_IRQE	Ethernet MAC TX frame status interrupt enable	
EMAC_MMC_RIRQS	Ethernet MAC MMC RX interrupt status	
EMAC_MMC_RIRQE	Ethernet MAC MMC RX interrupt enable	
EMAC_MMC_TIRQS	Ethernet MAC MMC TX interrupt status	
EMAC_MMC_TIRQE	Ethernet MAC MMC TX interrupt enable	
MAC Management Counter Register Group		
EMAC_MMC_CTL	MAC management counters control	For a list of the MMC counter registers, see Table 22-10 .

Table 22-10. MAC Management Counter Registers

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_RXC_OK (FramesReceivedOK) 30.3.1.1.5	Holds a count of frames that are successfully received. This does not include frames received with frame-too-long, FCS, length or alignment errors, or frames lost due to internal MAC sublayer (DMA/FIFO) errors. This also excludes frames with frame-too-short errors, or frames that do not pass the address filter as indicated by the receive frame accepted status bit. Such frames are not considered to be received by the station, and are not considered errors.
EMAC_RXC_FCS (FrameCheckSequenceErrors) 30.3.1.1.6	Holds a count of receive frames that are an integral number of octets in length and do not pass the FCS check. This does not include frames received with frame-too-long or frame-too-short (frame fragment) errors. This also excludes frames with frame-too-short errors, or which do not pass the address filter.
EMAC_RXC_ALIGN (AlignmentErrors) 30.3.1.1.7	Holds a count of frames that are not an integral number of octets in length and do not pass the FCS check. This counter is incremented when the receive status is reported as alignment error. This also excludes frames with frame-too-short errors, or which do not pass the address filter.
EMAC_RXC_OCTET (OctetsReceivedOK) 30.3.1.1.14	Holds a count of data and padding octets in frames that are successfully received. This does not include octets in frames received with frame-too-long, FCS, length or alignment errors, or frames lost due to internal MAC sublayer errors. This also excludes frames with frame-too-short errors, or which do not pass the address filter.
EMAC_RXC_DMAOVF (FramesLostDueToIntMAC RcvError) 30.3.1.1.15	Holds a count of frames that would otherwise be received by the station, but could not be accepted due to an internal MAC sublayer receive error. If this counter is incremented, then none of the other receive counters are incremented. This counts frames truncated during DMA transfer to memory, as indicated by the DMA overrun status bit.
EMAC_RXC_UNICST (UnicastFramesReceivedOK) No IEEE reference	Holds a count of frames counted by the EMAC_RXC_OK register that are not counted by the EMAC_RXC_MULTI or the EMAC_RXC_BROAD register.

Ethernet MAC Register Definitions

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_RXC_MULTI (MulticastFramesReceivedOK) 30.3.1.1.21	Holds a count of frames that are successfully received and are directed to an active non-broadcast group address. This does not include frames received with frame-too-long, FCS, length or alignment errors, or frames lost due to internal MAC sublayer error. This also excludes frames with frame-too-short errors, or that do not pass the address filter.
EMAC_RXC_BROAD (BroadcastFramesReceivedOK) 30.3.1.1.22	Holds a count of frames that are successfully received and are directed to the broadcast group address. This does not include frames received with frame-too-long, FCS, length or alignment errors, or frames lost due to internal MAC sublayer error. This also excludes frames with frame-too-short errors, or that do not pass the address filter.
EMAC_RXC_LNERRI (InRangeLengthErrors) 30.3.1.1.23	Holds a count of frames with a length/type field value between the minimum unpadded MAC client data size and the maximum allowed MAC client data size, inclusive, that does not match the number of MAC client data octets received. The counter also increments when a frame has a length/type field value less than the minimum allowed unpadded MAC client data size and the number of MAC client data octets received is greater than the minimum unpadded MAC client data size. This also excludes frames with frame-too-short errors (less than the minimum unpadded MAC client data size), or that do not pass the address filter.
EMAC_RXC_LNERRO (OutOfRangeLengthField) 30.3.1.1.24	Holds a count of frames with a Length field value greater than the maximum allowed LLC data size. This also excludes frames with frame-too-short errors, or that do not pass the address filter.
EMAC_RXC_LONG (FrameTooLongErrors) 30.3.1.1.25	Holds a count of frames received that exceed the maximum permitted frame size. This counter is incremented when the status of a frame reception is “frame too long.” This also excludes frames with frame-too-short errors, or that do not pass the address filter.

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_RXC_MACCTL (MACControlFramesReceived) 30.3.3.4	Holds a count of MAC control frames passed by the MAC sublayer to the MAC control sublayer. This counter is incremented upon receiving a valid frame with a Length/Type field value equal to 88-08. While the control frame may be received by the Ethernet MAC and yet not be delivered to the MAC client by DMA, depending on the state of the PCF bit, the control frame is still counted by this counter.
EMAC_RXC_OPCODE (UnsupportedOpCodesReceived) 30.3.3.5	Holds a count of MAC control frames received that contain an opcode that is not supported by the device. This counter is incremented when a receive frame function call returns a valid frame with a length/type field value equal to the reserved type, and with an opcode for a function that is not supported by the device. Only opcode 00-01 (pause) is supported by the Ethernet MAC.
EMAC_RXC_PAUSE (PAUSEMACCtrlFramesReceived) 30.3.4.3	Holds a count of MAC control frames passed by the MAC sublayer to the MAC control sublayer. This counter is incremented when a receive frame function call returns a valid frame with both a length/type field value equal to 88-08 and an opcode indicating the pause operation (00-01). This counter does not include or exclude frames on the basis of address, even though pause frames are required to contain the MAC control pause multicast address.
EMAC_RXC_ALLFRM (FramesReceivedAll) No IEEE reference	Holds a count of all frames or frame fragments detected by the Ethernet MAC, regardless of errors and regardless of address, except for DMA overrun frames.
EMAC_RXC_ALLOCT (OctetsReceivedAll) No IEEE reference	Holds a count of all octets in frames or frame fragments detected by the Ethernet MAC, regardless of errors and regardless of address, except for DMA overrun frames.
EMAC_RXC_TYPED (TypedFramesReceived) No IEEE reference	Holds a count of all frames received with a length/type field greater than or equal to 0x600. This does not include frames received with frame-too-long, frame-too-short, FCS, length or alignment errors, frames lost due to internal MAC sublayer error, or that do not pass the address filter.

Ethernet MAC Register Definitions

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_RXC_SHORT (FramesLenLt64Received) No IEEE reference	Holds a count of all frame fragments detected with frame-too-short errors (length < 64 bytes), regardless of address filtering or of any other errors in the frame.
EMAC_RXC_EQ64 (FramesLenEq64Received) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length of exactly 64 bytes.
EMAC_RXC_LT128 (FramesLen65_127Received) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length between 65 and 127 bytes, inclusive.
EMAC_RXC_LT256 (FramesLen128_255Received) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length between 128 and 255 bytes, inclusive.
EMAC_RXC_LT512 (FramesLen256_511Received) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length between 256 and 511 bytes, inclusive.
EMAC_RXC_LT1024 (FramesLen512_1023Received) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length between 512 and 1023 bytes, inclusive.
EMAC_RXC_GE1024 (FramesLen1024_MaxReceived) No IEEE reference	Holds a count of all good frames (with status receiveOK) that have a length greater than or equal to 1024 bytes. This does not include frames with a frame-too-long error.
EMAC_TXC_OK (FramesTransmittedOK) 30.3.1.1.2	Holds a count of frames that are successfully transmitted. This counter is incremented when the transmit status is reported as transmit OK.
EMAC_TXC_1COL (SingleCollisionFrames) 30.3.1.1.3	Holds a count of frames that are involved in a single collision and are subsequently transmitted successfully. This counter is incremented when the result of a transmission is reported as transmit OK and the attempt value is 2.

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_TXC_GT1COL (MultipleCollisionFrames) 30.3.1.1.4	Holds a count of frames that are involved in more than one collision and are subsequently transmitted successfully. This counter is incremented when the transmit status is reported as transmit OK and the value of the attempts variable is greater than 2 and less than or equal to 16.
EMAC_TXC_OCTET (OctetsTransmittedOK) 30.3.1.1.8	Holds a count of data and padding octets in frames that are successfully transmitted. This counter is incremented when the transmit status is reported as transmit OK.
EMAC_TXC_DEFER (FramesWithDeferredXmissions) 30.3.1.1.9	Holds a count of frames whose transmission was delayed on its first attempt because the medium was busy (that is, at the start of frame, CRS is asserted, or was previously asserted within the minimum interframe gap). Frames involved in any collisions are not counted.
EMAC_TXC_LATECL (LateCollisions) 30.3.1.1.10	Holds a count of times that a collision has been detected later than one slot time from the start of the frame transmission. A late collision is counted twice, both as a collision and as a late collision. This counter is incremented when the number of late collisions detected in transmission of any one frame is nonzero.
EMAC_TXC_XS_COL (FramesAbortedDueToXSColls) 30.3.1.1.11	Holds a count of frames that are not transmitted successfully due to excessive collisions. This counter is incremented when the number of attempts equals 16 during a transmission. Note this does not include frames that are successfully transmitted on the last possible attempt.
EMAC_TXC_DMAUND (FramesLostDueToIntMACXmit Error) 30.3.1.1.12	Holds a count of frames that would otherwise be transmitted by the station, but could not be sent due to an internal MAC sublayer transmit error. If this counter is incremented, then none of the other transmit counters are incremented. This counts frames whose transmission is interrupted by incomplete DMA transfer from memory, as indicated by the DMA underrun status bit.
EMAC_TXC_CRSEERR (CarrierSenseErrors) 30.3.1.1.13	Holds a count of the number of times that carrier sense was not asserted or was deasserted during the transmission of a frame without collision.

Ethernet MAC Register Definitions

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_TXC_UNICST (UnicastFramesXmittedOK) No IEEE reference	Holds a count of frames counted by the EMAC_TXC_OK register that are not counted by the EMAC_TXC_MULTI or the EMAC_TXC_BROAD register.
EMAC_TXC_MULTI (MulticastFramesXmittedOK) 30.3.1.1.18	Holds a count of frames that are successfully transmitted to a group destination address other than broadcast.
EMAC_TXC_BROAD (BroadcastFramesXmittedOK) 30.3.1.1.19	Holds a count of frames that are successfully transmitted to the broadcast address as indicated by the transmit status of OK.
EMAC_TXC_XS_DFR (FramesWithExcessiveDeferral) 30.3.1.1.20	Holds a count of frames that deferred for an excessive period of time. This counter can only be incremented once per LLC transmission.
EMAC_TXC_MACCTL (MACControlFramesTransmitted) 30.3.3.3	Holds a count of MAC control frames passed to the MAC sublayer for transmission. Note this counter is incremented only when a MAC pause frame is generated by writing to the EMAC_FLC register. The counter is not incremented for frames transmitted via the normal DMA mechanism which happen to contain valid MAC pause data.
EMAC_TXC_ALLFRM (FramesTransmittedAll) No IEEE reference	Holds a count of all frames whose transmission has been attempted, regardless of success. Each frame is counted only once, regardless of the number of retry attempts.
EMAC_TXC_ALLOCT (OctetsTransmittedAll) No IEEE reference	Holds a count of all octets in all frames whose transmission has been attempted, regardless of success. Each frame's length is counted only once, regardless of the number of retry attempts.
EMAC_TXC_EQ64 (FramesLenEq64Transmitted) No IEEE reference	Holds a count of all frames with status transmit OK that have a length of exactly 64 bytes.
EMAC_TXC_LT128 (FramesLen65_127Transmitted) No IEEE reference	Holds a count of all frames transmitted with status transmit OK that have a length between 65 and 127 bytes, inclusive.

Table 22-10. MAC Management Counter Registers (Continued)

Register Name (IEEE Name) IEEE 802.3 Reference	Description
EMAC_TXC_LT256 (FramesLen128_255Transmitted) No IEEE reference	Holds a count of all frames transmitted with status transmit OK that have a length between 128 and 225 bytes, inclusive.
EMAC_TXC_LT512 (FramesLen256_511Transmitted) No IEEE reference	Holds a count of all frames transmitted with status transmit OK that have a length between 256 and 511 bytes, inclusive.
EMAC_TXC_LT1024 (FramesLen512_1023Transmitted) No IEEE reference	Holds a count of all frames transmitted with status transmit OK that have a length between 512 and 1023 bytes, inclusive.
EMAC_TXC_GE1024 (FramesLen1024_MaxTransmitted) No IEEE reference	Holds a count of all frames transmitted with status transmit OK that have a length greater than or equal to 1024 bytes but not greater than the maximum frame size.
EMAC_TXC_ABORT (TxAbortedFrames) No IEEE reference	Holds a count of all frames attempted that were not successfully transmitted with status of transmit OK.

Control-Status Register Group

This set of registers is used by the application software to configure and monitor the functionality of the MAC block.

Ethernet MAC Register Definitions

MAC Operating Mode (EMAC_OPMODE) Register

The EMAC_OPMODE register, shown in Figure 22-13, controls the address filtering and collision response characteristics of the Ethernet controller in both the RX and TX modes.

MAC Operating Mode Register (EMAC_OPMODE)

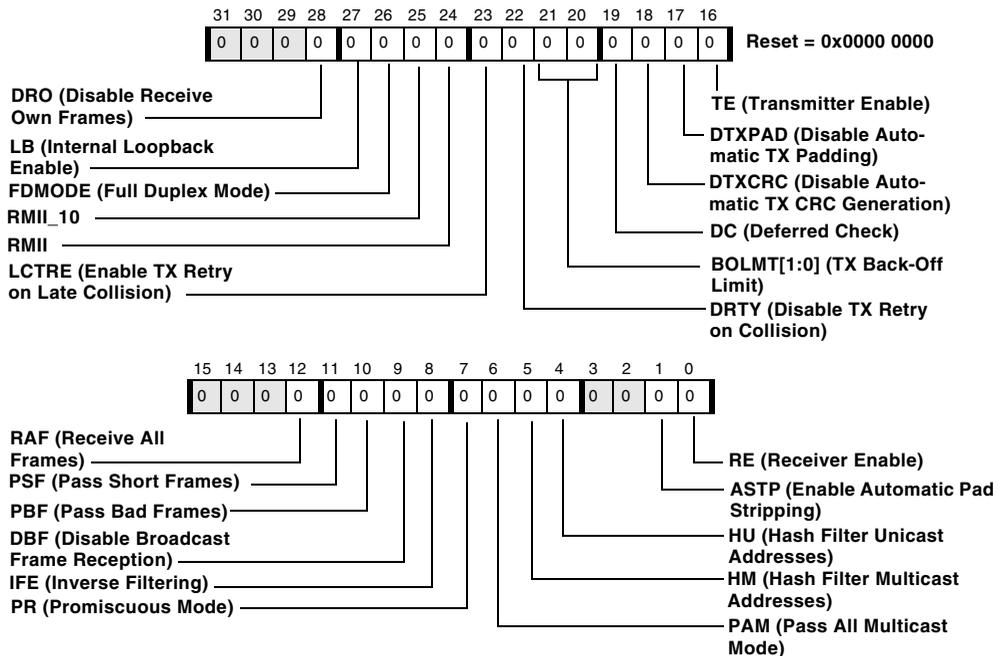


Figure 22-13. EMAC_OPMODE Register

Additional information for the `EMAC_OPMODE` register bits includes:

- **Disable receive own frames** (`DRO`)

When set in half-duplex mode, this bit blocks all frames transmitted by the MAC from being read into the receive path. This bit should be reset when the MAC is operating in full-duplex mode, MII mode only.

0 – Receive own frames enabled

1 – Receive own frames disabled

- **Internal loopback enable** (`LB`)

When internal loopback is enabled, the frames transmitted by the MAC are internally redirected to the receive MAC port. Loopback operation is supported in MII mode; loopback is not supported in RMI mode. During loopback, the external MII port is inactive, the RX pins are ignored and the TX pins are set to `TXEN = 0`, `TXD = 1111`.

0 – Internal loopback not enabled

1 – Internal loopback enabled

- **Full duplex mode** (`FDMODE`)

0 – Half duplex mode selected

1 – Full duplex mode selected

- **RMII port speed selector** (`RMII_10`)

When the interface is configured for RMII operation, software must query the PHY after any automatic negotiation to determine the link speed, and set the RMII port speed selector accordingly. This is because in RMII mode, the `REFCLK` input is always a

Ethernet MAC Register Definitions

constant speed regardless of link speed. In MII mode, by contrast, the PHY decreases the speed of the $RXCLK$ and $TXCLK$ to 2.5 MHz when the link speed is 10M bits.

0 – Speed for RMII port is 100M bits

1 – Speed for RMII port is 10M bits

- **RMII mode** ($RMII$)

This bit is used to select which interface, RMII or MII, is used by the MAC to transfer data to and from the external PHY. Note that MII and RMII modes use slightly different sets of package pins. Program different values into the port configuration register(s) accordingly.

0 – MII mode

1 – RMII mode

- **Enable TX retry on late collision** ($LCTRE$)

0 – TX retry on late collision not enabled

1 – TX retry on late collision enabled

- **Disable TX retry on collision** ($DRTY$)

0 – TX retry on collision not disabled

1 – TX retry on collision disabled

- **TX back-off limit** (BOLMT[1:0])

This field sets an upper bound on the random back-off interval time before the MAC resends a packet in the event of a collision. The bound can be set to 1, 15, 255, or 1023 slot times (1 slot time = 128 MII clock cycles). Thus, varying levels of aggressiveness with regard to packet re-transmission can be selected.

00 – The number of bits is 10 and the maximum back-off time is 1023 slots (relaxed, standard-compliant behavior).

01 – The number of bits is 8 and the maximum back-off time is 255 slots.

10 – The number of bits is 4 and the maximum back-off time is 15 slots.

11 – The number of bits is 1 and the maximum back-off time is 1 slot (aggressive)

- **Deferral check** (DC)

In half-duplex operation, a frame whose transmission defers to incoming traffic for longer than two maximum-length frame times is considered to have been excessively deferred. This time is $(2 \times 1518 \times 2) = 6072$ MII clocks. See IEEE 802.3 section 5.2.4.1 for more information.

0 – The MAC cannot abort transmission of frames due to excessive deferral.

1 – Enables the MAC to abort transmission of frames that encounter excessive deferral.

Ethernet MAC Register Definitions

- **Disable automatic TX CRC generation** (DTXCRC)
 - 0 – Automatic TX CRC generation is enabled. Four CRC bytes are appended to the frame data.
 - 1 – Automatic TX CRC generation is disabled.
- **Disable automatic TX padding** (DTXPAD)
 - 0 – Automatic TX padding is enabled. Pad bytes with value 0 are appended to the data, followed by the CRC, so that the minimum frame size is 64 bytes.
 - 1 – Automatic TX padding of frames shorter than 64 bytes is disabled.
- **Transmitter enable** (TE)

The MAC transmitter is reset when TE is 0. A rising (0 to 1) transition on TE causes the EMAC_TX_STAT register and the EMAC_TX_STKY register to be reset. TE and RE may be enabled either individually or together in either MII or RMII mode.
- **Receive all frames** (RAF)
 - 0 – Does not override filters.
 - 1 – Overrides the address and frame filters and causes all frames or frame fragments to be transferred to memory by DMA.
- **Pass short frame** (PSF)
 - 0 – The frame filter rejects frames with frame-too-short errors (runt frames, or frames with total length less than 64 bytes not including preamble).
 - 1 – Short frames are not rejected by the frame filter.

- **Pass bad frames (PBF)**
 - 0 – The frame filter rejects frames with FCS errors, alignment errors, length errors, frame-too-long errors, and DMA overrun errors.
 - 1 – Pass bad frames enabled.
- **Disable broadcast frame reception (DBF)**
 - 0 – Broadcast frame reception not disabled.
 - 1 – Removes the broadcast address (all 1s) from the set of addresses passed by the address filter, overriding promiscuous mode.
- **Inverse filtering (IFE)**
 - 0 – Inverse filtering not enabled
 - 1 – Removes the MAC address programmed in the `EMAC_ADDRHI` and `EMAC_ADDRLO` registers from the set of addresses passed by the address filter, overriding `PR` (promiscuous) and `HU` (hash unicast) modes. The effect is to block reception of a specific destination address.
- **Promiscuous mode (PR)**
 - 0 – Promiscuous mode not enabled
 - 1 – Promiscuous mode enabled, the address filter accepts all addresses.
- **Pass all multicast mode (PAM)**
 - 0 – Do not pass all multicast frames.
 - 1 – All multicast frames are added to the set of addresses passed by the address filter.

Ethernet MAC Register Definitions

- **Hash filter multicast addresses** (HM)

0 – Does not add multicast addresses that match the hash table to the set of addresses passed by the address filter.

1 – Adds multicast addresses that match the hash table to the set of addresses passed by the address filter.

- **Hash filter unicast addresses** (HU)

0 – Does not add unicast addresses that match the hash table to the set of addresses passed by the address filter.

1 – Adds unicast addresses that match the hash table to the set of addresses passed by the address filter.

- **Automatic pad stripping enable** (ASTP)

A received frame contains pad bytes if it is in IEEE format (the length/type field contains a length value $< 0x600$) and if the length value is less than 46 (corresponding to a frame whose total length including header and FCS is less than 64 bytes). If $ASTP = 1$, both the pad and the FCS bytes are removed from the received data.

0 – Automatic pad stripping is not enabled.

1 – Automatic pad stripping is enabled.

- **Receiver enable** (RE)

The MAC transmitter is reset when RE is 0. A rising (0 to 1) transition on RE causes the EMAC_RX_STAT register and the EMAC_RX_STKY register to be reset. RE and TE may be enabled either individually or together in either MII or RMII mode.

MAC Address Low (EMAC_ADDRLO) Register

The EMAC_ADDRLO register, shown in [Figure 22-14](#), holds the low part of the unique 48-bit station address of the MAC hardware. Writes to this register must be performed while the MAC receive and transmit paths are both disabled. The byte order of address transfer is lowest significant byte first and lowest significant bit first on the MII. Thus EMAC_ADDRLO[3:0] is the first nibble transferred and EMAC_ADDRHI[15:12] is the last nibble.

For example, the address 00:12:34:56:78:9A (where 00 is transferred first and 9A is transferred last) would be programmed as:

EMAC_ADDRLO = 0x56341200

EMAC_ADDRHI = 0x00009A78

MAC Address Low Register (EMAC_ADDRLO)

R/W, except cannot be written if RX or TX is enabled in the EMAC_OPMODE register.

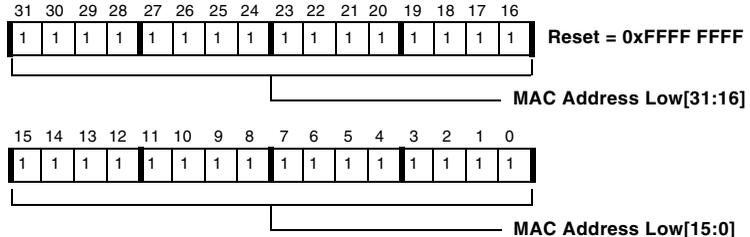


Figure 22-14. EMAC_ADDRLO Register

Ethernet MAC Register Definitions

MAC Address High Register (EMAC_ADDRHI) Register

The EMAC_ADDRHI register, shown in Figure 22-15, holds the high part of the unique 48-bit station address of the MAC hardware. Writes to this register must be performed while the MAC receive and transmit paths are both disabled.

MAC Address High Register (EMAC_ADDRHI)

R/W, except cannot be written if RX or TX is enabled in the EMAC_OPMODE register.

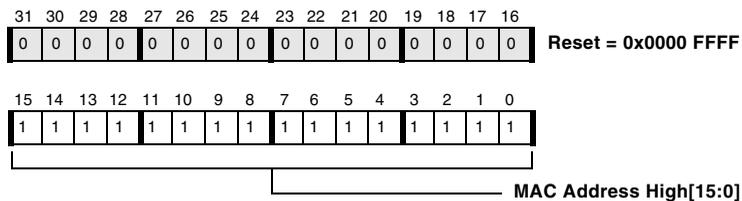


Figure 22-15. EMAC_ADDRHI Register

MAC Multicast Hash Table High (EMAC_HASHHI) and Low (EMAC_HASHLO) Registers

The EMAC_HASHLO register holds the values for bins 31–0 of the multicast hash table. The EMAC_HASHHI register holds the values for bins 63–32 of the multicast hash table. See Figure 22-16 and Figure 22-17.

The 64-bit multicast table is used for multicast frame address filtering. A cyclic redundancy check (CRC) based hash table scheme is used. After the destination address (6th byte) of the frame is received, the state of the CRC-32 checksum unit is sampled. This CRC-32 unit implements the IEEE 802.3 CRC algorithm used in validating the FCS field of the frame. The 6 most significant bits from this state identify one of 64 hash bins representing the frame's destination address. These 6 bits are then used to index into the two hash table registers and extract the corresponding hash bin enable bit. The most significant bit of this value determines the register to be used (high/low) while the other five bits determine the bit position within the register. A CRC value of 000000 selects bit 0 of the

EMAC_HASHLO register and a CRC value of 111111 selects bit 31 of the EMAC_HASHHI register.

MAC Multicast Hash Table Low Register (EMAC_HASHLO)

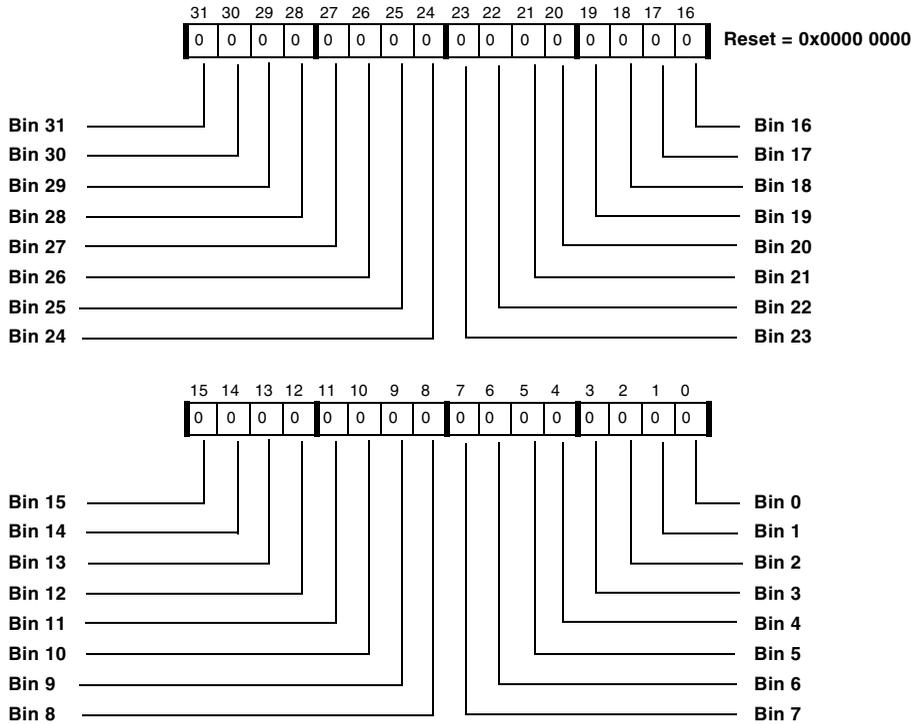


Figure 22-16. EMAC_HASHLO Register

Ethernet MAC Register Definitions

MAC Multicast Hash Table High Register (EMAC_HASHHI)

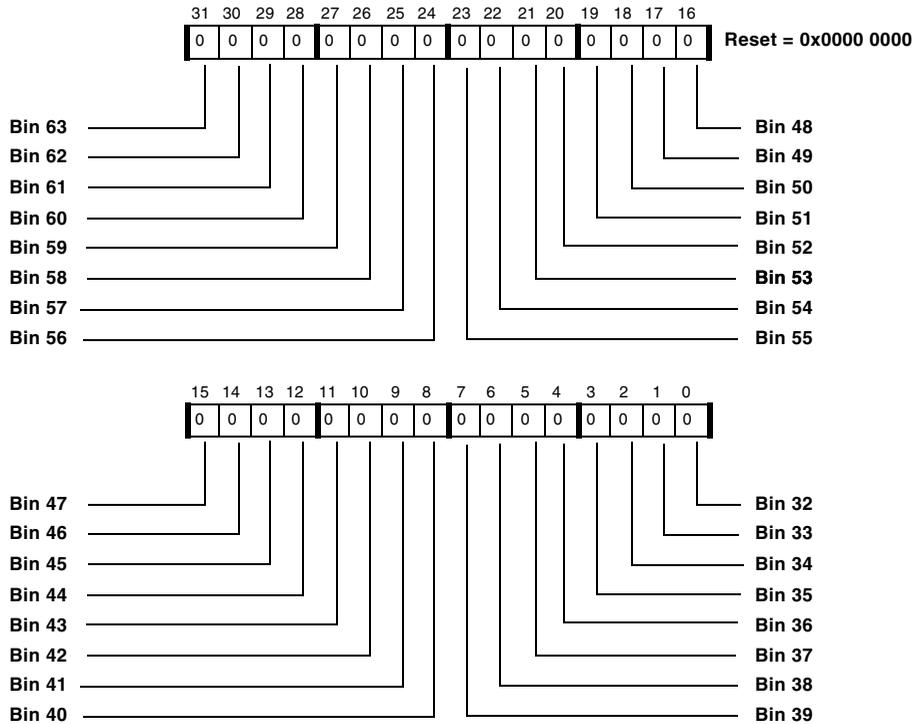


Figure 22-17. EMAC_HASHHI Register

If the corresponding bit in the hash table register is set, the multicast frame is accepted. Otherwise, it is rejected. If the PM bit in the EMAC_OPMODE register is set, all multicast frames are accepted regardless of the hash values.

For example, consider the calculation of the hash bin for the MAC address 01.23.45.67.89.AB. The CRC algorithm uses an LFSR with the prime generator polynomial specified in IEEE 802.3 Sec 3.2.8:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The bits of the MAC address are fed in left-most byte first, least significant bit first, in this sequence (left to right):

1000 0000 1100 0100 1010 0010 1110 0110 1001 0001 1101 0101

The 32-bit CRC register is initialized to all 1s. Then each input bit is processed as follows: first, the register is shifted left one place, shifting in a zero and shifting out the former MSB. The bit just shifted out is XORed with the current input bit, yielding the feedback bit. If this feedback bit is a 1, then the shift register contents are XORed with the generator polynomial value:

0x04C1 1DB7 = 0000 0100 1100 0001 0001 1101 1011 0111

Following this procedure, the CRC-32 for the MAC address is calculated. See [Table 22-11](#).

Table 22-11. CRC-32 Calculation

Bit Number	Input Bit	MSB Bit	Feedback Bit	Next CRC Shift Register
Start				1111 1111 1111 1111 1111 1111 1111 1111
0	1	1	0	1111 1111 1111 1111 1111 1111 1111 1110
1	0	1	1	1111 1011 0011 1110 1110 0010 0100 1011
2	0	1	1	1111 0010 1011 1100 1101 1001 0010 0001
3	0	1	1	1110 0001 1011 1000 1010 1111 1111 0101
4	0	1	1	1100 0111 1011 0000 0100 0010 0101 1101
5	0	1	1	1000 1011 1010 0001 1001 1001 0000 1101
6	0	1	1	0001 0011 1000 0010 0010 1111 1010 1101

Ethernet MAC Register Definitions

Table 22-11. CRC-32 Calculation (Continued)

Bit Number	Input Bit	MSB Bit	Feedback Bit	Next CRC Shift Register
7	0	0	0	0010 0111 0000 0100 0101 1111 0101 1010
...				
46	0	1	1	1101 0011 1001 0111 1111 0100 0100 1001
47	1	1	0	1010 0111 0010 1111 1110 1000 1001 0010

The resulting six MSBs are 101001 = 0x29 = 41 decimal. The hash bin enable bit for this address is then bit 41 – 32 = 9 of the EMAC_HASHHI register.

MAC Station Management Address (EMAC_STAADD) Register

The EMAC_STAADD register, shown in Figure 22-18, controls the transactions between the MII management (MIM) block and the registers on the external PHY. These transactions are used to appropriately configure the PHY and monitor its performance.

MAC Station Management Address Register (EMAC_STAADD)

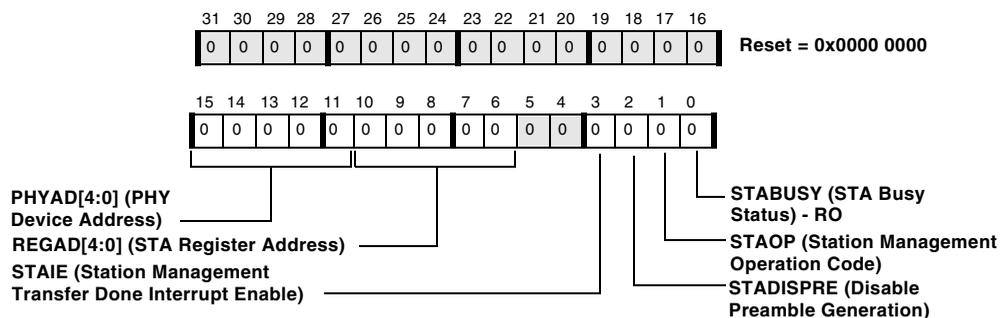


Figure 22-18. EMAC_STAADD Register

Additional information for the `EMAC_STAADD` register bits includes:

- **Station management transfer done interrupt enable** (`STAIE`)
 - 0 – Interrupt not enabled.
 - 1 – Enables an Ethernet event interrupt at the completion of a station management register access (when `STABUSY` changes from 1 to 0).
- **Disable preamble generation** (`STADISPRE`)
 - 0 – Preamble generation for station management transfers not disabled.
 - 1 – Preamble generation (32 ones) for station management transfers disabled.
- **Station management operation code** (`STAOP`)
 - 0 – Read
 - 1 – Write
- **STA busy status** (`STABUSY`)

This bit should be set by the application software in order to initiate a station management register access. This bit is automatically cleared when the access is complete. The MAC ignores new transfer requests made while the serial interface is busy. Writes to the STA address or data registers are discarded if `STABUSY` is 1.

 - 0 – No operation.
 - 1 – Initiate a station management register access across MDC/MDIO.

Ethernet MAC Register Definitions

MAC Station Management Data (EMAC_STADAT) Register

The EMAC_STADAT register, shown in [Figure 22-19](#), contains either the data to be written to the PHY register specified in the EMAC_STAADD register, or the data read from the PHY register whose address is specified in the EMAC_STAADD register.

MAC Station Management Data Register (EMAC_STADAT)

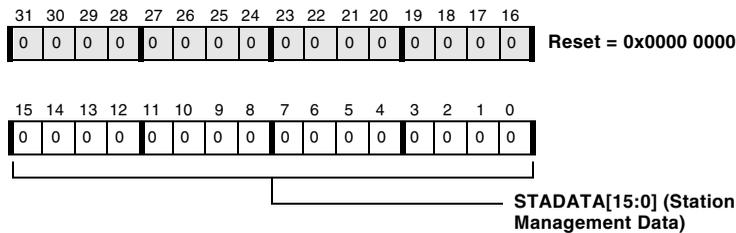


Figure 22-19. EMAC_STADAT Register

MAC Flow Control (EMAC_FLC) Register

The EMAC_FLC register, shown in [Figure 22-20](#), controls the generation and reception of control frames by the MAC. The control frame fields are selected as specified in the IEEE 802.3 specification. When flow control is enabled, the MAC acts upon MAC control pause frames received without errors. When an error-free MAC control pause frame is received (with length/type = MacControl = 88-08 and with opcode = pause = 00-01), the transmitter defers starting new frames for the number of slot times specified by the pause time field in the control frame.

The MAC can also generate and transmit a MAC control pause frame when the EMAC_FLC register is written with FLCBUSY = 1 and FLCPAUSE equal to the number of slot times of deferral being requested.

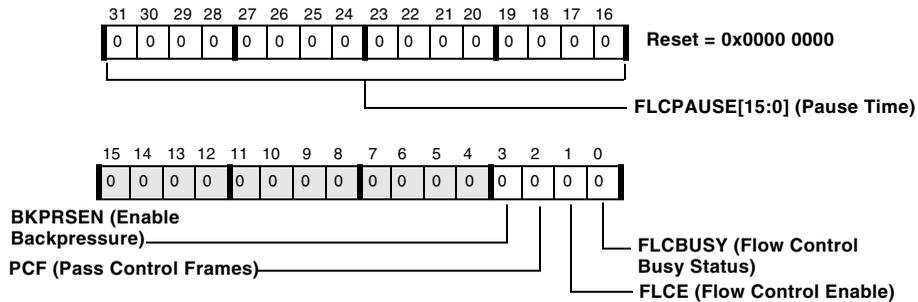
MAC Flow Control Register (EMAC_FLC)

Figure 22-20. EMAC_FLC Register

Additional information for the EMAC_FLC register bits includes:

- **Pause time** (FLCPAUSE)

The number of slot times for which the transmission of new frames is deferred.

- **Enable back pressure** (BKPRSEN)

Available only in half-duplex mode, this bit can be used as a form of flow control.

0 – Transmit and receive function is normal.

1 – Prevents frame reception by colliding with (continuously transmitting a jam pattern during) every incoming frame.

Ethernet MAC Register Definitions

- **Pass control frames (PCF)**

When cleared, the PCF bit causes the frame filter to reject all control frames (frames with length/type field equal to 88-08). When cleared, error-free pause control frames are still interpreted (if enabled by FLCE) but are not delivered via DMA.

0 – Do not pass control frames.

1 – Pass control frames.

- **Flow control enable (FLCE)**

When set, this bit enables interpretation of MAC control pause frames that are received without errors.

0 – Flow control not enabled.

1 – Flow control enabled.

- **FLC busy status (FLCBUSY)**

Setting this bit triggers the MAC to send a control frame. The MAC automatically clears the FLCBUSY bit once the control frame has been transferred onto the physical medium. Writes to the EMAC_FLC register are discarded if FLCBUSY is 1.

0 – No operation.

1 – Initiate sending flow control frame.

MAC VLAN1 Tag (EMAC_VLAN1) and MAC VLAN2 Tag (EMAC_VLAN2) Registers

The EMAC_VLAN1 register, shown in [Figure 22-21](#), and the EMAC_VLAN2 register, shown in [Figure 22-22](#), contain the tag fields used to identify VLAN frames. The MAC compares the 13th and 14th bytes of the incoming frame field to the values contained in these registers, so that the 13th frame byte is compared to the most significant byte of the registers and the

14th frame byte is compared to the least significant byte of the registers. If a match is found, the appropriate bit is set in the `EMAC_RX_STAT` register. The legal length of the frame is then increased from 1518 bytes to either 1522 bytes in the case of a VLAN1 match or 1538 bytes for a VLAN2 match.

MAC VLAN1 Tag Register (`EMAC_VLAN1`)

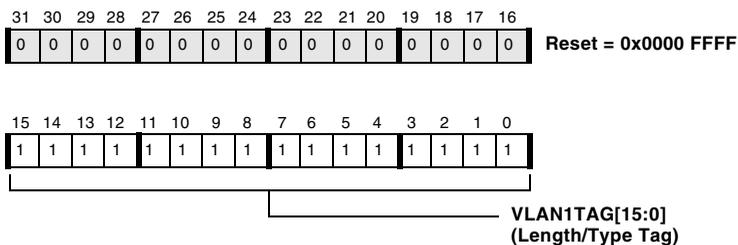


Figure 22-21. `EMAC_VLAN1` Register

MAC VLAN2 Tag Register (`EMAC_VLAN2`)

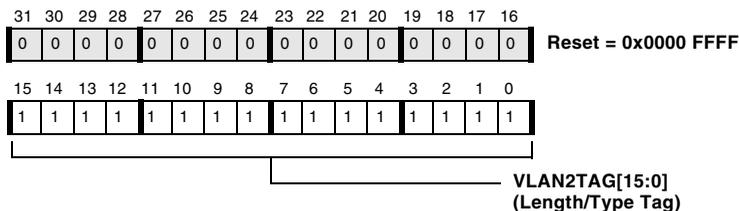


Figure 22-22. `EMAC_VLAN2` Register

MAC Wakeup Frame Control and Status (`EMAC_WKUP_CTL`) Register

The `EMAC_WKUP_CTL` register, shown in [Figure 22-23](#), contains data pertaining to the MAC's remote wakeup status and capabilities. A write to the `EMAC_WKUP_CTL` register causes an update into the receive clock domain

Ethernet MAC Register Definitions

of all the wakeup filter registers. Changes to these other registers do not affect the operation of the MAC until the `EMAC_WKUP_CTL` register is written. For this reason, it is recommended that the wakeup filters be programmed by writing all of the other registers first, and writing the `EMAC_WKUP_CTL` register last.

MAC Wakeup Frame Control and Status Register (`EMAC_WKUP_CTL`)

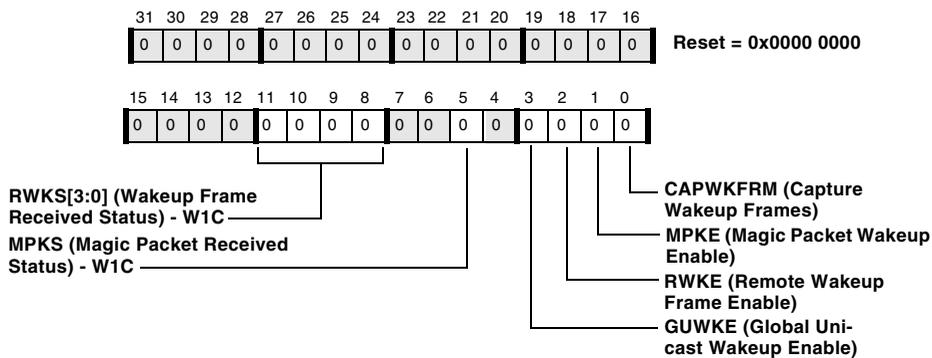


Figure 22-23. `EMAC_WKUP_CTL` Register

Additional information for the `EMAC_WKUP_CTL` register bits includes:

- **Wakeup frame received status (`RWKS`)**

These four frame status bits flag the receipt of wakeup frames corresponding to the respective wakeup frame filters.

- **Magic packet received status (`MPKS`)**

This bit is set by the MAC when it receives the magic packet received wakeup call. The MAC then resumes operation in the normal powered-up mode.

0 – Magic packet not received.

1 – Magic packet received.

- **Global unicast wake enable** (G UWKE)

When set, configures the MAC to wake up from the power-down mode on receipt of a global unicast frame. Such a frame has the MAC address [1:0] bits cleared.

0 – Global unicast wake not enabled.

1 – Global unicast wake enabled.

- **Remote wakeup frame enable** (RWKE)

When set, this bit enables the remote wakeup frame power-down mode.

0 – Remote wakeup frame not enabled.

1 – Remote wakeup frame enabled.

- **Magic packet wakeup enable** (MPKE)

When set, this bit enables the magic packet wakeup power-down mode.

0 – Magic packet wakeup not enabled.

1 – Magic packet wakeup enabled.

- **Capture wakeup frames** (CAPWKFRM)

0 – The RX DMA pathway is disabled when MPKE or RWKE is set.

1 – RX frames are delivered via DMA while in power-down mode (when either MPKE or RWKE is set).

Ethernet MAC Register Definitions

MAC Wakeup Frame0 Byte Mask (EMAC_WKUP_FFMSK0)

MAC Wakeup Frame1 Byte Mask (EMAC_WKUP_FFMSK1)

MAC Wakeup Frame2 Byte Mask (EMAC_WKUP_FFMSK2)

MAC Wakeup Frame3 Byte Mask (EMAC_WKUP_FFMSK3)

Registers

The EMAC_WKUP_FFMSK0, EMAC_WKUP_FFMSK1, EMAC_WKUP_FFMSK2, and EMAC_WKUP_FFMSK3 registers (see [Figure 22-24](#) through [Figure 22-27](#)) are a part of the mechanism used to select which bytes in a received frame are used for CRC computation. Each bit in these registers functions as a byte enable. If a bit *i* is set, then the byte (offset + *i*) is used for CRC computation, where offset is contained in the EMAC_WKUP_FFOFF register.

For example, to identify a wakeup packet containing the byte sequence (0x80, 0x81, 0x82) in bytes 14, 15, and 17, the EMAC_WKUP_FFOFF register should be set to 14 and the byte mask should be set to 0x000B. This byte mask has bits 0, 1, and 3 set, so that bytes 14+0, 14+1, and 14+3 are selected.

MAC Wakeup Frame0 Byte Mask Register (EMAC_WKUP_FFMSK0)

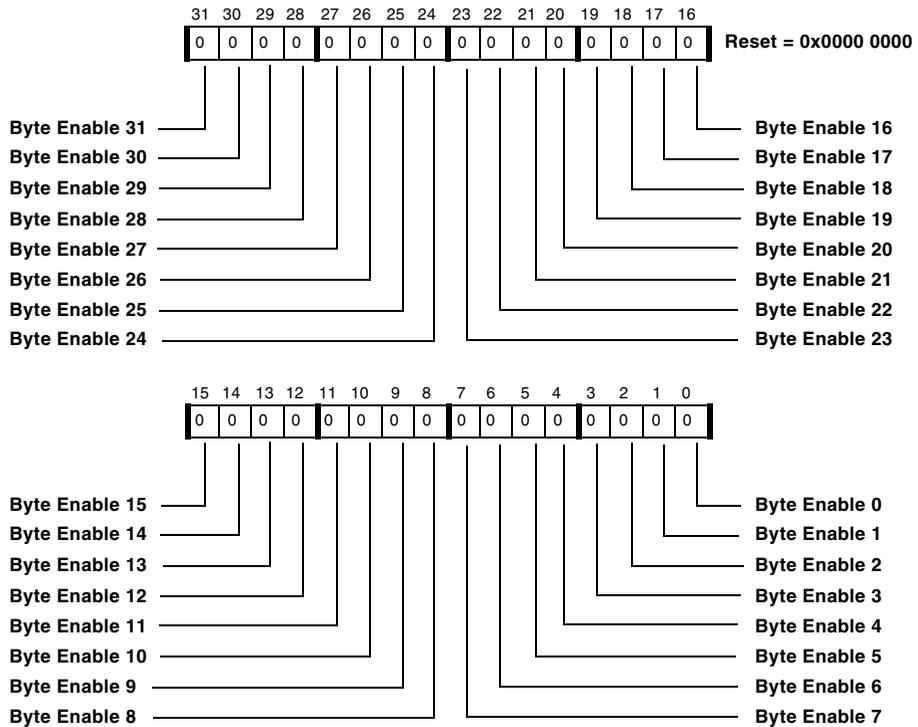


Figure 22-24. EMAC_WKUP_FFMSK0 Register

Ethernet MAC Register Definitions

MAC Wakeup Frame1 Byte Mask Register (EMAC_WKUP_FFMSK1)

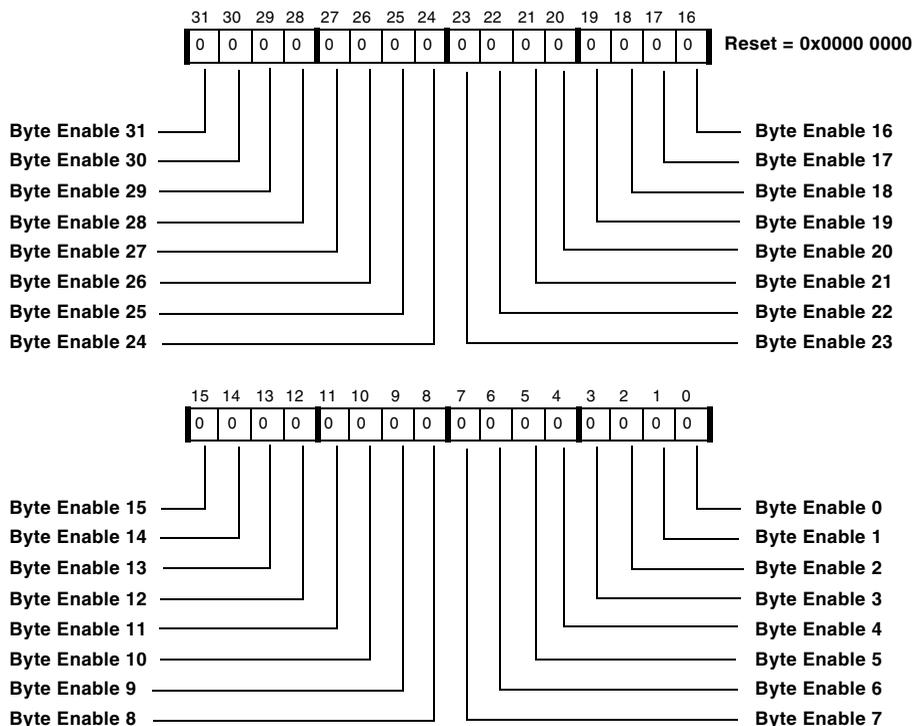


Figure 22-25. EMAC_WKUP_FFMSK1 Register

MAC Wakeup Frame2 Byte Mask Register (EMAC_WKUP_FFMSK2)

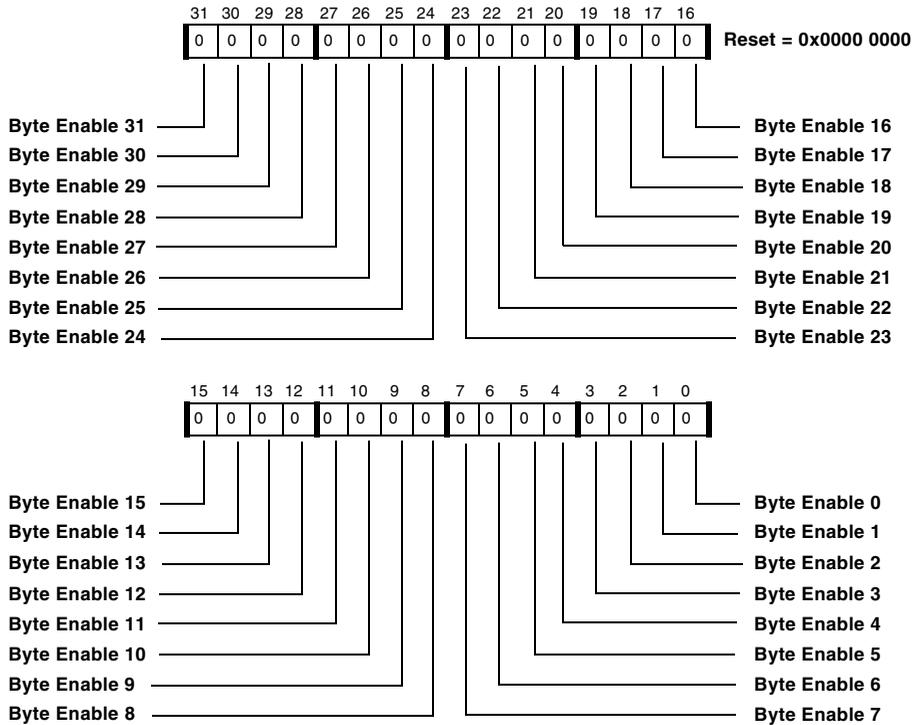


Figure 22-26. EMAC_WKUP_FFMSK2 Register

MAC Wakeup Frame Filter Commands (EMAC_WKUP_FFCMD) Register

The EMAC_WKUP_FFCMD register, shown in Figure 22-28, regulates which of the four frame filter registers are enabled and if so, whether they are configured for unicast or multicast address filtering.

MAC Wakeup Frame Filter Commands Register (EMAC_WKUP_FFCMD)

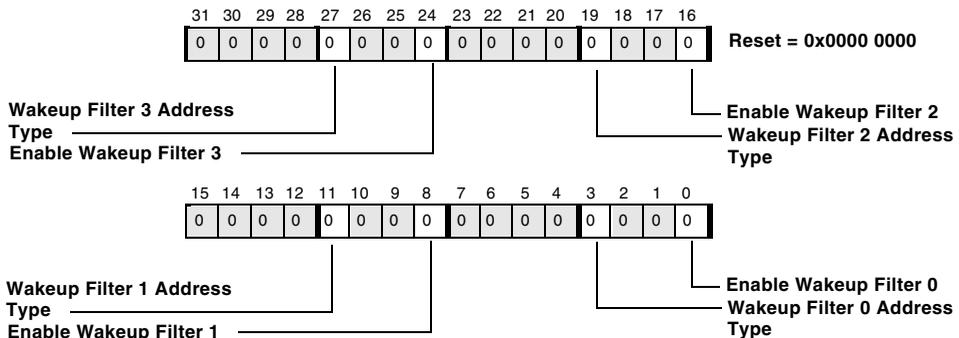


Figure 22-28. EMAC_WKUP_FFCMD Register

Additional information for the EMAC_WKUP_FFCMD register bits includes:

- **Wakeup filter 3 address type**
 - 0 – Unicast
 - 1 – Multicast
- **Enable wakeup filter 3**
 - 0 – Wakeup filter 3 not enabled.
 - 1 – Wakeup filter 3 enabled.

Ethernet MAC Register Definitions

- **Wakeup filter 2 address type**
 - 0 – Unicast
 - 1 – Multicast
- **Enable wakeup filter 2**
 - 0 – Wakeup filter 2 not enabled.
 - 1 – Wakeup filter 2 enabled.
- **Wakeup filter 1 address type**
 - 0 – Unicast
 - 1 – Multicast
- **Enable wakeup filter 1**
 - 0 – Wakeup filter 1 not enabled.
 - 1 – Wakeup filter 1 enabled.
- **Wakeup filter 0 address type**
 - 0 – Unicast
 - 1 – Multicast
- **Enable wakeup filter 0**
 - 0 – Wakeup filter 0 not enabled.
 - 1 – Wakeup filter 0 enabled.

Ethernet MAC Wakeup Frame Filter Offsets (EMAC_WKUP_FFOFF) Register

The EMAC_WKUP_FFOFF register, shown in [Figure 22-29](#), contains the byte offsets for CRC computation to be performed on potential wakeup frames.

Ethernet MAC Wakeup Frame Filter Offsets Register (EMAC_WKUP_FFOFF)

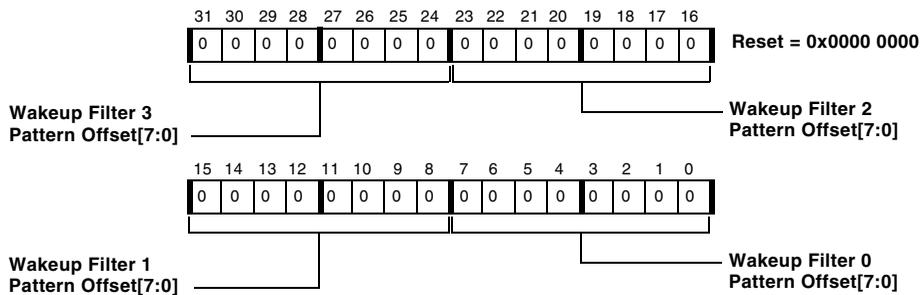


Figure 22-29. EMAC_WKUP_FFOFF Register

MAC Wakeup Frame Filter CRC0/1 (EMAC_WKUP_FFCRC0) and CRC2/3 (EMAC_WKUP_FFCRC1) Registers

The EMAC_WKUP_FFCRC0 register, shown in [Figure 22-30](#), and the EMAC_WKUP_FFCRC1 register, shown in [Figure 22-31](#), should be loaded with the results of the CRC computations for the relevant wakeup frame bytes. See [“Remote Wake-up Filters”](#) on page 22-34.

Ethernet MAC Register Definitions

MAC Wakeup Frame Filter CRC0/1 Register (EMAC_WKUP_FFCRC0)

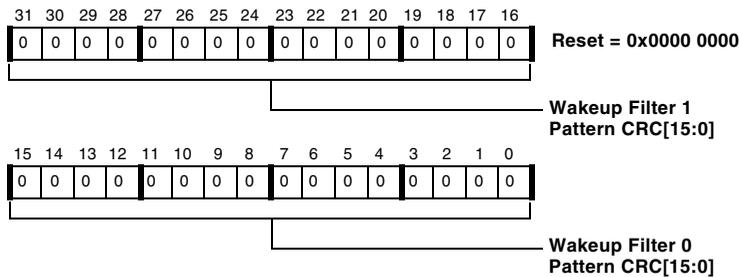


Figure 22-30. EMAC_WKUP_FFCRC0 Register

MAC Wakeup Frame Filter CRC2/3 Register (EMAC_WKUP_FFCRC1)

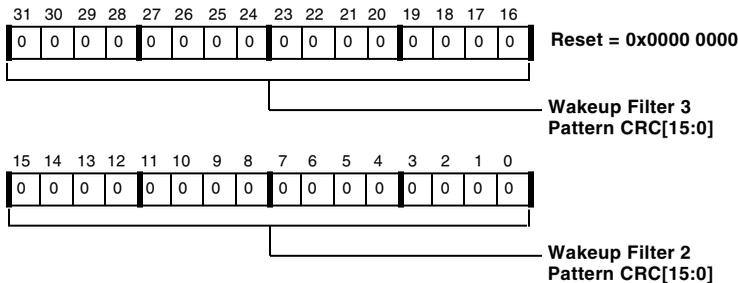


Figure 22-31. EMAC_WKUP_FFCRC1 Register

System Interface Register Group

The SIF block registers control and monitor the MAC's interactions with the Blackfin processor peripheral subsystem and the external PHY. The SIF block has several frame status registers whose bit descriptions can be found in [“Ethernet MAC Frame Status Registers”](#) on page 22-92.

MAC System Control (EMAC_SYSCTL) Register

The EMAC_SYSCTL register, shown in Figure 22-32, is used to set up MAC controls.

MAC System Control Register (EMAC_SYSCTL)

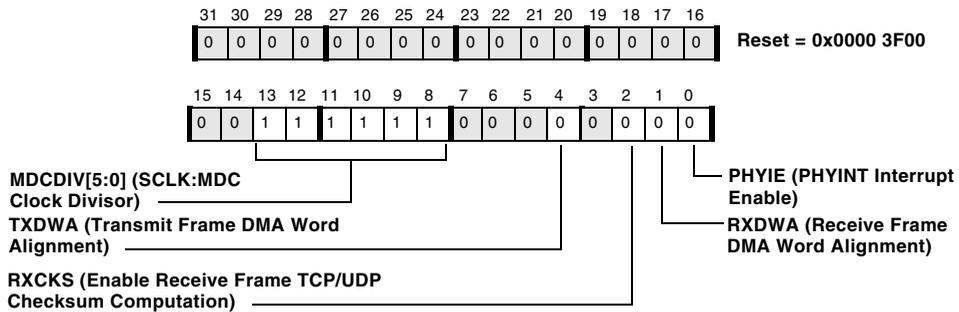


Figure 22-32. EMAC_SYSCTL Register

Additional information for the EMAC_SYSCTL register bits includes:

- **SCLK:MDC clock divisor** (MDCDIV[5:0])

This field contains the clock divisor that determines the ratio between the Blackfin system clock (SCLK) and the MAC data clock (MDC). The 6-bit ratio N determines the MDC rate as:

$$\text{MDC} = \text{SCLK} / (2 \times (N + 1)).$$

- **Transmit frame DMA word alignment** (TXDWA)

This bit determines whether outgoing frame data is aligned on odd or even 16-bit boundaries in memory.

0 – Odd word alignment.

1 – Even word alignment.

Ethernet MAC Register Definitions

- **Enable receive frame TCP/UDP checksum computation** (RXCKS)
0 – Receive frame TCP/UDP checksum computation not enabled.
1 – TCP/UDP checksum computation on received frames enabled.
- **Receive frame DMA word alignment** (RXDWA)
This bit determines whether incoming frames are aligned on odd or even 16-bit boundaries in memory.
0 – Even word alignment.
1 – Odd word alignment.
- **PHYINT interrupt enable** (PHYIE)
0 – PHYINT interrupt not enabled.
1 – PHYINT interrupt enabled.

MAC System Status (EMAC_SYSTAT) Register

The EMAC_SYSTAT register, shown in [Figure 22-33](#), contains a range of interrupt status bits that signal the occurrence of significant Ethernet events to the application. Detailed descriptions of the functionality can be found in the section entitled “[Ethernet Event Interrupts](#)” on [page 22-38](#).

Additional information for the EMAC_SYSTAT register bits includes:

- **Station management transfer done interrupt status** ()
This bit is set when a station management transfer on MDC/MDIO has completed, provided the STAIE interrupt enable control bit is set in the EMAC_STAADD register.

MAC System Status Register (EMAC_SYSTAT)

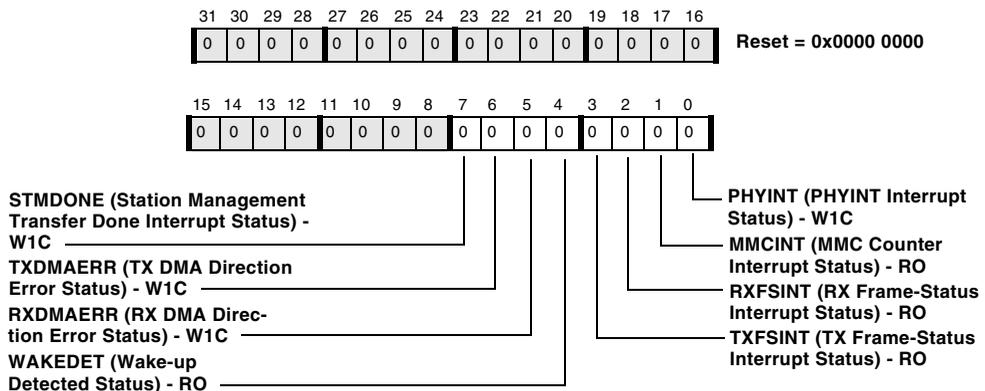


Figure 22-33. EMAC_SYSTAT Register

- **TX DMA direction error status (TXDMAERR)**

This bit is set if a TX data or status DMA request is granted by the DMA channel with transfer in the wrong direction. Data should be memory-read, status should be memory-write. This interrupt is non-maskable in the Ethernet MAC.

- **RX DMA direction error status (RXDMAERR)**

This bit is set if an RX data or status DMA request is granted by the DMA channel with transfer in the wrong (memory-read) direction. This interrupt is non-maskable in the Ethernet MAC.

- **Wakeup detected status (WAKEDET)**

To clear this bit, write 1 to the `EMAC_WKUP_CTL` register.

0 – Wakeup not detected.

1 – Wakeup detected.

Ethernet MAC Register Definitions

- **TX frame-status interrupt status** (TXFSINT)
To clear this bit, write 1s to the EMAC_RX_STKY register bits.
0 – TX frame-status interrupt has not occurred.
1 – TX frame-status interrupt has occurred.
- **RX frame-status interrupt status** (RXFSINT)
To clear this bit, write 1s to the EMAC_RX_STKY register bits.
0 – RX frame-status interrupt has not occurred.
1 – RX frame-status interrupt has occurred.
- **MMC counter interrupt status** (MMCINT)
To clear this bit, write 1 to the EMAC_MMC_RIRQS or EMAC_MMC_TIRQS register.
0 – MMC counter interrupt has not occurred.
1 – MMC counter interrupt has occurred.
- **PHYINT interrupt status** (PHYINT)
0 – PHYINT interrupt has not occurred.
1 – PHYINT interrupt has occurred.

Ethernet MAC Frame Status Registers

The Ethernet MAC frame status registers keep track of the status of each frame received or transmitted, as well as the status of MMC interrupts.

Ethernet MAC RX Current Frame Status (EMAC_RX_STAT) Register

The `EMAC_RX_STAT` register, shown in [Figure 22-34](#), tells the status of the most recently completed receive frame, including type of error for cases where an error occurs. When the receive complete bit is set, exactly one of bits 13 through 20 is 1. Bits 13 through 20 indicate the receive status as defined in IEEE 802.3, section 4.3.2. In case of multiple errors, errors are prioritized in the order listed in [Table 22-4 on page 22-20](#). Bits 18 and 19 identify frames which are not considered received by the station and also are not considered errors. (See section 4.1.2.1.2 and section 4.2.4.2.2 of IEEE 802.3.) Bit 20 identifies frames damaged within the MAC sublayer.

Note if the `PB` (pass bad frames) bit is 0, then delivery via DMA of frames with status bits 14 through 18 or 20 is cancelled. The DMA buffer is reused for the next frame. If the `PR` (promiscuous) bit is 0, then frames with bit 19 set are not delivered (the DMA is never initiated).

Additional information for the `EMAC_RX_STAT` register bits includes:

- **Receive frame accepted** (`RX_ACCEPT`)
 - 0 – Receive frame not accepted.
 - 1 – The receive frame was accepted, based on the address filter result and the frame filtering modes in the `EMAC_OPMODE` register. Note that this does not imply a status of receiveOK. If the `RA` (receive all) control bit is 0, then the only frames delivered by DMA are the frames whose receive frame accepted status bit is 1.
- **VLAN2 frame** (`RX_VLAN2`)
 - 0 – The frame does not meet those conditions.
 - 1 – The frame is a valid tagged frame with a length/type field matching the `EMAC_VLAN2` register, and with status of receiveOK.

Ethernet MAC Register Definitions

Ethernet MAC RX Current Frame Status Register (EMAC_RX_STAT)

All bits in this register are RO.

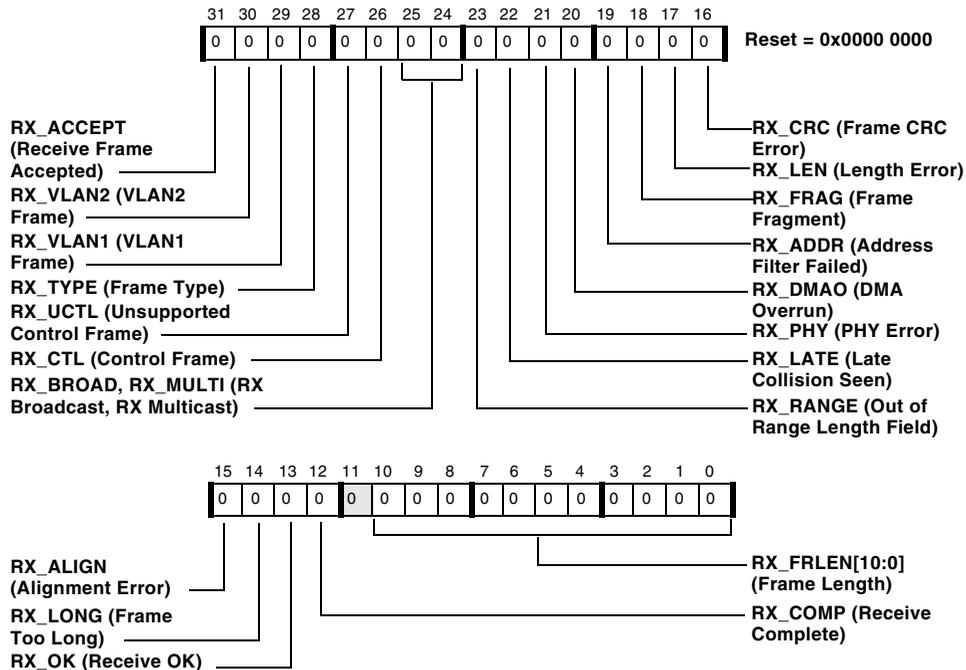


Figure 22-34. EMAC_RX_STAT Register

- VLAN1 frame (RX_VLAN1)**
 - 0 – The frame does not meet those conditions.
 - 1 – The frame is a valid tagged frame with a length/type field matching the EMAC_VLAN1 register, and with status of receiveOK.
- Frame type (RX_TYPE)**
 - 0 – The frame is not of that type.
 - 1 – The frame is a valid typed frame, with status of receiveOK and with a length/type field greater than or equal to 0x600.

- **Unsupported control frame** (RX_UCTL)
 - 0 – The frame does not meet those conditions.
 - 1 – The frame is a valid MAC control frame (with status of receiveOK and with a length/type field equal to 802.3_MAC_Control, 88-08), but does not contain the pause opcode, or is not 64 bits in length, or is received in half-duplex mode.
- **Control frame** (RX_CTL)
 - 1 – The frame is a valid MAC control frame in full duplex mode with status of receiveOK, with a length/type field equal to MAC_Control, 88-08, with length of 64 bytes, and with a MAC control opcode field equal to the pause opcode (00-01).
 - 0 – The frame does not meet those conditions.
- **RX broadcast, RX multicast** (RX_BROAD, RX_MULTII)
 - 00 – Unicast address
 - 01 – Group address
 - 10 – Broadcast address
 - 11 – Illegal
- **Out of range length field** (RX_RANGE)
 - 0 – The frame's length was not out of range.
 - 1 – The frame's length/type field was consistent with the length interpretation (<1536 = 0x600) but was greater than the maximum allowable frame size in bytes, as indicated by the frame too long bit).

Ethernet MAC Register Definitions

- **Late collision seen** (RX_LATE)
 - 0 – Late collision not detected.
 - 1 – A collision was detected after the first 64 bytes of the packet.
- **PHY error** (RX_PHY)
 - 0 – No PHY error.
 - 1 – RX_ER was asserted at some time during the frame. This condition always causes the FCS check to fail.
- **DMA overrun** (RX_DMAO)
 - 0 – No DMA overrun.
 - 1 – The received frame was truncated due to failure of the FIFO/DMA channel to continuously store data during DMA transfer to memory.
- **Address filter failed** (RX_ADDR)
 - 0 – Address did not fail.
 - 1 – The destination address did not pass the address filters specified by the station MAC address, the EMAC_HASHLO/EMAC_HASHHI registers, and the filter modes in the EMAC_OPMODE register.
- **Frame fragment** (RX_FRAG)
 - 0 – Frame length was at least 64 bytes.
 - 1 – Frame length was less than the minimum frame size (64 bytes).

- **Length error** (RX_LEN)
 - 0 – No frame length error.
 - 1 – The frame’s length/type field does not match the length of received data and is consistent with the length interpretation (< 0x600), although the frame had no “frame too long” errors and had a valid FCS.
- **Frame CRC error** (RX_CRC)
 - 0 – No frame CRC error.
 - 1 – The frame failed FCS validation, but had neither a “frame too long” error nor a partial number of octets. Note if RX_ER is asserted by the PHY during frame reception, the FCS validation will fail.
- **Alignment error** (RX_ALIGN)
 - 0 – No alignment error.
 - 1 – The frame ended with a partial octet and failed RCS validation, but had no frame too long error.
- **Frame too long** (RX_LONG)
 - 0 – Frame is not too long.
 - 1 – The number of octets received is greater than the maximum Ethernet frame size. Maximum frame size is 1522 bytes for a frame whose length/type field matches the EMAC_VLAN1 register, 1538 bytes for a frame whose length/type field matches the EMAC_VLAN2 register, or 1518 for all other frames. The frame data delivered by DMA is truncated to 1556 (0x614) bytes in all cases.

Ethernet MAC Register Definitions

- **Receive OK** (RX_OK)

0 – A receive error occurred.

1 – There was no receive error.

- **Receive complete** (RX_COMP)

This bit is cleared on reset and when the MAC RX is enabled (RE changes from 0 to 1). Frames that fail the address filter or the frame filter are not delivered by DMA, unless overridden by the RA (receive all) control bit. Note that in the RX frame status buffer written to memory by DMA, the receive complete bit is always 1. This bit acts as a semaphore, indicating that DMA of the frame has completed.

0 – The first RX frame is not yet complete.

1 – The first RX frame is complete.

- **Frame length** (RX_FRLLEN)

The number of bytes in the frame. If the ASTP bit is set, the pad and FCS are not included in the length.

Ethernet MAC RX Sticky Frame Status (EMAC_RX_STKY) Register

The EMAC_RX_STKY register, shown in Figure 22-35, accumulates state across multiple frames, unless software clears it after every frame.

Ethernet MAC RX Sticky Frame Status Register (EMAC_RX_STKY)

All bits in this register are W1C.

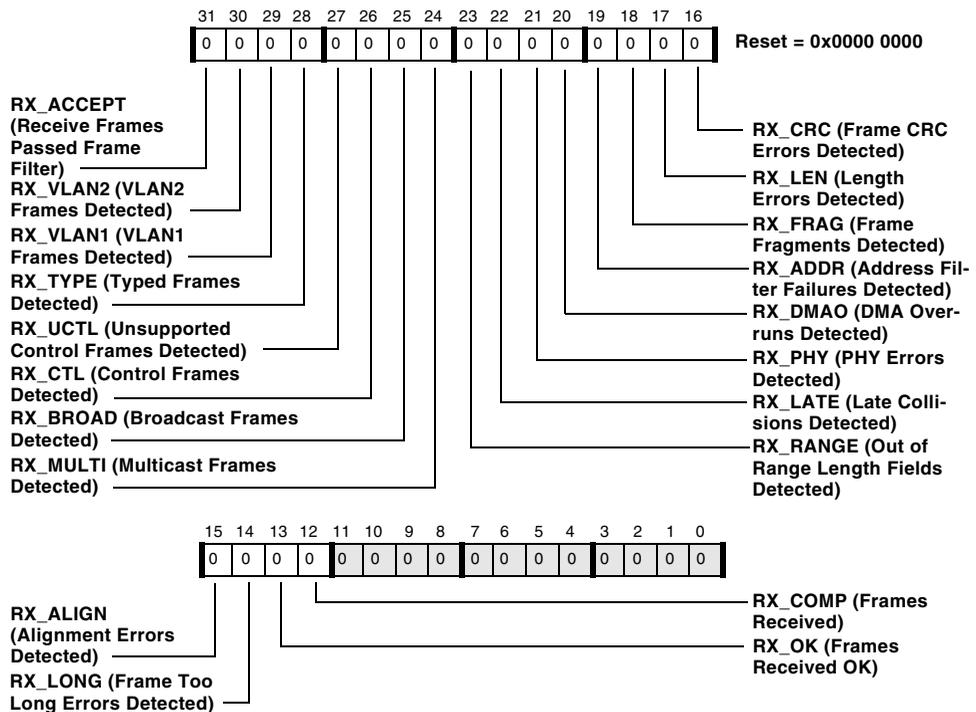


Figure 22-35. EMAC_RX_STKY Register

Ethernet MAC Register Definitions

Additional information for the `EMAC_RX_STKY` register bits includes:

- **Receive frames passed frame filter** (`RX_ACCEPT`)
 - 0 – No receive frames passed the frame filter.
 - 1 – At least one receive frame passed the frame filter.
- **VLAN2 frames detected** (`RX_VLAN2`)
 - 0 – No VLAN2 frames were detected.
 - 1 – At least one VLAN2 frame was detected.
- **VLAN1 frames detected** (`RX_VLAN1`)
 - 0 – No VLAN1 frames were detected.
 - 1 – At least one VLAN1 frame was detected.
- **Typed frames detected** (`RX_TYPE`)
 - 0 – No typed frames were detected.
 - 1 – At least one typed frame was detected.
- **Unsupported control frames detected** (`RX_UCTL`)
 - 0 – No unsupported control frames were detected.
 - 1 – At least one unsupported control frame was detected.
- **Control frames detected** (`RX_CTL`)
 - 0 – No control frames were detected.
 - 1 – At least one control frame was detected.

- **Broadcast frames detected** (RX_BROAD)
 - 0 – No broadcast frames were detected.
 - 1 – At least one broadcast frame was detected.
- **Multicast frames detected** (RX_MULTI)
 - 0 – No multicast frames were detected.
 - 1 – At least one multicast frame was detected.
- **Out of range length fields detected** (RX_RANGE)
 - 0 – No out of range length fields were detected.
 - 1 – At least one out of range length field was detected.
- **Late collisions detected** (RX_LATE)
 - 0 – No late collisions were detected.
 - 1 – At least one collision was detected after the first 64 bytes of the packet.
- **PHY errors detected** (RX_PHY)
 - 0 – No PHY errors were detected.
 - 1 – At least one PHY error was detected.
- **DMA overruns detected** (RX_DMA0)
 - 0 – No DMA overruns were detected.
 - 1 – At least one DMA overrun was detected.
- **Address filter failures detected** (RX_ADDR)
 - 0 – No address filter failures were detected.
 - 1 – At least one address filter failure was detected.

Ethernet MAC Register Definitions

- **Frame fragments detected** (RX_FRAG)
 - 0 – No frame fragments were detected.
 - 1 – At least one frame fragment was detected.
- **Length errors detected** (RX_LEN)
 - 0 – No length errors were detected.
 - 1 – At least one length error was detected.
- **Frame CRC errors detected** (RX_CRC)
 - 0 – No frame CRC errors were detected.
 - 1 – At least one CRC error was detected.
- **Alignment errors detected** (RX_ALIGN)
 - 0 – No alignment errors were detected.
 - 1 – At least one alignment error was detected.
- **Frame too long errors detected** (RX_LONG)
 - 0 – No frame too long errors were detected.
 - 1 – At least one frame too long error was detected.

- **Frames received OK** (RX_OK)

This bit can be used to generate an interrupt on the next RX frame.

0 – No good frames have been received.

1 – At least one frame has been received OK.

- **Frames received** (RX_COMP)

0 – No frames were received.

1 – At least one frame (good or bad) was received.

Ethernet MAC Register Definitions

Ethernet MAC RX Frame Status Interrupt Enable (EMAC_RX_IRQE) Register

The EMAC_RX_IRQE register, shown in Figure 22-36, enables the frame status interrupts.

Ethernet MAC RX Frame Status Interrupt Enable Register (EMAC_RX_IRQE)

For all bits, 1 = Interrupt enabled, 0 = Interrupt not enabled.

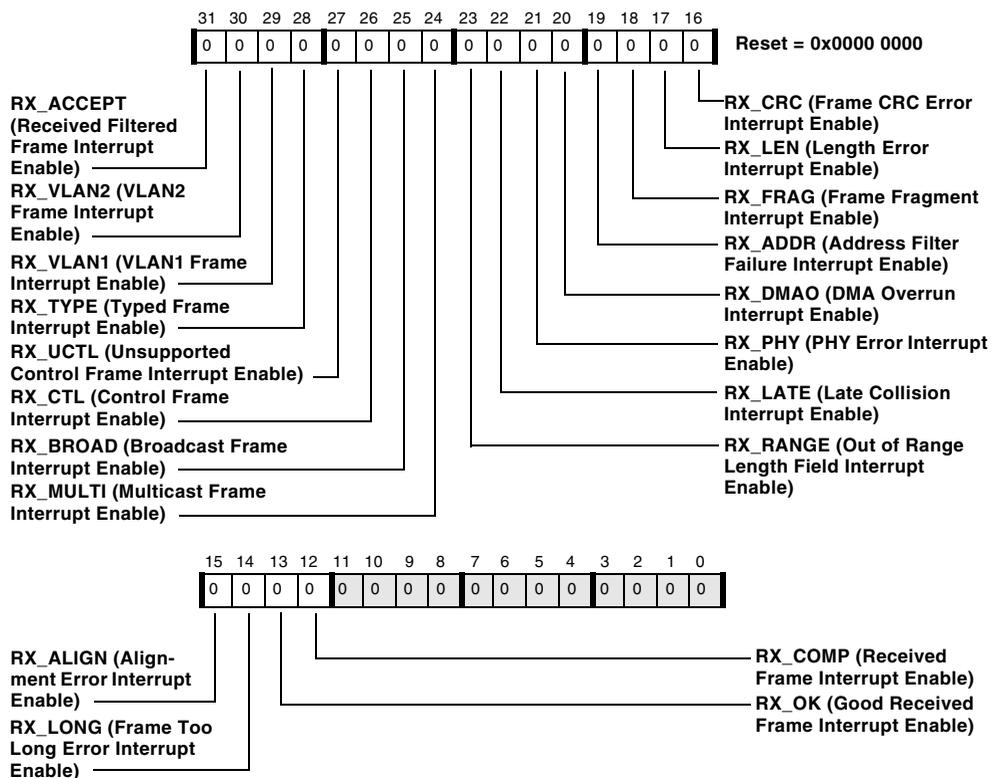


Figure 22-36. EMAC_RX_IRQE Register

Ethernet MAC TX Current Frame Status (EMAC_TX_STAT) Register

The EMAC_TX_STAT register, shown in Figure 22-37, tells the status of the most recently completed transmit frame, including type of error for cases where an error occurred. When the transmit complete bit is set, exactly one of bits 2, 3, 4, 13, or 14 is 1. Bits 1 through 3 indicate the transmit status as defined in IEEE 802.3, section 4.3.2.

Ethernet MAC TX Current Frame Status Register (EMAC_TX_STAT)

All bits in this register are RO.

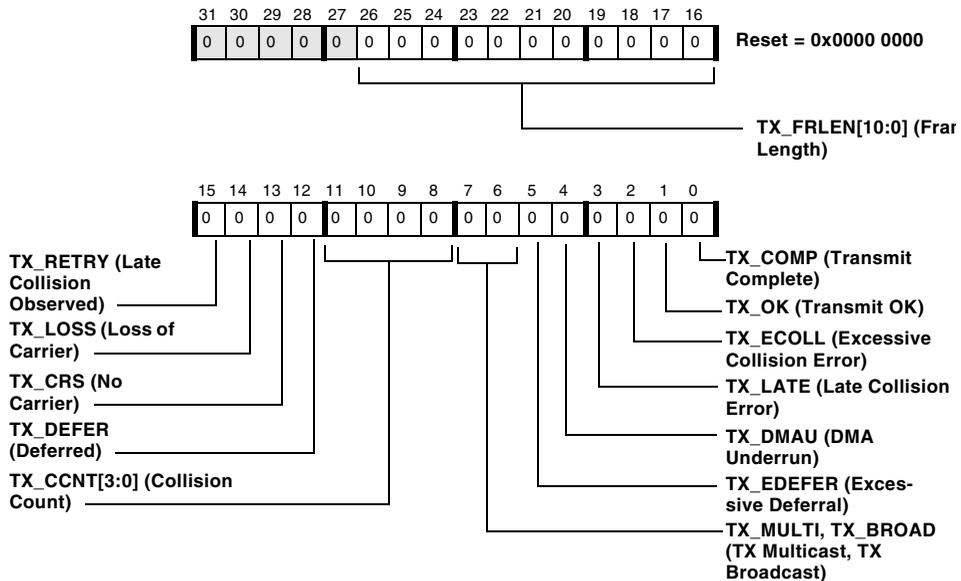


Figure 22-37. EMAC_TX_STAT Register

Additional information for the EMAC_TX_STAT register bits includes:

- **TX frame length** (TX_FRLLEN)

This field contains the length of the transmit frame in bytes.

Ethernet MAC Register Definitions

- **Late collision observed** (TX_RETRY)
 - 0 – No late collision occurred.
 - 1 – A late collision occurred, but the frame transmission was successful after retry.
- **Loss of carrier** (TX_LOSS)
 - 0 – No loss of carrier occurred.
 - 1 – The carrier sense transitioned from asserted to deasserted at some time during the frame transmission. Half-duplex only, MII mode only.
- **No carrier** (TX_CRS)
 - 0 – CRS was asserted.
 - 1 – Carrier sense (CRS) was not asserted at any time during frame transmission. Half-duplex only.
- **Deferred** (TX_DEFER)
 - 0 – Transmission not deferred.
 - 1 – The transmission was deferred in half-duplex mode because the medium was initially occupied (CRS was asserted) at the time the frame was ready to transmit (after the initial frame data was transferred by DMA to the MAC). Note the deferred status bit should be expected to be 1 on frames that have been retried after early collisions, since the MAC can restart the frame immediately after a collision using data available in its local FIFO. Since the MAC does not need to wait for DMA, the frame data is typically ready for retransmission before TXEN and CRS have deasserted from the prior attempt. Half-duplex only.

- **Collision count** (TX_CCNT)

This field contains the number of collisions that occurred during frame transmission.

- **TX broadcast, TX multicast** (TX_BROAD, TX_MULTII)

00 – Unicast address

01 – Broadcast address

10 – Group address

11 – Illegal

- **Excessive deferral** (TX_EDEFER)

0 – Excessive deferral did not occur.

1 – The frame transmission was deferred for more than 24,288 bit times or 6072 TX clocks:

$$\text{MaxDeferTime} = 2 \times (\text{MaxUntaggedFrameSize} \times 8) \text{ bits}$$

If the deferral check (DC) bit in the EMAC_OPMODE register is 1, frame transmission is aborted upon excessive deferral, and both the excessive deferral and excessive collision error status bits are set.

- **DMA underrun** (TX_DMAU)

0 – No DMA underrun.

1 – The frame transmission was interrupted by a failure of the FIFO/DMA channel to continuously supply frame data after the start of transmission on the MII/RMII.

Ethernet MAC Register Definitions

- **Late collision error** (TX_LATE)
 - 0 – No late collision error.
 - 1 – Frame transmission failed because a collision occurred after the end of the collision window (512 bit times) and the LCRTE bit was clear, disabling frame transmission retry.
- **Excessive collision error** (TX_ECOLL)
 - 0 – No excessive collision error.
 - 1 – Frame transmission failed because too many (16) attempts were interrupted by collisions, or because the frame was deferred for more than the maximum deferral time while the deferral check (DC) control bit was set.
- **Transmit OK** (TX_OK)
 - 0 – A transmit error occurred.
 - 1 – There was no transmit error.
- **Transmit complete** (TX_COMP)
 - This bit is cleared on reset and when the MAC TX is enabled (TE changes from 0 to 1). In the TX DMA status buffer, this bit is always set to 1 on every status word written via DMA. This bit thus acts as a semaphore, indicating to software that processing of this descriptor pair has been completed.
 - 0 – The first TX frame is not yet complete.
 - 1 – The first TX frame is complete.

Ethernet MAC TX Sticky Frame Status (EMAC_TX_STKY) Register

The EMAC_TX_STKY register, shown in [Figure 22-38](#), accumulates state across multiple frames, unless software clears it after every frame.

Ethernet MAC TX Sticky Frame Status Register (EMAC_TX_STKY)

All bits in this register are W1C.

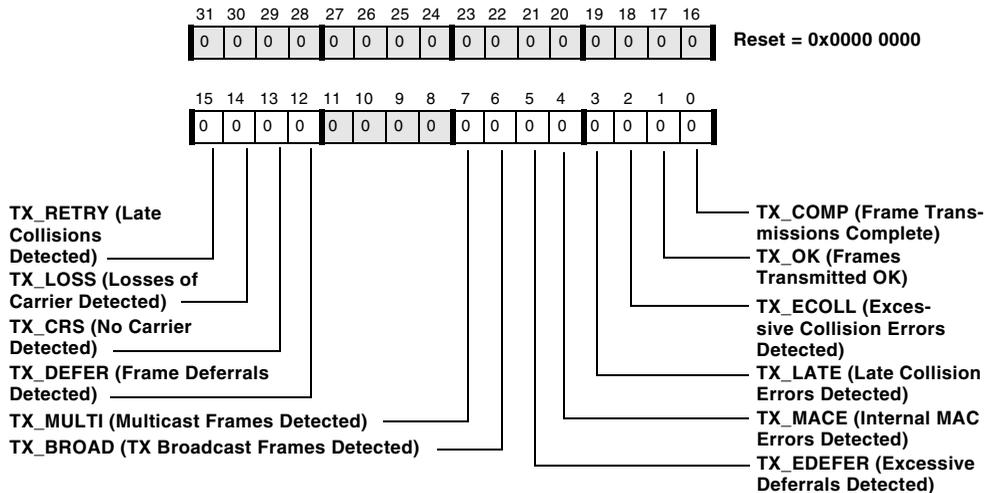


Figure 22-38. EMAC_TX_STKY Register

Additional information for the EMAC_TX_STKY register bits includes:

- **Late collisions detected (TX_RETRY)**
 - 0 – No late collisions were detected.
 - 1 – At least one late collision was detected on frames successfully transmitted after retry.

Ethernet MAC Register Definitions

- **Losses of carrier detected** (TX_LOSS)
 - 0 – No losses of carrier were detected.
 - 1 – At least one loss of carrier was detected
- **No carrier detected** (TX_CRIS)
 - 0 – No instances of no carrier were detected.
 - 1 – At least one occasion of no carrier was detected.
- **Frame deferrals detected** (TX_DEFER)
 - 0 – No frame deferrals were detected.
 - 1 – At least one frame deferral was detected.
- **TX multicast frames detected** (TX_MULTII)
 - 0 – No multicast frames were detected.
 - 1 – At least one multicast frame was detected.
- **TX broadcast frames detected** (TX_BROAD)
 - 0 – No broadcast frames were detected.
 - 1 – At least one broadcast frame was detected.
- **Excessive deferrals detected** (TX_EDEFER)
 - 0 – No excessive deferrals were detected.
 - 1 – At least one excessive deferral was detected.
- **Internal MAC errors detected** (TX_MACE)
 - 0 – No internal MAC errors were detected.
 - 1 – At least one internal MAC error was detected.

- **Late collision errors detected** (TX_LATE)
 - 0 – No late collision errors were detected.
 - 1 – At least one late collision error was detected.
- **Excessive collision errors detected** (TX_ECOLL)
 - 0 – No excessive collision errors were detected.
 - 1 – At least one excessive collision error detected.
- **Frame transmissions complete** (TX_COMP)
 - 0 – No frames have been transmitted.
 - 1 – At least one frame was transmitted.
- **Frames transmitted OK** (TX_OK)

This bit can be used to generate an interrupt at the completion of each TX frame.

 - 0 – No good frames have been transmitted.
 - 1 – At least one frame has been transmitted OK.

Ethernet MAC Register Definitions

Ethernet MAC TX Frame Status Interrupt Enable (EMAC_TX_IRQE) Register

The EMAC_TX_IRQE register, shown in [Figure 22-39](#), is used to enable TX frame status interrupts.

Ethernet MAC TX Frame Status Interrupt Enable Register (EMAC_TX_IRQE)

For all bits, 1 = Interrupt enabled, 0 = Interrupt not enabled.

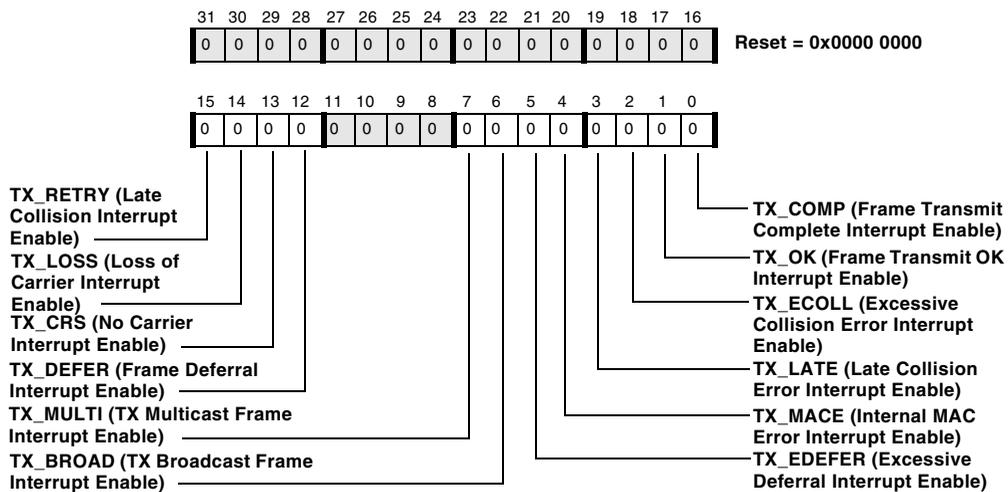


Figure 22-39. EMAC_TX_IRQE Register

Ethernet MAC MMC RX Interrupt Status (EMAC_MMC_RIRQS) Register

The EMAC_MMC_RIRQS register, shown in [Figure 22-40](#), indicates which of the receive MAC management counters have incremented past one-half of maximum range. Each bit is set from 0 to 1 when the corresponding counter increments from a value less than 0x8000 0000 to a value greater than or equal to 0x8000 0000 (regardless of the state of the EMAC_MMC_RIRQE register). Bits in this register are cleared by writing a 1; writing zero has no effect. For more information, see [“MAC Management Counters” on page 22-42](#).

Ethernet MAC MMC RX Interrupt Status Register (EMAC_MMC_RIRQS)

All bits are W1C. For all bits, 1 = Interrupt occurred, 0 = Interrupt did not occur.

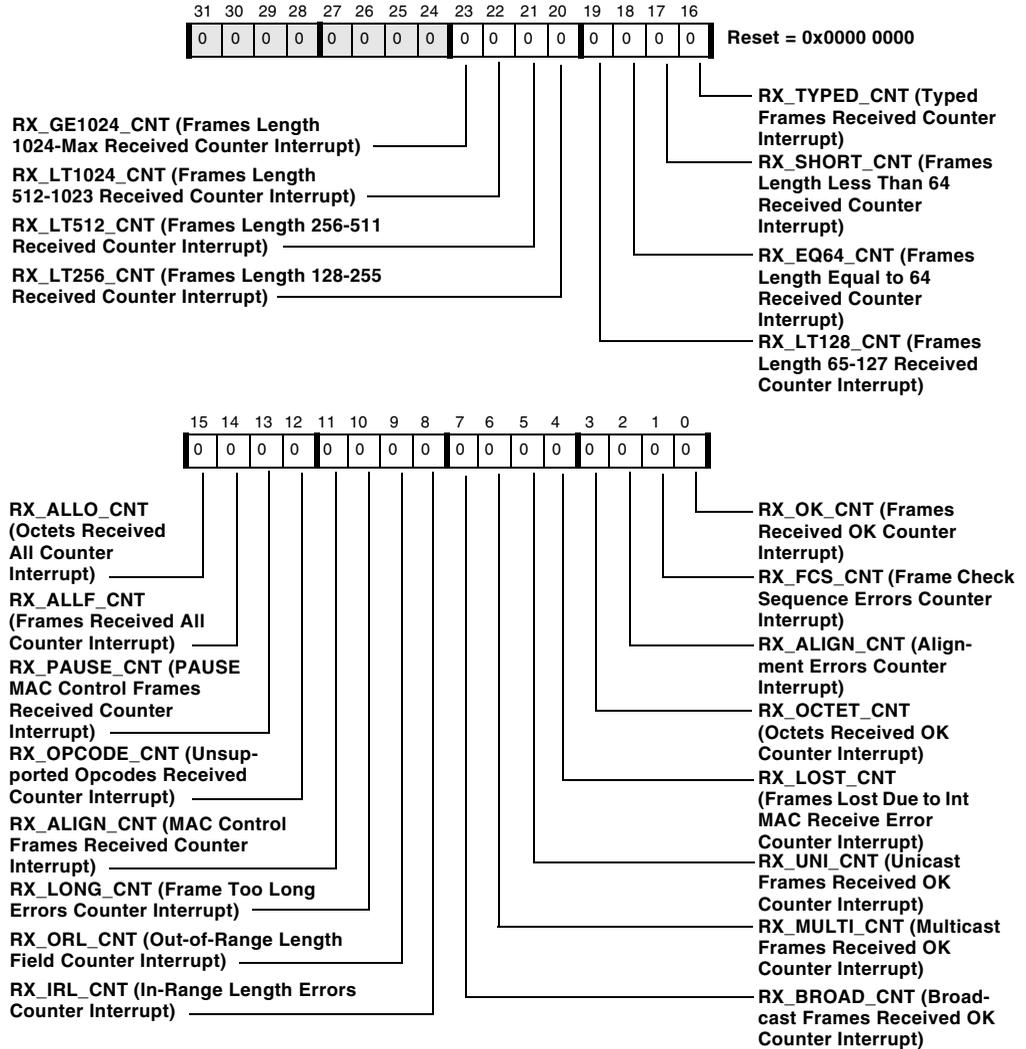


Figure 22-40. EMAC_MMC_RIRQS Register

Ethernet MAC Register Definitions

Ethernet MAC MMC RX Interrupt Enable (EMAC_MMC_RIRQE) Register

The `EMAC_MMC_RIRQE` register, shown in [Figure 22-41](#), indicates which of the receive MAC management counters are enabled to signal an `MMCINT` interrupt when they increment past one-half of maximum range.

If a given counter's interrupt is not enabled, and that counter passes `0x8000 0000`, then the counter's interrupt status bit is set to 1 but this does not cause the `MMCINT` interrupt to be signalled. If the corresponding interrupt enable bit is later written to 1, the `MMCINT` Ethernet event interrupt is signalled immediately.

Ethernet MAC MMC TX Interrupt Status (EMAC_MMC_TIRQS) Register

The `EMAC_MMC_TIRQS` register, shown in [Figure 22-42](#), indicates which of the transmit MAC management counters have incremented past one-half of maximum range. Each bit is set from 0 to 1 when the corresponding counter increments from a value less than `0x8000 0000` to a value greater than or equal to `0x8000 0000` (regardless of the state of the `EMAC_MMC_TIRQE` register). Bits in this register are cleared by writing a 1; writing zero has no effect. For more information, see “[MAC Management Counters](#)” on [page 22-42](#).

Ethernet MAC MMC RX Interrupt Enable Register (EMAC_MMC_RIRQE)

For all bits, 1 = Interrupt enabled, 0 = Interrupt not enabled.

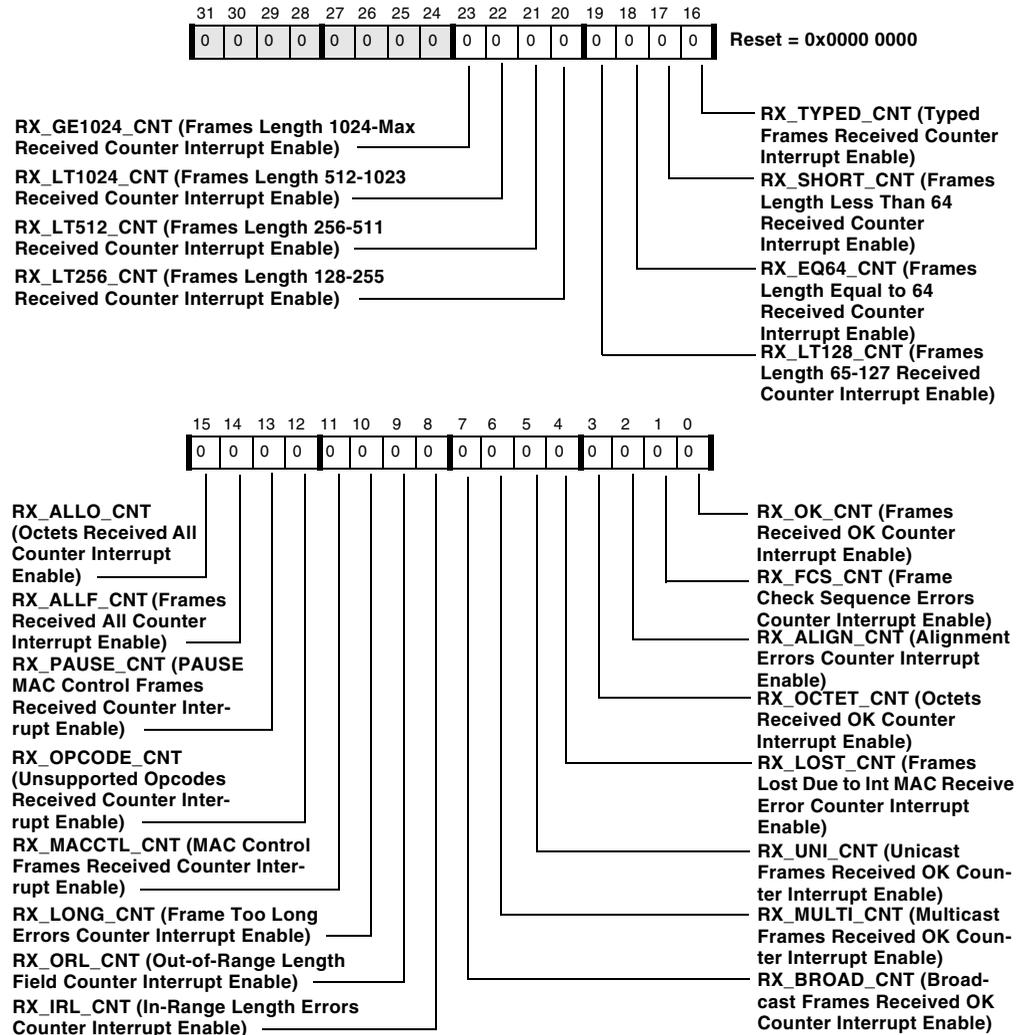


Figure 22-41. EMAC_MMC_RIRQE Register

Ethernet MAC Register Definitions

Ethernet MAC MMC TX Interrupt Status Register (EMAC_MMC_TIRQS)

All bits are W1C. For all bits, 1 = Interrupt occurred, 0 = Interrupt did not occur.

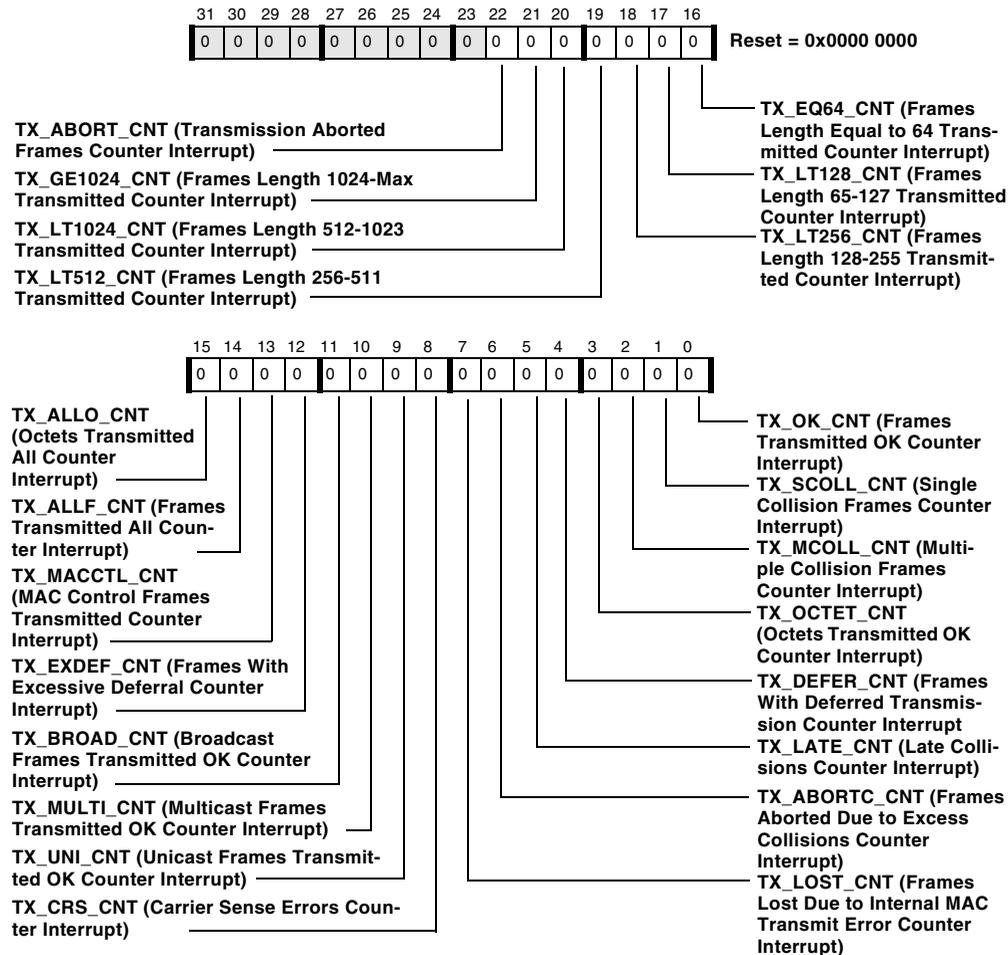


Figure 22-42. EMAC_MMC_TIRQS Register

Ethernet MAC MMC TX Interrupt Enable (EMAC_MMC_TIRQE) Register

The `EMAC_MMC_TIRQE` register, shown in [Figure 22-43](#), indicates which of the transmit MAC management counters are enabled to signal an `MMCINT` interrupt when they increment past one-half of maximum range.

If a given counter's interrupt is not enabled, and that counter passes `0x8000 0000`, then the counter's interrupt status bit is set to 1 but this does not cause the `MMCINT` interrupt to be signalled. If the corresponding interrupt enable bit is later written to 1, the `MMCINT` Ethernet event interrupt is signalled immediately.

MAC Management Counter Registers

The MAC Management Counter (MMC) block register group consists of a number of 32-bit unsigned counter registers that gather statistical data regarding the operation of the MAC. The MAC management counter registers update automatically at the completion of frame transmit and receive, whenever the `MMCE` bit in the `EMAC_MMC_CTL` register is set. Counters contain a 32-bit unsigned value, and may be configured to saturate at `0xFFFF FFFF` (`CROLL = 0`) or to wrap around to zero (`CROLL = 1`). Counters cannot be written directly, but can be collectively reset to zero by writing 1 to the `RSTC` bit, or they can be programmed for clear-on-read behavior by setting `CCOR` to 1. The reset value for all MMC registers is `0x0000 0000`. See [Table 22-10 on page 22-53](#) for more information.

Each of these counters can be set up to generate interrupts when they reach half of the maximum unsigned 32-bit value. This functionality is described in detail in the section entitled “[Ethernet Event Interrupts](#)” on [page 22-38](#).

Ethernet MAC Register Definitions

Ethernet MAC MMC TX Interrupt Enable Register (EMAC_MMC_TIRQE)

For all bits, 1 = Interrupt enabled, 0 = Interrupt not enabled.

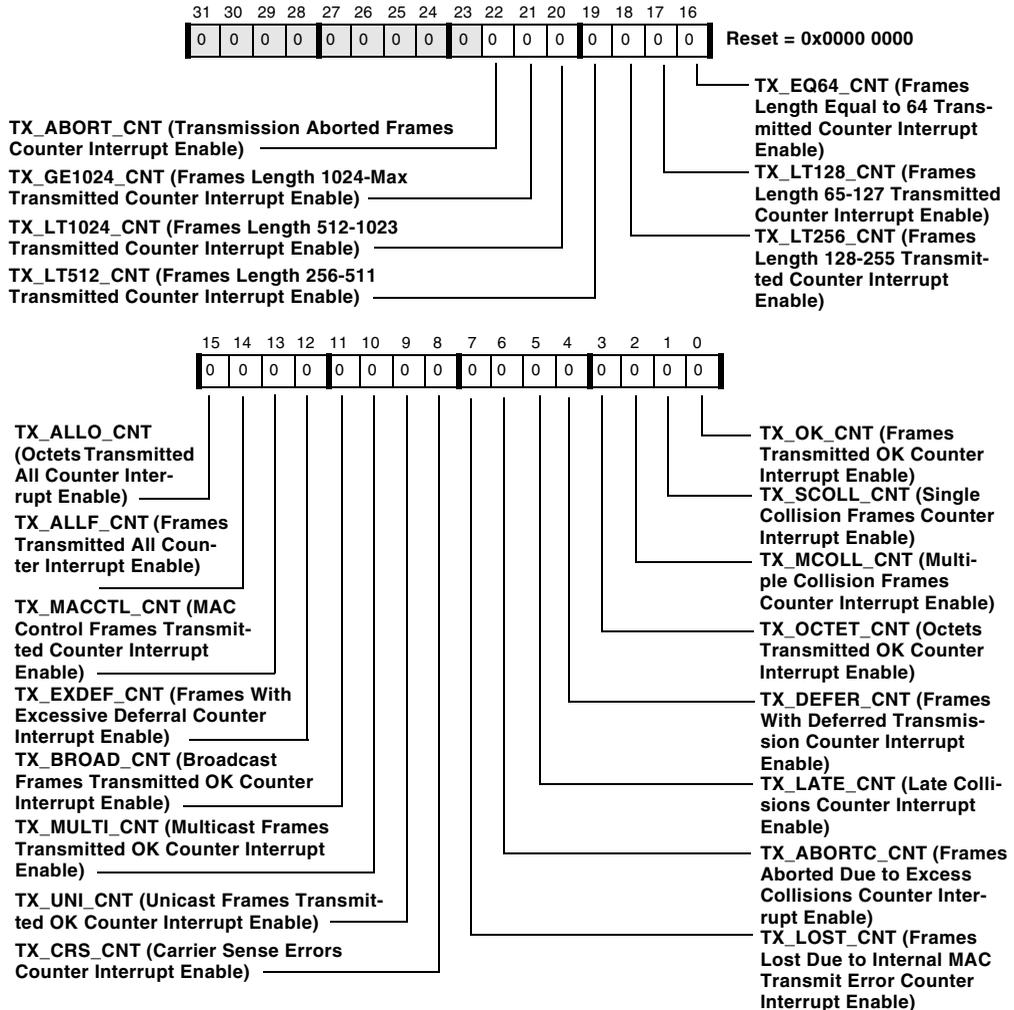


Figure 22-43. EMAC_MMC_TIRQE Register

MAC Management Counters Control (EMAC_MMC_CTL) Register

The EMAC_MMC_CTL register, shown in [Figure 22-44](#), is used to globally configure all MMC counter registers.

MAC Management Counters Control Register (EMAC_MMC_CTL)

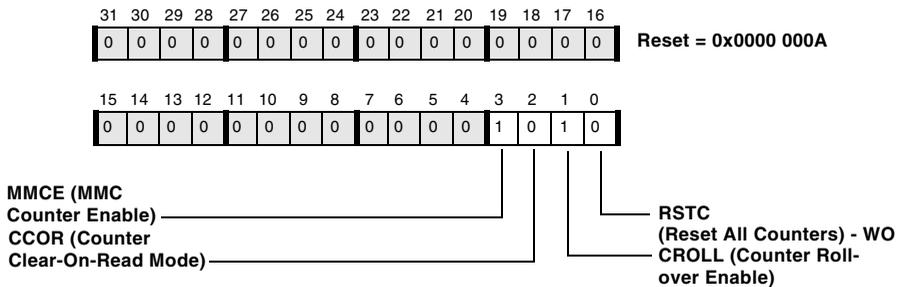


Figure 22-44. EMAC_MMC_CTL Register

Additional information for the EMAC_MMC_CTL register bits includes:

- **MMC counter enable** (MMCE)

Setting this bit turns on all the MMC counters, which update on every frame transmission or reception.

0 – MMC counters are not enabled. Counters retain their values but are not updated.

1 – MMC counters are enabled.

- **Counter clear-on-read mode** (CCOR)

0 – Counters are not in clear-on-read mode. Reads do not affect counter contents.

1 – Counters are in clear-on-read mode. The contents of each counter is reset each time it is read by the application.

Programming Examples

- **Counter rollover enable** (CROLL)

0 – Counter rollover is not enabled. All MMC registers saturate upon reaching 0xFFFF FFFF.

1 – Counter rollover is enabled. This causes all MMC counters to wrap around to zero when the count exceeds the maximum 32-bit value of 0xFFFF FFFF.

- **Reset all counters** (RSTC)

Writing a 1 to this bit at any time globally resets all MMC counters.

0 – Do not reset all counters.

1 – Globally clear all MMC counters.

Programming Examples

This section gives a general overview of the functionality of an Ethernet MAC driver. All necessary steps for reproducing and understanding this interface are explained with code listings and accompanying text. These code listings are similar to the driver model supported by VisualDSP++ and are mainly written in C. Data transfers over the MAC with DMA are explained in [Figure 22-5 on page 22-11](#) and [Figure 22-7 on page 22-23](#), which show receive and transmit DMA operations. Please examine these figures carefully—the code listings reproduce this kind of “linked list” in the form of C structures. Also provided are code listings that describe accessing an external PHY via the station management (MIM) block. All macros that are not explained in this section can be found in the `cdefBF5xx.h` and `defBF5xx.h` header files of VisualDSP++.

The code examples in this section ([Listing 22-1](#) through [Listing 22-9](#)) show basic functions and structures. The management counter registers and the interrupt settings are advanced functions and are not covered

here. There are many counter registers which are accessible by polling of the appropriate register or using interrupt service routines. The EMAC_SYSCTL and EMAC_SYSTAT register should be used to configure the Ethernet MAC interrupt capabilities. See [Figure 22-12 on page 22-38](#) for a detailed description of the MAC interrupts.

Ethernet Structures

Listing 22-1. Type Definition

```
// type definitions
typedef unsigned long int u32;
typedef unsigned short int u16;
typedef unsigned char u8;
typedef volatile u32 reg32;
typedef volatile u16 reg16;
```

The type definitions are placed here to help with reading of the following code.

Listing 22-2. DMA Configuration

```
typedef struct ADI_DMA_CONFIG_REG {
    u16 b_DMAEN:1;      /* 0      Enabled */
    u16 b_WNR:1;       /* 1      Direction */
    u16 b_WDSIZE:2;    /* 2:3    Transfer word size */
    u16 b_DMA2D:1;     /* 4      DMA mode */
    u16 b_SYNC:1;     /* 5      Retain FIFO */
    u16 b_DI_SEL:1;    /* 6      Data interrupt timing select */
    u16 b_DI_EN:1;     /* 7      Data interrupt enabled */
    u16 b_NDSIZE:4;    /* 8:11   Flex descriptor size */
    u16 b_FLOW:3;      /* 12:14  Flow */
} ADI_DMA_CONFIG_REG;
```

Programming Examples

A convenient way to handle the DMA properties in a “linked list” is to use structures, because each set should be assigned to the appropriate DMA descriptor. [Listing 22-3](#) shows a structure used to manage DMA descriptors. Before jumping to the next descriptor, like 1A-1B-2A-2B-1C in [Figure 22-5 on page 22-11](#) and [Figure 22-7 on page 22-23](#), the structure `ADI_DMA_CONFIG_REG` immediately loads to the DMA register before starting its DMA transfer.

Listing 22-3. DMA Descriptor

```
typedef struct dma_descriptor {
    struct dma_descriptor*  NEXT_DESC_PTR;
    u32                     START_ADDR;
    ADI_DMA_CONFIG_REG      CONFIG;
} DMA_DESCRIPTOR;
```

The structure shown in [Listing 22-3](#) shows how it is possible to create a “linked list” of DMAs. The `START_ADDR` points to the data and the `ADI_DMA_CONFIG_REG` structure (shown in [Listing 22-2](#)) holds all the necessary settings.

Structures like these are convenient for handling Ethernet streams, because they allow the programmer to simply call members of the structure instead of extracting meaningful items through array offsets. This structure, shown in [Listing 22-4](#), is mirrored in the Ethernet MAC header with additional NoBytes.

Listing 22-4. Ethernet Frame Buffer

```
typedef struct adi_ether_frame_buffer {
    u16  NoBytes; /* the no. of following bytes */
    u8   Dest[6]; /* destination MAC address */
    u8   Srce[6]; /* source MAC address */
    u16  LTfield; /* length/type field */
}
```

```

    u8    Data[0]; /* payload bytes */
} ADI_ETHER_FRAME_BUFFER;

```

The `ADI_ETHER_BUFFER` structure in [Listing 22-5, Top Level Structure](#), covers all the above structures and shows the general framework as described in [Figure 22-5 on page 22-11](#) and [Figure 22-7 on page 22-23](#). The two `Dma[2]` structures are needed for descriptors 1A, 1B and 2A, 2B. The pointer `*frmData` represents the payload of the frame, which has a specific number of bytes (as dictated by the `NoBytes` structure member). This is relevant only in transmit mode—in receive mode the driver will not touch this `NoBytes` variable. To ease programming by keeping the transmit and receive structures the same, the MAC can pad the first 16-bit word (that is, the data corresponding to the `NoBytes` structure member) with zeros if the `RXDWA` bit in `EMAC_SYSCTL` is 1. The `*pNext` and `*pPrev` pointers are necessary for creating a “linked list.” The `IPHdrChksum` and `IPPayloadChksum` are available in case the Ethernet MAC is set to calculate this. See the `RXCKS` bit in the `EMAC_SYSCTL` register (shown in [Figure 22-32 on page 22-89](#)). These two variables are relevant only in receive mode of the Ethernet MAC. The `StatusWord` variable holds the `EMAC_RX_STAT` register value in receive mode and holds the `EMAC_TX_STAT` register value in transmit mode.

Listing 22-5. Top Level Structure

```

typedef struct adi_ether_buffer {
    DMA_DESCRIPTOR Dma[2]; /* first for the frame, second for the
status */
    ADI_ETHER_FRAME_BUFFER *FrmData; /* pointer to data */
    struct adi_ether_buffer *pNext; /* next buffer */
    struct adi_ether_buffer *pPrev; /* prev buffer */
    u16 IPHdrChksum; /* the IP header checksum */
    u16 IPPayloadChksum; /* the IP header and payload checksum */
    u32 StatusWord; /* the frame status word */
} ADI_ETHER_BUFFER;

```

MAC Address Setup

Write `EMAC_ADDRLO` and `EMAC_ADDRHI` in the initialization routine of the Ethernet MAC, as shown in [Listing 22-6](#). The Ethernet MAC address is a unique number and may not be used twice. See the IEEE Std. 802.3-2002 specification for further information.

Listing 22-6. MAC Address Setup

```
// MAC address
u8 SrcAddr[6] = {0x5A,0xD4,0x9A,0x48,0xDE,0xAC};

// function
void SetupMacAddr(u8 *MACaddr)
{
    *pEMAC_ADDRLO = *(u32 *)&MACaddr[0];
    *pEMAC_ADDRHI = *(u16 *)&MACaddr[4];
}

// function call
SetupMacAddr(SrcAddr);
```

PHY Control Routines

The `EMAC_STAAD` register provides the option of either polling the `STABUSY` bit or getting an interrupt during each MIM block access. The function in [Listing 22-7](#) polls the `STABUSY` bit and should be placed after each read or write command to the PHY register.

Listing 22-7. Poll MIM Block

```

//
/* Wait until the previous MDC/MDIO transaction has completed */
//
void PollMdcDone(void)
{
    /* poll the STABUSY bit */
    while(*pEMAC_STAADD & STABUSY)
}

```

Shown in [Listing 22-8](#), the `SET_PHYAD` and `SET_REGAD` macros shift the `PHYAddr` and `RegAddr` values to the appropriate field within the `EMAC_STAADD` register. The other macros `STAOP`, `STAIE`, and `STABUSY`, also set bits in the `EMAC_STAADD` register. Use of the `STAOP` macro controls the read and write transfer of the MIM block.

Listing 22-8. Write Access to the PHY

```

//
/* Write an off-chip register in a PHY through the MDC/MDIO port
*/
//
void WrPHYReg(u16 PHYAddr, u16 RegAddr, u16 Data)
{
    PollMdcDone();
    *pEMAC_STADAT = Data;
    *pEMAC_STAADD = SET_PHYAD(PHYAddr) |\
                    SET_REGAD(RegAddr) |\
                    STAOP | STABUSY;
}

```

The data in the `EMAC_STADAT` register is immediately shifted out after a write to the `EMAC_STAADD` register. See [Figure 22-4 on page 22-9](#).

Programming Examples

The function in [Listing 22-9](#) shows how PHY data is read over the MIM function block of the MAC. First, the STABUSY bit of the EMAC_STAADD register will be polled until no other function is using the MIM block. The PHY address and register address is sent over the MIM block. Then, the STABUSY bit is polled again, before the data is finally read through the EMAC_STADAT register.

Listing 22-9. Read Access to the PHY

```
//
/* Read an off-chip register in a PHY through the MDC/MDIO port
*/
//
u16 RdPHYReg(u16 PHYAddr, u16 RegAddr)
{
    u16 Data;
    PollMdcDone();

    *pEMAC_STAADD = SET_PHYAD(PHYAddr) |\
                    SET_REGAD(RegAddr) |\
                    STABUSY;

    PollMdcDone();
    Data = (u16)*pEMAC_STADAT;

    return Data;
}
```

A complete PHY initialization also requires the initialization of the station management clock, which is described in detail in the section “[MII Station Management](#)” on page 22-47. The three PHY functions included in this section (write, read, and poll) and the initialization routine of the station management clock are the minimum requirements for setup and control of any PHYs.

Unique Information for the ADSP-BF51x Processor

None

Unique Information for the ADSP-BF51x Processor

23 IEEE 1588 PTP ENGINE

This chapter describes the IEEE 1588 engine module, PTP_TSYNC (Precision Time Protocol Time Synchronization), for ADSP-BF518 processors. Following an overview and list of key features is an introduction to the IEEE 1588 protocol and a description of the module's operation. The chapter concludes with consolidated register definitions and programming models. This chapter includes the following sections:

- [“PTP_TSYNC Overview”](#)
- [“General Operation” on page 23-2](#)
- [“PTP_TSYNC Module Registers” on page 23-11](#)
- [“PTP_TSYNC Module Programming Model” on page 23-37](#)

PTP_TSYNC Overview

The IEEE 1588 engine module (PTP_TSYNC) provides hardware assistance for the implementation of the IEEE 1588 standard on Blackfin processors. It supports IEEE Std. 1588-2002 and IEEE Std. 1588-2008.

Features

The PTP_TSYNC module includes these features:

- Support for IEEE 1588-2002 and IEEE 1588-2008
- Addend clock adjustment

General Operation

- Programmable PTP message detection
- Pulse-per-second (PPS) signal output
- Alarm features
- Three different clock input sources, and a clock output with a programmable clock divider
- Auxiliary GPIO/event triggered timestamp
- Timestamp overrun indication

General Operation

Multiple clocks of a distributed system may drift apart due to a variety of reasons like clock intrinsic characteristics, ambient temperature, power supply voltage etc. To serve the applications requiring synchronized clocks, a periodic correction mechanism can be applied. The IEEE 1588 standard defines a protocol to perform this correction. The protocol aids by synchronizing individual clocks (Slave Clock) in the system to a Master Clock which is elected to be the least sensitive to the drifts mentioned above, and hence has the best precision. In accordance with this protocol, individual clocks in the system exchange a set of timing messages to select the Master Clock and to keep Slave Clock synchronized to the Master.

The IEEE 1588-2002 standard defines four types of messages: Sync, Follow_Up, Delay_Req, and Delay_Resp. Of these, the Sync, Follow_Up, and Delay_Resp messages are sent from the Master Clock to Slave Clocks in the system while the Delay_Req is sent from the Slave to the Master. The IEEE 1588-2008 standard defines three additional messages to make the system more tolerant to faults; these messages (Pdelay_Req, Pdelay_Resp, and Pdelay_Resp_Follow_Up) are exchanged between adjacent nodes in the system.

With the timing information embedded in the protocol messages, the Slave uses software to calculate the difference between the Master's Clock and its own clock and dynamically adjusts its local clock to synchronize itself to the Master Clock.

Two aspects of the clock must be adjusted. One is the absolute clock time with respect to the protocol epoch, and the other is the frequency of the clock. An IEEE 1588 compliant clock must have the capability to perform both two types of adjustment.

The IEEE 1588 protocol can be implemented on a variety of communication technologies, such as IEEE 802.3(Ethernet) and IEEE 802.11(WLAN). The underlying communication path is responsible for conveying the IEEE 1588 messages among all the clocks. The protocol requires a symmetric communication path delay, meaning the forward and backward path delay between two nodes should be same.

In the protocol messages, some types (including Sync, Delay_Req, and Delay_Resp) are event messages and are required to be timestamped when they arrive or depart a node. This timestamp information will be inserted into messages or used locally by a node for adjustment calculation. To achieve a better Slave-to-Master synchronization performance, the timestamping point should be as close to the communication path as possible. This is where the hardware assistance is required. Due to the indeterministic delay of a node's software system, the software is unable to capture an accurate timestamp when the message is sent or received. However, the hardware is capable of monitoring the signal on the communication media and therefore getting accurate message arrival/departure time.

PTP_TSYNC Module Description of Operation

The PTP_TSYNC (Precision Time Protocol Time Synchronization Engine) module, closely integrated with the Ethernet EMAC module (see [Chapter 22, “Ethernet MAC”](#) for details about the EMAC module), provides hardware assistance to implement both the IEEE 1588-2002 and IEEE 1588-2008 standard on Ethernet (IEEE 802.3). It takes one input clock signal as its PTP clock and automatically does hardware clock adjustment and hardware event message timestamping under software control. With proper settings, it can also output a pulse-per-second (PPS) signal, trigger an alarm interrupt, and timestamp a general-purpose external signal.

[Figure 23-1](#) shows the block diagram of the PTP_TSYNC module.

Clock Source Selection

The PTP_TSYNC module can take any of three input clock sources as its input clock—processor system clock (SCLK), ethernet network clock (MII/RMII TXCLK), or external clock (ECLK).

However, if the clock source is TXCLK or ECLK, the maximum supported frequency is $SCLK/2$.

The selected clock source can be driven out after a clock divider at the processor’s GPIO pin PG[13] given the proper pin-muxing settings. The clock can be driven out without any divider ratio, with a divide-by-2 setting, or with a divide-by-4 setting. However, the maximum supported output clock rate is 33 MHz.

Clock Adjustment

The PTP_TSYNC module applies the Accumulator Addend Method to its input clock to generate the adjustable PTP clock to track the Master Clock. As shown in [Figure 23-1](#), it uses two 32-bit registers (addend and

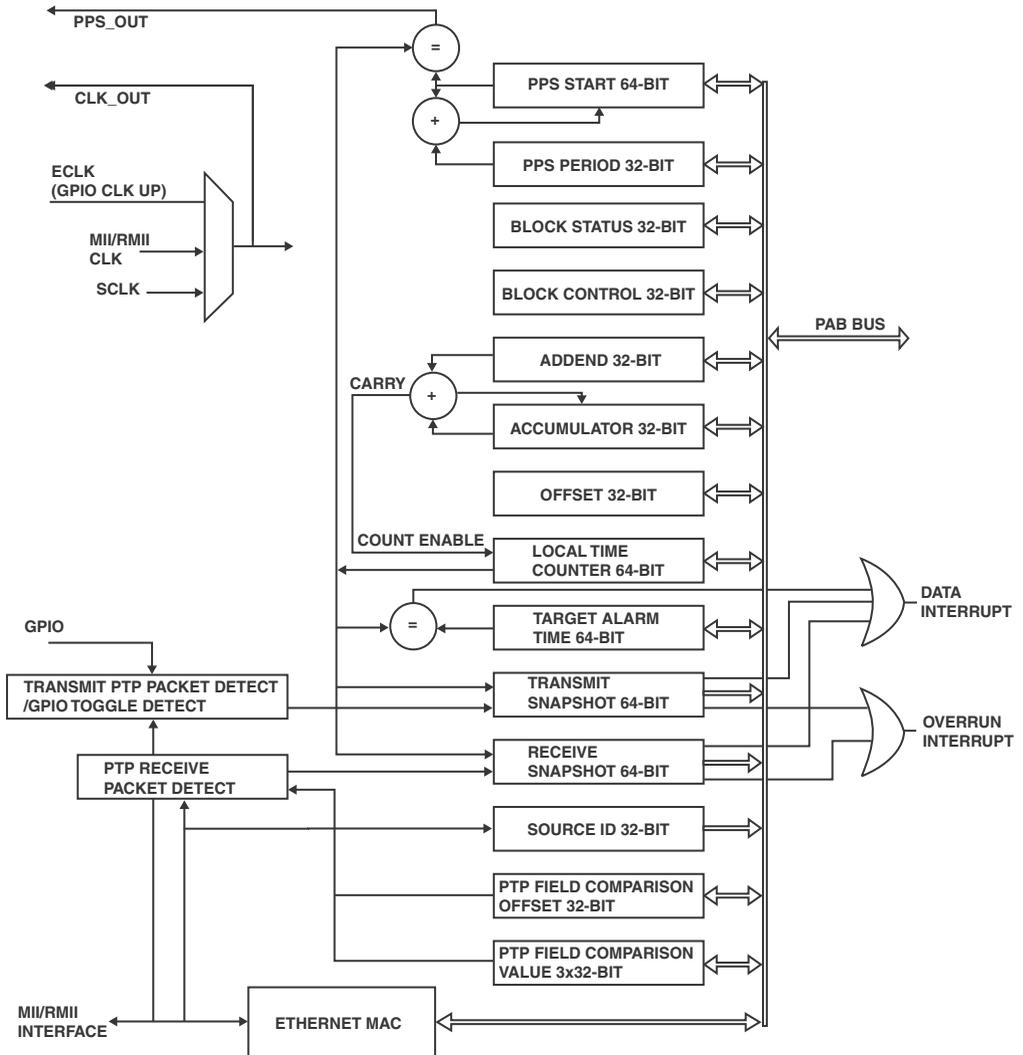


Figure 23-1. PTP_TSYNC Module Block Diagram

accumulator) and one 64-bit counter (local time counter), which stores the adjusted PTP clock time. The content of the addend register is added to the accumulator register at each PTP_TSYNC input clock. In turn, the

General Operation

carry bit of the addition increments the local time counter. The rate by which the local time counter increments is adjusted by changing the addend register value, resulting in the modification of the PTP clock rate. The PTP clock frequency can be calculated by

$$f_{PTP} = \frac{f_{in}}{\left(\frac{2^{32}}{A}\right)}$$

where:

f_{PTP} is the frequency of the adjusted PTP clock

f_{in} is the input clock frequency

A is the value of the addend register

Table 23-1 shows examples of addend settings.

Table 23-1. Examples of Addend Settings

f_{in} (MHz)	A	f_{PTP} (MHz)
80	0xA000 0000	50
40	0x4000 0000	10
33	0x9B26 C9B2	20

In addition to the adjustment of PTP clock's rate, its absolute time, which is the absolute value of local timer counter, can also be adjusted to be the same as the Master Clock's absolute time. The offset between local time and the master time can be programmed into the 32-bit offset register (EMAC_PTP_OFFSET) to compensate for the absolute time difference.

Event Message (Timestamping)

The PTP_TSYNC module automatically monitors all received and transmitted IEEE 1588 event messages on the Ethernet. If an event message is

detected, the module takes a snapshot of the local time counter register and stores its value, which is actually the current local time, to the 64-bit transmit snapshot register (EMAC_PTP_TXSNAPHI and EMAC_PTP_TXSNAPLO) or receive snapshot register (EMAC_PTP_RXSNAPHI and EMAC_PTP_RXSNAPLO). The timestamping is done at the EMAC module's MII/RMII interface when the module sees the Start of Frame of an event message packet. This interface is the closest possible place to the physical Ethernet transmission medium, providing the best timing accuracy.

The PTP_TSYNC modules can be programmed to detect the following types of messages:

- Sync messages sent by the Master and received by a Slave
- Delay_Req messages sent by a Slave and received by the Master
- Pdelay_Req messages sent and received (IEEE 1588-2008 only)
- Pdelay_Resp messages sent and received (IEEE 1588-2008 only)

Transmit Packet Detection

When transmitting an Ethernet packet, the PTP_TSYNC module requires a DMA_Length_Word field associated with the packet. The lower 12 bits of this field are the length of the packet payload in bytes and the higher 4 bits are the timestamping enable field. A value of 0x1 enables the PTP_TSYNC module to timestamp the packet, and a value of 0x0 disables timestamping. When transmitting a PTP sync message, Delay_Req message, or a Pdelay_Req and Pdelay_Resp (for IEEE 1588-2008), this field must be set to 0x1 to enable timestamping. When the packet is sent out, the corresponding timestamp value is saved in the transmit snapshot register for software to retrieve when it is needed.

Receive Packet Detection

The PTP_TSYNC module monitors the received packets on the Ethernet for PTP event messages. It uses up to five fields in the Ethernet packets to determine whether a packet is an event message and requires to be time-stamped. These five fields are described in [Table 23-2](#), and the positions and values of each of these fields within a packet are programmable. The IEEE 1588-2002 implementation over Ethernet IP layer provides default settings for these fields. The PTP_TSYNC module can also be programmed to select which of the five fields are to be used for packet type identification. A packet is identified as a valid event message and time-stamped only when all the selected fields of an incoming packet match the preset match values. The timestamp is saved in the receive snapshot register for software to retrieve when it is needed.

Table 23-2. Default Settings of Matching Fields

Field	Default octet positions after Start of Frame (octet number starts from zero)	Default Match Value of the field	Significance of the default value (IEEE 1588 V1 frame over Ethernet/IPV4/UDP)
MAC frame type	12, 13	0x0800	MAC frame type = IP
IP version	14	0x45	IP version = v4
Layer 4 protocol	23	0x11	Layer 4 protocol = PTP
UDP destination port	36, 37	0x013F	UDP destination port = PTP Event Port
PTP control	74	0x00 or 0x01	Message type: Sync and Delay_Req

For VLAN frames, all field offsets in [Table 23-2](#) are automatically incremented by the PTP_TSYNC module. Specifically, offsets are incremented by 4 for VLAN1 frames and are incremented by 8 for VLAN2 frames.

When the PTP_TSYNC module detects a PTP event message, in addition to the timestamp, it also saves a 32-bit ID of the message into the source

ID snapshot register (EMAC_PTP_ID_SNAP). The ID allows the software to associate a timestamp with the specific message. The 32-bit ID is extracted from the incoming message, based on the offset programmed in the source ID offset register (EMAC_PTP_ID_OFF).

A timestamp overrun may occur when multiple event messages are received and the software has not read the previous event message timestamp from the receive snapshot register before it is overwritten by the timestamp of a second event message. The PTP_TSYNC module provides two options to deal with this situation. One option is timestamp lock, whereby the snapshot register is locked when one timestamp is captured and the PTP_SYNC module does not overwrite it with following message timestamps until the software clears the status bit. If subsequent timestamps are locked out, the module also sets the overrun bit in the interrupt status register (EMAC_PTP_ISTAT) and raises an interrupt. The other option is for the receive snapshot register to be written whenever an event message is detected; this overrun condition is also flagged by the module, setting the overrun bit in EMAC_PTP_ISTAT.

Alarm

The PTP_TSYNC module provides alarm functionality by triggering an alarm at a preset time. It can generate an interrupt and set a bit in the control status register (EMAC_PTP_ISTAT) when the local time (that is, the local time counter register value) matches the value of the target alarm time register. Once an alarm has occurred, if another alarm is needed, the software must clear the status bit and reprogram the target alarm register to a future value. The alarm time is represented in absolute units, not relative units. For example, if the software needs to generate an alarm after 1 second, it must read the current time value, add the number corresponding to 1 second, and write the result back to the target alarm time register.

Pulse-Per-Second (PPS)

Pulse-per-second (PPS) is another physical representation of local time. It is composed of a train of pulses, where each pulse is synchronized to a second transition of local time. In other words, every pulse corresponds to a second boundary. PPS can be used as another synchronization method or to monitor the synchronization performance between clocks. With proper configuration, the PTP_TSYNC module can be programmed to generate this signal at GPIO pin PG[12]. The configuration involves three registers: EMAC_PTP_PPS_STARTHI, EMAC_PTP_PPS_STARTLO, and EMAC_PTP_PPS_PERIOD. The EMAC_PTP_PPS_STARTHI and EMAC_PTP_PPS_STARTLO registers together constitute a 64-bit PPS start time value. If the PPS function is enabled, when the local clock time reaches the PPS start time, a PPS pulse is generated. The EMAC_PTP_PPS_PERIOD value is then automatically added to the PPS start time to create a new PPS start time value for the next pulse. To generate the PPS signal, the PPS start time must be programmed to be equal to the next second value with PPS period set to one second.

Auxiliary Snapshot

In addition to performing timestamping for event messages, the PTP_TSYNC module can also be programmed to timestamp a GPIO toggle. With proper configuration, a low-to-high level change at GPIO PG[14] triggers the module to capture a timestamp and save it in the transmit snapshot register. A corresponding interrupt is raised, and the status bit is set in the control status register. Note that the auxiliary snapshot and the transmit snapshot are mutually exclusive; only one may be enabled at a time.

PTP_TSYNC Module Registers

Table 23-3 summarizes the registers of the PTP_TSYNC module together with their functions, memory-mapped addresses, and access.

Table 23-3. PTP_TSYNC Module Registers

Register Name	Function	Register Address	Register Access
EMAC_PTP_CTL	Control register on page 23-13	0xFFC0 30A0	RW
EMAC_PTP_IE	Interrupt enable register on page 23-15	0xFFC0 30A4	RW
EMAC_PTP_ISTAT	Interrupt status register on page 23-17	0xFFC0 30A8	RW1C
EMAC_PTP_FOFF	Message filter offset register on page 23-18	0xFFC0 30AC	RW
EMAC_PTP_FV1	Message filter value register 1 on page 23-20	0xFFC0 30B0	RW
EMAC_PTP_FV2	Message filter value register 2 on page 23-21	0xFFC0 30B4	RW
EMAC_PTP_FV3	Message filter value register 3 on page 23-21	0xFFC0 30B8	RW
EMAC_PTP_ADDEND	Addend register on page 23-23	0xFFC0 30BC	RW
EMAC_PTP_ACCR	Accumulator register on page 23-24	0xFFC0 30C0	RW
EMAC_PTP_OFFSET	Time offset register on page 23-24	0xFFC0 30C4	RW
EMAC_PTP_TIMELO	Local clock time low on page 23-25	0xFFC0 30C8	RW
EMAC_PTP_TIMEHI	Local clock time high on page 23-26	0xFFC0 30CC	RW
EMAC_PTP_RXSNAPLO	Receive snapshot low register on page 23-27	0xFFC0 30D0	RO

PTP_TSYNC Module Registers

Table 23-3. PTP_TSYNC Module Registers (Continued)

Register Name	Function	Register Address	Register Access
EMAC_PTP_RXSNAPHI	Receive snapshot high register on page 23-28	0xFFC0 30D4	RO
EMAC_PTP_TXSNAPLO	Transmit snapshot low register on page 23-29	0xFFC0 30D8	RO
EMAC_PTP_TXSNAPHI	Transmit snapshot high register on page 23-30	0xFFC0 30DC	RO
EMAC_PTP_ALARMLO	Target alarm time low register on page 23-31	0xFFC0 30E0	RW
EMAC_PTP_ALARMHI	Target alarm time high register on page 23-32	0xFFC0 30E4	RW
EMAC_PTP_ID_OFF	Source ID offset register on page 23-33	0xFFC0 30E8	RW
EMAC_PTP_ID_SNAP	Source ID snapshot register on page 23-34	0xFFC0 30EC	RW
EMAC_PTP_PPS_STARTLO	PPS start low register on page 23-35	0xFFC0 30F0	RW
EMAC_PTP_PPS_STARTHI	PPS start high register on page 23-36	0xFFC0 30F4	RW
EMAC_PTP_PPS_PERIOD	PPS period register on page 23-37	0xFFC0 30F8	RW

Control Register (EMAC_PTP_CTL)

The EMAC_PTP_CTL register, shown in Figure 23-2, controls the overall operation of the PTP_TSYNC module. The function of each bit is described in Table 23-4.

Control Register (EMAC_PTP_CTL)

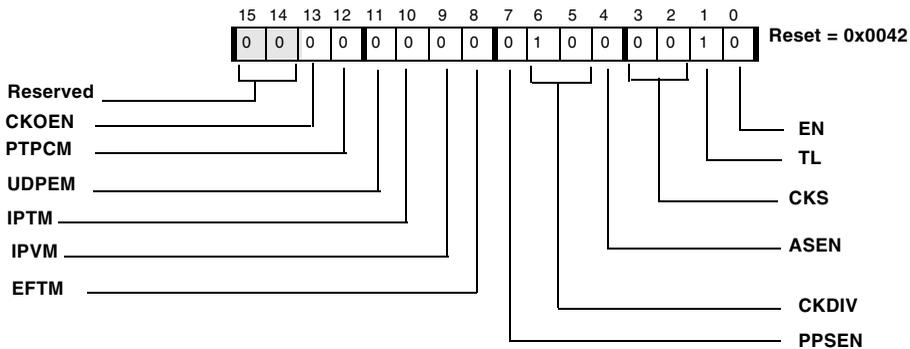


Figure 23-2. Control Register

Table 23-4. EMAC_PTP_CTL Register

Bits	Name	Description	Reset Value	Access
0	EN	0 = Disable the PTP_TSYNC module. 1 = Enable the PTP_TSYNC module.	0	RW
1	TL	Timestamp lock control 0 = The snapshot register is not locked when an event message is detected. The timestamp of any detected event message will overwrite the snapshot register. 1 = The snapshot register is locked when an event message is detected until the software has cleared the corresponding bit in the interrupt status register.	1	RW

PTP_TSYNC Module Registers

Table 23-4. EMAC_PTP_CTL Register (Continued)

Bits	Name	Description	Reset Value	Access
3:2	CKS	PTP clock source 00 = Clock source is peripheral clock (SCLK). 01 = Clock source is TX MII clock (TXCLK). 10 = Clock source is an external clock from GPOI (ECLK). 11 = Reserved clock source selection.	00	RW
4	ASEN	Auxiliary snapshot control 0 = Disable auxiliary snapshot. 1 = Enable auxiliary snapshot.	0	RW
6:5	CKDIV	Divider for the selected PTP_CLK output: 00 = The selected PTP_CLK is driven out without any division. 01 = The selected PTP_CLK is divided by 2 before driving out. 10 = The selected PTP_CLK is divided by 4 before driving out. 11 = Reserved.	10	RW
7	PPSEN	Pulse-per-second (PPS) control 0 = PPS feature is disabled. 1 = PPS feature is enabled.	0	RW
8	EFTM	Compare mask of the Ethernet frame type (EFT) field of EMAC_PTP_FV1 register 0 = Enables compare of the frame type field. 1 = Masks compare of the frame type field.	0	RW
9	IPVM	Compare mask of the IP version (IPV) field of the EMAC_PTP_FV1 register 0 = Enables compare of the IP version field. 1 = Masks compare of the IP version field.	0	RW
10	IPTM	Compare mask of the IP type frame (IPT) field of the EMAC_PTP_FV1 register. 0 = Enables compare of the IP frame type field. 1 = Masks compare of the IP frame type field.	0	RW
11	UDPEM	Compare mask of the UDP event port (UDP_EV) field of the EMAC_PTP_FV2 register. 0 = Enables compare of the UDP event port field. 1 = Masks compare of the UDP event port field.	0	RW

Table 23-4. EMAC_PTP_CTL Register (Continued)

Bits	Name	Description	Reset Value	Access
12	PTPCM	Compare mask of the PTP control field the EMAC_PTP_FV3 register. 0 = Enables compare of the PTP control field. 1 = Masks compare of the PTP control field.	0	RW
13	CKOEN	Clock output control 1 = The output of the clock selected by CKS is enabled on pin PG[13]. The pin mux must be programmed accordingly. 0 = Clock output is disabled.	0	RW
15:14	Reserved		0x0	RV

Interrupt Enable Register (EMAC_PTP_IE)

The EMAC_PTP_IE register, shown in [Figure 23-3](#), controls the interrupt sources that are enabled to generate a core interrupt. The function of each bit is described in [Table 23-5](#).

Interrupt Enable Register (EMAC_PTP_IE)

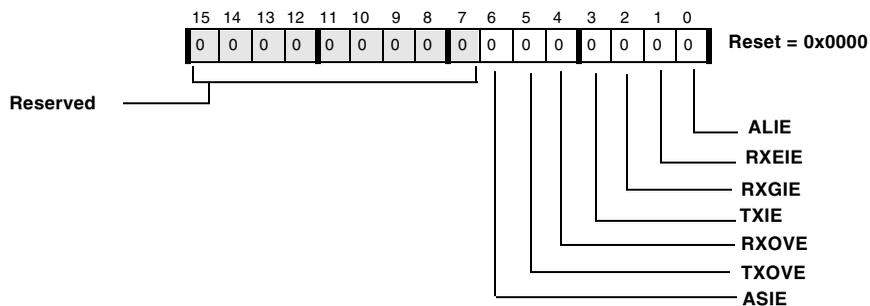


Figure 23-3. Interrupt Enable Register

PTP_TSYNC Module Registers

Table 23-5. EMAC_PTP_IE Register

Bits	Name	Description	Reset Value	Access
0	ALIE	Alarm interrupt enable 0: Disable the alarm interrupt. 1: Enable the alarm interrupt.	0	RW
1	RXEIE	Receive event interrupt enable: 1: Interrupt is enabled. A core interrupt is raised when an event message is detected and time stamped. 0: Interrupt is disabled.	0	RW
2	RXGIE	Receive general interrupt enable: 1: Interrupt is enabled. A core interrupt is raised when a general IEEE 1588 message is detected and timestamped. 0: Interrupt is disabled	0	RW
3	TXIE	Transmit interrupt enable 1: Interrupt is enabled. A core interrupt is raised when a transmit event message is detected and timestamped. 0: Interrupt is disabled	0	RW
4	RXOVE	Receive overrun error interrupt enable 1: Interrupt is enabled. A core interrupt is raised when a receive timestamp overrun condition occurs 0: Interrupt is disabled	0	RW
5	TXOVE	Transmit overrun error interrupt enable 1: Interrupt is enabled. A core interrupt is raised when a transmit timestamp overrun condition occurs 0: Interrupt is disabled.	0	RW
6	ASIE	Auxiliary snapshot interrupt enable 1: Interrupt is enabled. A core interrupt is raised when the auxiliary snapshot is captured. This bit is valid only when the auxiliary snapshot is enabled in the control register. 0: Interrupt is disabled	0	RW
15:7	Reserved	Reserved	0x0	RV

Interrupt Status Register (EMAC_PTP_ISTAT)

The EMAC_PTP_ISTAT register, shown in Figure 23-4, indicates the interrupts that have occurred among the enabled interrupt sources. The function of each bit is described in Table 23-6.

Interrupt Status Register (EMAC_PTP_ISTAT)

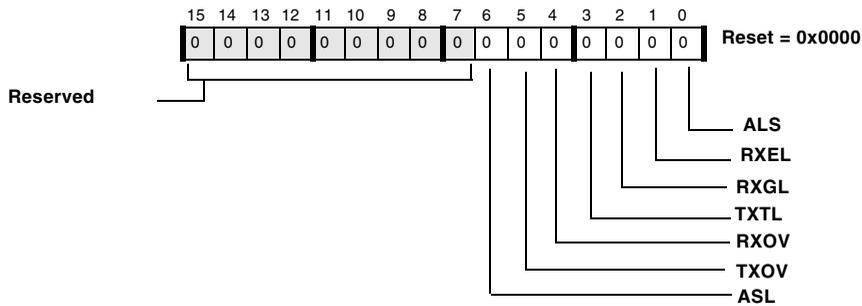


Figure 23-4. Interrupt Status Register

Table 23-6. EMAC_PTP_ISTAT Register

Bits	Name	Description	Reset Value	Access
0	ALS	Alarm status 1: Alarm is triggered. 0: Alarm is not triggered.	0	RW1C
1	RXEL	Receive event interrupt status 1: A receive timestamp is captured in the receive snapshot register. It is cleared when the software reads the receive snapshot register. 0: No receive timestamp is captured.	0	RO
2	RXGL	Receive general interrupt status 1: A general IEEE 1588 message is received. Software must write a “1” to clear this bit. 0: No general IEEE 1588 message is received.	0	RW1C

PTP_TSYNC Module Registers

Table 23-6. EMAC_PTP_ISTAT Register (Continued)

Bits	Name	Description	Reset Value	Access
3	TXTL	Transmit snapshot status: 1: A transmit timestamp has been captured in the transmit snapshot register. It is cleared when the software reads the transmit snapshot register. 0: No transmit timestamp is captured.	0	RO
4	RXOV	Receive snapshot overrun status 1: Receive snapshot overrun occurs. 0: No receive snapshot overrun occurs.	0	RW1C
5	TXOV	Transmit snapshot overrun status: 1: Transmit snapshot overrun occurs. 0: No transmit snapshot overrun occurs.	0	RW1C
6	ASL	Auxiliary snapshot interrupt status: 1: Auxiliary snapshot has been captured in the transmit snapshot register. It is cleared when the software reads the transmit snapshot register. 0: No auxiliary timestamp is captured.	0	RO
15:7	Reserved	Reserved	0x0	RV

Message Filter Offset Register (EMAC_PTP_FOFF)

The EMAC_PTP_FOFF register, shown in [Figure 23-5](#), specifies the location offsets of all the matching fields within an Ethernet frame for detecting PTP event messages. The function of each bit is described in [Table 23-7](#). All offset values are with respect to the start-of-frame (SOF) of Ethernet frames and are the number of octets.

Message Filter Offset Register (EMAC_PTP_FOFF)

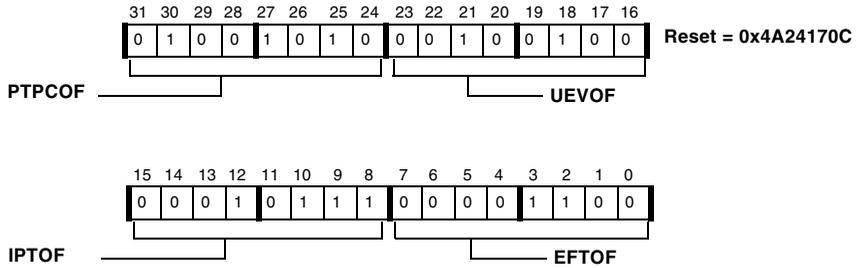


Figure 23-5. Message Filter Offset Register

Table 23-7. EMAC_PTP_FOFF Register

Bits	Name	Description	Reset Value	Access
7:0	EFTOF	Offset of the Ethernet frame type (EFT) field.	0x0C	RW
15:7	IPTOF	Offset of the IP frame type (IPT) field.	0x17	RW
23:16	UEVOF	Offset of the UDP event port (UDP_EVP) field.	0x24	RW
31:24	PTPCOF	Offset of the PTP control (PTPC) field.	0x4A	RW

Message Filter Value Register 1 (EMAC_PTP_FV1)

The EMAC_PTP_FV1 register, shown in Figure 23-6, specifies message filter matching values for three matching fields (Ethernet frame type, IP version, and IP type). The function of each bit is described in Table 23-8.

Message Filter Value Register 1 (EMAC_PTP_FV1)

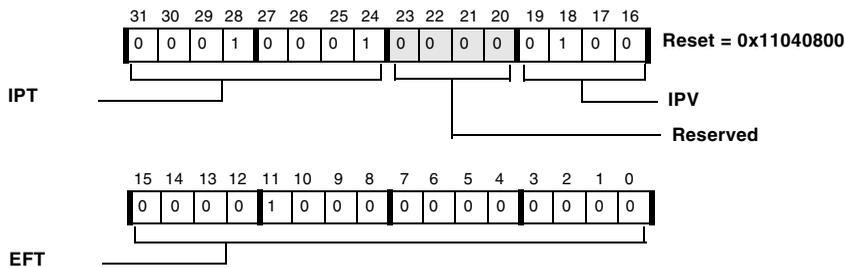


Figure 23-6. Message Filter Value Register 1

Table 23-8. EMAC_PTP_FV1 Register

Bits	Name	Description	Reset Value	Access
15:0	EFT	Ethernet frame type (EFT) matching value	0x0800	RW
19:16	IPV	IP version (IPV) matching value	0x4	RW
23:20	Reserved	Reserved	0x0	RV
31:24	IPT	IP type (IPT) matching value	0x11	RW

Message Filter Value Register 2 (EMAC_PTP_FV2)

The EMAC_PTP_FV2 register, shown in [Figure 23-7](#), specifies message filter matching values for two matching fields (UDP event port and UDP general port). The function of each bit is described in [Table 23-9](#).

Message Filter Value Register 2 (EMAC_PTP_FV2)

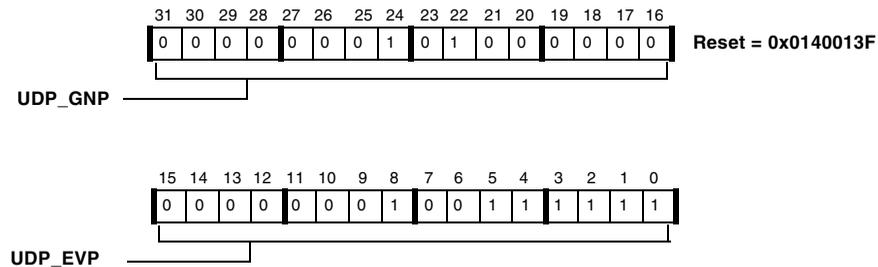


Figure 23-7. Message Filter Value Register 2

Table 23-9. EMAC_PTP_FV2 Register

Bits	Name	Description	Reset Value	Access
15:0	UDP_EVP	UDP event port (UDP_EVP) matching value	0x013F	RW
31:16	UDP_GNP	UDP general port (UDP_GNP) matching value	0x0140	RW

Message Filter Value Register 3 (EMAC_PTP_FV3)

The EMAC_PTP_FV3 register, shown in [Figure 23-8](#), controls how the PTP control field is masked when the PTP_TSYNC module detects event messages and general messages. Each bit in this register corresponds to one PTP control field value.

For event messages, clearing bit N (N = 0 through 15) enables the detection and timestamping of event messages with the PTP control field equal to N.

PTP_TSYNC Module Registers

For general messages, clearing bit N (N = 16 through 31) enables the detection of general messages with the PTP control field equal to (N-16).

The function of each bit is described in [Table 23-10](#).

Message Filter Value Register 3 (EMAC_PTP_FV3)

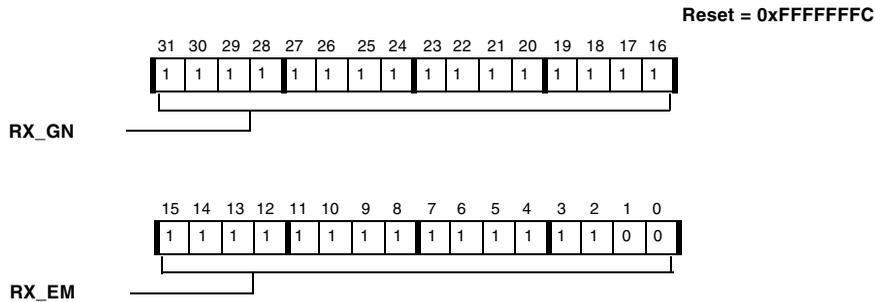


Figure 23-8. Message Filter Value Register 3

Table 23-10. EMAC_PTP_FV3 Register

Bits	Name	Description	Reset Value	Access
15:0	RX_EM	Mask for PTP Control field of event messages	0xFFFC	RW
31:16	RX_GN	Mask for PTP Control field of general messages	0xFFFF	RW

Addend Register (EMAC_PTP_ADDEND)

The EMAC_PTP_ADDEND register, shown in Figure 23-9, specifies the added value for the local clock adjustment. The function of each bit is described in Table 23-11.

Message Filter Value Register 3 (EMAC_PTP_FV3)

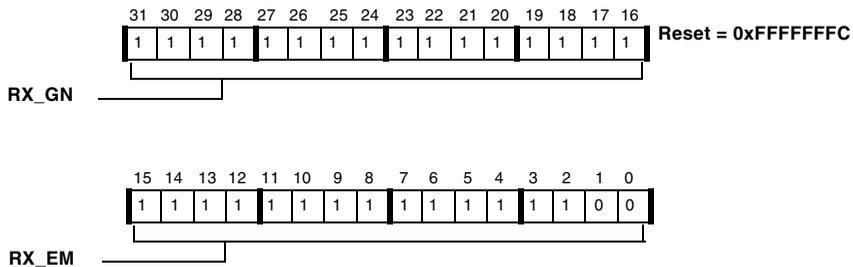


Figure 23-9. Addend Register

Table 23-11. EMAC_PTP_ADDEND Register

Bits	Name	Description	Reset Value	Access
31:0	ADDEND	Addend value for local clock adjustment	0x0	RW

Accumulator Register (EMAC_PTP_ACCR)

The EMAC_PTP_ACCR register, shown in [Figure 23-10](#), holds the accumulator value for the local clock adjustment. The function of each bit is described in [Table 23-12](#).

Accumulator Register (EMAC_PTP_ACCR)

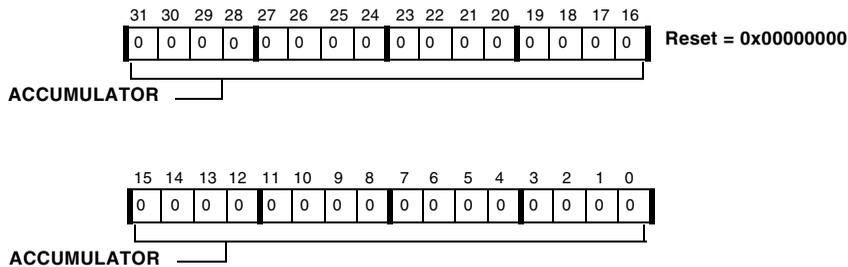


Figure 23-10. Accumulator Register

Table 23-12. EMAC_PTP_ACCR Register

Bits	Name	Description	Reset Value	Access
31:0	ACCUMULATOR	Accumulator value for local clock adjustment	0x0	RW

Time Offset Register (EMAC_PTP_OFFSET)

The EMAC_PTP_OFFSET register, shown in [Figure 23-11](#), is the offset (in two’s complement format) used by the PTP_TSYNC module to correct the local time. A write to this register triggers the PTP_TSYNC module to update the local clock time with the new value of (current local clock time + offset). The function of each bit is described in [Table 23-13](#).

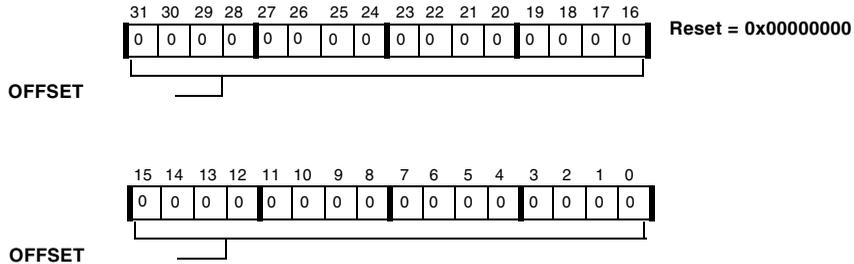
Time Offset Register (EMAC_PTP_OFFSET)

Figure 23-11. Time Offset Register

Table 23-13. EMAC_PTP_OFFSET Register

Bits	Name	Description	Reset Value	Access
31:0	OFFSET	Offset to the local time	0x0	RW

Local Clock Time Low Register (EMAC_PTP_TIMELO)

Two registers (EMAC_PTP_TIMELO and EMAC_PTP_TIMEHI), shown in [Figure 23-12](#) and in [Figure 23-13](#), hold the local adjusted clock time, represented as a 64-bit value spread across two 32-bit registers.

When reading the local time, the lower register (EMAC_PTP_TIMELO) must be read before reading the higher register (EMAC_PTP_TIMEHI). Similarly, when writing the local time, the lower register (EMAC_PTP_TIMELO) must be written before writing the higher register (EMAC_PTP_TIMEHI). This programming sequence is required to maintain synchronization between the two registers.

PTP_TSYNC Module Registers

Local Clock Time Low Register (EMAC_PTP_TIMELO)

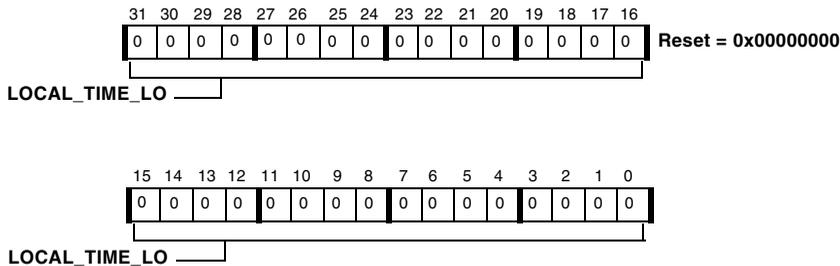


Figure 23-12. Local Clock Time Low Register

Table 23-14. EMAC_PTP_TIMELO Register

Bits	Name	Description	Reset Value	Access
31:0	LOCAL_TIME_LO	Lower 32 bits of the local clock time	0x0	RW

Local Clock Time High Register (EMAC_PTP_TIMEHI)

Two registers (EMAC_PTP_TIMELO and EMAC_PTP_TIMEHI), shown in [Figure 23-12](#) and in [Figure 23-13](#), hold the local adjusted clock time, represented as a 64-bit value spread across two 32-bit registers.

When reading the local time, the lower register (EMAC_PTP_TIMELO) must be read before reading the higher register (EMAC_PTP_TIMEHI). Similarly, when writing the local time, the lower register (EMAC_PTP_TIMELO) must be written before writing the higher register (EMAC_PTP_TIMEHI). This programming sequence is required to maintain synchronization between the two registers.

Local Clock Time High Register (EMAC_PTP_TIMEHI)

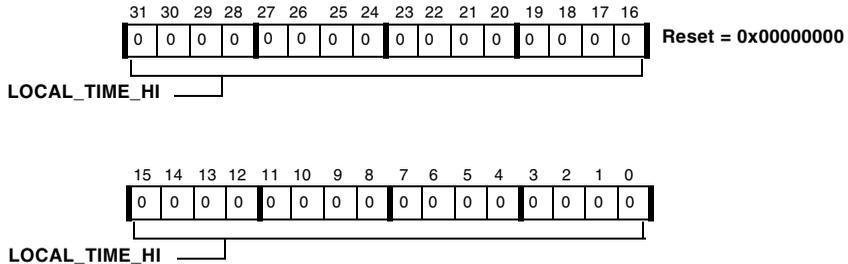


Figure 23-13. Local Clock Time High Register

Table 23-15. EMAC_PTP_TIMEHI Register

Bits	Name	Description	Reset Value	Access
31:0	LOCAL_TIME_HI	Higher 32 bits of the local clock time	0x0	RW

Receive Snapshot Low Register (EMAC_PTP_RXSNAPLO)

Two registers (EMAC_PTP_RXSNAPLO and EMAC_PTP_RXSNAPHI), shown in [Figure 23-14](#) and in [Figure 23-15](#), hold the timestamp of a detected incoming event message. It is a snapshot of the local time clock when the start of frame of the event message is recognized by the PTP_TSYNC module. It is a 64-bit value, spreading across two 32-bit registers. A read of this pair of registers clears the RXTL bit of the interrupt status register (EMAC_PTP_ISTAT).

PTP_TSYNC Module Registers

Receive Snapshot Low Register (EMAC_PTP_RXSNAPLO)

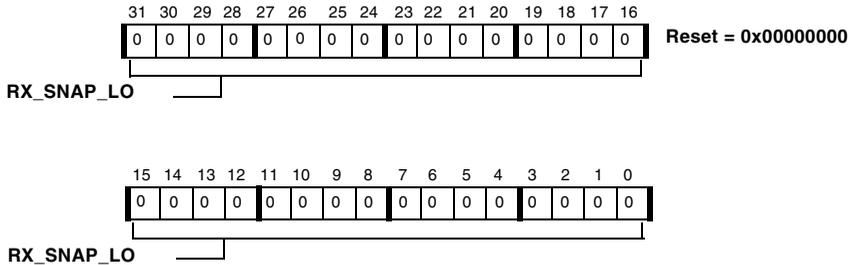


Figure 23-14. Receive Snapshot Low Register

Table 23-16. EMAC_PTP_RXSNAPLO Register

Bits	Name	Description	Reset Value	Access
31:0	RX_SNAP_LO	Lower 32 bits of the receive snapshot of local clock time	0x0	RW

Receive Snapshot High Register (EMAC_PTP_RXSNAPHI)

Two registers (EMAC_PTP_RXSNAPLO and EMAC_PTP_RXSNAPHI), shown in [Figure 23-14](#) and in [Figure 23-15](#), hold the timestamp of a detected incoming event message. It is a snapshot of the local time clock when the start of frame of the event message is recognized by the PTP_TSYNC module. It is a 64-bit value, which spreads across two 32-bit registers. A read of this pair of registers clears the RXTL bit of the interrupt status register (EMAC_PTP_ISTAT).

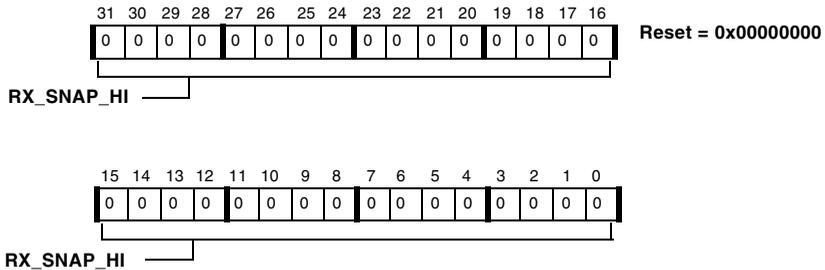
Receive Snapshot High Register (EMAC_PTP_RXSNAPHI)

Figure 23-15. Receive Snapshot High Register

Table 23-17. EMAC_PTP_RXSNAPHI Register

Bits	Name	Description	Reset Value	Access
31:0	RX_SNAP_HI	Higher 32 bits of the receive snapshot of local clock time	0x0	RW

Transmit Snapshot Low Register (EMAC_PTP_TXSNAPLO)

Two registers (EMAC_PTP_TXSNAPLO and EMAC_PTP_TXSNAPHI), shown in [Figure 23-16](#) and in [Figure 23-17](#), hold the timestamp of a detected outgoing event message. It is a snapshot of the local time clock when the start of frame of the event message is recognized by the PTP_TSYNC module. It is a 64-bit value, which spreads across two 32-bit registers. A read of this pair of registers clears the TXTL bit of the interrupt status register (EMAC_PTP_ISTAT).

PTP_TSYNC Module Registers

Transmit Snapshot Low Register (EMAC_PTP_TXSNAPLO)

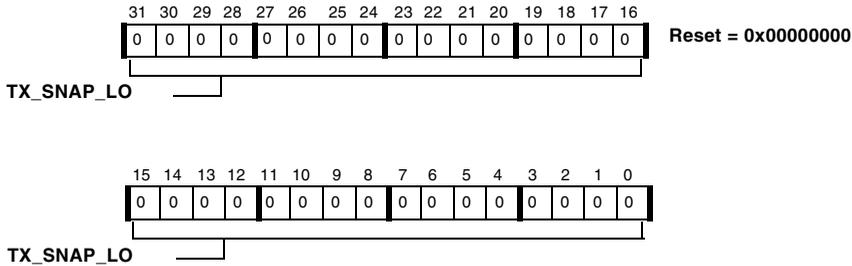


Figure 23-16. Transmit Snapshot Low Register

Table 23-18. EMAC_PTP_TXSNAPLO Register

Bits	Name	Description	Reset Value	Access
31:0	TX_SNAP_LO	Lower 32 bits of the transmit snapshot of local clock time	0x0	RW

Transmit Snapshot High Register (EMAC_PTP_TXSNAPHI)

Two registers (EMAC_PTP_TXSNAPLO and EMAC_PTP_TXSNAPHI), shown in [Figure 23-16](#) and in [Figure 23-17](#), hold the timestamp of a detected outgoing event message. It is a snapshot of the local time clock when the start of frame of the event message is recognized by the PTP_TSYNC module. It is a 64-bit value, which spreads across two 32-bit registers. A read of this pair of registers clears the TXTL bit of the interrupt status register (EMAC_PTP_ISTAT).

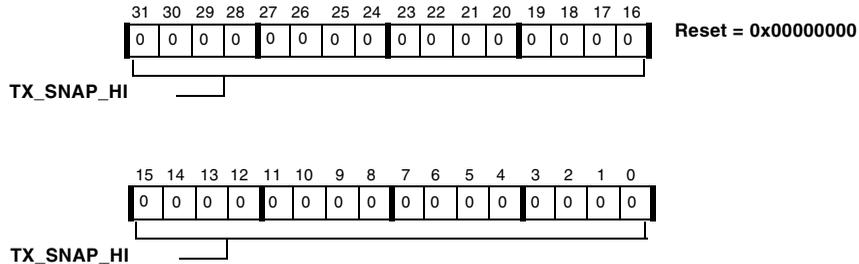
Transmit Snapshot High Register (EMAC_PTP_TXSNAPHI)

Figure 23-17. Transmit Snapshot High Register

Table 23-19. EMAC_PTP_TXSNAPHI Register

Bits	Name	Description	Reset Value	Access
31:0	TX_SNAP_HI	Higher 32 bits of the transmit snapshot of local clock time	0x0	RW

Target Alarm Time Low Register (EMAC_PTP_ALARMLO)

Two registers (EMAC_PTP_ALARMLO and EMAC_PTP_ALARMHI), shown in [Figure 23-18](#) and in [Figure 23-19](#), hold the target alarm time. When the local clock time reaches the target alarm time, the alarm is triggered and an interrupt is raised and the corresponding alarm bit in EMC_PTP_ISTAT is set if the alarm function is enabled. The target alarm time is a 64-bit value, which spreads across two 32-bit registers.

PTP_TSYNC Module Registers

Target Alarm Time Low Register (EMAC_PTP_ALARMLO)

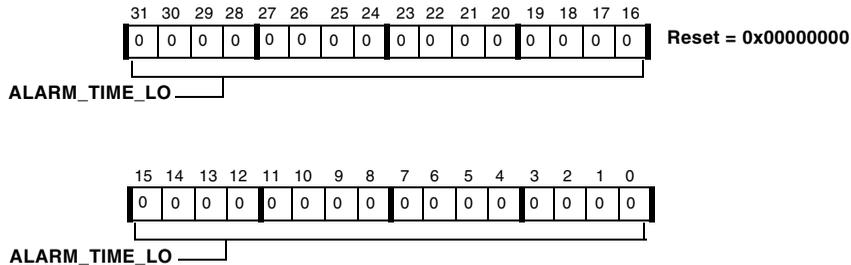


Figure 23-18. Target Alarm Time Low Register

Table 23-20. EMAC_PTP_ALARMLO Register

Bits	Name	Description	Reset Value	Access
31:0	ALARM_TIME_LO	Lower 32 bits of the target alarm time	0x0	RW

Target Alarm Time High Register (EMAC_PTP_ALARMHI)

Two registers (EMAC_PTP_ALARMLO and EMAC_PTP_ALARMHI), shown in [Figure 23-18](#) and in [Figure 23-19](#), hold the target alarm time. When the local clock time reaches the target alarm time, the alarm is triggered and an interrupt is raised and the corresponding alarm bit in EMC_PTP_ISTAT is set if the alarm function is enabled. The target alarm time is a 64-bit value, and spreads across two 32-bit registers.

Target Alarm Time High Register (EMAC_PTP_ALARMHI)

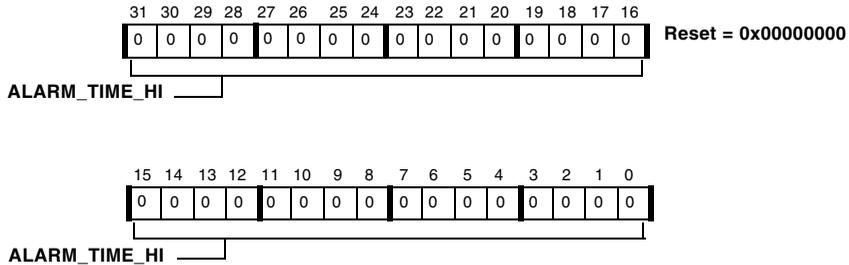


Figure 23-19. Target Alarm Time High Register

Table 23-21. EMAC_PTP_ALARMHI Register

Bits	Name	Description	Reset Value	Access
31:0	ALARM_TIME_HI	Higher 32 bits of the target alarm time	0x0	RW

Source ID Offset Register (EMAC_PTP_ID_OFF)

The EMAC_PTP_ID_OFF register, shown in [Figure 23-20](#), specifies the location offset of the source ID of an IEEE 1588 message within an Ethernet packet. The PTP_TSYNC module uses this value to extract the source ID of a detected incoming message. The function of each bit is described in [Table 23-22](#).

Source ID Offset Register (EMAC_PTP_ID_OFF)

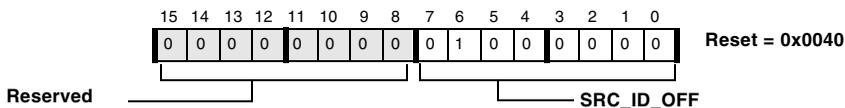


Figure 23-20. Source ID Offset Register

PTP_TSYNC Module Registers

Table 23-22. EMAC_PTP_ID_OFF Register

Bits	Name	Description	Reset Value	Access
7:0	SRC_ID_OFF	Source ID offset	0x40	RW
15:8	Reserved	Reserved	0x0	RV

Source ID Snapshot Register (EMAC_PTP_ID_SNAP)

The EMAC_PTP_ID_SNAP register, shown in [Figure 23-21](#), holds the source ID extracted from a timestamped event message. It is used by software to correlate a timestamp to the received message to which it belongs in case multiple incoming messages are received in the software's buffer. The function of each bit is described in [Table 23-23](#).

Source ID Snapshot Register (EMAC_PTP_ID_SNAP)

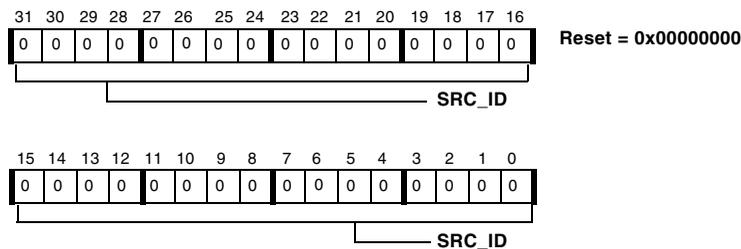


Figure 23-21. Source ID Snapshot Register

Table 23-23. EMAC_PTP_ID_SNAP Register

Bits	Name	Description	Reset Value	Access
31:0	SRC_ID	Source ID	0x0	RW

PPS Start Low Register (EMAC_PTP_PPS_STARTLO)

Two registers (EMAC_PTP_PPS_STARTLO and EMAC_PTP_PPS_STARTHI), shown in Figure 23-22 and Figure 23-23, define the PPS start time value, which controls how the pulse-per-second (PPS) signal is generated.

PPS Start Low Register (EMAC_PTP_PPS_STARTLO)

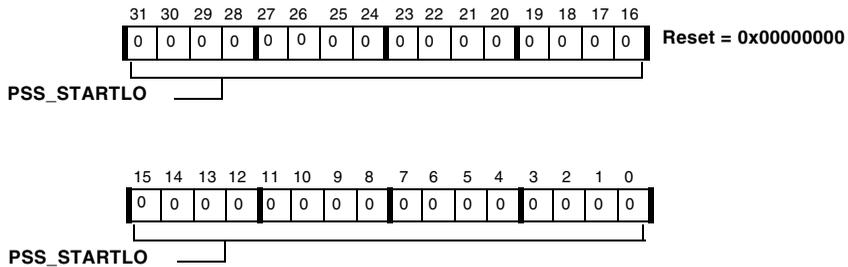


Figure 23-22. PPS Start Low Register

Table 23-24. EMAC_PTP_PPS_STARTLO Register

Bits	Name	Description	Reset Value	Access
31:0	PPS_STARTLO	Lower 32 bits of the PPS start time	0x0	RW

PPS Start High Register (EMAC_PTP_PPS_STARTHI)

Two registers (EMAC_PTP_PPS_STARTLO and EMAC_PTP_PPS_STARTHI), shown in Figure 23-22 and Figure 23-23, define the PPS start time value, which controls how the pulse-per-second (PPS) signal is generated.

PSS Start High Register (EMAC_PTP_PSS_STARTHI)

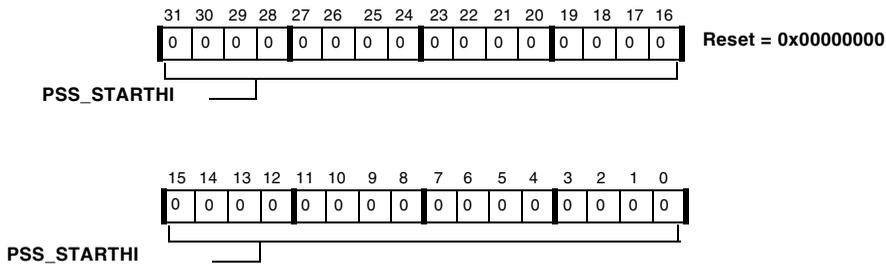


Figure 23-23. PSS Start High Register

Table 23-25. EMAC_PTP_PPS_STARTHI Register

Bits	Name	Description	Reset Value	Access
31:0	PPS_STARTHI	Higher 32 bits of the PPS start time	0x0	RW

PPS Period Register (EMAC_PTP_PPS_PERIOD)

The EMAC_PTP_PPS_PERIOD register, shown in [Figure 23-24](#), specifies the period of PPS signal, which is in number of adjusted local clock cycles. It is added to the PPS start time after each PPS pulse.

PSS Period Register (EMAC_PTP_PSS_PERIOD)

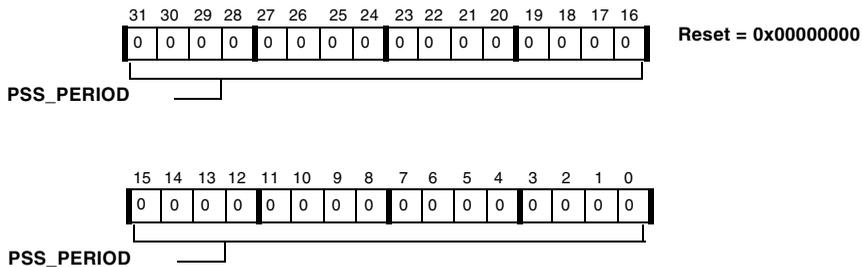


Figure 23-24. PSS Period Register

Table 23-26. EMAC_PTP_PPS_PERIOD Register

Bits	Name	Description	Reset Value	Access
31:0	PPS_PERIOD	PPS signal period	0x0	RW

PTP_TSYNC Module Programming Model

The following sections describe how to configure the PTP_TSYNC module for typical uses:

- [“IEEE 1588-2002 Implementation Over IP/UDP”](#) on page 23-38
- [“IEEE 1588-2008 Implementation Over IP/UDP”](#) on page 23-39

PTP_TSYNC Module Programming Model

- [“IEEE 1588-2008 Implementation Over MAC Layer” on page 23-39](#)
- [“Pulse-Per-Second \(PPS\) Signal Generation” on page 23-40](#)

IEEE 1588-2002 Implementation Over IP/UDP

For IEEE 1588-2002 V1 implementation over IP/UDP, most of the default values of the registers can be left unchanged. The detailed programming steps are:

1. Clear the five comparison mask bits (bits[12:8]) in `EMAC_PTP_CTL` to enable all the field matches.
2. Keep the default values of the `EMAC_PTP_FOFF` register.
3. Keep the default values of the `EMAC_PTP_FV1` and `EMAC_PTP_FV2` registers.
4. Depending on the event messages to be timestamped, clear the appropriate bits among bits[15:0] of the `EMAC_PTP_FV3` register. The default value (0xFFFC) allows the timestamping of both received Sync messages and Delay_Req messages.
5. Depending on the general messages to be detected, clear the appropriate bits among bits[31:16] of the `EMAC_PTP_FV3` register. The default value (0xFFFF) allows no detection of received general messages.

IEEE 1588-2008 Implementation Over IP/UDP

The detailed programming steps are:

1. Clear all five comparison mask bits (bits[12:8]) in the `EMAC_PTP_CTL` register to enable all the field matches.
2. Keep the default values of the `EMAC_PTP_FOFF` register, except set the `PTPCOF` field to `0x2A`.
3. Keep the default values of the `EMAC_PTP_FV1` and the `EMAC_PTP_FV2` registers.
4. Depending on the event messages to be timestamped, clear the appropriate bits among bits[15:0] of the `EMAC_PTP_FV3` register. The default value (`0xFFFC`) allows the timestamping of both received Sync messages and Delay_Req messages. To allow the timestamping of Pdelay_Req and Pdelay_Resp, set the value to `0xFFF0`.
5. Depending on the general messages to be detected, clear the appropriate bits among the bits[31:16] of the `EMAC_PTP_FV3` register. The default value (`0xFFFF`) allows no detection of received general messages.

IEEE 1588-2008 Implementation Over MAC Layer

The detailed programming steps are:

1. Clear bits 8 and 12 of the `EMAC_PTP_CTL` register to enable only the `EFTM` and `PTPCM` field comparison.
2. Keep the default values of all the fields of the `EMAC_PTP_FOFF` register, except set the `PTPCOF` field to `0x0E`.
3. Program bits [15:0] of the `EMAC_PTP_FV1` register to `0x88F7`, which corresponds to PTP messages on the MAC layer.

PTP_TSYNC Module Programming Model

4. Keep the default values for all other fields of the `EMAC_PTP_FV1` and `EMAC_PTP_FV2` registers.
5. Depending on the event messages to be timestamped, clear the appropriate bits among bits[15:0] of the `EMAC_PTP_FV3` register. The default value (0xFFFC) allows the timestamping of both received Sync messages and Delay_Req messages.

To allow the timestamping of Pdelay_Req and Pdelay_Resp messages, set the value to 0xFFF0.

6. Depending on the general messages to be detected, clear the appropriate bits among bits[31:16] of the `EMAC_PTP_FV3` register. The default value (0xFFFF) allows no detection of received general messages.

Pulse-Per-Second (PPS) Signal Generation

The detailed programming steps are:

1. Enable the `PPSEN` bit of the `EMAC_PTP_CTL` register to enable the PPS function.
2. Program the number of adjusted local clock cycles corresponding to 1 second in the `EMAC_PTP_PPS_PERIOD` register.
3. Program the time of the first pulse in the `EMAC_PTP_PPS_STARTLO` and `EMAC_PTP_PPS_STARTHI` registers.
4. After the first pulse, if the position of the next pulse must be advanced or delayed, re-program `EMAC_PTP_PPS_STARTLO` and `EMAC_PTP_PPS_STARTHI` registers accordingly.
5. If the period of the pulses needs to be changed, modify the `EMAC_PTP_PPS_PERIOD` register value accordingly.

24 REAL-TIME CLOCK

This chapter describes the real-time clock (RTC). Following an overview and list of key features is a description of operation. The chapter concludes with a programming model, consolidated register definitions, and programming examples.

Specific Information for the ADSP-BF51x

For RTC interrupt vector assignments, refer to [Table 5-3 on page 5-19](#) in [Chapter 5, “System Interrupts”](#).

For a list of MMR addresses for the RTC, refer to [Appendix A, “System MMR Assignments”](#).

RTC behavior for the ADSP-BF51x that differs from the general information in this chapter can be found at the end of this chapter in the section [“Unique Information for the ADSP-BF51x Processor” on page 24-27](#).

Overview

The RTC provides a set of digital watch features to the processor, including time of day, alarm, and stopwatch countdown. It is typically used to implement either a real-time watch or a life counter, which counts the elapsed time since the last system reset.

The RTC watch features are clocked by a 32.768 kHz crystal external to the processor. The RTC uses dedicated power supply pins and is

Overview

independent of any reset, which enables it to maintain functionality even when the rest of the processor is powered down.

The RTC input clock is divided down to a 1 Hz signal by a prescaler, which can be bypassed. When bypassed, the RTC is clocked at the 32.768 kHz crystal rate. In normal operation, the prescaler is enabled.

The primary function of the RTC is to maintain an accurate day count and time of day. The RTC accomplishes this by means of four counters:

- 60 second counter
- 60 minute counter
- 24 hour counter
- 32768 day counter

The RTC increments the 60 second counter once per second and increments the other three counters when appropriate. The 32768 day counter is incremented each day at midnight (0 hours, 0 minutes, 0 seconds). Interrupts can be issued periodically, either every second, every minute, every hour, or every day. Each of these interrupts can be independently controlled.

The RTC provides two alarm features, programmed with the `RTC_ALARM` register. The first is a time of day alarm (hour, minute, and second). When the alarm interrupt is enabled, the RTC generates an interrupt each day at the time specified. The second alarm feature allows the application to specify a day as well as a time. When the day alarm interrupt is enabled, the RTC generates an interrupt on the day and time specified. The alarm interrupt and day alarm interrupt can be enabled or disabled independently.

The RTC provides a stopwatch function that acts as a countdown timer. The application can program a second count into the RTC stopwatch count register (`RTC_SWCNT`). When the stopwatch interrupt is enabled and

the specified number of seconds have elapsed, the RTC generates an interrupt.

Interface Overview

The RTC external interface consists of two clock pins, which together with the external components form the reference clock circuit for the RTC. The RTC interfaces internally to the processor system through the peripheral access bus (PAB), and through the interrupt interface to the system interrupt controller (SIC).

The RTC has dedicated power supply pins that power the clock functions at all times, including when the core power supply is turned off.

[Figure 24-1](#) provides a block diagram of the RTC.

Description of Operation

The following sections describe the operation of the RTC.

RTC Clock Requirements

The RTC timer is clocked by a 32.768 kHz crystal external to the processor. The RTC system memory mapped registers (MMRs) are clocked by this crystal. When the prescaler is disabled, the RTC MMRs are clocked at the 32.768 kHz crystal frequency. When the prescaler is enabled, the RTC MMRs are clocked at the 1 Hz rate.

There is no way to disable the RTC counters using software. If a given system does not require the RTC functionality, then it may be disabled with hardware tie offs. Tie the `RTXI` and `RTCGND` pins to `EGND`, tie the `RTCVD` pin to `EVDD`, and leave the `RTX0` pin unconnected. Additionally, writing `RTC_PREN` to “0” saves a small amount of power.

Description of Operation

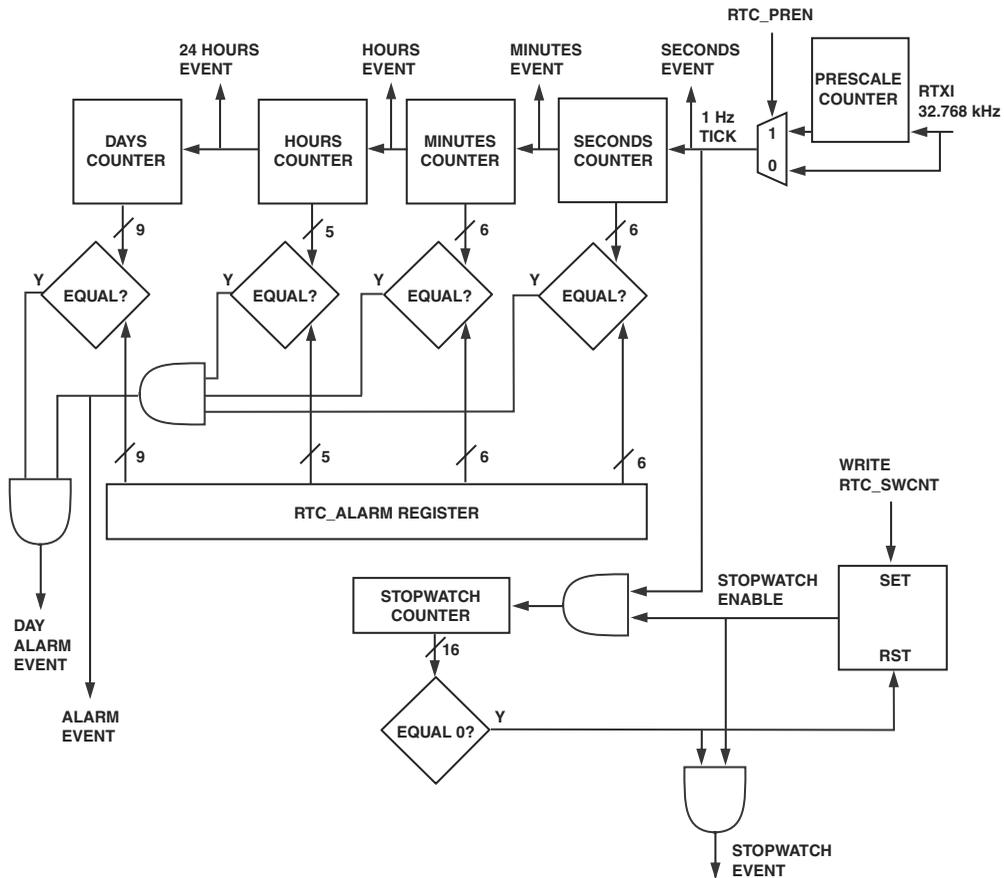


Figure 24-1. RTC Block Diagram

Prescaler Enable

The single active bit of the RTC prescaler enable register (`RTC_PREN`) is written using a synchronization path. Clearing of the bit is synchronized to the 32.768 kHz clock. This faster synchronization allows the module to be put into high speed mode (bypassing the prescaler) without waiting the full one second for the write to complete that would be necessary if the module were already running with the prescaler enabled. When this bit is

cleared, the prescaler is disabled, and the RTC runs at the 32.768 kHz crystal frequency.

When setting the `RTC_PREN` bit, the first positive edge of the 1 Hz clock occurs 1 to 2 cycles of the 32.768 kHz clock after the prescaler is enabled. The write complete status/interrupt works as usual when enabling or disabling the prescale counter. The new RTC clock rate is in effect before the write complete status is set. In order for the RTC to operate at the proper rate, software must set the prescaler enable bit after initial powerup. When this bit is set, the prescaler is enabled, and the RTC runs at a frequency of 1 Hz.

Write `RTC_PREN` and then wait for the write complete event before programming the other registers. It is safe to write `RTC_PREN` to “1” every time the processor boots. The first time sets the bit, and subsequent writes have no effect, as no state is changed.

 Do not disable the prescaler by clearing the bit in the `RTC_PREN` register without making sure that there are no writes to RTC MMRs in progress. Do not switch between fast and slow mode during normal operation by setting and clearing this bit, as this disrupts the accurate tracking of real time by the counters. To avoid these potential errors, initialize the RTC during startup using `RTC_PREN` and do not dynamically alter the state of the prescaler during normal operation.

Running without the prescaler enabled is provided primarily as a test mode. All functionality works, just 32,768 times as fast. Typical software should never program `RTC_PREN` to “0”. The only reason to do so is to synchronize the 1 Hz tick to a more precise external event, as the 1 Hz tick predictably occurs a few `RTXI` cycles after a 0-to-1 transition of `RTC_PREN`.

Use the following sequence to achieve synchronization to within 100 ms.

1. Write `RTC_PREN` to “0”.
2. Wait for the write to complete.

RTC Programming Model

3. Wait for the external event.
4. Write `RTC_PREN` to “1”.
5. Wait for the write to complete.
6. Reprogram the time into `RTC_STAT`.

RTC Programming Model

The RTC programming model consists of a set of system MMRs. Software can configure the RTC and can determine the status of the RTC through reads and writes to these registers. The RTC interrupt control register (`RTC_ICTL`) and the RTC interrupt status register (`RTC_ISTAT`) provide RTC interrupt management capability.

Note that software cannot disable the RTC counting function. However, all RTC interrupts can be disabled, or masked. At reset, all interrupts are disabled. The RTC state can be read via the system MMR status registers at any time.

The primary RTC functionality, shown in [Figure 24-1 on page 24-4](#), consists of registers and counters that are powered by an independent RTC supply (V_{DDRTC}). This logic is never reset; it comes up in an indeterminate state when V_{DDRTC} is first powered on.

The RTC also contains logic powered by the same internal V_{dd} as the processor core and other peripherals. This logic contains some control functionality, holding registers for PAB write data, and prefetched PAB read data shadow registers for each of the five V_{DDRTC} -powered registers. This logic is reset by the same system reset and clocked by the same `SCLK` as the other peripherals.

[Figure 24-2](#) shows the connections between the V_{DDRTC} -powered RTC MMRs and their corresponding V_{DDINT} -powered write holding registers and read shadow registers. In the figure, “REG” means each of the

RTC_STAT, RTC_ALARM, RTC_SWCNT, RTC_ICTL, and RTC_PREN registers. The RTC_ISTAT register connects only to the PAB.

The rising edge of the 1 Hz RTC clock is the “1 Hz tick”. Software can synchronize to the 1 Hz tick by waiting for the seconds event flag to set or by waiting for the seconds interrupt (if enabled).

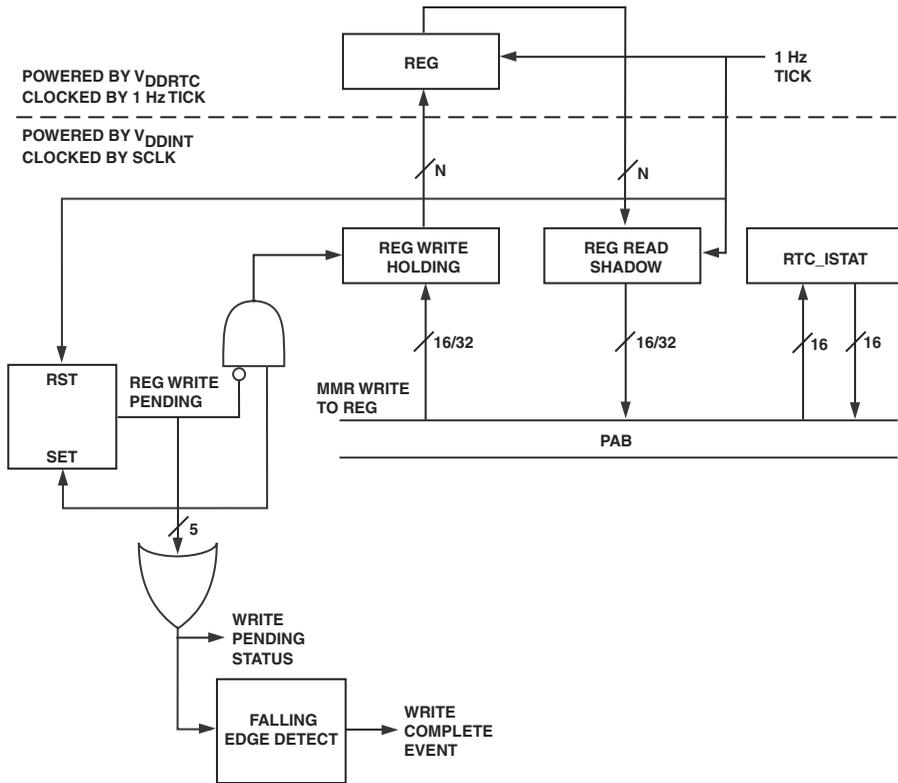


Figure 24-2. RTC Register Architecture

Register Writes

Writes to all RTC MMRs except `RTC_ISTAT` are saved in write holding registers and then are synchronized to the RTC 1 Hz clock. The write pending status bit in `RTC_ISTAT` indicates the progress of the write. The write pending status bit is set when a write is initiated and is cleared when all writes are complete. The falling edge of the write pending status bit causes the write complete flag in `RTC_ISTAT` to be set. This flag can be configured in `RTC_ICTL` to cause an interrupt. Software does not have to wait for writes to one RTC MMR to complete before writing to another RTC MMR. The write pending status bit is set if any writes are in progress, and the write complete flag is set only when all writes are complete.

-  Any writes in progress when peripherals are reset are aborted. Do not stop `SCLK` (for example, by entering deep sleep mode) or remove internal V_{dd} power until all RTC writes have completed.
-  Do not attempt another write to the same register without waiting for the previous write to complete. Subsequent writes to the same register are ignored if the previous write is not complete.
-  Reading a register that has been written before the write complete flag in `RTC_ISTAT` is set will return the old value. Always check the write pending status bit in `RTC_ISTAT` before attempting a read or write.

Write Latency

Writes to the RTC MMRs are synchronized to the 1 Hz RTC clock. When setting the time of day, do not factor in the delay when writing to the RTC MMRs. The most accurate method of setting the RTC is to monitor the seconds (1 Hz) event flag or to program an interrupt for this event and then write the current time to `RTC_STAT` in the interrupt service routine (ISR). The new value is inserted ahead of the incrementer. Hardware adds one second to the written value (with appropriate carries into

minutes, hours and days) and loads the incremented value at the next 1 Hz tick, when it represents the then-current time.

Writes posted at any time are properly synchronized to the 1 Hz clock. Writes complete at the rising edge of the 1 Hz clock. A write posted just before the 1 Hz tick may not be completed until the 1 Hz tick one second later. Any write posted in the first 990 ms after a 1 Hz tick completes on the next 1 Hz tick, but the simplest, most predictable and recommended technique is to only post writes to `RTC_STAT`, `RTC_ALARM`, `RTC_SWCNT`, `RTC_ICTL`, or `RTC_PREN` immediately after a seconds interrupt or event. All five registers may be written in the same second.

W1C bits in the `RTC_ISTAT` register take effect immediately.

Register Reads

There is no latency when reading RTC MMRs, as the values come from the read shadow registers. These shadow registers are updated and ready for reading by the time any RTC interrupts or event flags for that second are asserted. Once the internal V_{dd} logic completes its initialization sequence after `SCLK` starts, there is no point in time when it is unsafe to read the RTC MMRs for synchronization reasons. They always return coherent values, although the values may be indeterminate.

Deep Sleep

When the dynamic power management mode is set to deep sleep, all clocks in the system (except `RTXI` and the RTC 1 Hz tick) are stopped. In this state, the V_{DDRTC} -powered counters continue to increment. The internal V_{dd} shadow registers are not updated, but neither can they be read.

During deep sleep mode, all bits in `RTC_ISTAT` are cleared. Events that occur during deep sleep are not recorded in `RTC_ISTAT`. The internal V_{dd} RTC control logic generates a virtual 1 Hz tick within one `RTXI` period

RTC Programming Model

(30.52 μ s) after `SCLK` restarts. This loads all shadow registers with up-to-date values and sets the seconds event flag. Other event flags may also be set. When the system wakes up from deep sleep, whether by an RTC event or a hardware reset, all of the RTC events that occurred during that second (and only that second) are reported in `RTC_ISTAT`.

When the system wakes up from deep sleep mode, software does not need to `W1C` the bits in `RTC_ISTAT`. All `W1C` bits are already cleared by hardware. The seconds event flag is set when the RTC internal V_{dd} logic has completed its restart sequence. Software should wait until the seconds event flag is set and then may begin reading or writing any RTC register.

Event Flags



The indeterminate values in the registers at power-up can cause event flags to set before the correct value is written into each of the registers. By catching the 1 Hz clock edge, the write to `RTC_STAT` can occur a full second before the write to `RTC_ALARM`. This would cause an extra second of delay between the validity of `RTC_STAT` and `RTC_ALARM`, if the value of the `RTC_ALARM` out of reset is the same as the value written to `RTC_STAT`. Therefore, wait for the writes to complete on these registers before using the flags and interrupts associated with their values.

The following is a list of flags along with the conditions under which they are valid:

- Seconds (1 Hz) event flag
Always set on the positive edge of the 1 Hz clock and after shadow registers have updated after waking from deep sleep. This is valid as long as the RTC 1 Hz clock is running. Use this flag or interrupt to validate the other flags.
- Write complete and write pending status
Always valid.

- **Minutes event flag**
Valid only after the second field in `RTC_STAT` is valid. Use the write complete and write pending status flags or interrupts to validate the `RTC_STAT` value before using this flag value or enabling the interrupt.
- **Hours event flag**
Valid only after the minute field in `RTC_STAT` is valid. Use the write complete and write pending status flags or interrupts to validate the `RTC_STAT` value before using this flag value or enabling the interrupt.
- **24 Hours event flag**
Valid only after the hour field in `RTC_STAT` is valid. Use the write complete and write pending status flags or interrupts to validate the `RTC_STAT` value before using this flag value or enabling the interrupt.
- **Stopwatch event flag**
Valid only after the `RTC_SWCNT` register is valid. Use the write complete and write pending status flags or interrupts to validate the `RTC_SWCNT` value before using this flag value or enabling the interrupt.
- **Alarm event and day alarm event flags**
Valid only after the `RTC_STAT` and `RTC_ALARM` registers are valid. Use the write complete and write pending status flags or interrupts to validate the `RTC_STAT` and `RTC_ALARM` values before using this flag value or enabling its interrupt.

RTC Programming Model

Writes posted together at the beginning of the same second take effect together at the next 1 Hz tick. The following sequence is safe and does not result in any spurious interrupts from a previous state.

1. Wait for 1 Hz tick.
2. Write 1s to clear the `RTC_ISTAT` flags for alarm, day alarm, stopwatch, and/or per interval.
3. Write new values for `RTC_STAT`, `RTC_ALARM`, and/or `RTC_SWCNT`.
4. Write new value for `RTC_ICTL` with alarm, day alarm, stopwatch, and/or per interval interrupts enabled.
5. Wait for 1 Hz tick.
6. New values have now taken effect simultaneously.

Setting Time of Day

The `RTC_STAT` register is used to read or write the current time. Reads return a 32-bit value that always reflects the current state of the days, hours, minutes, and seconds counters. Reads and writes must be 32-bit transactions; attempted 16-bit transactions result in an MMR error. Reads always return a coherent 32-bit value. The hours, minutes, and seconds fields are usually set to match the real time of day. The day counter value is incremented every day at midnight to record how many days have elapsed since it was last modified. Its value does not correspond to a particular calendar day. The 15-bit day counter provides a range of 89 years, 260 or 261 days (depending on leap years) before it overflows.

After the 1 Hz tick, program `RTC_STAT` with the current time. At the next 1 Hz tick, `RTC_STAT` takes on the new, incremented value. For example:

1. Wait for 1 Hz tick.
2. Read `RTC_STAT`, get 10:45:30.

3. Write `RTC_STAT` to current time, 13:10:59.
4. Read `RTC_STAT`, still get old time 10:45:30.
5. Wait for 1 Hz tick.
6. Read `RTC_STAT`, get new current time, 13:11:00.

Using the Stopwatch

The `RTC_SWCNT` register contains the countdown value for the stopwatch. The stopwatch counts down seconds from the programmed value and generates an interrupt (if enabled) when the count reaches “0”. The counter stops counting at this point and does not resume counting until a new value is written to `RTC_SWCNT`. Once running, the counter may be overwritten with a new value. This allows the stopwatch to be used as a watchdog timer with a precision of one second.

The stopwatch can be programmed to any value between 0 and $(2^{16} - 1)$ seconds, which is a range of 18 hours, 12 minutes, and 15 seconds.

Typically, software should wait for a 1 Hz tick, then write `RTC_SWCNT`. One second later, `RTC_SWCNT` changes to the new value and begins decrementing. Because the register write occupies nearly one second, the time from writing a value of N until the stopwatch interrupt is nearly $N + 1$ seconds. To produce an exact delay, software can compensate by writing $N - 1$ to get a delay of nearly N seconds. This implies that you cannot achieve a delay of 1 second with the stopwatch. Writing a value of “1” immediately after a 1 Hz tick results in a stopwatch interrupt nearly two seconds later. To wait one second, software should just wait for the next 1 Hz tick.

The `RTC_SWCNT` register is not reset. After initial powerup, it may be running. When the stopwatch is not used, writing it to “0” to force it to stop saves a small amount of power.

Interrupts

The RTC can provide interrupts at several programmable intervals:

- Per second, minute, hour, and day—based on increments to the respective counters in `RTC_STAT`
- On countdown from a programmable value—value in `RTC_SWCNT` transitions to “0” or is written with “0” by software (whether it was previously running or already stopped with a count of “0”)
- Daily at a specific time—all fields of `RTC_ALARM` must match `RTC_STAT` except the day field
- On a specific day and time—all fields of `RTC_ALARM` register must match `RTC_STAT`

The RTC can be programmed to provide an interrupt at the completion of all pending writes to any of the 1 Hz registers (`RTC_STAT`, `RTC_ALARM`, `RTC_SWCNT`, `RTC_ICTL`, and `RTC_PREN`). The eight RTC interrupt events can be individually masked or enabled by the `RTC_ICTL` register. The seconds interrupt is generated on each 1 Hz clock tick, if enabled. The minutes interrupt is generated at the 1 Hz clock tick that advances the seconds counter from 59 to 0. The hour interrupt is generated at the 1 Hz clock tick that advances the minute counter from 59 to 0. The 24 hour interrupt occurs once per 24 hour period at the 1 Hz clock tick that advances the time to midnight (00:00:00). Any of these interrupts can generate a wakeup request to the processor, if enabled. All implemented bits are read/write.

This register is only partially cleared at reset, so some events may appear to be enabled initially. However, the RTC interrupt and the RTC wakeup to the PLL are handled specially and are masked (forced low) until after the first write to the `RTC_ICTL` register is complete. Therefore, all interrupts act as if they were disabled at system reset (as if all bits of `RTC_ICTL` were zero), even though some bits of `RTC_ICTL` may read as nonzero. If no RTC interrupts are needed immediately after reset, it is recommended to write

RTC_ICTL to 0x0000 so that later read-modify-write accesses function as intended.

Interrupt status can be determined by reading the RTC_ISTAT register. All bits in RTC_ISTAT are sticky. Once set by the corresponding event, each bit remains set until cleared by a software write to this register. Event flags are always set; they are not masked by the interrupt enable bits in RTC_ICTL. Values are cleared by writing a “1” to the respective bit location, except for the write pending status bit, which is read-only. Writes of “0” to any bit of the register have no effect. This register is cleared at reset and during deep sleep.

The RTC interrupt is set whenever an event latched into the RTC_ISTAT register is enabled in the RTC_ICTL register. The pending RTC interrupt is cleared whenever all enabled and set bits in RTC_ISTAT are cleared, or when all bits in RTC_ICTL corresponding to pending events are cleared.

As shown in [Figure 24-3](#), the RTC generates an interrupt request (IRQ) to the processor core for event handling and wakeup from a sleep state. The RTC generates a separate signal for wakeup from a deep sleep or from an internal V_{dd} power-off state. The deep sleep wakeup signal is asserted at the 1 Hz tick when any RTC interval event enabled in RTC_ICTL occurs. The assertion of the deep sleep wakeup signal causes the processor core clock (CCLK) and the system clock (SCLK) to restart. Any enabled event that asserts the RTC deep sleep wakeup signal also causes the RTC IRQ to assert once SCLK restarts.

RTC Programming Model

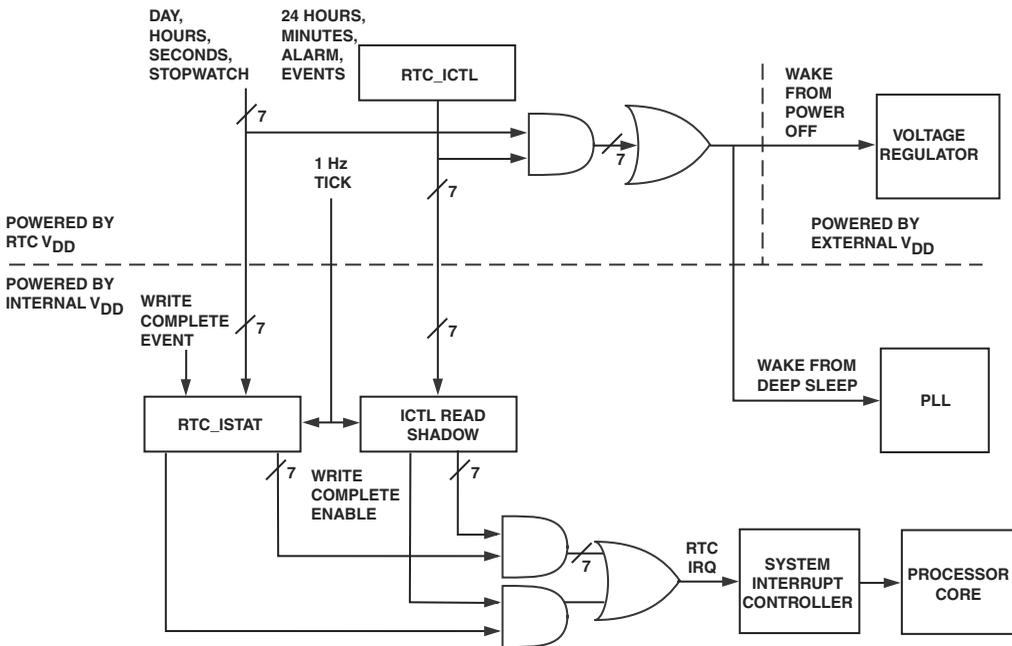


Figure 24-3. RTC Interrupt Structure

State Transitions Summary

Table 24-1 shows how each RTC MMR is affected by the system states. The phase locked loop (PLL) states are defined in the *Dynamic Power Management* chapter. “No power” means none of the processor power supply pins are connected to a source of energy. “Off” means the processor core, peripherals, and memory are not powered (internal V_{dd} is off), while the RTC is still powered and running. External V_{dd} may still be powered. Registers described as “as written” are holding the last value software wrote to the register. If the register has not been written since V_{DDRTC} power was applied, then the state is indeterminate (for all bits of `RTC_STAT`, `RTC_ALARM`, and `RTC_SWCNT`, and for some bits of `RTC_ISTAT`, `RTC_PREN`, and `RTC_ICTL`).

Table 24-1. Effect of States on RTC MMRs

RTC V _{dd}	IV _{dd}	System State	RTC_ICTL	RTC_ISTAT	RTC_STAT RTC_SWCNT	RTC_ALARM RTC_PREN
Off	Off	No power	X	X	X	X
On	On	Reset	As written	0	Counting	As written
On	On	Full on	As written	Events	Counting	As written
On	On	Sleep	As written	Events	Counting	As written
On	On	Active	As written	Events	Counting	As written
On	On	Deep sleep	As written	0	Counting	As written
On	Off	Off	As written	X	Counting	As written

Table 24-2 summarizes software’s responsibilities with respect to the RTC at various system state transition events.

Table 24-2. RTC System State Transition Events

At This Event:	Execute This Sequence:
Power on from no power	Write RTC_PREN = 1. Wait for write complete. Write RTC_STAT to current time. Write RTC_ALARM, if needed. Write RTC_SWCNT. Write RTC_ISTAT to clear any pending RTC events. Write RTC_ICTL to enable any desired RTC interrupts or to disable all RTC interrupts.
Full on after reset or Full on after power on from off	Wait for seconds event, or write RTC_PREN = 1 and wait for write complete. Write RTC_ISTAT to clear any pending RTC events. Write RTC_ICTL to enable any desired RTC interrupts or to disable all RTC interrupts. Read RTC MMRs as required.
Wake from deep sleep	Wait for seconds event flag to set. Write RTC_ISTAT to acknowledge RTC deep sleep wakeup. Read RTC MMRs as required. The PLL state is now active. Transition to full on as needed.

Register Definitions

Table 24-2. RTC System State Transition Events (Continued)

At This Event:	Execute This Sequence:
Wake from sleep	If wakeup came from RTC, seconds event flag will be set. In this case, write RTC_ISTAT to acknowledge RTC wakeup IRQ. Always, read RTC MMRs as required.
Before going to sleep	If wakeup by RTC is desired: Write RTC_ALARM and/or RTC_SWCNT as needed to schedule a wakeup event. Write RTC_ICTL to enable the desired RTC interrupt sources for wakeup. Wait for write complete. Enable RTC for wakeup in the system interrupt wakeup enable register (SIC_IWR).
Before going to deep sleep	Write RTC_ALARM and/or RTC_SWCNT as needed to schedule a wakeup event. Write RTC_ICTL to enable the desired RTC event sources for deep sleep wakeup. Wait for write complete.
Before going to off	Write RTC_ALARM and/or RTC_SWCNT as needed to schedule a wakeup event. Write RTC_ICTL to enable any desired RTC event sources for powerup wakeup. Wait for write complete. Set the wake bit in the voltage regulator control register (VR_CTL).

Register Definitions

The following sections contain the register definitions. [Figure 24-4](#) through [Figure 24-9](#) on page 24-22 illustrate the registers.

Table 24-3 shows the functions of the RTC registers.

Table 24-3. RTC Register Mapping

Register Name	Function	Notes
RTC_STAT	RTC status register	Holds time of day
RTC_ICTL	RTC interrupt control register	Bits 14:7 are reserved
RTC_ISTAT	RTC interrupt status register	Bits 13:7 are reserved
RTC_SWCNT	RTC stopwatch count register	Undefined at reset
RTC_ALARM	RTC alarm register	Undefined at reset
RTC_PREN	Prescaler enable register	Always set PREN = 1 for 1 Hz ticks

Register Definitions

RTC Status (RTC_STAT) Register

RTC Status Register (RTC_STAT)

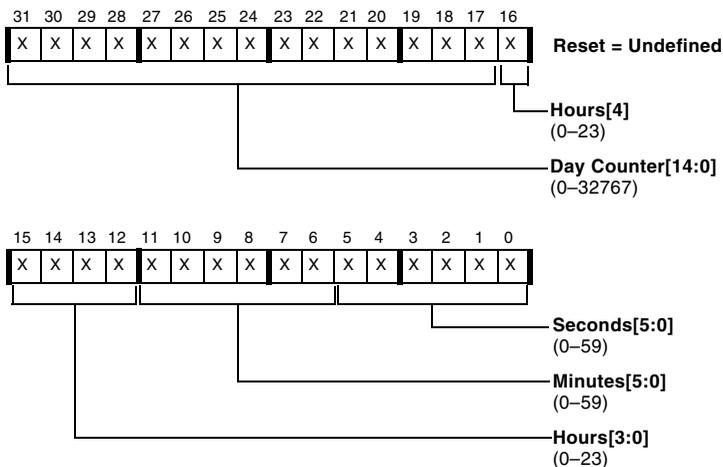


Figure 24-4. RTC Status Register

RTC Interrupt Control (RTC_ICTL) Register

RTC Interrupt Control Register (RTC_ICTL)

0 – Interrupt disabled, 1 – Interrupt enabled

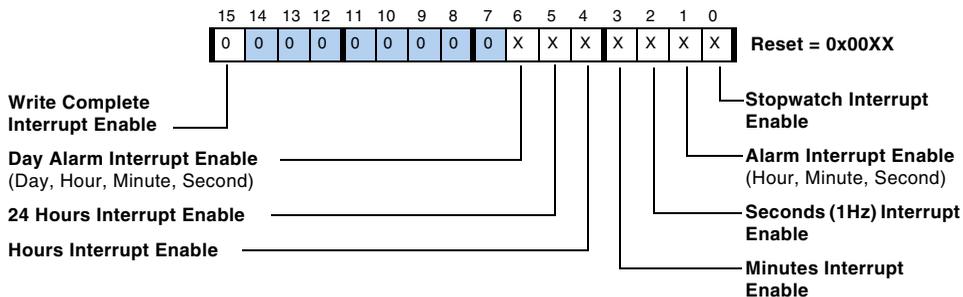


Figure 24-5. RTC Interrupt Control Register

RTC Interrupt Status (RTC_ISTAT) Register

RTC Interrupt Status Register (RTC_ISTAT)

All bits are write-1-to-clear, except bit 14

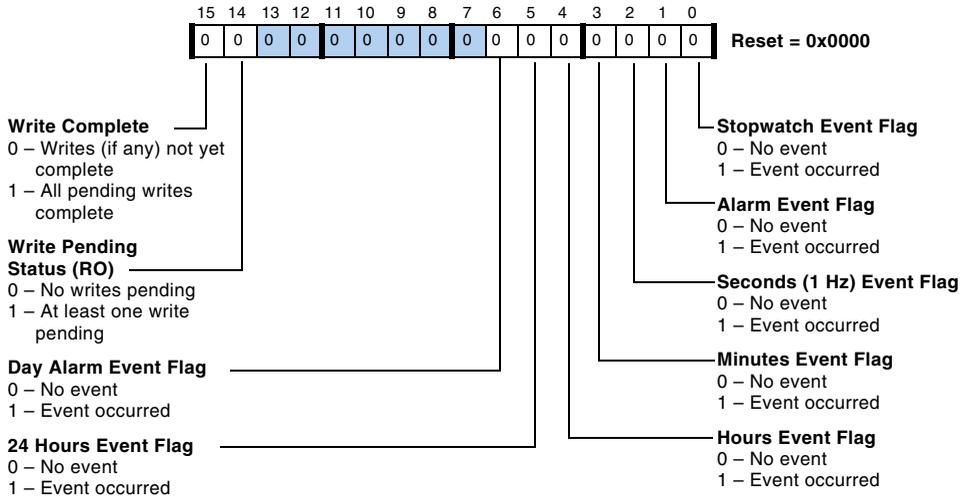


Figure 24-6. RTC Interrupt Status Register

RTC Stopwatch Count (RTC_SWCNT) Register

RTC Stopwatch Count Register (RTC_SWCNT)

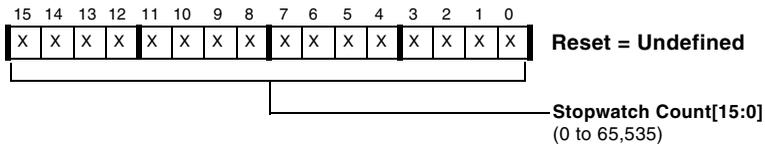


Figure 24-7. RTC Stopwatch Count Register

Register Definitions

RTC Alarm (RTC_ALARM) Register

RTC Alarm Register (RTC_ALARM)

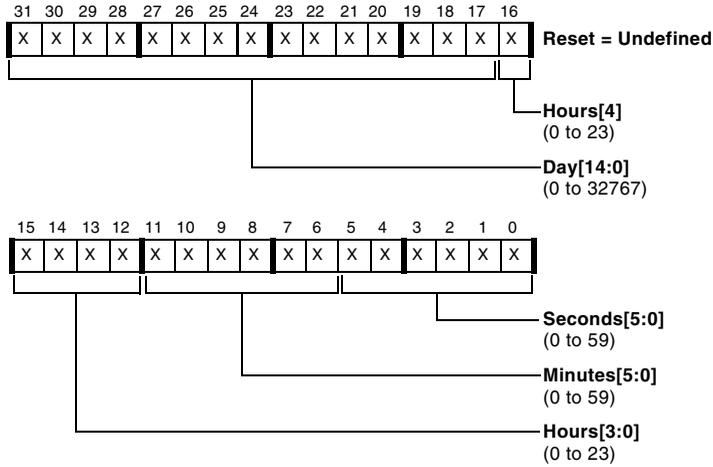


Figure 24-8. RTC Alarm Register

RTC Prescaler Enable (RTC_PREN) Register

RTC Prescaler Enable Register (RTC_PREN)

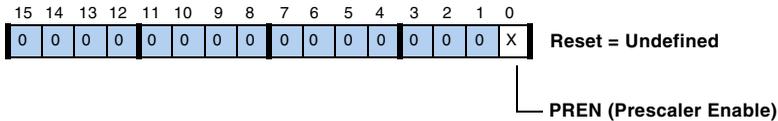


Figure 24-9. RTC Prescaler Enable Register

Programming Examples

The following RTC code examples show how to enable the RTC prescaler, how to set up a stopwatch event to take the RTC out of deep sleep mode, and how to use the RTC alarm to exit hibernate state. Each of these code examples assumes that the appropriate header file is included in the source code (for instance, `#include <defBF525.h>` for ADSP-BF525 projects).

Enable RTC Prescaler

[Listing 24-1](#) properly enables the prescaler and clears any pending interrupts.

Listing 24-1. Enabling the RTC Prescaler

```
RTC_Initialization:
    P0.H = HI(RTC_PREN);
    P0.L = LO(RTC_PREN);
    R0=PREN(Z); /* enable prescaler for 1 Hz ticks */
    W[P0] = R0.L;

    P0.L = LO(RTC_ISTAT);
    R0 = 0x807F(Z);
    W[P0] = R0.L; /* clear any pending interrupts */

    R0 = WRITE_COMPLETE(Z); /* mask for WRITE-COMPLETE bit */
Poll_WC:    R1 = W[P0](Z);
            R1 = R1 & R0; /* wait for Write Complete */
            CC = AZ;
            IF CC JUMP Poll_WC;

RTS;
```

RTC Stopwatch For Exiting Deep Sleep Mode

[Listing 24-2](#) sets up the RTC to utilize the stopwatch feature to come out of deep sleep mode. This code assumes that the `_RTC_Interrupt` label is properly registered as the ISR destination for the real-time clock event, the RTC interrupt is enabled in both `IMASK` and `SIC_IMASK`, and that the RTC prescaler has already been enabled properly.

Listing 24-2. RTC Stopwatch Interrupt to Exit Deep Sleep

```
/* RTC Wake-Up Interrupt To Be Used With Deep Sleep Code */
_RTC_Interrupt:
    PO.H = HI(PLL_CTL);
    PO.L = LO(PLL_CTL);
    R0 = W[P0](Z);
    BITCLR (R0, BITPOS(BYPASS));
    W[P0] = R0; /* If BYPASS Set, Must Clear It */

    IDLE; /* Must go to IDLE for PLL changes to be effected */

    R0 = 0x807F(Z);
    PO.H = HI(RTC_ISTAT);
    PO.L = LO(RTC_ISTAT);
    W[P0] = R7; /* clear pending RTC IRQs */

    R0 = WRITE_COMPLETE(Z); /* mask for WRITE-COMPLETE bit */
Poll_WC_IRQ:    R1 = W[P0](Z);
                R1 = R1 & R0; /* wait for Write Complete */
                CC = AZ;
                IF CC JUMP Poll_WC_IRQ;

    RTI;
```

```

Deep_Sleep_Code:
    P0.H = HI(RTC_SWCNT);
    P0.L = LO(RTC_SWCNT);
    R1 = 0x0010(Z); /* set stop-watch to 16 seconds */
    W[P0] = R1.L; /* will produce ~15 second delay */

    P0.L = LO(RTC_ICTL);
    R1 = STOPWATCH(Z);
    W[P0] = R1.L; /* enable Stop-Watch interrupt */
    P0.L = LO(RTC_ISTAT);
    R1 = 0x807F(Z);
    W[P0] = R1.L; /* clear any pending RTC interrupts */

    R0 = WRITE_COMPLETE(Z); /* mask for WRITE-COMPLETE bit */
Poll_WC1:    R1 = W[P0](Z);
             R1 = R1 & R0; /* wait for Write Complete */
             CC = AZ;
             IF CC JUMP Poll_WC1;

/* RTC now running with correct stop-watch count and interrupts
*/
    P0.H = HI(PLL_CTL);
    P0.L = LO(PLL_CTL);
    R0 = W[P0](Z);
    BITSET (R0, BITPOS(PDWN)); /* set PDWN To Go To Deep Sleep */
    W[P0] = R0.L; /* Issue Command for Deep Sleep */

    CLI R0; /* Perform PLL Programming Sequence */
    IDLE;
    STI R0; /* In Deep Sleep When Idle Exits */

    RTS;

```

RTC Alarm to Come Out of Hibernate State

[Listing 24-3](#) sets up the RTC to utilize the alarm feature to come out of hibernate state. This code assumes that the prescaler has already been properly enabled.

Listing 24-3. Setting RTC Alarm to Exit Hibernate State

```
Hibernate_Code:
    PO.H = HI(RTC_ALARM);
    PO.L = LO(RTC_ALARM);
    R0 = 0x0010(Z); /* set alarm to 16 seconds from now */
    W[PO] = R0.L;

    PO.L = LO(RTC_STAT);
    R0 = 0; /* Clear RTC Status to Start Counting at 0 */
    W[PO] = R0.L;

    PO.L = LO(RTC_ICTL);
    R0 = ALARM(Z);
    W[PO] = R0.L; /* enable Alarm interrupt */

    PO.L = LO(RTC_ISTAT);
    R0 = 0x807F(Z);
    W[PO] = R0.L; /* clear any pending RTC interrupts */

    R0 = WRITE_COMPLETE(Z);
Poll_WC1:  R1 = W[PO](Z);
           R1 = R1 & R0; /* wait for Write Complete */
           CC = AZ;
           IF CC JUMP Poll_WC1;
```

```
/* RTC now running with correct RTC status */
GoToHibernate:
    PO.H = HI(VR_CTL);
    PO.L = LO(VR_CTL);
    RO = W[P0](Z);
    BITCLR(RO, 0); /* Clear FREQ (bits 0 and 1) to */
    BITCLR(RO, 1); /* go to Hibernate State      */
    BITSET(RO, BITPOS(WAKE)); /* Enable RTC Wakeup */
    W[P0] = RO.L;

    CLI RO; /* Use PLL programming sequence to */
    IDLE; /* make VR_CTL changes take effect */
    RTS; /* Should Never Execute This */
```

Unique Information for the ADSP-BF51x Processor

None.

Unique Information for the ADSP-BF51x Processor

25 SECURITY

This chapter describes the security features and functionality of the ADSP-BF51x Blackfin processor. Following an overview and a list of key features are a description of operation and functional modes of operation.

This chapter includes the following sections:

- [“Overview”](#)
- [“Features” on page 25-4](#)
- [“Description of Operation” on page 25-6](#)
- [“Programming Model” on page 25-32](#)
- [“Security Registers” on page 25-45](#)

The intention of the chapter is to describe security features of the ADSP-BF51x Blackfin processor and how they can be used to facilitate a secure system. It is beyond the scope of this chapter to fully describe various ways to implement secure systems or to describe security protocols and primitives in any great detail.

Overview

Lockbox™ Secure Technology for Analog Devices Blackfin processors is comprised of a mix of hardware and software mechanisms designed to prevent unauthorized accesses and allow trusted code to execute on the processor. Throughout the rest of this chapter, the terms Blackfin Lockbox secure technology and Lockbox will be used interchangeably.



The developer's decision to use security features is completely optional. No security features are enabled by default. The developer can choose to never implement security features in their application if it is so desired. The Blackfin will always power up/boot in Open Mode with no security features or restrictions enabled.

Blackfin Lockbox secure technology allows users to:

- Safeguard as little as a single function, as much as a complete system, or anything in-between.
- Uniquely identify each processor by a Unique Chip ID.
- Utilize secure key storage provided by non-volatile, write-protectable One Time Programmable (OTP) memory.
- Perform digital signature authentication using elliptic curve cryptography (ECC) and secure one-way hash (SHA-1) algorithms implemented in firmware.
- Keep secret information in secure OTP Memory.
- Use any encryption algorithm to protect code or other assets.
- Ensure data integrity through digital signature authentication.
- Safeguard confidentiality by encrypting any or all of the system from core IP (code security) to data integrity.

These features in combination provide the following benefits.

- **Authenticity/Origin verification**—Lockbox secure technology allows verification of a code image against its associated digital signature, and provides for a process to identify entities and data origins.
- **Integrity**—Developers can use a digital signature authentication process to ensure that the message or the content of the storage media has not been altered in any way. If either the message or digital signature was altered, Lockbox fails during the authentication process.
- **Confidentiality**—Cryptographic encryption/decryption supports situations that require the ability to prevent unauthorized users from seeing and using designated files and streams. Methods for ensuring confidentiality are supported by the secure processing environment (Secure Mode) and secure memory.
- **Renewability**—System components can be updated to enhance security.

The Unique Chip ID enables end users to identify each Blackfin processor and hence each OEM device in which the processor resides.

Features

This Lockbox feature can be used in support of revocation and renewability of licenses in case of security violations in digital rights management systems, for example:

- **Unique Chip ID**—In combination with a trusted DRM agent (sourced by the OEM), this feature enables developers to implement renewability in DRM systems.
- **Unique Chip ID**—Provides capability to identify each OEM device and “blacklist” devices to remove them from a system.
- **Prevention of Mass Copying**—Lockbox supports cryptographic encryption/decryption algorithms for situations when confidentiality is required. The Unique Chip ID can also be utilized to “bind” the processor to one specific boot source/device and can be used to facilitate antitheft schemes and prevent OEM device cloning.

The ADSP-BF51x Blackfin processors featuring Lockbox secure technology provide security features that enable developer’s applications to use secure protocols consisting of code authentication and execution of code within a secure environment. Together these features protect secure memory spaces and restrict control of security features to authenticated developer code.

Features

Lockbox is comprised of a combination of hardware and software elements. These elements are:

- **OTP Memory.** An array of non-volatile write-protectable memory that can be programmed by the developer only one time. Half of the array is public (accessible in any mode) and the other half is

private (only accessible in Secure Mode). For more information on OTP memory, refer to [Chapter 3, “One-Time Programmable Memory”](#).

- **Secured System Switches.** Programmable bitfields in the Secured System Switches MMR to disable and enable different methods of memory access in support of a secured environment. Some of these protection mechanisms include disabling DMA access to L1 memory and disabling ADI JTAG instructions from the ICE port.
- **Secure Mode Control.** This involves the Secure State Machine hardware required to support a transition from an unsecured state of operation (Open Mode), through an authentication state (Secure Entry Mode), and finally to a secured state (Secure Mode) where secrets are accessible.
- **Firmware.** Code that resides in the on-chip ROM and performs digital signature authentication. Having the code that performs the digital signature authentication in ROM ensures integrity of the code.
- **User callable cryptographic ciphers.** In addition to the control code that resides in the on-chip ROM used for authentication, the SHA-1 cryptographic function is user-callable. The API is documented in [“Programming Model” on page 25-32](#).
- **Unique Chip ID.** Each ADSP-BF51x Blackfin processor has a 128-bit unique chip identification value stored in public OTP memory. The Unique Chip ID is programmed and write protected before a processor leaves the Analog Devices factory. It is always located at the same OTP page address.



The 128-bit Unique Chip ID value can be read but cannot be modified by the developer or end user. A total of 64K bits of OTP memory is available to the developer if additional user-defined ID values are desired. These IDs can be stored in either public or

Description of Operation

private areas of OTP memory depending on application requirements. Refer to [Chapter 3, “One-Time Programmable Memory”](#) for details.

Description of Operation

Blackfin Lockbox technology is based upon the concept of authentication of digital signatures using standards-based algorithms and provides a secure processing environment in which to execute code and access protected assets.

Digital signatures are created using a public-key signature algorithm, the Elliptic Curve Cryptography (ECC) public-key cipher, and a secure one-way hash algorithm, SHA-1. A public-key algorithm actually uses two different keys; the public key and the private key (called a key pair). The private key is known only to its owner and is not stored on-chip, while the public key can be available to anyone and is stored in the public OTP memory region on-chip. Public-key algorithms, such as ECC, are designed so that if one key is used for encryption, the other is necessary for decryption. Furthermore, the encryption key cannot be reasonably calculated from the decryption key. In a digital signature authentication scheme like Lockbox, the private key is used to generate the signature and the corresponding public key is used to validate the signature. Each ADSP-BF51x Blackfin processor has an on-chip ROM that contains firmware with the Elliptic Curve Cryptography (ECC) and SHA-1 algorithms. These are called to verify the digital signatures (ECDSA¹).

JTAG emulation and test features are disabled in hardware, and certain memory access restrictions are enabled during verification of the digital signature. Once the signature is authenticated, the access restrictions are still in effect and can only be controlled by the authenticated user code.

¹ ECDSA implementation on the ADSP-BF51x Blackfin products only supports the Koblitz curve.

Secure State Machine

The ADSP-BF51x processor includes a Secure State Machine to handle the different protection configurations of the processor depending on the security situation. The machine states are “Open Mode”, “Secure Entry Mode”, and “Secure Mode” (See [Figure 25-1](#)). The following sections describe these machine states.

The state of the Secure State Machine can be identified by reading `SECURE_STATUS[1:0]` bits. The bit values in the upper right of the states shown in [Figure 25-1](#) correspond to the bit values in `SECURE_STATUS[1:0]`.

For more information on the `SECURE_STATUS` MMR, see “[Security Registers](#)” on page 25-45.

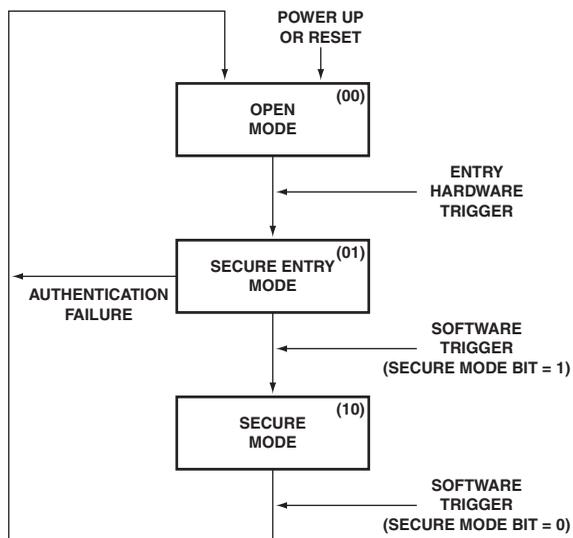


Figure 25-1. Secure State Machine Modes

Description of Operation

Open Mode

This is the default operating state of the processor, in which no restrictions are present except restricted access to the Private OTP memory area. The processor powers up and boots in Open Mode. This is the default state upon power up and after processor reset. No Lockbox security features or protection mechanisms are enabled in this state.

The state flow illustrated in [Figure 25-1](#) shows that the Secure State Machine can only transition from Open Mode into Secure Entry Mode, and there is no direct path from Open Mode into Secure Mode.

Secure Entry Mode

The on-chip ROM firmware performs the authentication process in this operating state. This mode is entered when NMI is active, and the program counter (PC) is vectored to the first address of the authentication firmware in the on-chip ROM. The program counter is monitored to ensure that it remains within the address range allocated to the Authentication firmware code. If the program counter vectors outside of the address range of the authorization code, authentication fails and the state returns to Open Mode. Any errors caught by firmware or hardware monitor will result in authentication failure and an abortion of the authentication process with the firmware exiting Secure Entry Mode and transitioning back to Open Mode. If authentication is successful, the firmware initiates the transition from Secure Entry Mode to Secure Mode.

In Secure Entry Mode, no DMA access is allowed to certain regions of internal SRAM, and JTAG emulation is disabled. The user should disable cache prior to initiating authentication. Interrupts are disabled by firmware prior to entry into Secure Mode. Interrupts are either re-enabled by dropping the interrupt level from NMI via the SESR arguments, or they are reenabled after authentication in the authenticated code after entry into Secure Mode. In addition, only the public area of OTP memory is accessible in this mode. For more information on memory access

restrictions within Secure Entry Mode, see [“Secure Entry Service Routine \(SESR\) API” on page 25-32](#).

State flow, illustrated in [Figure 25-1](#), shows that the Secure State Machine can only transition from Secure Entry Mode to Secure Mode upon successful digital signature authentication. A transition from Secure Entry Mode back into Open Mode can occur if digital signature authentication fails or if the authentication process is aborted due to an error observed by the firmware. Such errors include illegal memory boundary conditions or jumps outside of the firmware range (for example, servicing an interrupt).

Secure Mode

Secure operating state in which trusted, authenticated code is allowed unrestricted access to the processor resources, execution of authenticated code occurs, decryption of sensitive information, etc. This is the only mode that allows access (reads and writes) to the private OTP memory space where secure data, such as secret keys, can be stored. Hence, the private area of OTP memory can be used to store confidential, secret information that only authorized authenticated code can access. Therefore, this is the only operating state in which users can securely run their own Blackfin implementation of any cryptographic cipher in which secret keys are used.

Only the code (or message) digitally signed by a trusted source and successfully passed through Lockbox’s authentication process can gain access to Secure Mode.

State flow illustrated in [Figure 25-1](#) shows that the Secure State Machine can only transition from Secure Mode back into Open Mode, and there is no direct path from Secure Mode into Secure Entry Mode. Exit from Secure Mode is implemented through software control by writing a “0” value to the `SECURE0` bit within the `SECURE_CONTROL` register.

Description of Operation



Assertion of reset or power cycling will also return the processor to the default Open Mode regardless of the state of operation when the reset or power cycle event occurred. See special handling of hardware reset in [“Reset Handling in Secure Mode” on page 25-20](#).

Access to private OTP memory is restricted in Open Mode and Secure Entry Mode regardless of whether or not other security features are enabled or disabled.

Secure Mode Control

[Figure 25-2](#) describes the inputs that control the secure state machine flow.

Hardware supports transition from an Open Mode of operation, through a Secure Entry Mode, to a Secure Mode where secrets are accessible.

Open Mode is characterized by being the default mode of processor upon power up/reset/boot, holding all secured system switches deactivated and protecting the private OTP memory area from access. The processor is open with all features being available with no restrictions (except for the private area of OTP memory).

Secure Entry Mode is characterized by executing firmware out of internal ROM memory to authenticate information loaded into on-chip memory. All secured system switches are activated. However, private OTP Memory is not accessible yet.

Secure Mode is entered only after a successful digital signature authentication process from Secure Entry Mode. It provides access to the private OTP memory area and makes secured system switches accessible to user (authenticated) code. This is the mode of operation in which to perform sensitive decryption or execution of trusted, authenticated code.

Authentication can only be requested and initiated while the processor is operating in Open Mode. If authentication is requested while the processor is operating in Secure Mode, the Secure State Machine will not

transition into Secure Entry Mode. Instead, the Secure State Machine will remain in Secure Mode.

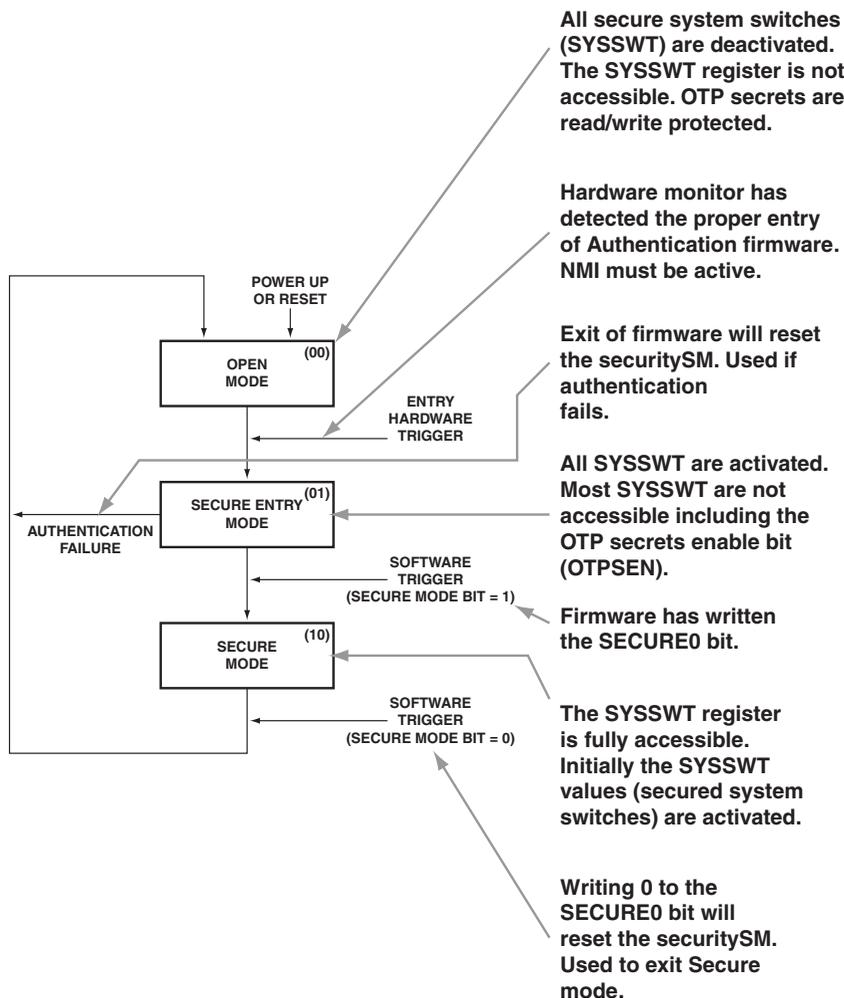


Figure 25-2. Secure Mode Control

Description of Operation



Please note that Open Mode, Secure Entry Mode and Secure Mode are states which pertain to the Secure State Machine. User Mode and Supervisor Mode are modes of operation which pertain to the core. The use of the term “mode” should not be confused and are not necessarily mutually exclusive. In Open Mode, the processor can operate in either User or Supervisor Mode. Since the firmware is entered when the NMI is being handled, Secure Entry Mode must start in Supervisor Mode. Finally, authenticated code executing in Secure Mode must be either operating at NMI interrupt level or the interrupt level that triggered the NMI.

Security Features

The following sections provide a functional description of the Security features.

Protection relies on the on-chip ROM code that includes Elliptic Curve Cryptography (ECC) and SHA-1 algorithms, applied towards verification of code authenticity using a digital signature. A processor has emulation and test features disabled in hardware as well as certain memory access restrictions upon entry into Secure Entry Mode (where authentication is performed) and maintained into Secure Mode. These functions can be controlled only by authenticated user application software executing in Secure Mode.

User code must request authentication by complying with two criteria: (1) asserting a Non-Maskable Interrupt (NMI) and (2) vector the Program Counter (PC) to the first executable address in the Secure Entry Service Routine (SESR) in firmware which resides in on-chip boot ROM.

During the authentication process, JTAG emulation is disabled, memory protection restrictions are enabled and interrupts are masked. The user has the option to pass arguments to the security firmware to control certain functionality during the authentication process. Refer to [“Secure Entry Service Routine \(SESR\) API” on page 25-32](#).

Digital Signature Authentication

Digital signatures are created off-chip (typically on a host computer) using the ECC algorithm and SHA-1, both of which are available in the public domain. In digital signature authentication, the private key generates the signature (off-chip), and the corresponding public key validates the signature (on-chip). The private key is known only to its owner and is not stored on-chip, while the public key can be available to anyone and is stored on-chip in OTP memory.

Lockbox uses standards-based cryptographic algorithms for digital signature authentication. ECDSA¹ is implemented in the Blackfin ADSP-BF51x processors. Digital signature validation on ADSP-BF51x utilizes Elliptic Curve Cryptography² (ECC) based on a binary field size of 163 bits and SHA-1³ secure one-way hash (which produces a 160-bit message digest).

In order to generate public/private key pairs or prepare digital signatures and apply them to application code, developers can use any method that complies with the Elliptic Curve Digital Signature Algorithm (ECDSA) specified in FIPS 186-2 with Change Notice 1 dated October 5, 2001, Digital Signature Standard (DSS). ECDSA is described in ANSI X9.62-1998. The Lockbox implementation in the ADSP-BF51x processors supports the following Koblitz curve, which is recommended in FIPS 186-2 for US Federal Government use:

1. m : 163 (degree of binary field)
2. a : 1

¹ ECDSA implementation on these Blackfin products only supports the Koblitz curve.

² These implementations are based on the Elliptic Curve Digital Signature Algorithm (ECDSA) specified in FIPS 186-2 with Change Notice 1 dated October 5, 2001, Digital Signature Standard (DSS) (<http://csrc.nist.gov/cryptval/dss.htm>), and specified in ANSI X9.62-1998.

³ SHA-1 is based on the publicly available standard for FIPS 180-2 (Secure Hash Signature Standard [SHS]) (FIPS PUB 180-2), <http://csrc.nist.gov/CryptoToolkit/tkhash.html>).

Description of Operation

3. b : 1 (a and b are the constants in the elliptic curve equation: $y^2 + xy = x^3 + ax + b$)
4. X_g : 2FE13C0537BBC11ACAA07D793DE4E6D5E5C94EEE8
5. Y_g : 289070FB05D38FF58321F2E800536D538CCDAA3D9 (X_g and Y_g define the base point G)
6. r : 4000000000000000000020108A2E0CC0D99F8A5EF (r is the order of the base point G)
7. T : 4 (T is the normal basis type)
8. $p(t)$: $t^{163} + t^7 + t^6 + t^3 + 1$ ($p(t)$ is the field polynomial)

The following steps summarize the Digital Signature Authentication process. Steps 1 to 3 correspond to the off-chip creation of a digital signature of a file or message. Steps 4 to 6 correspond to the on-chip digital signature authentication. These steps are preceded by generation of a key pair (Private Key and Public Key) and the programming of the Public Key in the Public OTP Memory.

1. A one-way hash of the file (message to be authenticated) is produced using SHA-1 off-chip (for example, using a host PC).
2. The hash is encrypted through ECC off-chip with the private key, thereby signing the file and completing the generation of the digital signature.
3. The file and the signed hash are stored on an external device such as Flash memory or a host device.
4. Upon transfer to the Blackfin processor's internal memory, a one-way hash of the file is calculated on-chip through SHA-1 (residing in the Blackfin on-chip boot ROM).

5. Using the ECC algorithm (residing in the Blackfin on-chip boot ROM), the Blackfin decrypts the signed hash with the user's public key stored in the Blackfin processor OTP memory.
6. The two hash results are then compared. If the signed hash matches the calculated hash, the signature is valid and the file is intact.

If the digital signature authentication process is successful, the Blackfin processor transitions from Secure Entry Mode to Secure Mode. At this time, all of the access restrictions mentioned will be in place. JTAG will be disabled and certain portions of on-chip SRAM memory are restricted from DMA access. The restrictions can be controlled once in Secure Mode by having the authenticated code modify the Secure System Switches (SECURE_SYSSWT) appropriate for use by the developer's application.



Encryption/decryption is only necessary when an application requires *confidentiality*. It is not always necessary to work with encrypted code to ensure code security. Authentication alone can be used when confidentiality is not required when ensuring tamper-proof code image and/or non-repudiation in a system. Thus, authentication safeguards code *integrity*.

Since the digital signature uniquely describes its corresponding code/message, the code/message itself does not have to be encrypted if *confidentiality* is not required. If the code/message is modified, either intentionally or inadvertently, authentication fails since the *integrity* of the code message has been compromised.

Digital Signature Authentication Performance Measurement

Authentication can be performed at any point during processor operation in Open Mode. It can be performed immediately upon boot or it can be performed any time after boot.

The algorithms used in the Lockbox firmware are highly optimized Blackfin code running from the on-chip boot ROM in the system clock

Description of Operation

domain. Firmware execution time for the digital signature authentication process is on the order of 40 million core clock cycles, depending upon the size of the digitally signed application code. This must be considered when architecting an application in order to allow a sufficient window of time in which authentication can proceed without requiring servicing of interrupts in the system.

The time it takes for authentication is dependent on several factors. These include the size of the message to be authenticated. This affects the amount of calculations done in the secure hash function (SHA-1). It also affects the DMA time required to move the message out of L1 data memory and place it into L1 code memory.

Protection Features

In order to establish a secure processing environment and protect the security of applications that establish trust and reach the privileged mode of operation, Lockbox implements access restrictions. These restrictions include disabling JTAG emulation and disabling DMA access to portions of on-chip SRAM memory. The memory access restrictions implemented in hardware on the Blackfin processor are not applied to off-chip memory. Therefore, external memory is always considered insecure and caching external memory while operating in Secure Mode represents a security risk.

Protection features include the following:

- Secure State Machine for implementing privileged states of operation in which access restrictions may be imposed to protect code and data.
 - Disable DMA access to L1 memory
 - These restrictions to memory areas are configurable (see [“Secure System Switch \(SECURE_SYSSWT\) Register” on page 25-45](#)).

- Protection of L1 regions of memory with DMA access controlled when in Secure Mode
- Disable ADI JTAG emulation from ICE port
- Divert hardware reset to NMI during Secure Mode operation to prevent “reset attack”
- Provide software control over hardware protection features accessible to trusted code operating in Secure Mode
- OTP memory for storage of customer programmable cipher keys, unique chip ID or a customer ID
- OTP write protection to protect programmed OTP memory locations from future tampering
- Private/Secret OTP memory region accessible only in Secure Mode
 - Store private key(s) for decryption of data or other validation
- A privileged mode (including firmware execution out of on-chip ROM) to perform code authentication

Description of Operation

Protection mechanisms are summarized [Table 25-1](#) for each state of the Secure State Machine along with the Secure System Switch register (SECURE_SYSSWT) that provides control over the protection feature.

Table 25-1. Secure State Machine

Secure State Machine	SECURE_SYSSWT	Description	Protected Memory Range
Open Mode (0x00000000)	The switches are involuntarily set with all controls OFF (unrestricted access)	No protection mechanisms or restrictions enabled	No restrictions ¹
Secure Entry (0x000704D9)	EMUDABL	Emulation Disable	Emulation disabled
	L1IDABLE	L1 Instruction Memory Disable 0xFFA00000—0xFFA07FFF SRAM	32 KB
	L1DADABL	L1 Data Bank A Memory Disable 0xFF800000—0xFF807FFF SRAM and SRAM/Cache	32 KB
	L1DBDABL	L1 Data Bank B Memory Disable 0xFF900000—0xFF901FFF SRAM	8 KB

Table 25-1. Secure State Machine (Continued)

Secure State Machine	SECURE_SYSSWT	Description	Protected Memory Range
Secure Mode (0x000704D9)	EMUDABL	Emulation Disable	User Configurable
	RSTDABL	RESET Disable	User Configurable
	LIIDABLE	L1 Instruction Memory Disable 0xFFA00000—0xFFA07FFF SRAM	0-32 KB
	L1DADABL	L1 Data Bank A Memory Disable 0xFF800000—0xFF807FFF SRAM and SRAM/Cache	0-32 KB
	L1DBDABL	L1 Data Bank B Memory Disable 0xFF900000—0xFF901FFF SRAM	0-32 KB

1 Private OTP is only accessible when operating in Secure Mode with OTPSEN bit set in SECURE_SYSSWT register

On-chip SRAM memory protection takes the form of DMA access restrictions only. There is no need to protect the on-chip SRAM from processor core access because, while operating in Secure Mode, the developer's authenticated code has full control over the processor core and execution of all core software instructions. It is the responsibility of the developer to take steps to avoid surrendering control of the Program Sequencer and the core to untrusted code execution.

Operating in Secure Mode

This section details Secure Mode operation, namely entering and exiting the mode.

Description of Operation

Entering Secure Mode

Upon successful digital signature authentication, the Secure State Machine transitions into Secure Mode. The same default protection features enabled in Secure Entry Mode are carried forward into Secure Mode. This includes JTAG emulation being disabled, and DMA access restrictions to memory and interrupts being masked. It is the responsibility of the authenticated code to manipulate or remove these restrictions as desired.

Exiting Secure Mode

Secure Mode provides a secure operating environment to execute sensitive code, run cryptographic ciphers, and process sensitive data. Upon exiting Secure Mode, the authenticated code should remove any sensitive code and data from memory because this sensitive information will still be accessible in Open Mode if it is not removed prior to exiting Secure Mode. Exit from Secure Mode is implemented through software control by writing a “0” value to the `SECURE0` bit within the `SECURE_CONTROL` register. Refer to [“Security Registers” on page 25-45](#) and [“Clearing Private Data” on page 25-21](#) for more information.

Reset Handling in Secure Mode

This section describes handling resets in Secure Mode.

Hardware Reset

Hardware reset is diverted to NMI when operating in Secure Mode only. When operating outside of Secure Mode, hardware reset behaves normally. This protection feature is configurable via the `RSTDABL` bit within the `SECURE_SYSSWT` register when operating within Secure Mode.

This is a protection feature to prevent malicious entities from attempting to assert hardware reset while sensitive code or data is present in the pro-

cessor's on-chip SRAM or in the processor's registers. A “reset attack” could take the following form: If hardware reset were left unprotected and reset was asserted while sensitive information were present on-chip, the processor would return to the default state of Open Mode with no protection features enabled and a malicious entity could gain access to the on-chip memory and registers, for example via JTAG emulation. In such a scenario assets intended to be protected could be compromised.

By diverting hardware reset to NMI while the processor operates in Secure Mode, servicing of hardware reset can be controlled and delayed in order to first implement a memory clean-up routine in software to purge sensitive information from internal memory and registers prior to servicing reset. At the completion of the memory clean-up, the processor can then be reset via software command and safely returned to Open Mode with no sensitive information available to be compromised.

By default, the SESR loads the address of a memory clean-up routine stored in the on-chip boot ROM into the NMI `EVT2` prior to transitioning from Secure Entry Mode into Secure Mode. See [“Clearing Private Data” on page 25-21](#) for more information.

Clearing Private Data

As part of the SESR firmware, there is a small routine stored in the on-chip boot ROM that clears the internal L1 data memory, generates a `RESET` event, and puts the processor into idle. Note that this firmware memory clear routine does not clear the contents of L1 Instruction memory or Data, Pointer, and DAG registers within the computational units. It is recommended that the user sets this routine as the new `EVT2` NMI vector once the user's authenticated application code is executing. This will prevent a malicious user from trying to reset the processor while it is operating in Secure Mode and then view the contents of internal memory when the processor returns to Open Mode after servicing `RESET`.

Description of Operation

 It is recommended that user software running in Secure Mode should also perform RAM clean-up prior to clearing the `SECURE0` Secure Mode bit and exiting Secure Mode via normal code execution within user's secure function. If sensitive code/data remains in on-chip RAM after exiting Secure Mode without wiping memory and register contents or cycling power to the processor, it is visible and accessible in Open Mode.

The memory clear routine in the on-chip boot ROM executes a watchdog `RESET` to reset the processor at the completion of the memory clear. The code also performs a clear of the `OTP_DATA0-3` registers which are used to hold data from OTP access reads (i.e. which could contain secret key or other sensitive data left by user code execution).

If a custom memory cleanup routine is part of an authenticated message, the user can use that routine instead of the one provided with the Lockbox firmware. The user can simply update `EVT2` in the event vector table to point to the start of the custom memory cleanup routine while operating in Secure Mode.

 It is strongly recommended that developers substitute their own custom memory clear routines if they require clearing of L1 instruction as the ROM memory clear routine will only clear the contents of L1 Data (Bank A and B) memory. The ROM memory clear routine will not protect instruction code from being exposed after reset is serviced or when the Secure State Machine transitions to Open Mode via other means.

Due to the fact that hardware reset is configured by default to be redirected to NMI when the processor is operating in Secure Mode, it is recommended that the user implements a watchdog reset within the `EVT2` NMI ISR in order to reset the processor. A Watchdog reset is implemented by writing a value `2'b00` in `WD0G_CTL[2:1]` and causes a complete core reset. The watchdog reset will not be redirected to the NMI pin as in the case of the external hardware reset and it will properly reset the proces-

sor. For more details of watchdog reset, refer to “[Software Resets](#)” on page 26-6“ in [Chapter 26, “System Reset and Booting”](#).

This “reset attack” protection scheme needs to protect only against hardware reset. Since it can be applied externally, the system developer typically has no control over reset in an embedded system. While operating in Secure Mode, the developer’s authenticated code has full control over the processor core and execution of all software instructions, so there is no need to protect against soft reset instructions. It is not recommended that the user’s secure application code implement a soft reset without first deleting sensitive information from memory and registers.

Public Key Requirements

A valid ECC public key must be a non-zero value and meet the following criteria:

Given the public key value shown here:

```
369368AF243193D001E39CE76BB1D5DA08A9BC0A6
```

```
15F7A90C841D4F1E1B005E70F167F6EF7CD2E251B
```

format in 32-bit little endian as follows:

```
8A9B C0A6
```

```
BB1D 5DA0
```

```
1E39 CE76
```

```
4319 3D00
```

```
6936 8AF2
```

```
0000 0003
```

```
CD2E 251B
```

Description of Operation

167F 6EF7

B005 E70F

41D4 F1E1

5F7A 90C8

0000 0001

The values should be stored in OTP pages 0x10, 0x11, 0x12 as follows, where 'L' denotes lower half of page (OTP page bits 63:0), 'H' denotes upper or high half of page (OTP page bits 127:64):

page: 0x010L: 0xBB1D 5DA0 8A9B C0A6,

page: 0x010H: 0x4319 3D00 1E39 CE76

page: 0x011L: 0x0000 0003 6936 8AF2

page: 0x011H: 0x167F 6EF7 CD2E 251B

page: 0x012L: 0x41D4 F1E1 B005 E70F

page: 0x012H: 0x0000 0001 5f7A 90C8

The general format takes the form of twelve (12) 32-bit words:

Word 1

Word 2

Word 3

Word 4

Word 5

Word 6

Word 7

Word 8

Word 9

Word 10

Word 11

Word 12

Stored into OTP pages in the following order (where 'L' denotes lower half of page, 'H' denotes upper or high half of page):

page: 0x010L:Word 2 Word 1

page: 0x010H:Word 4 Word 3

page: 0x011L:Word 6 Word 5

page: 0x011H:Word 8 Word 7

page: 0x012L:Word 10 Word 9

page: 0x012H:Word 12 Word 11

Storing Public Cipher Key in Public OTP

In order to make use of security features, the user must first store an ECC public key in the Blackfin processor public region of OTP memory pages 0x10, 0x11, and 0x12 as specified in the Firmware's Secure Entry Service Routine (SESR) API and the OTP memory map (see [“Secure Entry Service Routine \(SESR\) API” on page 25-32](#)). If no ECC public key is stored in this area of OTP, digital signature authentication cannot be successfully completed and no Lockbox security features can be enabled. For more information see [Chapter 3, “One-Time Programmable Memory”](#).

Description of Operation

Cryptographic Ciphers

Lockbox uses SHA-1 and ECC to implement ECDSA as part of the authentication process to enter into Secure Mode. These ciphers reside in the firmware in the on-chip boot ROM. The SHA-1 cipher is user-callable in Open Mode or in Secure Mode. The API is documented in [“Programming Model” on page 25-32](#). Note that ECC is not user-callable and is only executed as part of firmware during the authentication process.

Keys

Although Lockbox uses an ECC public key for digital signature authentication and has private OTP memory to store private keys for other cryptographic algorithms, Lockbox does not implement key management. Lockbox does not implement key generation, nor does it implement key exchanges natively in the Blackfin hardware.

In order to use Lockbox, an ECDSA key pair must be generated. The private key is used off-chip (typically on a host PC) to sign the message. The public key is placed in the public OTP memory where it is used to authenticate the signed message. Lockbox is only part of a full cryptosystem. It is the responsibility of the user to develop the other parts of the cryptosystem necessary for the intended application.

Debug Functionality

The processor is fully compatible with the IEEE 1149.1 standard, also known as the Joint Test Action Group (JTAG) standard. Full details of the JTAG standard can be found in the document IEEE Standard Test Access Port and Boundary-Scan Architecture, ISBN 1-55937-350-4.

ADSP-BF51x debug functionality has some modified behavior dependent upon the access privileges associated with the state of the Secure State Machine operating mode. This is to ensure that sensitive information and processing performed within Secure Entry Mode and Secure Mode will

not be compromised via JTAG. Furthermore, public JTAG instructions necessary for system test and debug (such as boundary scan and bypass mode) remain in effect regardless of the state of the Secure State Machine and are not hindered by the ADSP-BF51x Secure Mode operation. This makes it possible for developers to debug their systems without interference from the Blackfin processor or its security features.

In compliance with the JTAG standard, ADSP-BF51x processors provide an Instruction Register (IR) that interprets 5-bit instruction codes to select the test mode that performs the desired test operation. The instruction register is five bits wide and accommodates up to 32 boundary-scan instructions. The instruction register holds both public and private instructions. The JTAG standard requires some of the public instructions; other public instructions are optional. Private instructions are reserved for the manufacturer's use.

All supported public and private JTAG instructions remain operational when operating in Open Mode. All supported public JTAG features remain operational and all private JTAG features are disabled when operating in Secure Entry Mode and Secure Mode. Refer to [Appendix B, “Test Features”](#) for full details of supported JTAG instructions.

By default, JTAG emulation is disabled when the processor enters Secure Entry Mode or Secure Mode. There is only one way to enter Secure Mode—through successful authentication of user code based on digital signature validation. Once the digital signature authentication process results in success, the user's trusted, authenticated code is given full control over the processor, including access to Secured System Switches register (SECURE_SYSSWT) that enables/disables various protection mechanisms, including JTAG emulation. The Secured System Switch register provides a setting that allows authenticated code to enable JTAG emulation either in a one-time secure session setting or in a “sticky” persistent manner that allows emulation to be enabled by default the next time the processor enters Secure Mode. These settings are cleared when reset is asserted or if processor core power is cycled. (See the `EMUOVR` and `EMUDABL`

Description of Operation

bits within the `SECURE_SYSSWT` Secure System Switches Register in “[Secure System Switch \(SECURE_SYSSWT\) Register](#)” on page 25-45.)

Two bits within the `SECURE_SYSSWT` Secure System Switches register control JTAG emulation; they are Emulation Disable (`EMUDABL`) and Emulation Override (`EMUOVR`). To enable JTAG emulation for the current session while operating within Secure Mode, `SECURE_SYSSWT` bit 0 (`EMUDABL`) must be set to 0. To enable JTAG emulation to remain “sticky” and persistently enabled for the current session and for all subsequent entries into Secure Mode until cleared by the user or until cleared via `RESET` or cycling power to the processor, `SECURE_SYSSWT` bit 0 (`EMUDABL`) must be set to 0 AND `SECURE_SYSSWT` bit 14 (`EMUOVR`) must be set to 1 simultaneously. See “[Secure System Switch \(SECURE_SYSSWT\) Register](#)” on page 25-45 for details.



The `EMUDABL` bit is writable only directly when in Secure Mode. `EMUOVR` can be written to a 0 at any time. `RESET` will clear `EMUOVR`. `EMUOVR` can be cleared by the user at any time and in any mode, including Open Mode, Secure Entry Mode, and Secure Mode. You do not have to operate in Secure Mode in order to clear `EMUOVR`.

The `EMUDABL` bit is writable only directly when in Secure Mode. `EMUOVR` can be written to a 0 at any time. This means if you are in Secure Mode and wish to remove the privilege of emulation override, you are allowed to clear `EMUOVR`. Or if you are operating in Open Mode and wish to remove emulation override, you can clear `EMUOVR`. In case of Secure Entry Mode, writing the `EMUOVR` bit to a 0 immediately blocks emulation (and the `EMUDABL` bit would read 0 immediately). While operating in Secure Entry Mode, the value of `EMUDABL` is the *not* of `EMUOVR`, that is, $EMUDABL = \overline{EMUOVR}$. While operating in Secure Mode, you can read or write the `EMUOVR` bit, which has no immediate affect since `EMUDABL` is in control at that point.

Upon setting `EMUDABL = 0` AND `EMUOVR = 1`, JTAG emulation remains active and enabled for the current session during Secure Mode operation AND for ALL subsequent entries into Secure Mode until `EMUOVR` is cleared

(set to 0) or until `RESET` or power cycle clears this setting. This is also known as “sticky” emulation setting.

If “sticky” emulation is enabled (`EMUDABL = 0 AND EMUOVR = 1`), JTAG emulation is active and enabled in all modes, i.e. Secure Entry, Secure Mode, as well as in Open Mode. The Secure State Machine can cycle through all modes of operation, and JTAG emulation will remain active and enabled in every mode with these settings in place until cleared by the user application code, or until `RESET` or power cycle clears the setting.

For example, a user creates code to be authenticated with a valid digital signature. The code and digital signature are loaded onto the Blackfin processor in Open Mode, Authentication is requested (JTAG emulation is disabled by default during Authentication in Secure Entry Mode), and the Authentication process is successful. The processor enters Secure Mode (JTAG emulation still is disabled by default) and control is given to the authenticated code. Authenticated code sets bits within the Secure System Switches to enable JTAG Emulation and sets the “sticky” bit to allow JTAG emulation to be enabled by default the next time the processor transitions into Secure Mode as well. Debug within Secure Mode can occur using emulation now. If a different set of trusted code must be loaded into the processor, the user can do so now without leaving Secure Mode, or the user can choose to exit Secure Mode and return back to Open Mode in order to authenticate another set of code or load test/problematic code. A new set of code and digital signature now can be loaded and authenticated. Upon entry into Secure Mode, JTAG emulation will be enabled by default due to the sticky bit setting in the Secure System Switches. Debug can be performed within Secure Mode without changes to problematic code.

One possible usage scenario for persistent (sticky) emulation might be as follows: a “final” production code that must run in Secure Mode is prepared. There seems to be an issue with the code, but emulation prevents working with it. You would take advantage of the `EMUOVR` bit within the `SECURE_SYSSWT` register by, first, performing a simple authentication of

Description of Operation

code that sets the `EMUOVR` bit in order to enable JTAG emulation within Secure Mode. From there you exit Secure Mode (write a value of “1” to the `SECURE0` bit in the `SECURE_CONTROL` register, but do not invoke any processor reset), and call the routine to debug. You would then set a breakpoint just after authentication. That way you can now step through your code using JTAG emulation and operate in Secure Mode.

 Digitally signed user code, which enables either single session or sticky JTAG emulation, must be treated as confidential by users in the same manner as private keys. If this code is allowed to fall outside of developer control or become public, it can be used to compromise a developer’s security.

In summation, in order to enable JTAG emulation during Secure Mode, the user must successfully perform the Authentication process at least one time, and then program the Secured System Switches while operating in Secure Mode to enable emulation.

Programming Examples

Listing 25-1. Assembly Code – Enable (“Sticky”) Persistent JTAG Emulation for Secure Mode Debug

```

#include <defBF518.h> /* ADSP-BF518 used as an example */

.section L1_code;
.align 4;
.global _secure_function;
_secure_function:

// required nops to account for
// SESR PC vector target+4 for overlay ID accommodation
nop;
nop;

P0.H = ((SECURE_SYSSWT) >> 16);
P0.L = ((SECURE_SYSSWT) & 0xFFFF);
R0 = [P0];
BITCLR(R0,0);
[P0] = R0;
SSYNC;

_secure_function.END:

```

Listing 25-2. C Code – Enable JTAG Emulation for Secure Mode Debug (single session)

```

#include <cdefBF518.h> /* ADSP-BF518 used as an example*/
#define ENABLE_JTAG_MASK 0xFFFFFFFF

void secure_function(void)
{

```

Programming Model

```
/* Enable JTAG */
*pSECURE_SYSSWT = ( *pSECURE_SYSSWT & ENABLE_JTAG_MASK );
ssync();

return;
}
```

Programming Model

This section contains the following procedures:

- [“Secure Entry Service Routine \(SESR\) API”](#)
- [“Starting Authentication” on page 25-33](#)
- [“Memory Configuration” on page 25-34](#)
- [“Secure Function and Secure Entry Service Routine Arguments” on page 25-36](#)
- [“Secure System Switch \(SECURE_SYSSWT\) Register” on page 25-45](#)
- [“Secure Control \(SECURE_CONTROL\) Register” on page 25-53](#)
- [“Secure Status \(SECURE_STATUS\) Register” on page 25-55](#)

Secure Entry Service Routine (SESR) API

This section describes the procedure to use Lockbox to authenticate a message. Memory configuration, input arguments and return codes are also described here.

In this chapter, the term “message” was widely used to describe the entity being digitally signed off-chip, and later authenticated on-chip by the SESR security firmware. “Message”, “secure function” (SF), and “secure application” are used interchangeably in this section and mean the same thing.

Starting Authentication

For an application to establish trust and reach the privileged mode of operation (for example, enter Secure Mode), the Secure State Machine has to transition from Open Mode, through Secure Entry Mode, to Secure Mode. In order to transition from Open Mode to Secure Entry Mode, NMI must be asserted and the program counter (PC) must vector to the beginning address of the firmware (SESR) at location 0xFFA1 4000.

This can be achieved by loading the beginning address of the SESR (0xFFA1 4000) as the NMI handler in the event vector table (EVT2). Then in supervisor mode, issue a raise 2; instruction. Similarly, NMI hardware pin may be asserted instead of issuing a software raise; instruction. Once the PC vectors to the SESR, while NMI assertion is sensed by the hardware, the Secure State Machine transitions into Secure Entry Mode.

Before actually going into Secure Entry Mode, the user will have to set up the memory environment. This includes specifying the arguments (described in this section) and moving the message to be authenticated into L1 data memory.

Memory Configuration

Figure 25-3 illustrates the Secure Entry Mode default memory configuration upon initiating authentication and entering the SESR.

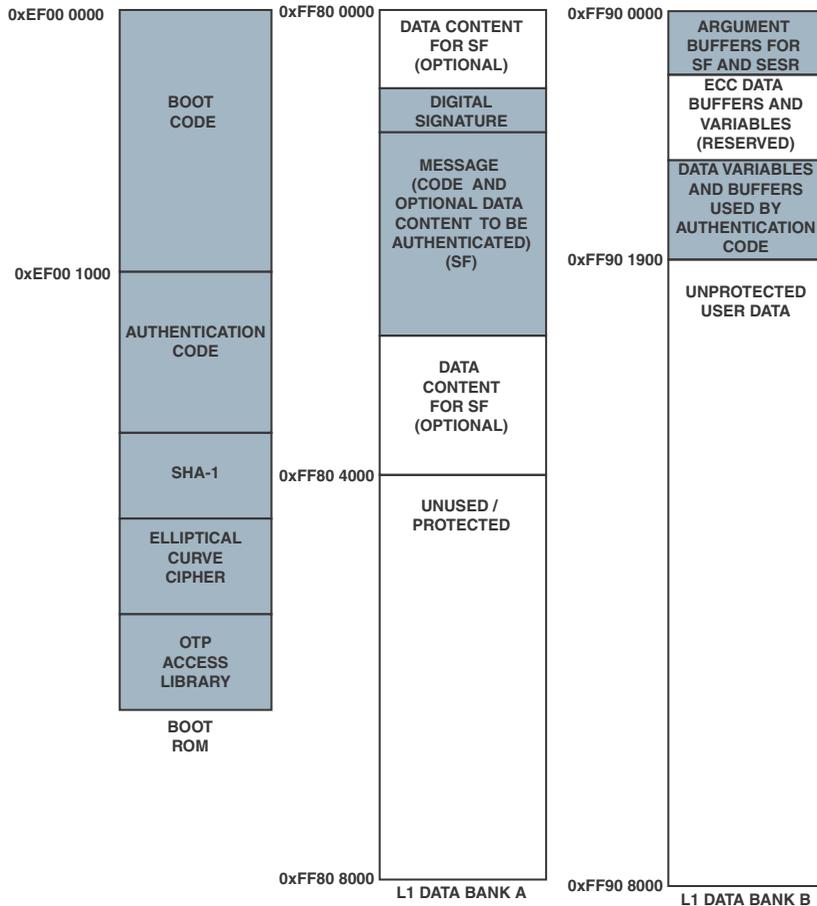


Figure 25-3. Memory Configuration for Authentication

Message Placement

The message must be placed in L1A for authentication. If the message (for example, code) is put into L1A for authentication, it must be DMA'd to L1 code space, where it can execute. It is the user's responsibility to provide the message in L1A memory for the SESR. If authentication is successful, the SESR then moves the message via DMA to the final destination according to the SESR arguments. No further action is required by the developer to perform this DMA as it is executed by the firmware.

Digital Signature

The digital signature is a pair of 163-bit integers. Each integer is padded to the nearest 32-bit word, resulting in 192 bits for each integer, resulting in a total size of 384 bits. The authentication firmware always expects the digital signature to be followed by the message. For example, if the message is placed in L1A data memory, and the digital signature starts at address 0xFF80 0000, the message must immediately follow the digital signature and be located at address 0xFF80 0030. The message and digital signature must be stored together contiguously in memory with the message always immediately following the digital signature.

Message Size Constraints

The maximum size of any message to be authenticated is limited by the size of on-chip memory in the Blackfin processor. When the Secure State Machine enters into Secure Entry Mode (authentication), certain portions of on-chip SRAM memory are protected from DMA accesses. These protected memory regions include L1A (32 KB) and L1B data memory (8 KB each and 32 KB of L1 code memory). This means that the maximum allowable message/code size that can be authenticated is 32 KB less 48 bytes for the digital signature when placed in L1A data memory.

Programming Model

Memory Usage

In data bank B of the L1 memory, the arguments for both the SESR and the secure function are stored beginning at address 0xFF90 0000. In addition, a portion of the L1B data memory is reserved for the firmware for scratch space. All memory above address 0xFF90 1900 is reserved for authentication. The user can either allocate this area of memory solely for Lockbox or save any data elsewhere in memory prior to starting authentication.



Any user information residing in the scratch space reserved area of L1 Data Bank B will be overwritten during the authentication process.

Memory Protection

This Secure Entry Mode default memory configuration with both protected and unprotected regions of on-chip SRAM is implemented in order to allow developers to initiate digital signature authentication at any time during Open Mode processor operation. If an application is already running on the processor, the unprotected memory regions can be used for placement of data buffers. When authentication occurs, access to these data buffers is not restricted, thus the application can be given higher precedence over the authentication process if necessary.

The Secure Entry Mode default memory protection configuration put into place upon initiating authentication cannot be modified by the developer. This is to ensure integrity of the secure processing environment during the authentication process and help prevent malicious tampering.

Secure Function and Secure Entry Service Routine Arguments

Prior to initiating the authentication, the arguments for both the SESR and the message (also known as Secure Function or SF) must be set up.

The arguments are stored in argument buffers stored in L1B data memory. Specifically, the arguments for the Secure Function are stored at the top of L1B data memory, at address 0xFF90 0000. There are 24 bytes allocated for the arguments for the secure function. Following the argument buffer for the Secure Function is the argument buffer for the SESR, at address 0xFF90 0018. For security reasons this authentication protocol accesses fixed locations for arguments. When the user starts executing the SF, it receives two arguments. The first argument (R0) contains the address of the SF argument buffer. The second argument (R1) holds the IMASK value before shut off interrupts.

Secure Function Arguments

When the message is successfully authenticated, the Program Counter will vector to the Secure Function with the first argument (R0) containing a pointer to top of L1B data memory. The second argument (R1) of the secure function is the IMASK value. This value is obtained when the SESR successfully authenticates the message. Before the message is transferred via DMA to its final target run location, interrupts are shut off so tampering cannot occur between the time of successful authentication and execution of the secure function. The prototype for the secure function is: `void secure_function(tSecureFunctionArgs *, unsigned short imask);`

The 24-byte Secure Function argument buffer is for the convenience of the user to be able to pass arguments to the Secure Function prior to starting authentication.

It will be the responsibility of the user's Secure Function responsibility to re-enable interrupts by using the saved IMASK value or by using a new IMASK value.

The 24-byte Secure Function argument buffer can be used in any aligned fashion. For example, it can be used to store six 32-bit words or twelve 16-bit words, or any combination of data types such as integers, shorts and characters, as long as the accesses are aligned.

Secure Entry Service Routine Arguments

The argument buffer for the SESR is shown in [Listing 25-3](#).

Listing 25-3. Argument Buffer for SESR

```
/* SESR argument structure. Expected to reside at address
0xFF900018 */
typedef struct SESR_args {

    unsigned short usFlags;        /* security firmware flags*/
    unsigned short usIRQMask;     /* interrupt mask*/
    unsigned long ulMessageSize; /* message length in bytes*/
    unsigned long ulSFEntryPoint; /* entry point of secure function*/
    unsigned long ulMessagePtr;  /* pointer to the buffer containing
the digital signature and message */
    unsigned long ulReserved1;   /* reserved*/
    unsigned long ulReserved2;   /* reserved*/
} tSESR_args;
```

usFlags

The first argument, `usFlags`, is a 16-bit flag that signals authentication what to do. [Figure 25-4](#) shows the meaning of the bits.

Bit 0 tells the authentication firmware whether or not to drop the interrupt level. To execute raise 2;, the Blackfin processor must operate in supervisor mode, in other words, operate at one of the interrupt levels. NMI must be asserted when authentication is initiated. The caller/user has the option to deassert NMI and drop back down to a lower interrupt level (the interrupt level in effect when NMI was asserted to initiate authentication) or continue authentication at NMI level.

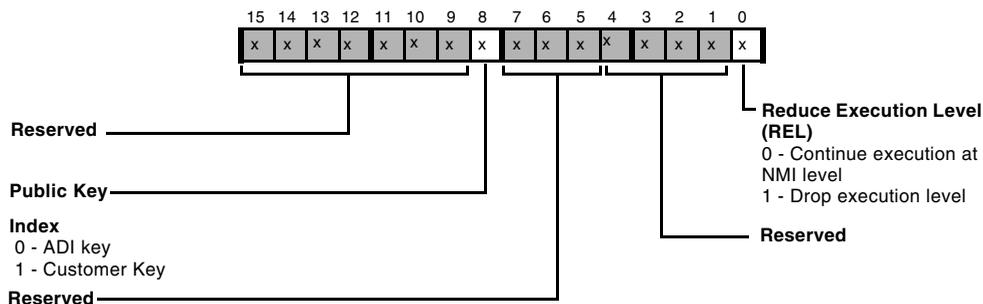


Figure 25-4. Bit Fields for Flags Argument

By lowering the interrupt level at which the authentication firmware executes, other interrupts can be serviced. Be aware that if another interrupt is serviced and the PC vectors out of the authentication firmware during authentication, the authentication process fails and returns an error code.

Bit 8 tells the firmware which public key is used for authentication. The OTP memory holds two public keys. One is programmed by Analog Devices for failure analysis purposes only, and the other is programmed by the developer.

usIRQMask

The `usIRQMask` argument is a 16-bit user-defined bitmask to be loaded into the lower 16 bits of the `IMASK` MMR if the execution level is to be lowered from NMI level. This argument allows the user to specify which, if any, interrupts will be allowed to be serviced should they occur during the time authentication occurs. Note that if any interrupt is serviced, the authentication process fails and returns an error code as mentioned above. For more information regarding `IMASK`, please refer to the *Blackfin Programming Reference* manual.

Programming Model

ulMessageSize

The `ulMessageSize` argument is a 32-bit non-negative integer that tells the SESR how big the message is, in bytes. The `ulMessageSize` must be a multiple of two, otherwise the SESR returns an error code.

ulSFEntryPoint

The `ulSFEntryPoint` argument is the final address that the message will be moved to and executed from L1 Instruction memory. Again, since the authentication firmware expects code as the first portion of the message, the address must be a multiple of four since instructions can be either 16-bit or 32-bit lengths. If the message consists of both code and data, it is the user's responsibility to move the data to the proper area of data memory for subsequent use within the application.

ulMessagePtr

The `ulMessagePtr` argument holds the address of where the digital signature and message is found in L1 Data memory.

Secure Message Execution

If the authentication of the digital signature is successful, the authentication firmware directly vectors the Program Counter to the Secure Function at its final target location, plus an offset of four bytes. The offset provides a location for the overlay ID if overlays are used with Lockbox. To return to the calling function, the authenticated message must execute `rtm`; if execution level was not signaled to be lowered in the authentication firmware. Otherwise, if the execution level was lowered, the Secure Function can return via `rts`.

To prevent tampering, interrupts and the watchdog timer are shut off near the end of successful authentication. It is the user's responsibility to re-enable the interrupts and the watchdog timer in the Secure Function if they are required in the user's application while operating in Secure Mode.

Return Codes

If for any reason an error occurs, the SESR returns an error code, and bit 7 in the SECURE_STAT MMR sets to indicate that register R0 contains a valid error code. [Table 25-2](#) lists a portion of the valid return codes.

Table 25-2. List of Return Codes From SESR

Return Codes	Value	Description
SECFW_SUCCESS	0	Success
SECFW_ERROR_INV_FLAGS	-1	“Flags” argument to firmware is invalid
SECFW_ERROR_INV_INTMASK	-2	IRQ mask specified is invalid
SECFW_ERROR_INV_CODESZ	-3	Code size specified is either non-positive or odd
SECFW_ERROR_OOB_CODE	-6	The message (Secure function) is too big and surpasses the protected region in L1A
SECFW_ERROR_BAD_EVT	-10	One of the ISR specified in the Event Vector table points inside the authentication firmware.
SECFW_ERROR_PUBKEY_ZERO	-11	Invalid public key of (0,0)
SECFW_ERROR_AUTH_FAILED	-12	Invalid message/signature pair
SECFW_ERROR_DMA	-15	MDMA error occurred during DMA transfer or the message to the final target vector.
SECFW_ERROR_DROPPING_INT_FAILED	-17	Could not drop interrupt level from NMI.
SECFW_ERROR_FUSE_READ_FAILED	-18	Error occurred while reading OTP memory.
SECFW_ERROR_TGTVECT_NONALIGNED	-19	Target vector is not 4 Byte aligned.
SECFW_ERROR_SECURE0_WRITE_FAILED	-20	Write to Secure0 bit failed. Secure State Machine might be blocking the write because ISR was taken.

Programming Model

Table 25-2. List of Return Codes From SESR (Continued)

Return Codes	Value	Description
SECFW_ERROR_SM_NOT_ENTERED	-21	Secure0 bit was written three times but secure mode was still not entered.
SECFW_ERROR_BAD_TGT_ADDR	-22	Target vector must be in L1 code space.
SECFW_ERROR_SF_TOO_BIG	-23	Message (Secure function) too big to fit at target location.

In addition to the return codes listed in [Table 25-2](#), a return value between -62 and -252 is also a valid error return code. These errors are from OTP accesses.

To decipher the error from an OTP access, there is an offset that must be added to the error code. The macro `OTP_READ_ERROR_OFFSET` (defined in VisualDSP++ header files with a value of -285) is added to the return value. The result is a bit mask. [Figure 25-5](#) shows the definition of the bit fields.

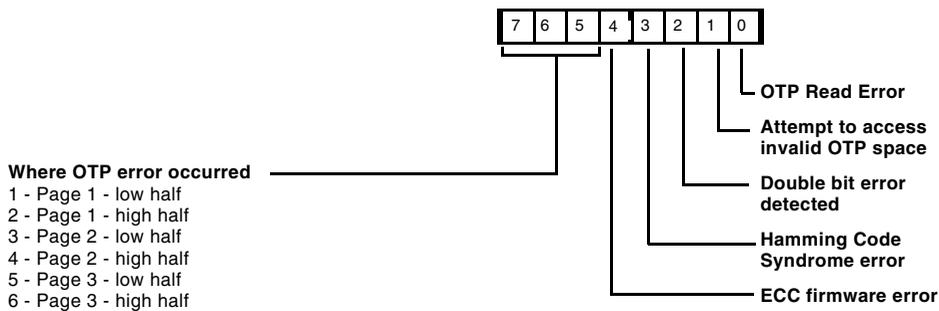


Figure 25-5. Bit Field Definition Return Value if OTP Error Occurred

SECURE HASH ALGORITHM (SHA-1) API

The ADSP-BF51x processor includes a software implementation of the Secure Hash Algorithm (SHA-1) in the on-chip boot ROM. This implementation of the SHA-1 hash algorithm is C-callable.

The following describes the application programming interface (API) for using SHA-1, including both data types and ROM routines.

ADI_SHA1 Data Type

```
typedef struct ADI_SHA1 {
    u8    *pInputMessage;
    u32   udMessageSize;
    u8    *pOutputDigest;
    u8    *pScratchBuffer;

} ADI_SHA1;
```

The SHA1 hash routine, `bfrom_Sha1Hash`, when provided with a reference to an object of type `ADI_SHA1`, hashes the `udMessageSize`-long message referenced by `pInputMessage`, and stores the hash value (also referred to as message digest) in the buffer referenced by `pOutputDigest`. The elements in an object of type `ADI_SHA1`, are shown in [Table 25-3](#).

Table 25-3. Elements in an Object of Type `ADI_SHA1`

<code>pInputMessage</code>	Pointer to the input buffer
<code>udMessageSize</code>	The size, in bytes, of the valid input data in <code>pInputMessage</code> .
<code>pOutputDigest</code>	Pointer to the output data buffer. After hashing, this buffer will contain the digest of the input message. The digest is 160-bits (<code>SHA1_HASH_SIZE</code> -bytes) long
<code>pScratchBuffer</code>	Pointer to a data buffer of size, <code>SHA1_SCRATCH_BUFFER_SIZE</code> -bytes, used by the SHA-1 module.

Programming Model

bfrom_Sha1Init ROM Routine

Entry address: Defined as `BFROM_SHA1_INIT` in the `bfrom.h` header file in the VisualDSP++ installation directory.

Arguments:

R0: Pointer to a buffer of size `SHA1_SCRATCH_BUFFER_SIZE`

C prototype:

```
void bfrom_Sha1Init (u8 *pScratchBuffer);
```

This function initializes some data elements in `pScratchBuffer`. It is called first before making any calls to `bfrom_Sha1Hash`.

bfrom_Sha1Hash ROM Routine

Entry address: Defined as `BFROM_SHA1_HASH` in the `bfrom.h` header file in the VisualDSP++ installation directory.

Arguments:

R0: Pointer to an object of type `ADI_SHA1`

C prototype:

```
void bfrom_Sha1Hash (ADI_SHA1 *pSha1);
```

This function performs the hash operation.

Security Registers

There are three registers for security mode control and status of the Secure State Machine states. These registers require privileged access depending upon the operating state of the processor.

Table 25-4. Security Registers

Register	Description	Size (Bits)	Memory-Mapped Address
SECURE_SYSSWT	Secure System Switches	32	0xFFC03620
SECURE_CONTROL	Secure Control	16	0xFFC03624
SECURE_STATUS	Secure Status	16	0xFFC03628

Secure System Switch (SECURE_SYSSWT) Register

The SECURE_SYSSWT register controls hardware that would otherwise allow a threat of attack to a secured system. Hardware is controlled voluntarily and involuntarily as follows.

- During Open Mode the switches are involuntarily set with all controls off (unrestricted access, with exception of access to OTP protected “secrets” area). OTP secrets are always protected and can only be accessed upon entry into Secure Mode.
- During Secure Entry Mode all switches are initially set with all controls on (restricted access); except that the OTP secrets control (OTPSSEN bit) is not accessible so access to the secrets OTP area remains restricted, and the RSTDABL bit remains deactivated (External Reset is allowed).
- During Secure Mode operation all switches are voluntary (initially set) and under the control of authenticated code. Therefore, restricted access controls can be reconfigured by authenticated user code. This includes the activation of Reset Disable (RSTDABL) bit.

Security Registers

The register, shown in [Figure 25-6](#) and [Figure 25-7](#), is 32-bits wide and requires 32-bit access. Limited write access to a few bits is allowed in Secure Entry mode, and full write access to all bits is allowed in Secure mode. No write access is allowed in Open Mode.

Secure System Switch Register (SECURE_SYSSWT) Bits 15:0

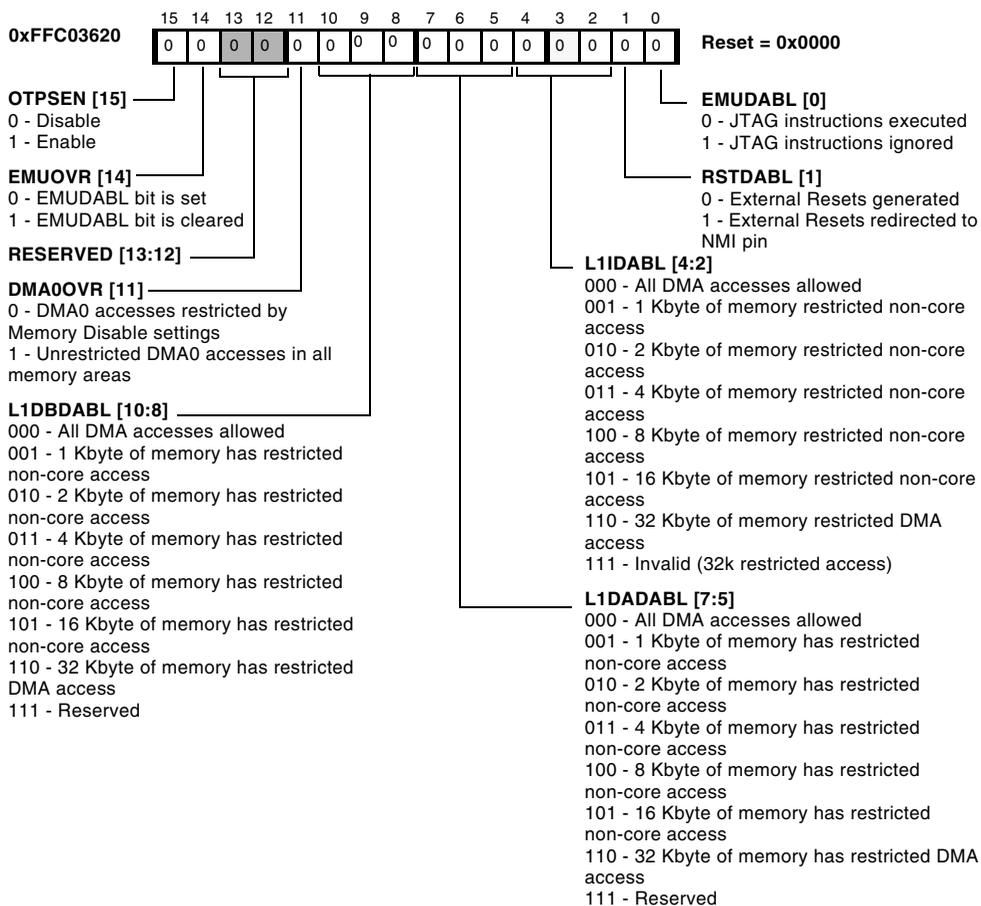


Figure 25-6. Secure System Switch Register, Bits 15:0

Security Registers

Secure System Switch Register (SECURE_SYSSWT) Bits 31:16

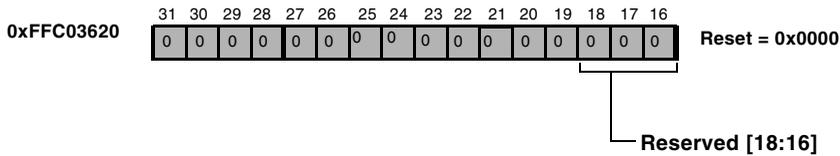


Figure 25-7. Secure System Switch Register, Bits 31:16

Table 25-5. Secure System Switch Register

Bit Position	Bit Name	Bit Description
		Reset = 0x0000 Secured Entry = 0x000704d9 Secure Mode = 0x000704db
0	EMUDABL	Emulation Disable. Upon Secured Entry EMUDABL's setting is based on the previous state of EMUOVR. Upon re-entering Open Mode, EMUDABL is cleared. This bit is always read accessible. This bit is write accessible only in Secure Mode. 0 - Analog Devices JTAG emulation instructions is recognized and executed. Once this bit is cleared while in Secure Mode it will not be set upon Secured Entry. This condition will remain until reset at which time it is cleared. This feature is used in security debug. 1 - Analog Devices JTAG emulation instructions are ignored. Standard emulation commands such as bypass is allowed.
1	RSTDABL	Reset Disable. This bit is not effected upon Secured Entry. This bit is set upon entering Secure Mode. Upon re-entering Open Mode, RSTDABL is cleared. This bit is always read accessible. This bit is write accessible only in Secure Mode. 0 - External Resets are generated and serviced normally. 1 - External Resets are redirected to the NMI pin. This avoids circumventing memory clean operations.

Table 25-5. Secure System Switch Register (Continued)

Bit Position	Bit Name	Bit Description
4:2	L1IDABL	<p>L1 Instruction Memory Disable.</p> <p>Upon Secured Entry L1IDABL is set to 0x6. Upon re-entering Unsecure Mode, L1IDABL is cleared. These bits are always read accessible. These bits are write accessible only in Secure Mode. In the event a DMA access is performed to a restricted memory area a DMA memory access error will occur resulting in a DMA_ERR interrupt and a clearing of DMA_RUN.</p> <p>000 - All DMA accesses are allowed to L1 Instruction areas.</p> <p>001 - 1 Kbyte of memory (0xFFA00000 - 0xFFA003FF) has restricted non core access</p> <p>010 - 2 Kbyte of memory (0xFFA00000 - 0xFFA007FF) has restricted non core access</p> <p>011 - 4 Kbyte of memory (0xFFA00000 - 0xFFA00FFF) has restricted non core access</p> <p>100 - 8 Kbyte of memory (0xFFA00000 - 0xFFA01FFF) has restricted non core access</p> <p>101 - 16 Kbyte of memory (0xFFA00000 - 0xFFA03FFF) has restricted non core access</p> <p>110 - 32 Kbyte of memory (0xFFA00000 - 0xFFA07FFF) has restricted DMA access. This is the initial setting upon entering Secured Entry.</p> <p>111 - Reserved</p>

Security Registers

Table 25-5. Secure System Switch Register (Continued)

Bit Position	Bit Name	Bit Description
7:5	L1DADABL	<p>L1 Data Bank A Memory Disable.</p> <p>Upon Secured Entry L1DADABL is set to 0x6. Upon re-entering Open Mode, L1DADABL is cleared. These bits are always read accessible. These bits are write accessible only in Secure Mode. In the event a DMA access is performed to a restricted memory area a DMA memory access error will occur resulting in a DMA_ERR interrupt and a clearing of DMA_RUN.</p> <p>000 - All DMA accesses are allowed to L1 data bank A areas.</p> <p>001 - 1 Kbyte of memory (0xFF800000 - 0xFF8003FF) has restricted non core access</p> <p>010 - 2 Kbyte of memory (0xFF800000 - 0xFF8007FF) has restricted non core access</p> <p>011 - 4 Kbyte of memory (0xFF800000 - 0xFF800FFF) has restricted non core access</p> <p>100 - 8 Kbyte of memory (0xFF800000 - 0xFF801FFF) has restricted non core access</p> <p>101 - 16 Kbyte of memory (0xFF800000 - 0xFF803FFF) has restricted non core access</p> <p>110 - 32 Kbyte of memory (0xFF800000 - 0xFF807FFF) has restricted DMA access. This is the initial setting upon entering Secured Entry.</p> <p>111 - Reserved</p>

Table 25-5. Secure System Switch Register (Continued)

Bit Position	Bit Name	Bit Description
10:8	L1DBDABL	<p>L1 Data Bank B Memory Disable.</p> <p>Upon Secured Entry L1DBDABL is set to 0x4 giving L1 Data Bank B 8 Kbyte of non core restricted access. Upon re-entering Open Mode, L1DBDABL is cleared. These bits are always read accessible. These bits are write accessible only in Secure Mode. In the event a DMA access is performed to a restricted memory area a DMA memory access error will occur resulting in a DMA_ERR interrupt and a clearing of DMA_RUN.</p> <p>000 - All DMA accesses are allowed to L1 data bank B areas. This is the initial setting upon entering Secured Entry.</p> <p>001 - 1 Kbyte of memory (0xFF900000 - 0xFF9003FF) has restricted non core access</p> <p>010 - 2 Kbyte of memory (0xFF900000 - 0xFF9007FF) has restricted non core access</p> <p>011 - 4 Kbyte of memory (0xFF900000 - 0xFF900FFF) has restricted non core access</p> <p>100 - 8 Kbyte of memory (0xFF900000 - 0xFF901FFF) has restricted non core access. This is the initial setting upon entering Secured Entry.</p> <p>101 - 16 Kbyte of memory (0xFF900000 - 0xFF903FFF) has restricted non core access</p> <p>110 - 32 Kbyte of memory (0xFF900000 - 0xFF907FFF) has restricted DMA access.</p> <p>111 - Reserved</p>
11	DMA0OVR	<p>DMA0 Memory Access Override</p> <p>Entering Secured Entry or Secure Mode does not effect this bit. Upon re-entering Open Mode, DMA0OVR is cleared. This bit is always read accessible. This bit is write accessible in both Secured Entry and Secure Mode.</p> <p>Controls DMA0 access to L1 Instruction, L1 Data and memory other than L1 regions. When clear access restrictions are based on Memory Disable settings within this register.</p> <p>0 - DMA0 accesses are restricted based on Memory Disable settings.</p> <p>1 - Unrestricted DMA0 accesses are allowed to all memory areas.</p>
12	Reserved	Reserved bit. This reserved bit always returns a “0” value on a read access. Writing this bit with any value has no effect.
13	Reserved	Reserved bit. This reserved bit always returns a “0” value on a read access. Writing this bit with any value has no effect.

Security Registers

Table 25-5. Secure System Switch Register (Continued)

Bit Position	Bit Name	Bit Description
14	EMUOVR	<p>Emulation Override</p> <p>This bit is always read accessible. This bit may be written with a 1 in Secure Mode only.</p> <p>This bit can be cleared in any mode (Open Mode, Secured Entry and Secure mode). Controls the value of EMUDABL upon Secured Entry.</p> <p>0 - Upon Secured Entry the EMUDABL bit is set.</p> <p>1 - Upon Secured Entry the EMUDABL bit is cleared. This bit can only be set when EMUDABL (bit-0) is written with a “0” while this bit (bit-14) is simultaneously written with a 1.</p>
15	OTPSEN	<p>OTP Secrets Enable.</p> <p>This bit can be read in all modes but is write accessible in Secure Mode only.</p> <p>0 - Read and Programming access of the “secured” OTP Fuse area is restricted. Accesses will result in an access error (FERROR)</p> <p>1 - Read and Programming access of the “secured” OTP Fuse area is allowed. If the corresponding program protection bit for an access is set, a program access is protected regardless of this bit's setting.</p>
18:16	Reserved	<p>Reserved. To ensure upward compatibility with future implementations, write back the value that is read from these bits.</p>

Secure Control (SECURE_CONTROL) Register

The SECURE_CONTROL register is used during Secure Entry Mode authentication. This register is used to establish Secure Mode transition and can be used at any time to exit from Secure Mode. The register, shown in Figure 25-8, is 16-bits wide and requires 16-bit access.

Secure Control Register (SECURE_CONTROL)

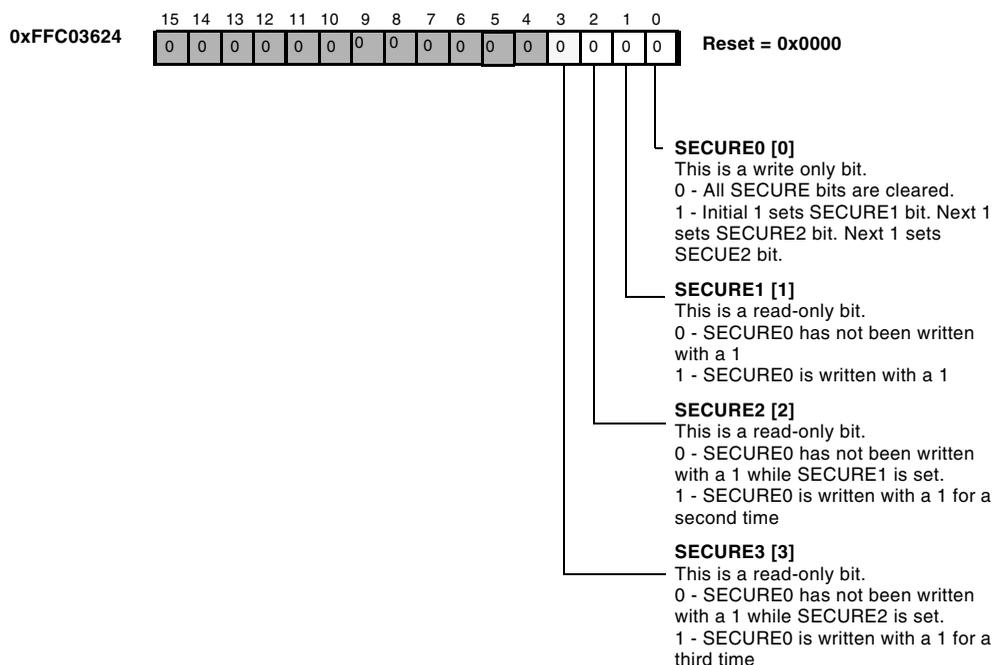


Figure 25-8. Secure Control Register

Security Registers

Table 25-6. Secure Control Register

Bit Position	Bit Name	Bit Description
		Reset = 0x0000
0	SECURE0	<p>SECURE 0</p> <p>This is a write only bit. A read always returns “0”. A 1 value can only be written to SECURE0 when in Secured Entry. The purpose of this control bit is to require 3 successive writes with a value of 1 to SECURE0 in order to enter Secure Mode.</p> <p>0 - When written with a “0” value, all SECURE bits within this register are cleared and Open Mode is entered. All SYSSWT bits are cleared with the exception of EMUOVR. If EMUOVR had been set by the user, it will remain set (until RESET is asserted or until it is written with a “0”).</p> <p>1 - Initially when written with a 1 value SECURE1 is set. With a subsequent 1 written SECURE2 is set. A subsequent 1 written will set SECURE3. Upon a set of SECURE3 Secure Mode is entered.</p>
1	SECURE1	<p>SECURE 1</p> <p>This is a read-only bit and indicates a successful write of SECURE0 with a data value of 1</p> <p>0 - SECURE0 has not been written with a 1 value</p> <p>1 - SECURE0 is written with a 1 value</p>
2	SECURE2	<p>SECURE 2</p> <p>This is a read-only bit and indicates two successful writes of SECURE0 with a data value of 1 has occurred</p> <p>0 - SECURE0 has not been written with a 1 value while SECURE1 was set.</p> <p>1 - SECURE0 is written with a 1 value for a second time.</p>
3	SECURE3	<p>SECURE 3</p> <p>This is a read-only bit and indicates three successful writes of SECURE0 with a data value of 1 has occurred.</p> <p>0 - SECURE0 has not been written with a 1 value while SECURE2 was set</p> <p>1 - SECURE0 is written with a 1 value for a third time. The part is currently in Secure Mode and the SYSSWT register is writable by Authenticated code.</p>

SECURE0 bit is user accessible and is used to exit from Secure Mode. Bits SECURE1, SECURE2, and SECURE3 are not user accessible and are accessed only by the firmware during the digital signature validation process.

Secure Status (SECURE_STATUS) Register

The SECURE_STATUS register provides information about the current secure state. This information can be used during security mode control as well as understanding why an authentication attempt has failed. The register, shown in Figure 25-9, is 16-bits wide and requires 16-bit access.

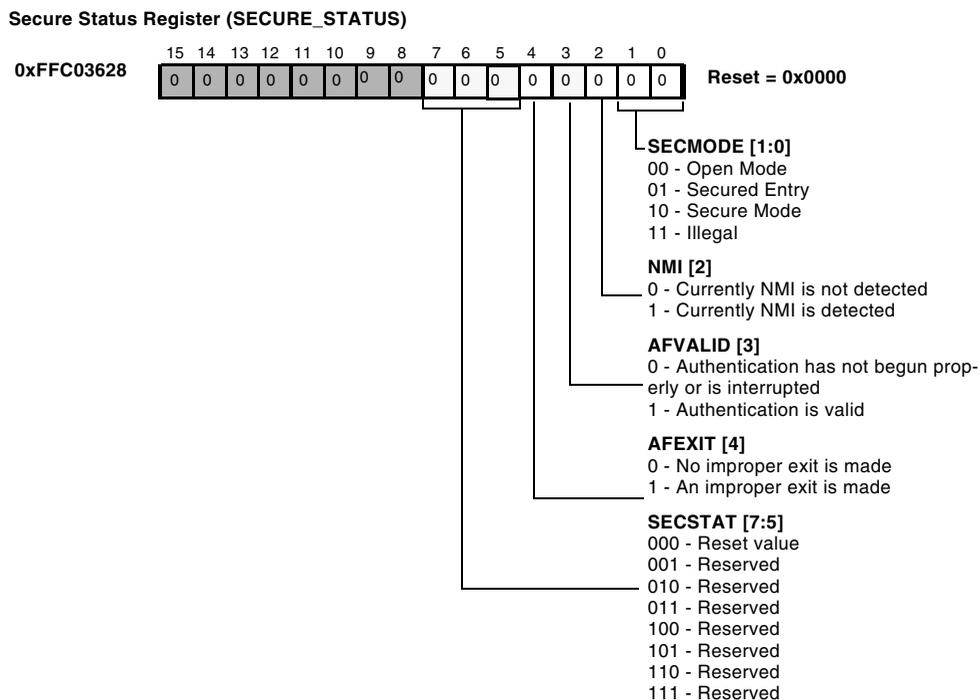


Figure 25-9. Secure Status Register

Security Registers

Table 25-7. Secure Status Register

Bit Position	Bit Name	Bit Description
		Reset = 0x0000
1:0	SECMODE	Secure Mode Control State These are read-only bits that reflects the current Secure Mode Control's state. 00 - Open Mode 01 - Secured Entry 10 - Secure Mode 11 - Illegal
2	NMI	This is a read-only bit that reflects the detection of NMI. 0 - Currently NMI is not detected. 1 - Currently NMI is detected.
3	AFVALID	Authentication Firmware Valid This is a read-only bit that reflects the state of the hardware monitor logic. If execution of authentication has begun properly and has had uninterrupted operation the authentication is considered valid. A valid authentication is required for Secured Entry and Secure Mode operation. 0 - Authentication has not begun properly or is interrupted. 1 - Authentication is valid and is progressing properly and uninterrupted.
4	AFEXIT	Authentication Firmware Exit This is a write one to clear status bit. In the event authentication has begun properly but has had an improper exit before completion, this bit is set. This can only occur on an exit from Secured Entry back to Open Mode. 0 - No improper exit is made while executing authentication firmware. 1 - An improper exit from authentication firmware is made.
7:5	SECSTAT	Secure Status These are some read write bits which is defined later. These are intended to pass a status back to the handler in the event an authentication has failed. 000 - Reset value 001 - Reserved 010 - Reserved 011 - Reserved 100 - Reserved 101 - Reserved 110 - Reserved 111 - Reserved

 Authentication Firmware Valid (AFVALID) is an input to the Secure State Machine and not an output control/status. AFVALID goes active based on hitting the correct Program Counter address.

Security Registers

26 SYSTEM RESET AND BOOTING

This document contains material that is subject to change without notice. The content of the boot ROM as well as hardware behavior may change across silicon revisions. See the anomaly list for differences between silicon revisions.

Overview

When the $\overline{\text{RESET}}$ input signal releases, the processor starts fetching and executing instructions from the on-chip boot ROM at address 0xEF00 0000.

The internal boot ROM includes a small boot kernel that loads application data from an external memory or host device. The application data is expected to be available in a well-defined format called the boot stream. A boot stream consists of multiple blocks of data and special commands that instruct the boot kernel how to initialize on-chip L1 memories as well as off-chip volatile memories.

The boot kernel processes the boot stream block-by-block until it is instructed by a special command to terminate the procedure and jump to the application's programmable start address, which traditionally is at 0xFFA0 0000 in on-chip L1 memory. This process is called "booting."

The processor features three dedicated input pins $\text{BMODE}[2:0]$ that select the booting mode. The boot kernel evaluates the BMODE pins and performs

booting from respective sources. [Table 26-1](#) describes the modes of the BMODE pins.

Table 26-1. Booting Modes

BMODE[2:0]	Boot Source	Description
000	No boot – idle	The processor does not boot. Rather, the boot kernel executes an IDLE instruction.
001	Boot from 8-bit or 16-bit external flash memory	The kernel boots from address 0x2000 0000 in asynchronous memory bank 0. The first byte of the boot stream contains further instructions whether the memory is eight or 16 bits wide.
010	Boot from internal serial SPI memory	The kernel boots from the on-chip, 24-bit addressable SPI flash memory via the SPI0 interface.
011	Boot from external serial SPI memory	After an initial device detection routine, the kernel boots from either 8-bit, 16-bit, 24-bit or 32-bit addressable SPI flash or EEPROM memory that connects to SPI0_SEL2.
100	Boot from SPI0 host	In this slave mode, the kernel expects the boot stream to be applied to SPI0 by an external host device.
111	Boot from UART0 host	In this slave mode, the kernel expects the boot stream to be applied to UART0 by an external host device. Prior to providing the boot stream, the host device is expected to send a 0x40 (ASCII '@') character that is examined by the kernel to adjust the bit rate.

Table 26-1. Booting Modes (Continued)

BMODE[2:0]	Boot Source	Description
110	Boot from SDRAM memory ¹	This mode provides a quick warm boot option. It requires the SDRAM controller to be programmed by the preboot routine based on OTP settings. The kernel starts booting from address 0x0000 0010.
101	Boot from on-chip OTP memory	This is the only stand-alone booting mode. It boots from the on-chip serial OTP memory. By default, the boot stream is expected to reside from OTP page 0x40 on. The start page can be altered by programming the OTP_START_PAGE field in OTP page PBS01H.

¹ This chapter uses the term SDRAM as a synonym for off-chip synchronous dynamic memory. For the ADSP-BF51x products, this includes the Mobile SDRAM standard.

Reset and Power-up

There is a subroutine in the boot kernel known as "preboot", which is executed prior to the boot mode being processed. This preboot routine can customize default values of MMR registers, such as the PLL and SDRAM controller registers. Furthermore, the SPI master mode can be customized. The preboot behavior is controlled through OTP programming.

To enable booting from volatile memories such as SDRAM, the SDRAM controller must be programmed *before* data can be loaded into the memory. Either the preboot or the initialization code mechanism can be used for this purpose.

Table 26-2 describes the six types of resets.

 Each type resets the core except for the System Software reset.

Table 26-2. 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 bits [15:4] of the SYSCR register. For more information, see “ System Reset Configuration (SYSCR) Register ” on page 26-73.
Wake up from hibernate state	Wake-up event as enabled in the VR_CTL register and reported by the PLL_STAT register.	Behaves as hardware reset except the WURESET bit in the SYSCR register is set. Booting can be performed conditionally on this event.
System software reset	Calling the <code>bfrom_SysControl()</code> routine with the SYCTRL_SYSRESET option triggers a system reset.	Resets only the peripherals, excluding the RTC (real time clock) block and most of the DPMC. The system software reset clears bits [15:13] and bits [11:4] of the SYSCR register, but not the WURESET bit. The core is not reset and a boot sequence is not triggered. Sequencing continues at the instruction after <code>bfrom_SysControl()</code> returns.
Watchdog timer reset	Programming the watchdog timer causes a watchdog timer reset.	Resets both the core and the peripherals, excluding the RTC block and most of the DPMC. (Because of the partial reset to the DPMC, the watchdog timer reset is not functional when the processor is in Sleep or Deep Sleep modes.) The SWRST or the SYSCR register can be read to determine whether the reset source was the watchdog timer.

Table 26-2. Resets (Continued)

Reset	Source	Result
Core double-fault reset	A core double fault occurs when an exception happens while the exception handler is executing. 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 core and the peripherals, excluding the RTC block and most of the DPMC. The SWRST or SYSCR registers can be read to determine whether the reset source was a core double-fault.
Software reset	This reset is caused by executing a RAISE 1 instruction or by setting the software reset (SYSRST) bit in the core debug control register (DBGCTL) through emulation software through the JTAG port. The DBGCTL register is not visible to the memory map.	Program executions vector to the 0xEF00 0000 address. The boot code executes an immediate system reset to ensure system consistency.

Hardware Reset

The processor chip reset is an asynchronous reset event. The $\overline{\text{RESET}}$ input pin must be deasserted after a specified asserted hold time to perform a hardware reset. For more information, see the product data sheet.

A hardware-initiated reset results in a system-wide reset that includes both core and peripherals. After the RESET pin is deasserted, 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 state of the BMODE pins.

The BMODE 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_{DDEXT} or GND. The pins and the corresponding bits in the SYSCR register configure the boot mode that is employed after hardware reset or system software reset. See the *Blackfin Processor Programming Reference* for further information.

Software Resets

A software reset may be initiated in three ways.

- By the watchdog timer, if appropriately configured
- Calling the `bfrom_SysControl()` API function residing in the on-chip ROM. For further information, see [Chapter 8, “Dynamic Power Management”](#).
- By the `RAISE 1` instruction

The watchdog timer resets both the core and the peripherals, as long as the processor is in Active or Full-On mode. A system software reset results in a reset of the peripherals without resetting the core and without initiating a booting sequence.

 In order to perform a system reset, the `bfrom_SysControl()` routine must be called while executing from L1 memory (either as cache or as SRAM). When L1 instruction memory is configured as cache, make sure the system reset sequence is read into the cache.

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 formatting the watchdog timer, the processor transitions into the boot mode sequence. The boot mode is configured by the state of the `BMODE` bit field in the `SYSCR` register.

A software reset is initiated by executing the `RAISE 1` instruction or setting the software reset (`SYSRST`) bit in the core debug control register (`DBGCTL`) (`DBGCTL` is not visible to the memory map) through emulation software through the JTAG port.

A software reset only affects the state of the core. The boot kernel immediately issues a system reset to keep consistency with the system domain.

Reset Vector

When reset releases, the processor starts fetching and executing instructions from address 0xEF00 0000. This is the address where the on-chip boot ROM resides.

On a hardware reset, the boot kernel initializes the `EVT1` register to 0xFFA0 0000. When the booting process completes, the boot kernel jumps to the location provided by the `EVT1` vector register. The content of the `EVT1` register is overwritten by the `TARGET ADDRESS` field of the first block of the applied boot stream. If the `BCODE` field of the `SYSCR` register is set to 1 (no boot option), the `EVT1` register is not modified by the boot kernel on software resets. Therefore, programs can control the reset vector for software resets through the `EVT1` register. This process is illustrated by the flow chart in [Figure 26-1](#).

The content of the `EVT1` register may be undefined in emulator sessions.

Servicing Reset Interrupts

The processor services a reset event like other interrupts. The reset interrupt has top priority. Only emulation events have higher priority. When coming out of reset, the processor is in supervisor mode and has full access to all system resources. The boot kernel can be seen as part of the reset service routine. It runs at the top interrupt priority level.

Even when the boot process has finished and the boot kernel passes control to the user application, the processor is still in the reset interrupt. To enter user mode, the reset service routine must initialize the `RETI` register and terminate with an `RTI` instruction.

For a programming example, see [“System Reset” on page 26-101](#).

[Listing 26-1](#) and [Listing 26-2 on page 26-102](#) show code examples that handle the reset event. See the *Blackfin Processor Programming Reference* for details on user and supervisor modes.

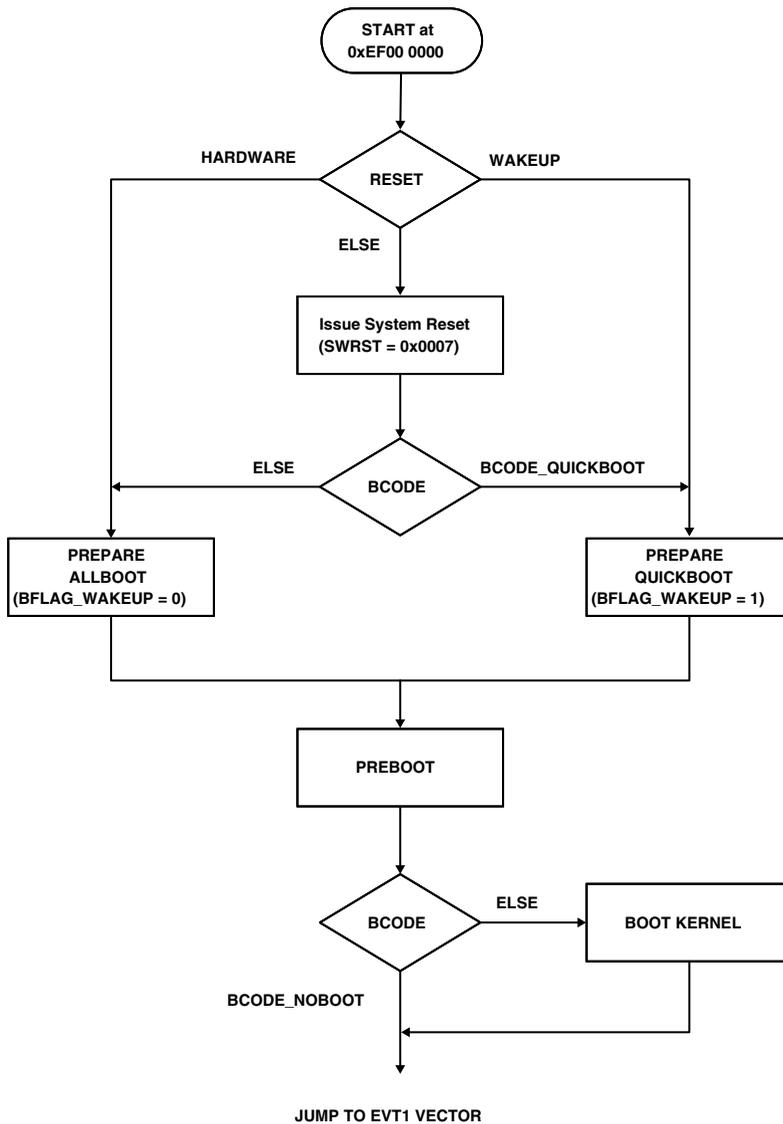


Figure 26-1. Global Boot Flow

Systems that do not work in an operating system environment may not enter user mode. Typically, the interrupt level needs to be degraded down to IVG15. [Listing 26-3](#) and [Listing 26-4 on page 26-103](#) show how this is accomplished.



Since the boot kernel is running at reset interrupt priority, NMI events, hardware errors and exceptions are not serviced at boot time. As soon as the reset service routine returns, the processor can service the events that occurred during the boot sequence. It is recommended that programs install NMI, hardware error, and exception handlers before leaving the reset service routine. This includes proper initialization of the respective event vector registers `EVTx`.

Preboot

After reset, the boot kernel residing in the on-chip boot ROM does not immediately start processing the boot stream. First it calls a subroutine called preboot, as shown in [Figure 26-2 on page 26-11](#) and [Figure 26-3 on page 26-12](#). The preboot routine customizes the default values of several system MMR registers based on user-configurable OTP (one-time programmable) memory. The following modules can be customized in this way.

- PLL and voltage regulator settings
- SDRAM controller settings
- Asynchronous EBIU settings

Some OTP bits customize the boot process:

- Bit rate of SPI boot modes
- Activation of SPI fast read mode
- Boot host wait (`HWAIT`) signal

Further OTP bits let the user disable certain features of the processor:

- Individual boot modes (for security reasons)

Finally, certain bits are already preset in the factory:

- Individual boot modes

Factory Page Settings (FPS)

The content of the boot ROM is identical across all ADSP-BF51x Blackfin processors. The factory settings prevent the boot ROM from accidentally accessing resources that are not present on a given processor, which would result in unpredictable behavior and/or hardware errors. The boot kernel goes to a safe idle state when the user configures the `BMODE` pins to a boot mode that is not available on a specific part.

For this purpose, the preboot routine always reads the `FPS01L` and `FPS01H` half pages from OTP memory.

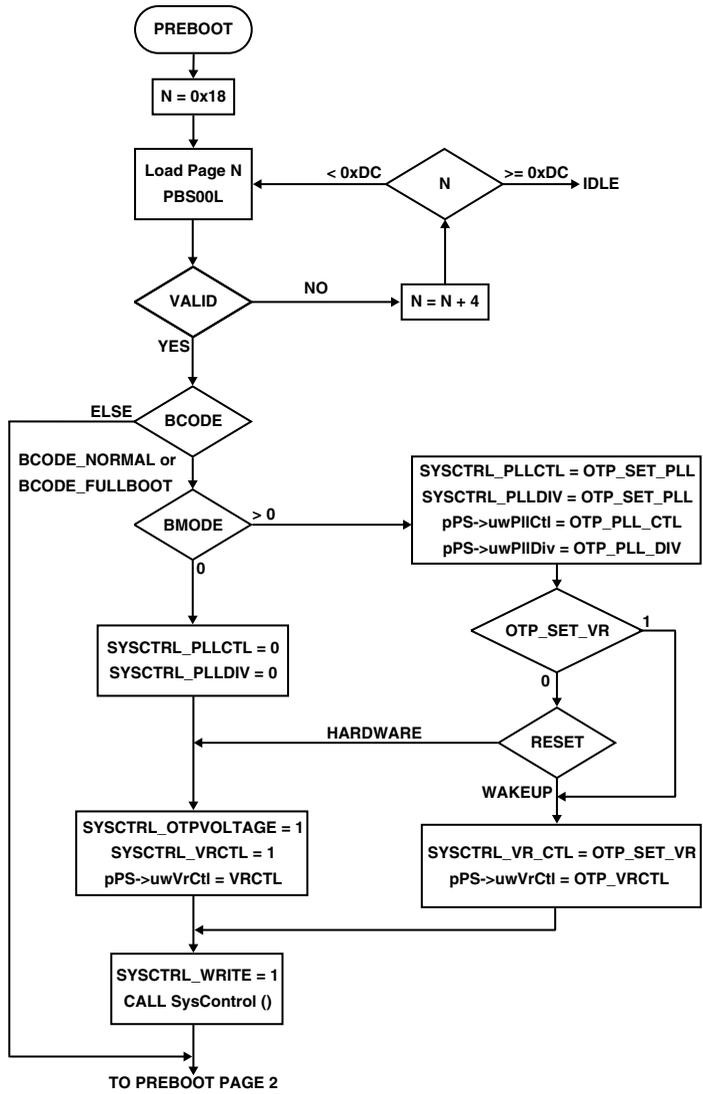


Figure 26-2. Preboot Flow 1 of 2

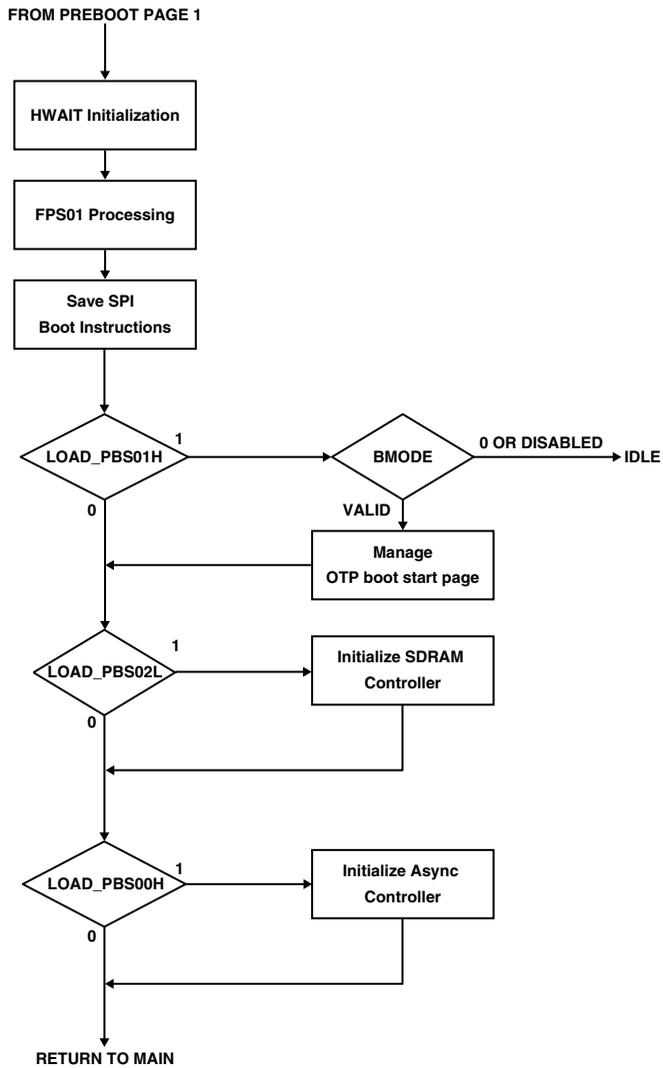


Figure 26-3. Preboot Flow 2 of 2

Preboot Page Settings (PBS)

Four OTP pages optionally enable the user to customize the behavior of the processor immediately after reset. These four pages (eight half pages) can be seen as one contiguous pre-boot settings (PBS) block. By default, the block spans OTP pages 0x18 to 0x1B. The OTP pages serve the following purposes:

- PBS00L (by default, on half page 0x18L, see [“Lower PBS00 Half Page” on page 26-78](#) for details)
 - PLL and voltage regulator settings
 - Boot customization
 - Instruction whether to load further half pages
- PBS00H (by default, on half page 0x18H, see [“Upper PBS00 Half Page” on page 26-81](#) for details)
 - Asynchronous EBIU register settings
- PBS01L (by default on half page 0x19L)
 - Reserved
- PBS01H (by default, on half page 0x19H, see [“Upper PBS01 Half Page” on page 26-82](#) for details)
 - Disabling of boot modes
 - OTP boot start page
- PBS02L (by default, on half page 0x1AL, see [“Lower PBS02 Half Page” on page 26-85](#) for details)
 - Synchronous EBIU register settings

- PBS02H (by default, on half page 0x1AH, see [“Upper PBS02 Half Page” on page 26-86](#) for details)

Reserved in current silicon revision. Do not use.

- PBS03L (by default, on half page 0x1BL, see [“Reserved Half Pages” on page 26-86](#) for details)

Reserved in current silicon revision. Do not use.

- PBS03H (by default, on half page 0x1BH, see [“Reserved Half Pages” on page 26-86](#) for details)

Reserved in current silicon revision. Do not use.

The preboot routine reads the main page `PBS00L` first. Since this page may instruct the preboot routine to alter the PLL settings, further pages may read more quickly. This page also instructs the preboot whether further OTP half pages have to be loaded and processed. By default, the `PBS00L` page reads all zeroes, and the preboot does not load further PBS pages.

Alternative PBS Pages

Especially during the development cycle, the user may fail to write the proper value to OTP memory and may make multiple attempts to get things right. Therefore, the `PBS00L` page provides a mechanism to invalidate the entire PBS block (consisting of pages (0x18, 0x19, 0x1A and 0x1B) and to use pages 0x1C to 0x1F instead. To do so, set the two `OTP_INVALID` bits (bits 62 and 63 on the `PBS00L` page). If both bits are set, the preboot routine disregards potential error codes returned by the `bfrom_OtpRead()` routine and continues processing from page 0x1C on. The active PBS block now spans the pages 0x1C to 0x1F. If the user wants to invalidate the second set of OTP pages as well, setting bits 62 and 63 on page 0x1C (which is the new `PBS00L` half page) instructs the preboot routine to continue at page 0x20, and so on.

Theoretically, this can be repeated up to page 0xD8L, if the pages are not required for other purposes. There are 49 chances to get things right, before a device may become useless. Note that every page that needs to be read by the preboot routine causes additional delay to the boot process.

Programming PBS Pages

Due to the need for error checking and correction (ECC), a 64-bit OTP half page must be written all at once. It is recommended that PBS pages be programmed only through the API function `bfrom_OtpWrite()`.

 If it is anticipated that the user is customizing the boot-related OTP pages for safety or security reasons, it is recommended that all PBS blocks be locked at production time to protect these pages from being tampered with in the field.

Reading OTP memory is subject to a potential failure rate. Since the preboot only accesses OTP memory through the `bfrom_OtpRead()` function, the ECC error correction is applied and the statistical failure rate is very low. However, the way the `PBS00L` page is tested for being invalid may at some point reduce the ECC reliability. To keep failure rates at a minimum, it is a good idea to duplicate the content of pages 0x18–0x1B on pages 0x1C–0x1F. For production parts, the final block should be followed by its exact copy to maintain the lowest failure rates. Then, even the unlikely case where one of the `OTP_INVALID` bits is read incorrectly would not cause the boot to fail.

Recovering From Misprogrammed PBS Pages

The preboot mechanism provides a powerful method to customize the chip to the needs of the user. However, as a downside, there are chances that invalid values programmed to the PBS pages prevent the processor from operating within required operating conditions. There is specific risk when the PLL and the voltage regulator are programmed with meaningless values during the development cycle.

In such cases, the boot mode `BMODE = b#000` helps. In this mode, the preboot routine does not attempt to read any of the user-programmable PBS pages, and the boot kernel does not try to boot any data. Rather, the processor is idled immediately after the `FPS` pages have been processed. Using the in-circuit emulator, the user then has the option to invalidate the actual PBS settings by overwriting both `OTP_INVALID` bits in the actual `PBS00L` with 1s.

For safety reasons, none of the boot modes, except the emulator, can get control over the processor when in this state.

Customizing Power Management

When the processor awakes with default PLL and voltage regulator settings, the preboot mechanism can be used to alter these settings to custom values before the boot process takes place. This is done by programming the OTP half page `PBS00L`.

If the `OTP_SET_PLL` bit is programmed to a 1, the value in the `OTP_PLL_DIV` bit field is copied into the `PLL_DIV` register, and the `OTP_PLL_CTL` bit field is copied into the `PLL_CTL` register, followed by the required `IDLE` instruction (if the contents of `PLL_CTL` are being altered).

If the `OTP_SET_VR` bit is programmed to a 1, the value in the `OTP_VR_CTL` bit field is copied into the `VR_CTL` register, followed by the required `IDLE` instruction (if the contents of `VR_CTL` are being altered).

The preboot mechanism invokes the `bfrom_SysControl()` routine to alter the PLL and the voltage regulator. The `bfrom_SysControl()` routine not only performs custom instructions, it also applies correction values from factory OTP pages `FPS01` and `FPS04`. See [Chapter 8, “Dynamic Power Management”](#) for details on the `bfrom_SysControl()` routine.

Customizing Booting Options

The OTP pages accessible by the preboot mechanism can also be used to customize some of the booting options. For example:

- SPI master boot mode operating frequency
- SPI master boot mode read operation mode
- Start page for OTP boot mode
- HWAIT signal behavior
- Disabling of unwanted boot modes

In SPI master boot mode, the `OTP_SPI_BAUD` register in the preboot half page `PBS00L` controls the value written to the `SPIO_BAUD` registers. By default, the clock divider value of 133 can be reduced in power-of-two steps. The table of values can be found in [“SPI Master Boot Modes” on page 26-60](#). The `OTP_SPI_BAUD` bit instructs the boot kernel to use the 0x0B SPI read command instead of the normal 0x03 read command when accessing the SPI memory device.

In OTP boot mode, the boot kernel normally assumes that the boot stream starts at OTP page 0x40L. The user can change this start page by programming the `OTP_START_PAGE` bit field in the preboot half page `PBS01H`.

The boot host wait (HWAIT) signal is available in all boot modes. If the `OTP_RESETOUT_HWAIT` bit in the preboot half page `PBS00L` is set, the boot kernel does not toggle HWAIT. Rather, it simply drives it to simulate a reset output signal.

If safety or security of an application is impacted by the existence of certain boot modes, the boot mode disable bits in preboot half page `PBS01H` can be used to disable unwanted boot modes. If a disabled boot mode is chosen by the `BMODE` pins, the boot kernel goes into a safe idle state and

stops processing. The half page `PBS01H` is only loaded when the `OTP_LOAD_PBS01H` bit in the `PBS00L` page is set.

Customizing the Asynchronous Port

The preboot half page `PBS00H` contains instructions to customize the asynchronous portion of the EBIU controller. This half page is only loaded and processed when the `OTP_LOAD_PBS00H` bit in the `PBS00L` half page is programmed to a 1.

The `OTP_EBIU_AMG` field is copied into the `EBIU_AMGCTL` register. While the lower bit controls the `CLKOUT` signal, the upper three `AMBEN` bits control which of the four asynchronous banks are enabled.

The preboot routine analyzes the three `AMBEN` bits and initializes the 16-bit portions (this routine is similar to the enabled banks in the `EBIU_AMBCTL0` and `EBIU_AMBCTL1` registers) with the value provided in the 16-bit `OTP_EBIU_AMBCTL` field. In this way, the bus timing of the asynchronous port can be customized prior to the boot process.

Customizing the Synchronous Port

Since many Blackfin applications require data and/or instruction code to be loaded into the SDRAM memory at boot time, the SDRAM controller must be initialized beforehand. This can be done by using either the [“Initialization Code” on page 26-35](#) or the preboot mechanism described here. For the SDRAM boot mode, only the preboot mechanism is valid.

The preboot half page `PBS02L` contains instructions to customize the MMR register of the SDRAM controller. This half page is only loaded and processed when the `OTP_LOAD_PBS02L` bit in the `PBS00L` half page is programmed to a “1”.

Basic Booting Process

Once the preboot routine returns, the boot kernel residing in the on-chip boot ROM starts processing the boot stream. The boot stream is either read from memory or received from a host processor. A boot stream represents the application data and is formatted in a special manner. The application data is segmented into multiple blocks of data. Each block begins with a block header. The header contains control words such as the destination address and data length information.

As [Figure 26-4](#) illustrates, the VisualDSP++ tools suite features a loader utility (`elfloader.exe`). The loader utility parses the input executable file (`.DXE`), segments the application data into multiple blocks, and creates the header information for each block. The output is stored in a loader file (`.LDR`). The loader file contains the boot stream and is made available to hardware by programming or burning it into non-volatile external memory. Refer to the *VisualDSP++ Loader Manual* for information on switches for loader files.

[Figure 26-5](#) shows the parallel or serial boot stream contained in a flash memory device. In host boot scenarios, the non-volatile memory more likely connects to the host processor rather than directly to the Blackfin processor. After reset, the headers are read and parsed by the on-chip boot ROM, and processed block-by-block. Payload data is copied to destination addresses, either in on-chip L1 memory or off-chip SRAM/SDRAM.

When the `BFLAG_INDIRECT` flag for any block is set, the boot kernel uses another memory block in L1 data bank B (by default, `0xFF90 7E00–0xFF90 7FFF`) for intermediate data storage. To avoid conflicts, the VisualDSP++ `elfloader` utility ensures this region is booted last.

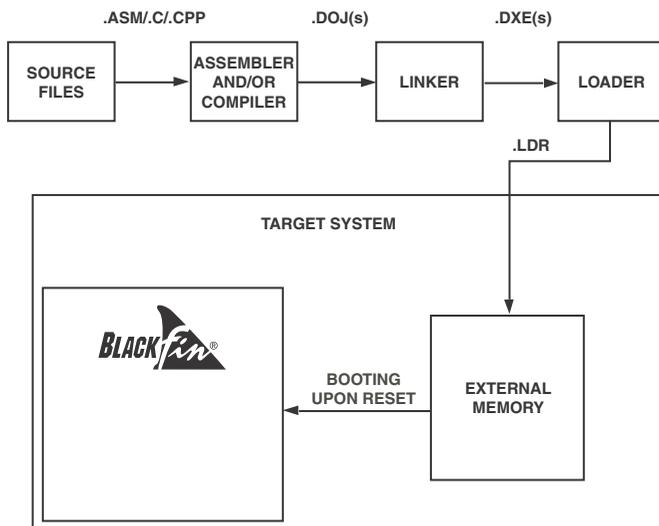


Figure 26-4. Project Flow for a Standalone System

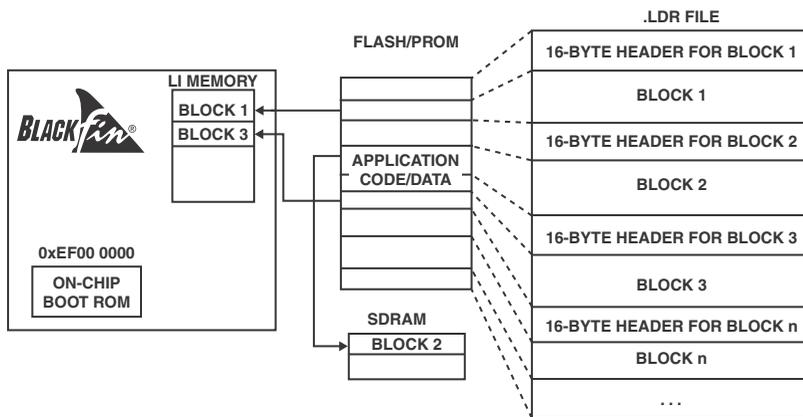


Figure 26-5. Booting Process

 Booting into scratchpad memory (0xFFB0 0000–0xFFB0 0FFF) is not supported. If booting to scratchpad memory is attempted, the processor hangs within the on-chip boot ROM. Similarly, booting into the upper 16 bytes of L1 data bank A (0xFF80 7FF0–0xFF80 7FFF by default) is not supported. These memory locations are used by the boot kernel for intermediate storage of block header information. These memory regions cannot be initialized at boot time. After booting, they can be used by the application during runtime.

The entire source code of the boot ROM is shipped with the VisualDSP++ tools installation. Refer to the source code for any additional questions not covered in this manual. Note that minor maintenance work may be done to the content of the boot ROM when silicon is updated.

Block Headers

A boot stream consists of multiple boot blocks, as shown in [Figure 26-5](#). Every block is headed by a 16-byte block header. However, every block does not necessarily have a payload, as shown in [Figure 26-6](#).

The 16 bytes of the block header are functionally grouped into four 32-bit words, the `BLOCK CODE`, the `TARGET ADDRESS`, the `BYTE COUNT`, and the `ARGUMENT` fields.

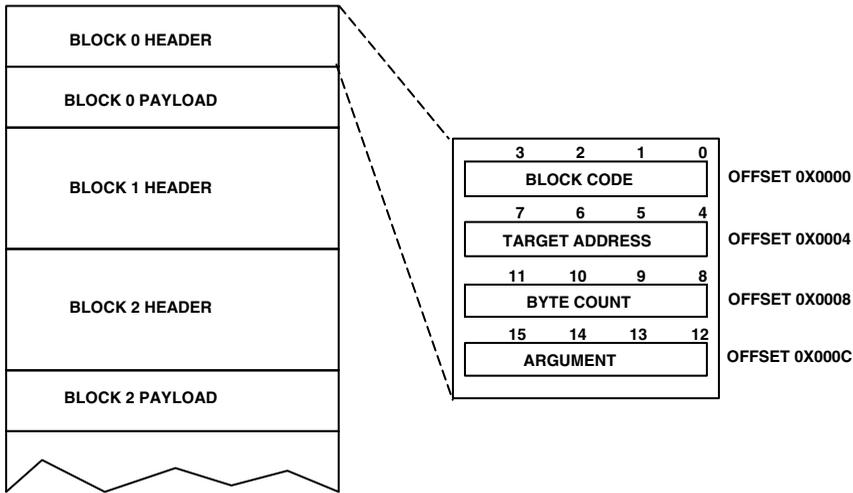


Figure 26-6. Boot Stream Headers

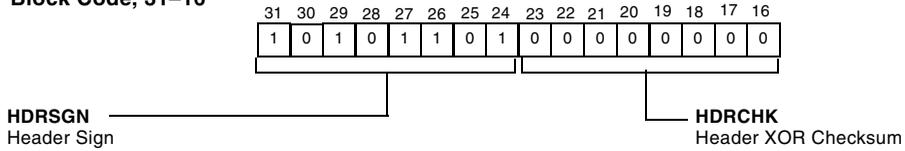
Block Code

The first 32-bit word is the `BLOCK CODE`. See [Figure 26-7](#).

DMA Code Field

The DMA code (`DMACODE`) field instructs the boot kernel whether to use 8-bit, 16-bit or 32-bit DMA and how to program the source modifier of a memory DMA. Particularly in case of memory boot modes, this field is interrogated by the boot kernel to differentiate the 8-bit, 16-bit, and 32-bit cases.

Block Code, 31–16



Block Code, 15–0

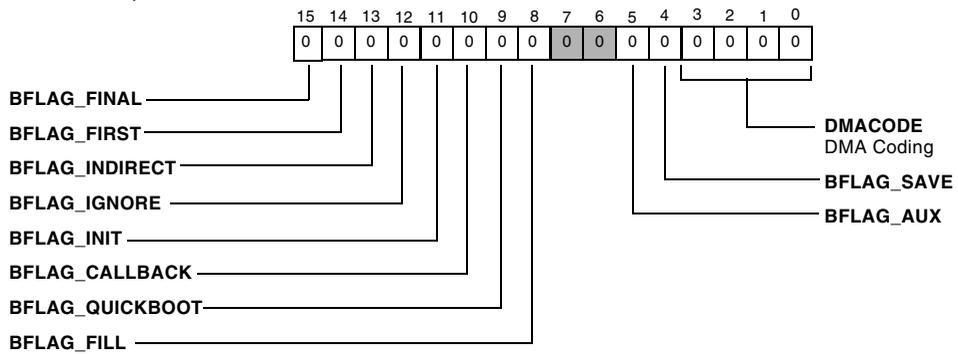


Figure 26-7. Block Code, 31–0

The boot kernel tests this field only on the first block and ignores the field in further blocks (See [Table 26-3](#)).

Table 26-3. Bus and DMA Width Coding

DMA Code	DMA Width	Source DMA Modify	Application
0	reserved ¹		
1	8-bit	1	Default 8-bit boot from 8-bit source ²
2	8-bit	2	Zero-padded 8-bit boot from 16-bit EBIU
3	8-bit	4	Zero-padded 8-bit boot from 32-bit EBIU ³
4	8-bit	8	Zero-padded 8-bit boot from 64-bit EBIU ⁴

Table 26-3. Bus and DMA Width Coding (Continued)

DMA Code	DMA Width	Source DMA Modify	Application
5	8-bit	16	Zero-padded 8-bit boot from 128-bit EBIU ⁴
6	16-bit	2	Default 16-bit boot from 16-bit source
7	16-bit	4	Zero-padded 16-bit boot from 32-bit EBIU ³
8	16-bit	8	Zero-padded 16-bit boot from 64-bit EBIU ⁴
9	16-bit	16	Zero-padded 16-bit boot from 128-bit EBIU ⁴
10	32-bit	4	Default 32-bit boot from 32-bit source ³
11	32-bit	8	Zero-padded 32-bit boot from 64-bit EBIU ⁴
12	32-bit	16	Zero-padded 32-bit boot from 128-bit EBIU ⁴
13	64-bit	8	Default 64-bit boot from 64-bit source ⁴
14	64-bit	16	Zero-padded 64-bit boot from 128-bit EBIU ⁴
15	128-bit	16	Default 128-bit boot from 128-bit source ⁴

- 1 Reserved to differentiate from ADSP-BF53x boot streams.
- 2 Used by all byte-wise serial boot modes.
- 3 Applicable only to memory boot modes and OTP mode. This code is expected by OTP boot mode.
- 4 Not supported by ADSP-BF51x Blackfin products.

Block Flags Field

Table 26-4. Block Flags

Bit	Name	Description
4	BFLAG_SAVE	Saves the memory of this block to off-chip memory in case of power failure or a hibernate request. This flag is not used by the on-chip boot kernel.
5	BFLAG_AUX	Nests special block types as required by special purpose second-stage loaders. This flag is not used by the on-chip boot kernel.
6	Reserved	
7	Reserved	
8	BFLAG_FILL	Tells the boot kernel to not process any payload data. Instead the target memory (specified by the TARGET ADDRESS and BYTE COUNT fields) is filled with the 32-bit value provided by the ARGUMENT word. The fill operation is always performed by 32-bit DMA; therefore target address and byte count must be divisible by four.
9	BFLAG_QUICKBOOT	Processes the block for full boot only. Does not process this block for a quick boot (warm boot).
10	BFLAG_CALLBACK	Calls a subfunction that may reside in on-chip or off-chip ROM or is loaded by an initcode in advance. Often used with the BFLAG_INDIRECT switch. If BFLAG_CALLBACK is set for any block, an initcode must register the callback function first. The function is called when either the entire block is loaded or the intermediate storage memory is full. The callback function can do advanced processing such as CRC checksum.
11	BFLAG_INIT	This flag causes the boot kernel to issue a CALL instruction to the target address of the boot block after the entire block is loaded. The initcode should return by an RTS instruction. It may or may not be overwritten by application data later in the boot process. If the code is loaded earlier or resides in ROM, the init block can be zero sized (no payload).

Table 26-4. Block Flags (Continued)

Bit	Name	Description
12	BFLAG_IGNORE	Indicates a block that is not booted into memory. It instructs the boot kernel to skip the number of bytes of the boot stream as specified by <code>BYTE_COUNT</code> . In master boot modes, the boot kernel simply modifies its source address pointer. In this case the <code>BYTE_COUNT</code> value can be seen as a 32-bit two's-complement offset value to be added to the source address pointer. In slave boot modes, the boot kernel actively loads and changes the payload of the block. In slave modes the byte count must be a positive value.
13	BFLAG_INDIRECT	Boots to an intermediate storage place, allowing for calling an optional callback function, before booting to the destination. This flag is used when the boot source does not have DMA support (TWI for example) and either the destination cannot be accessed by the core (L1 instruction SRAM) or cannot be efficiently accessed by the core (SDRAM or RAM). This flag is also used when <code>CALLBACK</code> requires access to data to calculate a checksum, or when performing tasks such as decryption or decompression.
14	BFLAG_FIRST	This flag, which is only set on the first block of a DXE, tells the boot kernel about the special nature of the <code>TARGET_ADDRESS</code> and the <code>ARGUMENT</code> fields. The <code>TARGET_ADDRESS</code> field holds the start address of the application. The <code>ARGUMENT</code> field holds the offset to the next DXE.
15	BFLAG_FINAL	This flag causes the boot kernel to pass control over to the application after the final block is processed. This flag is usually set on the last block of a DXE unless multiple DXEs are merged.

The `BFLAG_FIRST` flag must not be combined with the `BFLAG_FILL` flag. The `BFLAG_FIRST` flag may be combined with the `BFLAG_IGNORE` flag to deposit special user data at the top of the boot stream. Note the special importance of the VisualDSP++ elfloader `-readall` switch.

Header Checksum Field

The header checksum (`HDRCHK`) field holds a simple XOR checksum of the other 31 bytes in the boot block header. The boot kernel jumps to the error routine if the result of an XOR operation across all 32 header bytes

(including the `HDRCHK` value) differs from zero. The default error routine is a simple `IDLE;` instruction. The user can overwrite the default error handler using the `initcode` mechanism.

Header Sign Field

The header signature (`HDRSGN`) byte always reads as `0xAD` and is used to verify whether the block pointer actually points to a valid block. The `HDRSGN` byte can also be used as a boot stream version control. For the `ADSP-BF54x`, `ADSP-BF52x` and `ADSP-BF51x` Blackfin processors, the byte always reads `0xAD`. The `ADSP-BF53x` boot streams always read `0xFF`. The `ADSP-BF561` boot streams always read `0xA0`.

Target Address

This 32-bit field holds the target address where the boot kernel loads the block payload data. When the `BFLAG_FILL` flag is set, the boot kernel fills the memory with the value stored in the `ARGUMENT` field starting at this address. If the `BFLAG_INIT` flag is set the kernel issues a `CALL(TARGET ADDRESS)` instruction after the optional payload is loaded.

If the `BFLAG_FIRST` flag is set, the `TARGET ADDRESS` field contains the start address of the application to which the boot kernel jumps at the end of the boot process. This address will also be stored in the `EVT1` register. By default the VisualDSP++ elfloader utility sets this value to `0xFFA0 0000` for compatibility with other Blackfin products.

The target address should be divisible by four, because the boot kernel uses 32-bit DMA for certain operations. The target address must point to valid on-chip or off-chip memory locations. When booting to external memories, the memory controller must first be set up by either the pre-boot or the `initcode` mechanism. When booting through peripherals that do not support DMA transfers, such as the OTP boot mode, the `BFLAG_INDIRECT` flag must be set if the target address points to L1 instruction memory. For performance reasons this is also recommended when booting to off-chip memories.

For the OTP boot mode, the VisualDSP++ elfloader utility manages the `BFLAG_INDIRECT` flag automatically. Refer to the *VisualDSP++ Loader and Utilities* reference guide for manual control of the flag.

 Booting to scratchpad memory is not supported. The scratchpad memory functions as a stack for the boot kernel. The L1 data memory locations `0xFF807FF0` to `0xFF807FFF` are used by the boot kernel and should not be overwritten by the application. The memory range used for intermediate storage as controlled by the `BFLAG_INDIRECT` switch should only be booted after the last `BFLAG_INDIRECT` bit is processed. By default the address range `0xFF907E00–0xFF907FFF` is used for intermediate storage.

For normal boot operation, the target address points to RAM memory. There are however a few exceptions where the target address can point to on-chip or off-chip ROM. For example a zero-sized `BFLAG_INIT` block would instruct the boot kernel to call a subroutine residing in ROM or flash memory. This method is used to activate the CRC32 feature.

Byte Count

This 32-bit field tells the boot kernel how many bytes to process. Normally, this is the size of the payload data of a boot block. If the `BFLAG_FILL` flag is set there is no payload. In this case the `BYTE_COUNT` field uses the value in its `ARGUMENT` field to tell the boot kernel how many bytes to process.

The byte count is a 32-bit value that should be divisible by four. Zero values are allowed in all block types. Most boot modes are based upon DMA operation which are only 16-bit words for Blackfin processors. The boot kernel may therefore start multiple DMA work units for large boot blocks. This enables a single block to fill to zero the entire SDRAM memory, for example, resulting in compact boot streams. The `HWAIT` signal may toggle for each work unit.

If the `BFLAG_IGNORE` flag is set, the byte count is used to redirect the boot source pointer to another memory location. In master boot modes, the byte count is a two's-complement (signed long integer) value. In slave boot modes, the value must be positive.

Argument

This 32-bit field is a user variable for most block types. The value is accessible by the initcode or the callback routine and can therefore be used for optional instructions to these routines. When the CRC32 feature is activated, the `ARGUMENT` field holds the checksum over the payload of the block.

When the `BFLAG_FILL` flag is set there is no payload. The argument contains the 32-bit fill value, which is most likely a zero.

If the `BFLAG_FIRST` flag is set, the argument contains the relative next-DXE pointer for multi-DXE applications. For single-DXE applications the field points to the next free boot source address after the current DXE's boot stream.

Boot Host Wait (HWAIT) Feedback Strobe

The `HWAIT` feedback strobe is a handshake signal that is used to hold off the host device from sending further data while the boot kernel is busy.

On ADSP-BF51x processors this feature is implemented by a GPIO that is toggled by the boot kernel as required. The `PG0` GPIO is used for this purpose.

The signal polarity of the `HWAIT` strobe is programmable by an external resistor in the 10 k Ω range.

A pull-up resistor instructs the `HWAIT` signal to be active high. In this case the host is permitted to send header and footer data when `HWAIT` is low,

but should pause while `HWAIT` is high. This is the mode used in SPI slave boot on other Blackfin products.

Similarly, a pull-down resistor programs active-low behavior.

 Note that the `HWAIT` signal is implemented slightly differently than on ADSP-BF53x Blackfin processors. In the ADSP-BF51x processors, the meaning of the pulling resistor is inverted and `HWAIT` is asserted by default during reset and preboot.

After preboot, the boot kernel first senses the polarity on the respective `HWAIT` pin. Then it enables the output driver but keeps the signal in its asserted state. The signal is not released until the boot kernel is ready for data, or when a receive DMA is started. As soon as the DMA completes, `HWAIT` becomes active again.

The boot host wait signal holds the host from booting in any slave boot mode and prevents it from being overrun with data. The `HWAIT` signal is, however, available in all boot modes.

In general the host device must interrogate the `HWAIT` signal before every word that is sent. This requirement can be relaxed for boot modes using on-chip peripherals that feature larger receive FIFOs. However, the host must not rely on the DMA FIFO since its content is cleared at the end of a DMA work unit.

While the `HWAIT` signal is only used for boot purposes, it may also play a significant role after booting. In slave boot modes, for example, the host device does not necessarily know whether the Blackfin processor is in an active mode or a power-down mode. For example, the `HWAIT` signal can be used to signal when the processor is in hibernate mode.

Using `HWAIT` as Reset Indicator

While the `HWAIT` signal is mandatory in some boot modes, it is optional in others. When not required for booting, the behavior of the `HWAIT` signal

can be changed by programming the `OTP_RESETOUT_HWAIT` bit in OTP page `PBS00L`.

If this bit is set, `HWAIT` does not toggle during the boot process. Rather, after page `PBS00L` is processed (and therefore the PLL has settled) the pre-boot routine first enables the `HWAIT` GPIO as an input and senses its state. Then `HWAIT` becomes an output and is driven to the invert of the state that is sensed. An external pulling resistor is required. If using a pull-up resistor, the `HWAIT` signal is driven low for the rest of the boot process (and beyond). If using a pull-down resistor, `HWAIT` is driven high.

With a pull-down resistor, this feature can be used to simulate an active-low reset output. When the processor is reset, or in hibernate, the GPIO is in a high impedance state and `HWAIT` is pulled low by the resistor. As soon as the processor recovers and has settled the PLL again, the `HWAIT` is driven high and can alert external circuits.

Boot Termination

After the successful download of the application into the bootable memory, the boot kernel passes control to the user application. By default this is performed by jumping to the vector stored in the `EVT1` register. The boot kernel provides options to execute an `RTS` instruction or a `RAISE 1` instruction instead. The default behavior can be changed by an `initcode` routine. The `EVT1` register is updated by the boot kernel when processing the `BFLAG_FIRST` block. See [“Servicing Reset Interrupts”](#) on page 26-7 to learn how the application can take control.

Before the boot kernel passes program control to the application it does some housekeeping. Most of the registers that were used are changed back to their default state but some register values may differ for individual boot modes. DMA configuration registers and primary register control registers (`UART0_LCR`, `SPIO_CTL`, etc.) are restored, while others are purposely not restored. For example `SPIO_BAUD`, `UART0_DLH` and `UART_DLL`

remain unchanged so that settings obtained during the booting process are not lost.

Single Block Boot Streams

The simplest boot stream consists of a single block header and one contiguous block of instructions. With appropriate flag instructions the boot kernel loads the block to the target address and immediately terminates by executing the loaded block.

Table 26-5 shows an example of a single block boot stream header that could be loaded from any serial boot mode. It places a 256-byte block of instructions at L1 instruction SRAM address 0xFFA0 0000. The flags `BFLAG_FIRST` and `BFLAG_FINAL` are both set at the same time. Advanced flags, such as `BFLAG_IGNORE`, `BFLAG_INIT`, `BFLAG_CALLBACK` and `BFLAG_FILL`, do not make sense in this context and should not be used.

Table 26-5. Header for a Single Block Boot Stream

Field	Value	Comments
BLOCK CODE	0xAD33 C001	0xAD00 0000 XORSUM BFLAG_FINAL BFLAG_FIRST (DMACODE & 0x1)
TARGET ADDRESS	0xFFA0 0000	Start address of block and application code
BYTE COUNT	0x0000 0100	256 bytes of code
ARGUMENT	0x0000 0100	Functions as next-DXE pointer in multi-DXE boot streams

With the `BFLAG_FIRST` flag set, the `ARGUMENT` field functions as the next-DXE pointer. This is a relative pointer to the next free source address or to the next DXE start address in a multi-DXE stream.

Direct Code Execution

Applications may want to avoid long booting times and start code execution directly from 16-bit flash or SDRAM memory. This feature is called

direct code execution. This is a special case of boot termination that replaces the no-boot/bypass mode in the ADSP-BF53x Blackfin processors.

An initial boot block header is needed for the processor to fetch and execute program code from the boot device as early as possible. The safety mechanisms of the block, such as the header signature and the XOR checksum, avoid unpredictable processor behavior due to the boot memory not being programmed with valid data yet. Rather than blindly executing code, the boot kernel first executes the preboot routine for system customization, then loads the first block header and checks it for consistency. If the block header is corrupted, the boot kernel goes into a safe idle state and does not start code execution.

If the initial block header checks good, the boot kernel interrogates the block flags. If the block has the `BFLAG_FINAL` flag set, the boot kernel immediately terminates and jumps directly to the address stored in the `EVT1` register. To cause the boot kernel to customize the `EVT1` register in advance, the initial blocks must also have the `BFLAG_FIRST` flag set. The `TARGET ADDRESS` field is then copied to the `EVT1` register. In this way, the `TARGET ADDRESS` field of the initial block defines the start address of the application.

For example in `BMODE = 001`, when the block header described in [Table 26-6](#) is placed at address `0x2000 0000`, the boot kernel is instructed to issue a `JUMP` command to address `0x2000 0020`.

The development tools must be instructed to link the above block to address `0x2000 0000` and the application code to address `0x2000 0020`.

An example shown in “[Direct Code Execution](#)” on page 26-111 illustrates how this is accomplished using the VisualDSP++ tools suite.

Table 26-6. Initial Header for Direct Code Execution in $BMODE = 001$

Field	Value	Comments
BLOCK CODE	0xAD7B D006	0xAD00 0000 XORSUM BFLAG_FINAL BFLAG_FIRST BFLAG_IGNORE (DMACODE & 0x6)
TARGET ADDRESS	0x2000 0020	Start address of application code
BYTE COUNT	0x0000 0010	Ignores 16 bytes to provide space for control data such as version code and build data. This is optional and can be zero.
ARGUMENT	0x0000 0010	Functions as next-DXE pointer in multi-DXE boot streams

Similarly for direct code execution in the SDRAM boot mode ($BMODE = 110$), an initial block as shown in [Table 26-7](#) has to be linked to address 0x0000 0010.

Table 26-7. Initial Header for Direct Code Execution in $BMODE = 110$

Field	Value	Comments
BLOCK CODE	0xAD5B D006	0xAD000000 XORSUM BFLAG_FINAL BFLAG_FIRST BFLAG_IGNORE (DMACODE & 0x6)
TARGET ADDRESS	0x0000 0020	Start address of application code
BYTE COUNT	0x0000 0000	No bubble for control data
ARGUMENT	0x0000 0000	Functions as next-DXE pointer in multi-DXE boot streams

For multi-DXE boot streams, [Figure 26-11 on page 26-54](#) shows a linked list of initial blocks that represent different applications.

Advanced Boot Techniques

This section details techniques for advanced booting.

Initialization Code

Initcode routines are subroutines that the boot kernel calls during the booting process. The user can customize and speed up the booting mechanisms using this feature. Traditionally, an initcode is used to set up system PLL, bit rates, wait states and the SDRAM controller. If executed early in the boot process, the boot time can be significantly reduced.

After the payload data is loaded for a specific boot block, if the `BFLAG_INIT` flag is set, the boot kernel issues a `CALL` instruction to the target address of the block.

On ADSP-BF51x Blackfin processors, initcode routines follow the C language calling convention so they can be coded in C language or assembly.

The expected prototype is

```
void initcode(ADI_BOOT_DATA* pBootStruct);
```

The VisualDSP++ header files define the `ADI_BOOT_INITCODE_FUNC` type:

```
typedef void ADI_BOOT_INITCODE_FUNC (ADI_BOOT_DATA* ) ;
```

Optionally, the initcode routine can interrogate the formatting structure and customize its own behavior or even manipulate the regular boot process. A pointer to the structure is passed in the `R0` register. Assembly coders must ensure that the routine returns to the boot kernel by a terminating `RTS` instruction.

Initcodes can rely on the validity of the stack, which resides in scratchpad memory. The `ADI_BOOT_DATA` structure resides on the stack. Rules for register usage conform to the compiler conventions. See the *VisualDSP++ C/C++ Compiler and Library Manual* for more information.

In the simple case, initcodes consist of a single instruction section and are represented by a single block within the boot stream. This block has the `BFLAG_INIT` bit set.

An init block can consist of multiple sections where multiple boot blocks represent the initcode within the boot stream. Only the last block has the `BFLAG_INIT` bit set.

The VisualDSP++ elfloader utility ensures that the last of these blocks vectors to the initcode entry address. The utility instructs the on-chip boot ROM to execute a `CALL` instruction to the given target address.

When the on-chip boot ROM detects a block with the `BFLAG_INIT` bit set, it boots the block into Blackfin memory and then executes it by issuing a `CALL` to its target address. For this reason, every initcode must be terminated by an `RTS` instruction to ensure that the processor vectors back to the on-chip boot ROM for the rest of the boot process.

Sometimes initcode boot blocks have no payload and the `BYTE COUNT` field is set to zero. Then the only purpose of the block may be to instruct the boot kernel to issue the `CALL` instruction.

Initcode routines can be very different in nature. They might reside in ROM or SRAM. They might be called once during the booting process or multiple times. They might be volatile and be overwritten by other boot blocks after executing, or they might be permanently available after boot time. The boot kernel has no knowledge of the nature of initcodes and has no restrictions in this regard. Refer to the *VisualDSP++ Loader and Utilities Manual* for how this feature is supported by the tools chain.

It is the user's responsibility to ensure that all code and data sections that are required by the initcode are present in memory by the time the initcode executes. Special attention is required if initcodes are written in C or C++ language. Ensure that the initcode does not contain calls to the runtime libraries. Do not assume that parts of the runtime environment, such as the heap are fully functional. Ensure that all runtime components are loaded and initialized before the initcode executes.

The VisualDSP++ elfloader utility provides two different mechanisms to support the initcode feature.

- The `-init initcode.dxe` command line switch
- The `-initcall address/symbol` command line switch

If enabled by the VisualDSP++ elfloader `-init initcode.DXE` command line switch, the initcode is added to the beginning of the boot stream. Here, `initcode.DXE` refers to the user-provided custom initialization executable—a separate VisualDSP++ project. [Figure 26-8](#) shows a boot stream example that performs the following steps.

1. Boot initcode into L1 memory.
2. Execute initcode.
3. Initcode initializes the SDRAM controller and returns.
4. Overwrite initcode with final application code.
5. Boot data/code into SDRAM.
6. Continue program execution with block n.

Although `initcode.DXE` files are built as VisualDSP++ projects, they differ from standard projects. Initcodes provide only a callable sub-function, so they look more like a library than an application. Nevertheless, unlike library files (`.DLB` file extension), the symbol addresses have already been resolved by the linker.

An initcode is always a heading for the regular application code. Consequently whether the initcode consists of one or multiple blocks, it is not terminated by a `BFLAG_FINAL` bit indicator—this would cause the boot ROM to terminate the boot process.

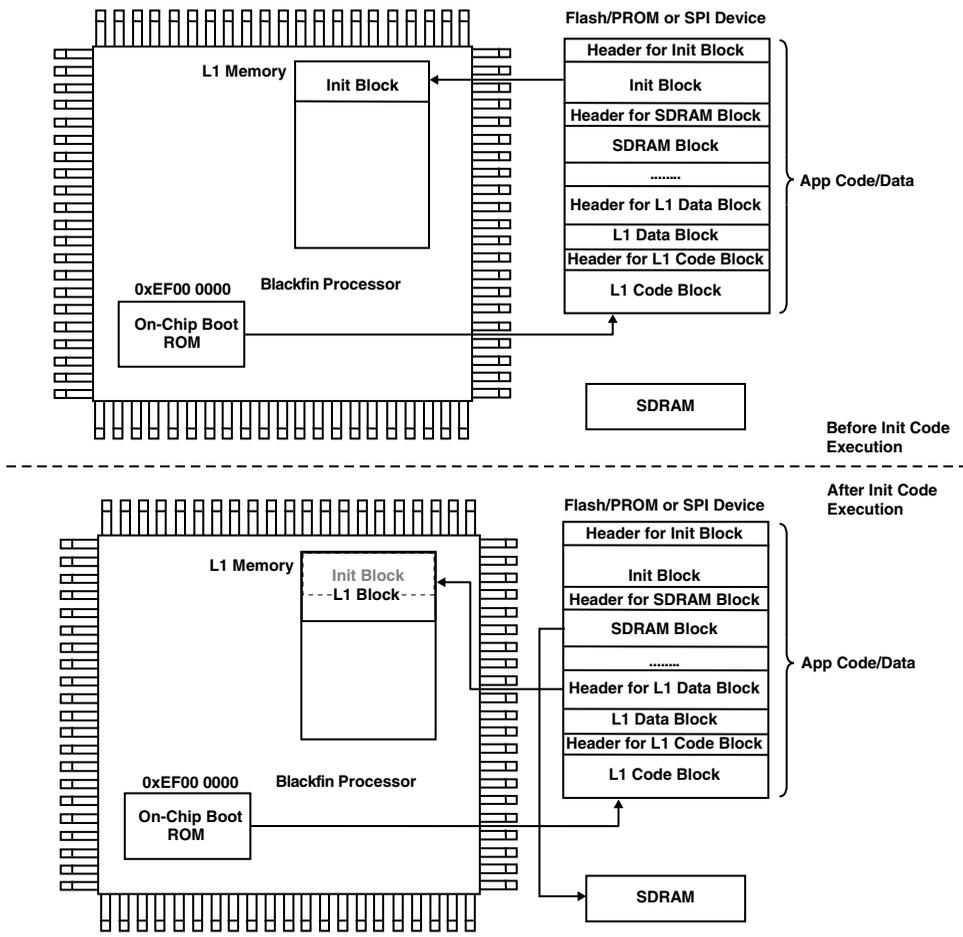


Figure 26-8. Initialization Code Execution/Boot

It is advantageous to have a clear separation between the initcode and the application by using the `-init` switch. If this separation is not needed, the `elfloader -initcall` command-line switch might be preferred. It enables fractions of the application code to be traded as initcode during the boot process. See the *VisualDSP++ Loader and Utilities Manual* for further details.

Initcode examples are shown in [“Programming Examples” on page 26-101](#).

Quick Boot

In some booting scenarios, not all memories need to be re-initialized. For example in a wake-up from hibernate state, off-chip SRAM might not be impacted if it was powered while the processor was in hibernate state. Dynamic RAM might also not be impacted if it was put into self-refresh mode before the processor powered down.

The ADSP-BF51x processor’s boot kernel can conditionally process boot blocks. The normal scenario is all boot, the shortened version is quick boot. It relies on the following primitives.

- The `SYSCR` register is read to determine what kind of boot is expected from the boot kernel. Refer to [Figure 26-24 on page 26-74](#).

The `WURESET` bit is used to distinguish between cold boot and warm boot situations and to identify wake-up from hibernate situations.

The `BCODE` bit field in the `SYSCR` register can overrule the native decision of the boot kernel for a software boot. See the flowchart in [Figure 26-1 on page 26-8](#).

- The `BFLAG_WAKEUP` bit in the `dFlag` word of the `ADI_BOOT_DATA` structure indicates that the final decision was to perform a quick boot. If the boot kernel is called from the application, then the application can control the boot kernel behavior by setting the `BFLAG_WAKEUP` flag accordingly. See the `dFlags` variable on [Figure 26-37 on page 26-91](#).
- The `BFLAG_QUICKBOOT` flag in the `BLOCK_CODE` word of the block header controls whether the current block is ignored for quick boot.

If both the global `BFLAG_WAKEUP` and the block-specific `BFLAG_QUICKBOOT` flags are set, the boot kernel ignores those blocks. But since the `BFLAG_INIT`, `BFLAG_CALLBACK`, `BFLAG_FINAL`, and `BFLAG_AUX` flags are internally cleared and the `BFLAG_IGNORE` flag is toggled, through double negation, the “ignore the ignore block” command instructs the boot kernel to process the block.

Although the `BFLAG_INIT` flag is suppressed in quick boot, the user may not want to combine the `BFLAG_INIT` flag with the `BFLAG_QUICKBOOT` flag. The initialization code can interrogate the `BFLAG_WAKEUP` flag and execute conditional instructions. For more information see [“Quickboot With Restore From SDRAM” on page 26-109](#).

Indirect Booting

The processor’s boot kernel provides a control mechanism to let blocks either boot directly to their final destination or load to an intermediate storage place, then copy the data to the final destination in a second step. This feature is motivated by the following requirements.

- Some boot modes do not use DMA. They load data by core instruction. The core cannot access some memories directly (for example L1 instruction SRAM), or is less efficient than the DMA in accessing some memories (for example, external SDRAM).
- In some advanced booting scenarios, the core needs to access the boot data during the booting process, for example in processing decompression, decryption and checksum algorithms at boot time. The indirect booting option helps speed-up and simplify such scenarios. Software accesses off-chip memory less efficiently and cannot access data directly if it resides in L1 instruction SRAM.

Indirect booting is not a global setting. Every boot block can control its own processing by the `BFLAG_INDIRECT` flag in the block header.

In general a boot block may not fit into the temporary storage memory so the boot kernel processes the block in multiple steps. The larger the temporary buffer, the faster the boot process. By default the L1 data memory region between 0xFF90 7E00 and 0xFF90 7FFF is used for intermediate storage. Initialization code can alter this region by modifying the `pTempBuffer` and `dTempByteCount` variables in the `ADI_BOOT_DATA` structure. The default region is at the upper end of a physical memory block. When increasing the `dTempByteCount` value, `pTempBuffer` also has to change.

Callback Routines

Callback routines, like initialization codes, are user-defined subroutines called by the boot kernel at boot time. The `BFLAG_CALLBACK` flag in the block header controls whether the callback routine is called for a specific block.

There are several differences between initcodes and callback routines. While the `BFLAG_INIT` flag causes the boot kernel to issue a `CALL` instruction to the target address of the specific boot block, the `BFLAG_CALLBACK` flag causes the boot kernel to issue a `CALL` instruction to the address held by the `pCallbackFunction` pointer in the `ADI_BOOT_DATA` structure. While a boot stream can have multiple individual initcodes, it can have just one callback routine. In the standard boot scenario, the callback routine has to be registered by an initcode prior to the first block that has the `BFLAG_CALLBACK` flag set.

The purpose of the callback routine is to apply standard processing to the block data. Typically, callback routines contain checksum, decryption, decompression, or hash algorithms. Checksum or hash words can be passed through the block header `ARGUMENT` field.

Since callback routines require access to the payload data of the boot blocks, the block data must be loaded before it can be processed. Unlike initcodes, a callback usually resides permanently in memory. If the block

is loaded to L1 instruction memory or off-chip memory, the `BFLAG_CALLBACK` flag is likely combined with the `BFLAG_INDIRECT` bit. The boot kernel performs these steps in the following order.

1. Data is loaded into the temporary buffer defined by the `pTempBuffer` variable.
2. The `CALL` to the `pCallbackFunction` is issued.
3. After the callback routine returns, the memory DMA copies data to the destination.

If a block does not fit into the temporary buffer, for example when the `BLOCK_COUNT` is greater than the `dTempByteCount` variable, the three steps are executed multiple times until all payload data is loaded and processed. The boot kernel passes the parameter `dCbFlags` to the callback routine to tell it that it is being invoked the first or the last time for a specific block. To store intermediate results across multiple calls the callback routine can use the `uwUserShort` and `dUserLong` variables in the `ADI_BOOT_DATA` structure.

Callback routines meet C language calling conventions for subroutines. The prototype is as follows.

```
s32 CallbackFunction (ADI_BOOT_DATA* pBootStruct,  
ADI_BOOT_BUFFER* pCallbackStruct, s32 dCbFlags);
```

The VisualDSP++ header file defines the `ADI_BOOT_CALLBACK_FUNC` type the following way:

```
typedef s32 ADI_BOOT_CALLBACK_FUNC (ADI_BOOT_DATA*,  
ADI_BOOT_BUFFER*, s32 );
```

The `pBootStruct` argument is passed in `R0` and points to the `ADI_BOOT_DATA` structure used by the boot kernel. These are handled by the `pTempBuffer` and `dTempByteCount` variables as well as the `pHeader` pointer to the `ARGUMENT` field. The callback routine may process the block further by modifying the `pTempBuffer` and `dTempByteCount` variables.

The `pCallbackStruct` structure passed in `R1` provides the address and length of the data buffer. When the `BFLAG_INDIRECT` flag is not set, the `pCallbackStruct` contains the target address and byte count of the boot block. If the `BFLAG_INDIRECT` flag is set, the `pCallbackStruct` contains a copy of the `pTempBuffer`. Depending on the size of the boot block and processing progress, the byte count provided by `pCallbackStruct` equals either `dTempByteCount` or the remainder of the byte count.

When the `BFLAG_INDIRECT` flag is set along with the `BFLAG_CALLBACK` flag, memory DMA is invoked by the boot kernel after the callback routine returns. This memory DMA relies on the `pCallbackStruct` structure not the global `pTempBuffer` and `dTempByteCount` variables.

The callback routine can control the source of the memory DMA by altering the content of the `pCallbackStruct` structure, as may be required if the callback routine performs data manipulation such as decompression.

The `dCbFlags` parameter passed in `R2` tells the callback routine whether it is invoked the first time (`CBFLAG_FIRST`) or whether it is called the last time (`CBFLAG_FINAL`) for a specific block. The `CBFLAG_DIRECT` flag indicates that the `BFLAG_INDIRECT` bit is not active so that the callback routine will only be called once per block. When the `CBFLAG_DIRECT` flag is set, the `CBFLAG_FIRST` and `CBFLAG_FINAL` flags are also set.

```
#define CBFLAG_FINAL      0x0008
#define CBFLAG_FIRST     0x0004
#define CBFLAG_DIRECT    0x0001
```

A callback routine also has a boolean return parameter in register `R0`. If the return value is non-zero, the subsequent memory DMA does not execute. When the `CBFLAG_DIRECT` flag is set, the return value has no effect.

Error Handler

While the default handler simply puts the processor into idle mode, an initcode routine can overwrite this pointer to create a customized error handler. The expected prototype is

```
void ErrorFunction (ADI_BOOT_DATA* pBootStruct, void  
*pFailingAddress);
```

Use an initcode to write the entry address of the error routine to the `pErrorFunction` pointer in the `ADI_BOOT_DATA` structure. The error handler has access to the boot structure and receives the instruction address that triggered the error.

CRC Checksum Calculation

The ADSP-BF51x Blackfin processors provide an initcode and a callback routine in ROM that can be used for CRC32 checksum generation during boot time. The checksum routine only verifies the payload data of the blocks. The block headers are already protected by the native XOR checksum mechanism.

Before boot blocks can be tagged with the `BFLAG_CALLBACK` flag to enable checksum calculation on the blocks, the boot stream must contain an initcode block with no payload data and with the CRC32 polynomial in the block header `ARGUMENT` word.

The initcode registers a proper CRC32 wrapper to the `pCallbackFunction` pointer. The registration principle is similar to the XOR checksum example shown in [“Programming Examples” on page 26-101](#).

Load Functions

All boot modes are processed by a common boot kernel algorithm. The major customization is done by a subroutine that must be registered to the

`pLoadFunction` pointer in the `ADI_BOOT_DATA` structure. Its simple prototype is as follows.

```
void LoadFunction (ADI_BOOT_DATA* pBootStruct);
```

The VisualDSP++ header files define the following type:

```
typedef void ADI_BOOT_LOAD_FUNC (ADI_BOOT_DATA* ) ;
```

For a few scenarios some of the flags in the `dFlags` word of the `ADI_BOOT_DATA` structure, such as `BFLAG_PERIPHERAL` and `BFLAG_SLAVE`, slightly modify the boot kernel algorithm.

The boot ROM contains several load functions. One performs a memory DMA for flash boot, others perform peripheral DMAs or load data from booting source by polling operation. The first is reused for fill operation and indirect booting as well.

In second-stage boot schemes, the user can create customized load functions or reuse the original `BFROM_PDMA` routine and modify the `pDmaControlRegister`, `pControlRegister` and `dControlValue` values in the `ADI_BOOT_DATA` structure. The `pDmaControlRegister` points to the `DMAX_CONFIG` or `MDMA_Dx_CONFIG` register. When the `BFLAG_SLAVE` flag is not set, the `pControlRegister` and `dControlValue` variables instruct the peripheral DMA routine to write the control value to the control register every time the DMA is started.

Load functions written by users must meet the following requirements.

- Protect against `dByteCount` values of zero.
- Multiple DMA work units are required if the `dByteCount` value is greater than 65536.
- The `pSource` and `pDestination` pointers must be properly updated.

In slave boot modes, the boot kernel uses the address of the `dArgument` field in the `pHeader` block as the destination for the required dummy

DMA's when payload data is consumed from `BFLAG_IGNORE` blocks. If the load function requires access to the block's `ARGUMENT` word, it should be read early in the function.

The most useful load functions `BFROM_MDMA` and `BFROM_PDMA` are accessible through the jump table. Others, do not have entries in the jump table. Their start address can be determined with the help of the hook routine when calling the respective `BFROM_SPIBOOT`, `BFROM_OTPBOOT` etc. functions. In this way they can be repurposed for runtime utilization.

Calling the Boot Kernel at Runtime

The boot kernel's primary purpose is to boot data to memory after power-up and reset cycles. However some of the routines used by the boot kernel might be of general value to the application. The boot ROM supports reuse of these routines as C-callable subroutines. Programs such as second-stage boot kernels, boot managers, and firmware update tools may call the function in the ROM at runtime. This could load entirely different applications or a fraction of an application, such as a code overlay or a coefficient array.

To call these boot kernel subroutines, the boot ROM provides an API at address `0xEF00 0000` in the form of a jump table.

When calling functions in the boot ROM, the user must ensure the presence of a valid stack following C language conventions. See the `VisualDSP++ Compiler` documentation for details.

Debugging the Boot Process

If the boot process fails, very little information can be gained by watching the chip from outside. In master boot modes, the interface signals can be observed. In slave boot modes only the `HWAIT` signal tells about the progress of the boot process.

However, by using the emulator, there are many possibilities for debugging the boot process. The entire source code of the boot kernel is provided with the VisualDSP++ installation. This includes the project executable (DXE) file. The `LOAD SYMBOLS` feature of the VisualDSP++ IDDE helps to navigate the program. Note that the content of the ROM might differ between silicon revisions. Hardware breakpoints and single-stepping capabilities are also available.

Table 26-8 shows program symbols that are of interest.

Table 26-8. Boot Kernel Symbols for Debug

Symbol	Comment
<code>_bootrom.assert.default</code>	If the program counter halts at the <code>IDLE</code> instruction at the <code>_bootrom.assert.default</code> address, either the boot kernel or the preboot has detected an error condition and will not continue the boot process. A misformatted boot stream, or invalid PBS settings are the most likely causes of such an error. The <code>RETS</code> register points to the failing routine. When stepping a couple of instructions further, there is a way to ignore the error and to continue the boot process by clearing the <code>>ASTAT</code> register while the emulator steps over the subsequent <code>IF CC JUMP 0</code> instruction.
<code>_bootrom.bootmenu</code>	If the emulator hits a hardware breakpoint at the <code>_bootrom.bootmenu</code> address, this indicates that the preboot returned properly. Otherwise the program may hang during preboot due to improper PBS settings or invalid boot modes.
<code>_bootrom.bootkernel.entry</code>	If the emulator hits a hardware breakpoint at the <code>_bootrom.bootkernel.entry</code> label, this indicates that device detection or autobaud returned properly.
<code>_bootrom.bootkernel.breakpoint</code>	This is a good address to place a hardware breakpoint. The boot kernel loads a new block header at this breakpoint. The block header can be watched at address <code>0xFF80 7FF0</code> or wherever the <code>pHeader</code> points to.

Table 26-8. Boot Kernel Symbols for Debug (Continued)

Symbol	Comment
<code>_bootrom.bootkernel.initcode</code>	All payload data of the current block is loaded by the time the program passes the <code>_bootrom.bootkernel.initcode</code> label. The boot kernel is about to interrogate the <code>BFLAG_INIT</code> flag. If set, the <code>initcode</code> can be debugged.
<code>_bootrom.bootkernel.exit</code>	Once the boot kernel arrives at the <code>_bootrom.bootkernel.exit</code> label, it detects a <code>BFLAG_FINAL</code> flag. After some housekeeping, it jumps to the <code>EVT1</code> vector.

The boot kernel also generates a circular log file in scratch pad memory. While the `pLogBuffer` and the `dLogByteCount` variables describe the location and dimension of the log buffer, the `pLogCurrent` points to the next free location in the buffer. The log file is updated whenever the kernel passes the `_bootrom.bootkernel.breakpoint` label.

At each pass, nine 32-bit words are written to the log file, as follows.

- block code word (`dBlockCode`) of the block header
- target address (`pTargetAddress`) of the block header
- byte count (`dByteCount`) of the block header
- argument word (`dArgument`) of the block header
- source pointer (`pSource`) of the boot stream
- block count (`dBlockCount`)
- internal copy of the `dBlockCode` word OR'ed with `dFlags`
- content of the `SEQSTAT` register
- `0xFFFF FFFA (-6)` constant

The ninth word is overwritten by the next entry set, so that `0xFFFF FFFA` always marks the last entry in the log file.

Most of the data structures used by the boot kernel reside on the stack in scratchpad memory. While executing the boot kernel routine (excluding subroutines), the P5 points to the ADI_BOOT_DATA structure. Type “(ADI_BOOT_DATA*) \$P5” in the VisualDSP++ expression window to see the structure content.

Boot Management

Blackfin processor hardware platforms may be required to run different software at different times. An example might be a system with at least one application and one in-the-field firmware upgrade utility. Other systems may have multiple applications, one starting then terminating, to be replaced by another application. Conditional booting is called boot management. Some applications may self-manage their booting rules, while others may have a separate application that controls the process, namely a boot manager.

In a master boot mode where the on-chip boot kernel loads the boot stream from memory, the boot manager is a piece of Blackfin software which decides at runtime what application is booted next. This may simply be based on the state of a GPIO input pin interrogated by the boot manager, or it may be the conclusion of complex system behavior.

Slave boot scenarios are different from master boot scenarios. In slave boot modes, the host masters boot management by setting the Blackfin processor to reset and then applying alternate boot data. Optionally, the host could alter the BMODE configuration pins, resulting in little impact to the Blackfin processor since the intelligence is provided by the host device.

Booting a Different Application

The boot ROM provides a set of user-callable functions that help to boot a new application (or a fraction of an application). Usually there is no

need for the boot manager to deal with the format details of the boot stream.

These functions are:

- `BFROM_MEMBOOT` discussed in “Flash Boot Modes” on page 26-57 and “SDRAM Boot Mode” on page 26-60
- `BFROM_SPIBOOT` discussed in “SPI Master Boot Modes” on page 26-60
- `BFROM_OTPBOOT` discussed in “OTP Boot Mode” on page 26-70

The user application, the boot manager application, or an initcode can call these functions to load the requested boot data. Using the `BFLAG_RETURN` flag the user can control whether the routine simply returns to the calling function or executes the loaded application immediately.

These ROM functions expect the start address of the requested boot stream as an argument. For `BFROM_MEMBOOT`, this is a Blackfin memory address, for `BFROM_SPIBOOT` it is a serial address. The SPI function can also accept the code for the GPIO pin that controls the device select strobe of the SPI memory.

Multi-DXE Boot Streams

If the start addresses of all the boot streams are predefined, the boot manager needs only to call the ROM functions directly. However since the addresses tend to vary from build to build they may have to be calculated at runtime.

In the world of the VisualDSP++ elfloader, a boot stream is always generated from a DXE file. It is therefore common to talk about multi-DXE or multi-application booting. When the elfloader utility accepts multiple DXE files on its command line, it generates a contiguous boot image by default. The second boot stream is appended immediately to the first one.

Since the utility updates the `ARGUMENT` field of all `BFLAG_FIRST` blocks, the `ARGUMENT` field of a `BFLAG_FIRST` block is called next-DXE pointer (NDP).

The next-DXE pointer of the first DXE boot stream points relatively to the start address of the second DXE boot stream. A multi-DXE boot image can be seen as a linked list of boot streams. The next-DXE pointer of the last DXE boot stream points relatively to the next free address. This is illustrated by an example shown in the next two figures. [Figure 26-9](#) shows a commented sketch as an example. [Figure 26-10](#) shows a screenshot of the Blackfin loader file viewer utility for the same example. The `LdrViewer` utility is not part of the VisualDSP++ tools suite. It is a third-party freeware product available on www.dolomitics.com.

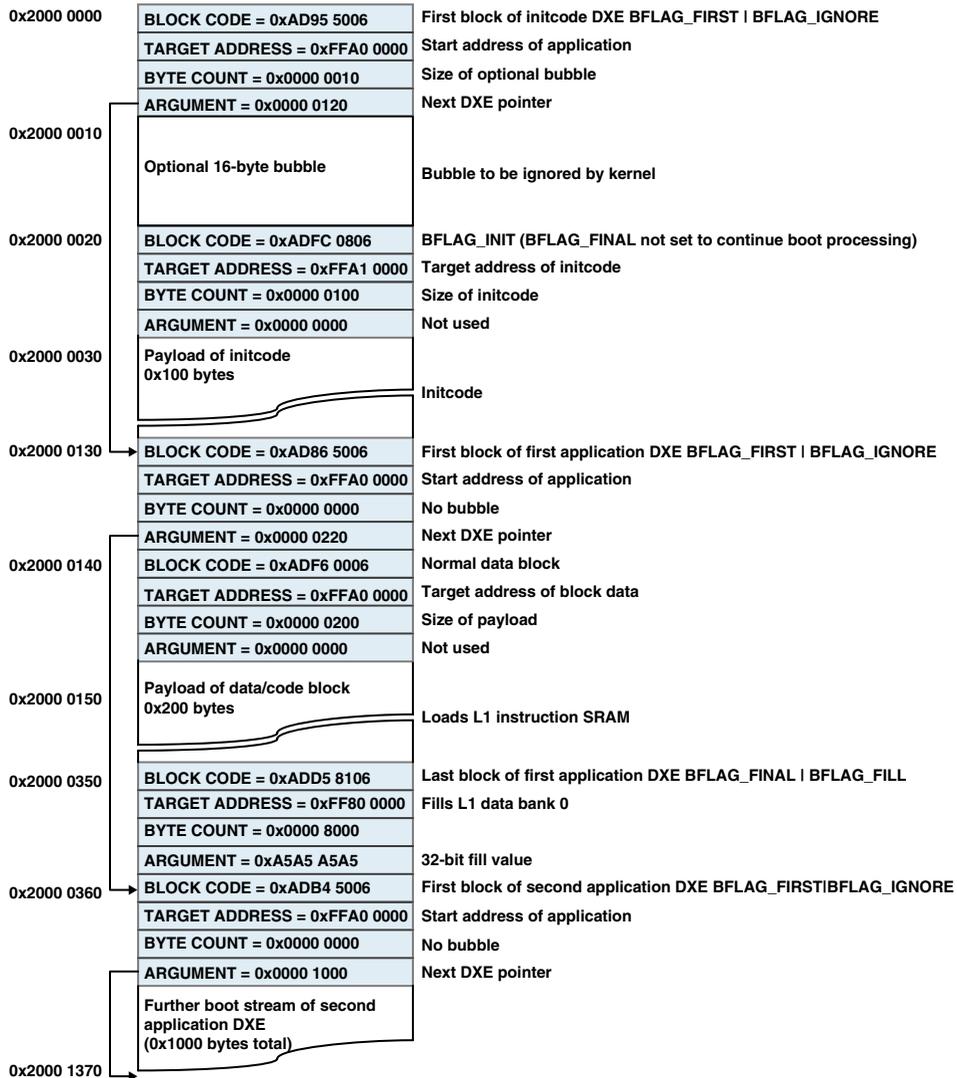


Figure 26-9. Multi-DXE Boot Stream Example for Flash Boot

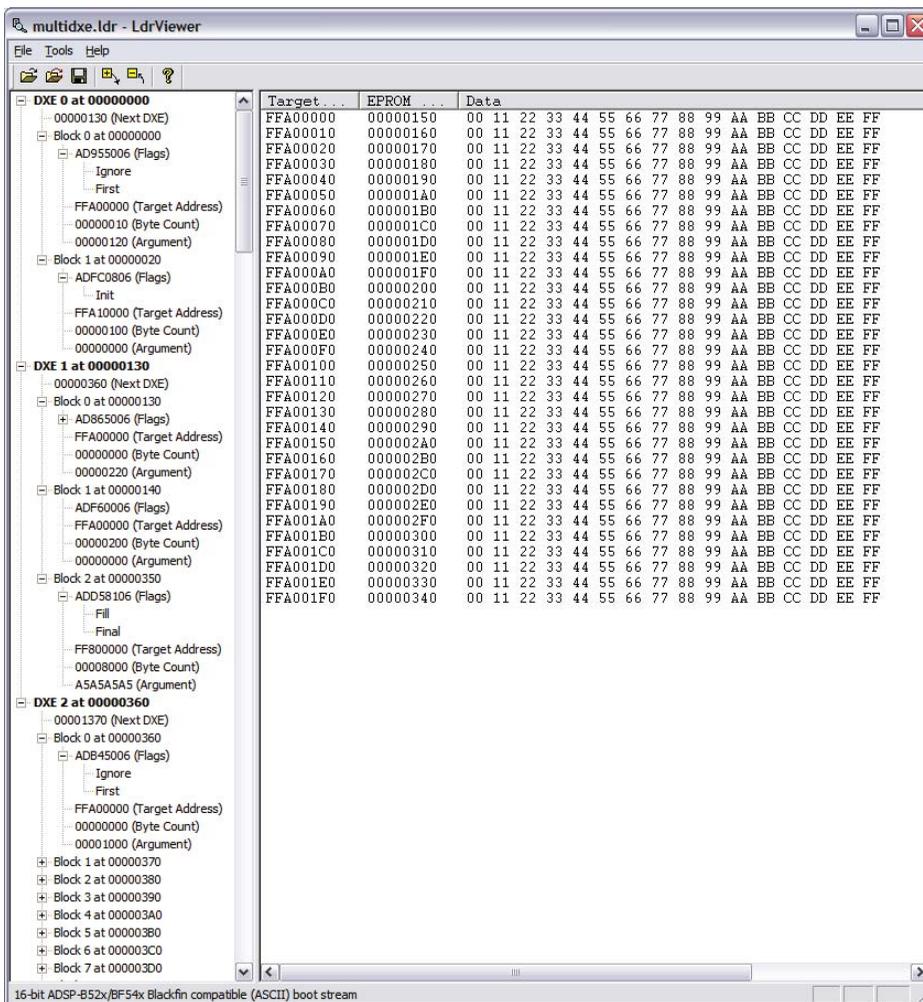


Figure 26-10. LdrViewer Screen Shot

Boot management principles are not only applicable to multi-DXE boot streams. The same scheme, as shown in [Figure 26-11](#), can be applied to direct code executions of multiple applications. See “[Direct Code Execution](#)” on [page 26-32](#) for more information. The example shows a linked list of initial block headers that instruct the boot kernel to terminate

immediately and to start code execution at the address provided by the TARGET ADDRESS field of the individual blocks. There is nothing in the boot ROM that prevents multi-DXE applications from mixing regular boot streams and direct code execution blocks.

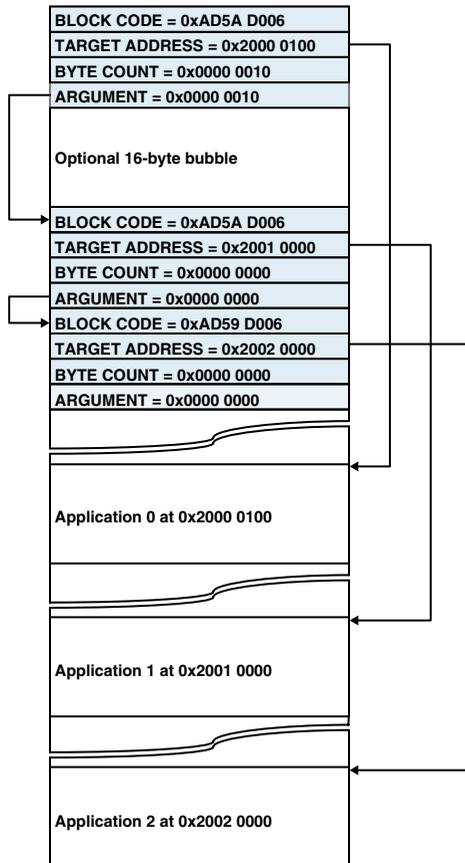


Figure 26-11. Multi-DXE Direct Code Execution Arrangement Example

Determining Boot Stream Start Addresses

The ROM functions `BFROM_MEMBOOT`, `BFROM_SPIBOOT`, etc. not only allow the application to boot a subroutine residing at a given start address, they also assist in walking through linked multi-DXE streams.

When the `BFLAG_NEXTDXE` bit in `dFlags` is set and these functions are called, the system does not boot but instead walks through the boot stream following the next-DXE pointers. The `dBlockCount` parameter can be used to specify the DXE of interest. The routines then return the start address of the requested DXE's boot stream.

Initialization Hook Routine

When the ROM functions `BFROM_MEMBOOT`, `BFROM_SPIBOOT`, etc. are called, they create an instance of the `ADI_BOOT_DATA` structure on the stack and fill the items with default values. If the `BFLAG_HOOK` is set, the boot kernel invokes a callback routine which was passed as the fourth argument of the ROM routines, after the default values have been filled. The hook routine can be used to overwrite the default values. Every hook routine should fit the prototype:

```
void hook (ADI_BOOT_DATA* pBS);
```

The VisualDSP++ header files define the `ADI_BOOT_HOOK_FUNC` type the following way:

```
typedef void ADI_BOOT_HOOK_FUNC (ADI_BOOT_DATA*);
```

The hook function also gives access to the DMA load function used by the respective boot mode, which can be used for general purposes at runtime. For example, in the `BFROM_SPIBOOT` case, an instance of the load function:

```
ADI_BOOT_LOAD_FUNC *pSpiLoadFunction;
```

can be initialized by equipping the hook function with the instruction:

```
pSpiLoadFunction = pBS->pLoadFunction;
```

Specific Boot Modes

This section discusses individual boot modes and the required hardware connections.

The boot modes differ in terms of the booting source— for example whether data is loaded through the SPI or the parallel interface. Boot modes can also be grouped into slave boot modes and master boot modes.

In slave boot modes, the Blackfin processor functions as a slave to any host device, which is typically another embedded processor, an FPGA device or even a desktop computer. Likely, the Blackfin processor $\overline{\text{RESET}}$ input is controlled by the host device. So, usually the host sets $\overline{\text{RESET}}$ first, then waits until the preboot routine terminates by sensing the HWAIT output, and finally provides the boot data.

If a Blackfin processor, configured to operate in any of the slave boot modes, awakens from hibernate, it cannot boot by its own control. A feedback mechanism has to be implemented at the system level to inform the host device whether the processor is in hibernate state or not. The HWAIT strobe is an important primitive in such systems.

In the master boot modes, the Blackfin processor usually does not need to be synchronized and can load the boot data by itself. Master modes typically read from memory. This can be parallel memory such as flash devices, or serial memory that is read through SPI interfaces.

Memory boot modes should also be differentiated from peripheral boot modes. Boot modes that load boot streams through memory DMA are referred to as memory boot mode, reading data from regular memory. Peripheral modes load boot data through peripherals such as UART. All memory boot modes are master modes. The boot source is typically non-volatile memory, such as a flash or EPROM device or even on-chip ROM. When supported by the system in warm boot scenarios, the boot source can also be SRAM or SDRAM.

Whether from the host (slave booting mode) or from memory (master booting mode), the boot source does not need to know about the structure of the boot stream.

No Boot Mode

When the `BMODE` pins are all tied low (`BMODE = 000`), the Blackfin processor does not boot. Instead it processes factory-programmed OTP pages, then executes an `IDLE` instruction, preventing it from executing any instructions provided by the regular boot source. The purpose of this mode is to bring the processor up to a clean state after reset.

This mode helps to recover from malicious OTP configuration since it prevents execution of the user-controllable portion of the preboot routine.

When connecting an emulator and starting a debug session, the processor awakens from an idle due to the emulation interrupt and can be debugged in the normal manner.

 The no boot mode is not the same as the bypass mode featured by the ADSP-BF53x Blackfin processor. To simulate that bypass mode feature using `BMODE = 001`, see [“Direct Code Execution” on page 26-32](#) and [“Direct Code Execution” on page 26-111](#).

Flash Boot Modes

These booting modes are intended to boot from flash or EEPROM memories or even from battery-buffered SRAMs. The flash boot modes are activated by `BMODE = 001`. Although this is a single `BMODE` setting, the ADSP-BF51x Blackfin products support various configurations.

- Boot from 8-bit asynchronous flash memory
- Boot from 16-bit asynchronous flash memory

By default, the boot kernel does not alter any EBIU registers. Therefore, traditional asynchronous flash is assumed and maximum wait states are applied. By programming OTP half pages `PBS00L` and `PBS00H`, the user has the option to instruct the preboot routine to alter the EBIU registers as desired. In this way, the EBIU can be preset to access the flash device in either page mode or burst mode. There are also options to customize bus settings, such as wait states and `ARDY` behavior.

After the preboot routine returns and `HWAIT` is deasserted the first time, the boot kernel loads an initial burst of four 16-bit words. Then it interrogates the `DMACODE` field in the byte loaded from the `0x2000 0000` address. For flash mode, the DMA options shown in [Table 26-9](#) are supported.

Table 26-9. DMA Options

DMACODE	DMA Width	Source Modify	Comment
1	8	1	Not recommended Provides ADSP-BF533 style 8-bit boot from 16-bit flash memory
2	8	2	8-bit MDMA boots from 8-bit flash mapped to lower byte of address bus.
6	16	2	16-bit MDMA boots from 16-bit flash
10	32	4	32-bit MDMA boots from 16-bit flash

The `DMACODE` field is filled by the `elfloader` utility based on boot mode, `-width` and `-dmawidth` settings. See the *VisualDSP++ Loader and Utility Manual* for details.

After the boot kernel has loaded and interpreted the first four 16-bit words, it continues loading the rest of the first block header and processes the boot stream.

Hardware configurations for the individual modes are shown in [Figure 26-12](#) and [Figure 26-13](#). The chip select is always controlled by the

$\overline{\text{AMS0}}$ strobe. This maps the boot stream to the Blackfin processor's address 0x2000 0000.

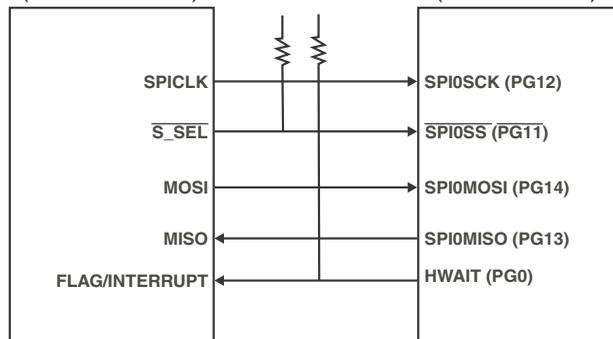


Figure 26-12. 8-Bit Flash Interconnection

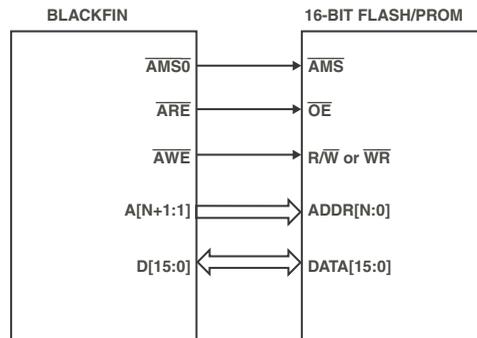


Figure 26-13. 16-Bit Flash Interconnection

Some flash devices provide write protection mechanisms, which can be activated during the power-up and reset cycles of the Blackfin processor. In the absence of such mechanisms, a pull-up resistor on the $\overline{\text{AMS0}}$ strobe prevents the chip select from floating when the state of the processor is unknown.

The boot mode $\text{BMODE} = 001$ can also be used to instruct the boot kernel to terminate immediately and directly execute code from the 16-bit flash

memory instead. Code execution from 8-bit flash memory is not supported. See [“Direct Code Execution” on page 26-32](#) for details.

SDRAM Boot Mode

From the boot kernel perspective, the SDRAM boot mode is just another memory boot mode like flash boot. The only differences are that the boot stream is expected at address 0x0000 0010 and the initial eight bytes are loaded by two 32-bit loads.

From the application point of view, SDRAM boot is a completely different scheme. Since SDRAM is volatile memory, $B_{MODE} = 110$ is not a valid setting when the processor and the memories have just been powered up. This mode can only be used as a dynamically applied B_{MODE} setting to install warm boot scenarios.

OTP programming is required to boot from SDRAM. Other boot modes can configure the SDRAM controller by execution of an initcode. But in the case of SDRAM boot, the initcode cannot be loaded without having the SDRAM controller already configured.

SDRAM boot is meaningful when the Blackfin processor is in hibernate state or is completely shut off for power savings while the SDRAM is kept alive in self-refresh mode.

Users who prefer to execute code out of SDRAM, rather than performing a boot from it, may refer to [“Direct Code Execution” on page 26-32](#) for details.

SPI Master Boot Modes

The ADSP-BF51x processors feature booting from on-chip as well as from off-chip SPI memories. The internal SPI boot mode ($B_{MODE} = 010$) boots from the on-chip SPI memory using the 24-bit addressing scheme on $SPI0$. An internal strobe signal controls the on-chip SPI memory via a virtual GPIO signal represented as $PH8$. This signal is not available off chip.

The external SPI boot mode ($B_{MODE} = 011$) boots from SPI memories connected to the $\overline{SPIOSEL2}$ interface. 8-, 16-, 24-, and 32-bit address words are supported. Standard SPI memories are read using either the standard 0x03 SPI read command or the 0x0B SPI fast read command.

i Unlike other Blackfin processors, the ADSP-BF51x Blackfin processors have no special support for DataFlash devices from Atmel. Nevertheless, DataFlash devices can be used for booting and are sold as standard 24-bit addressable SPI memories. They also support the fast read mode. If used for booting, DataFlash memory must be programmed in the power-of-2 page mode.

For booting, the SPI memory is connected as shown in [Figure 26-14](#).

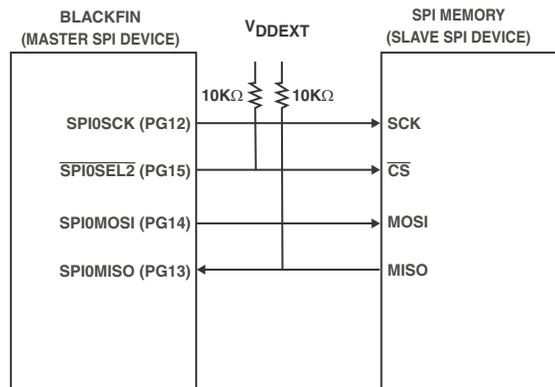


Figure 26-14. Blackfin to SPI Memory Connections

The pull-up resistor on the MISO line is required for automatic device detection. The pull-up resistor on the $\overline{SPIOSEL2}$ line ensures that the memory is in a known state when the Blackfin GPIO is in a high-impedance state (for example, during reset). A pull-down resistor on the SPIO_SCK line displays cleaner oscilloscope plots during debugging.

For SPI master boot, the SPE, MSTR and SZ bits are set in the SPIO_CTL register. For details see [Chapter 17, “SPI-Compatible Port Controller”](#). With

TIMOD = 2, the receive DMA mode is selected. Clearing both the CPOL and CPHA bits results in SPI mode 0. The boot kernel does not allow SPI0 hardware to control the $\overline{\text{SPIOSEL2}}$ pin. Instead, this pin is toggled in GPIO mode by software. Initialization code is allowed to manipulate the `uwSsel` variable in the `ADI_BOOT_DATA` structure to extend the boot mechanism to a second SPI memory connected to another GPIO pin.

By default, the boot kernel sets the `SPIO_BAUD` register to a value of 133, resulting in a bit rate of $\text{SCLK}/266$. This default value can be altered by programming the 4-bit `OTP_SPI_BAUD` field in OTP page `PBS00L` to one of the values in [Table 26-10](#).

Table 26-10. Bit Rate

OTP_SPI_BAUD	SPIO_BAUD	Bit Rate
b#0000	133	$\text{SCLK}/(2 \times 133)$
b#0001	Reserved	
b#0010	2	$\text{SCLK}/(2 \times 2)$
b#0011	4	$\text{SCLK}/(2 \times 4)$
b#0100	8	$\text{SCLK}/(2 \times 8)$
b#0101	16	$\text{SCLK}/(2 \times 16)$
b#0110	32	$\text{SCLK}/(2 \times 32)$
b#0111	64	$\text{SCLK}/(2 \times 64)$

Similarly, the boot kernel uses the standard `0x03` SPI read command, by default. Programming the `OTP_SPI_FASTREAD` bit in OTP page `PBS00L` enables the fast read mode where the boot kernel uses the `0x0B` read command instead and transmits a dummy zero byte after the address bytes.

The `OTP_SPI_FASTREAD` and `OTP_SPI_BAUD` values apply both to internal and external SPI master boot modes.

SPI Device Detection Routine

Since $BMODE = 011$ supports booting from various SPI memories, the boot kernel automatically detects what type of memory is connected. To determine whether the SPI memory device requires an 8-, 16-, 24- or 32-bit addressing scheme, the boot kernel performs a device detection sequence prior to booting. The $MISO$ signal requires a pull-up resistor, since the routine relies on the fact that memories do not drive their data outputs unless the right number of address bytes are received.

Boot from internal SPI memory does not perform any device detection. Internal SPI flash is known to require 24-bit addressing. Since the $MISO$ signal is not pulled up internal, device detection is not recommended at all in $BMODE = 010$.

Initially, the boot kernel transmits a read command (either $0x03$ or $0x0B$) on the $MOSI$ line, which is immediately followed by two zero bytes. Once the transmission is finished, the boot kernel interrogates the data received on the $MISO$ line. If it does not equal $0xFF$ (usually a $DMACODE$ value of $0x01$ is expected), then an 8-bit addressable device is assumed.

If the received value equals $0xFF$, it is assumed that the memory device has not driven its data output yet and that the $0xFF$ value is due to the pull-up resistor. Thus, another zero byte is transmitted and the received data is tested again. If it differs from $0xFF$, either a 16-bit addressable device (standard mode) or an 8-bit addressable device (fast read mode) is assumed.

If the value still equals $0xFF$, device detection continues. Device detection aborts immediately if a byte different than $0xFF$ is received. The boot kernel continues with normal boot operation and it re-issues a read command to read from address 0 again. The first block header is loaded by two read sequences, further block headers and block payload fields are loaded by separate read sequences.

Figure 26-15 illustrates how individual devices would behave.

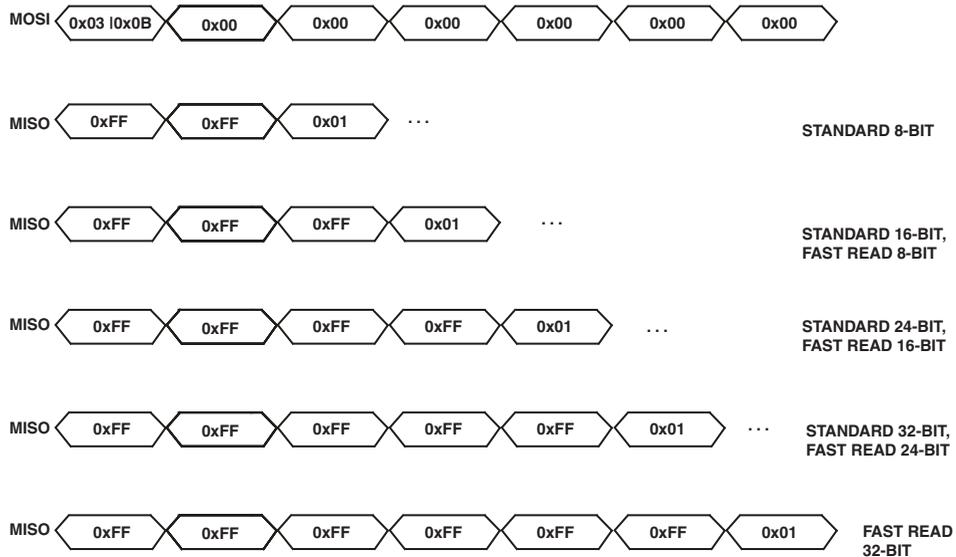


Figure 26-15. SPI Device Detection Principle

Figure 26-16 shows the initial signaling when a 24-bit addressable SPI memory is connected in SPI master boot mode. After $\overline{\text{RESET}}$ releases and preboot has processed relevant OTP pages, a 0x03 command is transmitted to the MOSI output, followed by a number of 0x00 bytes. The 24-bit addressable memory device returns a first data byte at the fourth zero byte. Then, the device detection has completed and the boot kernel re-issues a 0x00 address to load the boot stream.

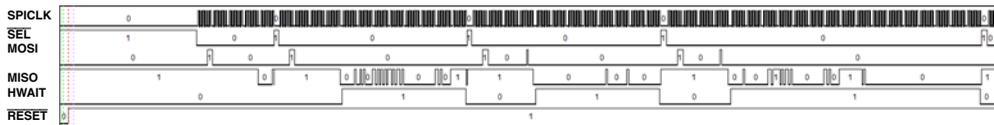


Figure 26-16. Typical SPI Master Boot Waveforms

SPI Slave Boot Mode

For SPI slave mode boot ($BMODE = 100$), the Blackfin processor is consuming boot data from an external SPI host device. SPI0 is configured as an SPI slave device. The hardware configuration is shown in Figure 26-17. As in all slave boot modes, the host device controls the Blackfin processor \overline{RESET} input.

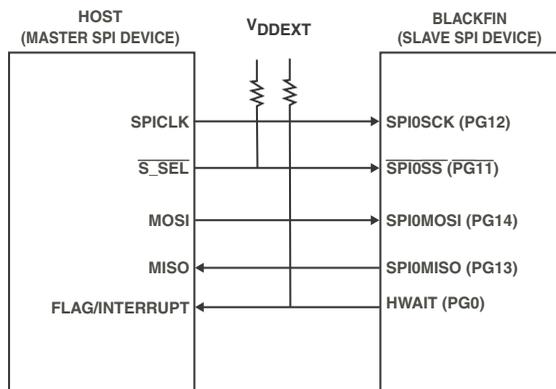


Figure 26-17. Connections Between Host (SPI Master) and Blackfin Processor (SPI Slave)

The host drives the SPI clock and is responsible for the timing. The host must provide an active-low chip select signal that connects to the $\overline{SPI0SS}$ input of the Blackfin processor. It can toggle with each byte transferred or

remain low during the entire procedure. 8-bit data is expected. The 16-bit mode is not supported.

In SPI slave boot mode, the boot kernel sets the `CPHA` bit and clears the `CPOL` bit in the `SPIO_CTL` register. Therefore the `MISO` pin is latched on the falling edge of the `MOSI` pin. For details see [Chapter 17, “SPI-Compatible Port Controller”](#).

In SPI slave boot mode, `HWAIT` functionality is critical. When high, the resistor shown in [Figure 26-17](#) programs `HWAIT` to hold off the host. `HWAIT` holds the host off while the Blackfin processor is in reset or executing the preboot. Once `HWAIT` turns inactive, the host can send boot data. The SPI module does not provide very large receive FIFOs, so the host must test the `HWAIT` signal for every byte. [Figure 26-19](#) illustrates the required program flow on the host side.

[Figure 26-18](#) shows the initial waveform for an SPI slave boot case. As soon as the Blackfin processor releases `HWAIT` after reset, the host device pulls the `SPI0SS` pin low and starts transmitting data. After the eighth data word has been received, the boot kernel asserts `HWAIT` again as it has to process the `DMACODE` field of the first block header. When the host detects the asserted `HWAIT` it gracefully finishes the transmission of the on-going word. Then, it pauses transmission until `HWAIT` releases again.

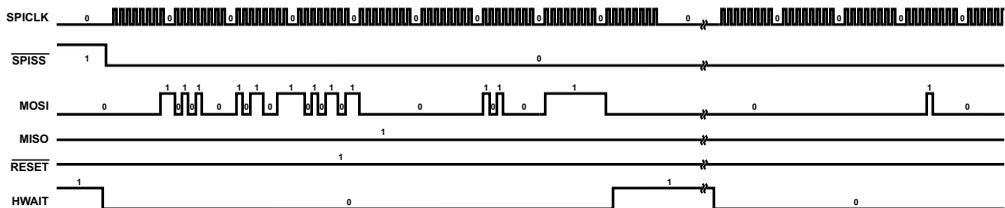


Figure 26-18. Typical SPI Slave Boot Waveforms

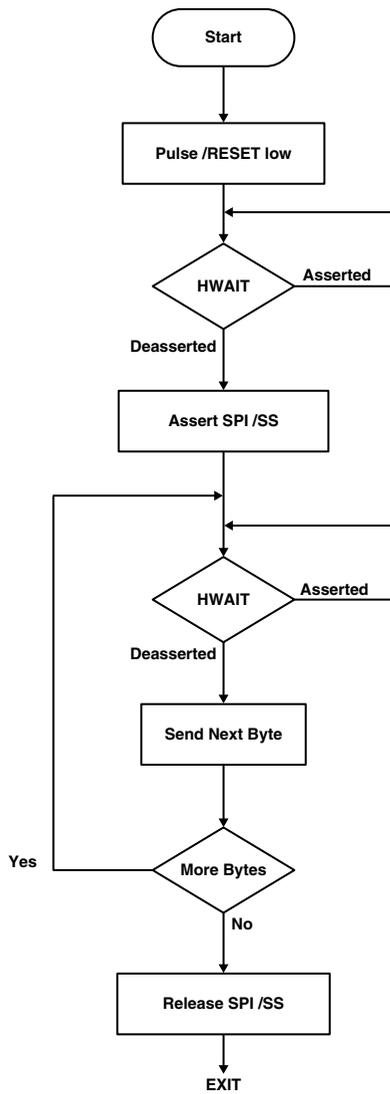


Figure 26-19. SPI Program Flow on Host Device

UART Slave Mode Boot

Figure 26-20 shows the interconnection required for booting. The figure does not show physical line drivers and level shifters that are typically required to meet the individual UART-compatible standards.

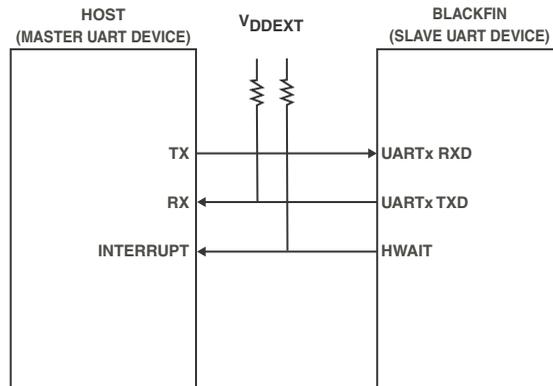


Figure 26-20. UART Slave Boot Mode Connections

For $B_{MODE} = 111$, the ADSP-BF51x processor consumes boot data from a UART host device connected to the UART0 interface on port G.

The host downloads programs formatted as boot streams using an auto-baud detection sequence. The host selects a bit rate within the UART clocking capabilities. To determine the bit rate when performing the auto-baud, the boot kernel expects an “@” character (0x40, eight data bits, one start bit, one stop bit, no parity bit) on the UART RXD input. The boot kernel acknowledges, and the host then downloads the boot stream. The acknowledgement consists of four bytes: 0xBF, UARTx_DLL, UARTx_DLH, 0x00. The host is requested to not send further bytes until it has received the complete acknowledge string. Once the 0x00 byte is received, the host can send the entire boot stream. The host should know the total byte count of the boot stream, but it is not required to have any knowledge about the content of the boot stream. Further information regarding auto-baud detection is given in “Autobaud Mode” on page 10-31.

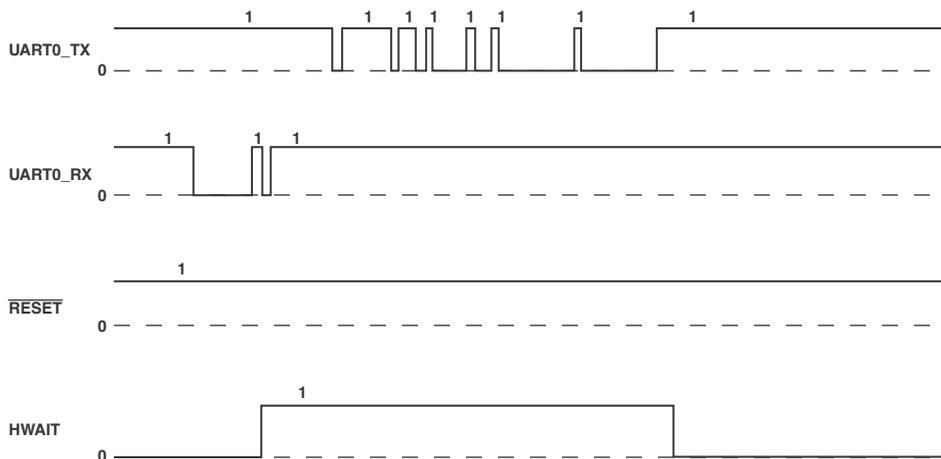


Figure 26-21. UART Autobaud Waveform

When the boot kernel is processing fill or initcode blocks it might require extra processing time and needs to hold the host off from sending more data. This is signalled with the `HWAIT` output. When equipped with a pull-up resistor the `HWAIT` signal imitates the behavior of an `RTS` output and could be connected to the `CTS` input of the booting host. The host is not allowed to send data until `HWAIT` turns inactive after a reset cycle. Therefore a pulling resistor on the `HWAIT` signal is required.

If the resistor pulls to ground, the host must pause transmission when `HWAIT` is low and is permitted to send when `HWAIT` is high. A pull-up resistor inverts the signal polarity of `HWAIT`. The host should test `HWAIT` at every transmitted byte.

[Figure 26-22](#) shows the initial case of the UART boot mode. As soon as `HWAIT` releases after reset, the boot kernel expects to receive a 0x40 byte for bit rate detection. After the bit rate is known, the UART is enabled and the kernel transmits for bytes.



Figure 26-22. UART Boot

i For UART boot, it is not obvious on how to change the PLL by an initcode routine. This is because the `UARTx_DLL` and `UARTx_DLH` registers have to be updated to keep the required bit rate constant after the `SCLK` frequency has changed. It must be ensured that the host does not send data while the PLL is changing. The initcode examples provided along with the VisualDSP++ tools installation demonstrate how this can be accomplished.

OTP Boot Mode

In the OTP boot mode (`BMODE = 101`), the boot kernel loads the boot stream from the on-chip OTP memory. OTP booting is a self-sufficient booting mechanism that does not require external boot memory or a host device.

By default the boot kernel starts loading the boot stream starting from OTP page `0x40`. This is in the public OTP region. The boot stream can occupy all pages up to OTP page `0xDF`, resulting in a boot stream length of up to 2560 bytes. The start address of the boot stream can be altered by programming the `OTP_START_PAGE` field in the `PBS01H` page. If there is no

conflict with the alternate preboot pages feature, the `OTP_START_PAGE` field can be set to `0x20`, resulting in a boot stream length of up to 3072 bytes.

In the current implementation, the OTP engine has no DMA support. Data is loaded and copied by core instructions. Nevertheless the `DMACODE` field should be set to `0xA`, indicating 32-bit operation. The boot kernel ensures proper operation at 32-bit granularity, but 64-bit alignment may help to reduce the number of OTP pages that have to be read during boot processing. Byte 0 of the boot stream is expected to be byte 0 of the lower 32-bit word of the lower `0x40` half page.

 In the OTP boot mode, the upper 512 bytes starting at address `0xFF903E00` either must not be used or must be booted last. The boot ROM code uses this space to temporarily hold the serial data which is then transferred to L1 instruction memory using DMA. All boot blocks that target the L1 instruction memory or external memories must have the `BFLAG_INDIRECT` bit set. Initcodes can alter the placement of the temporary buffer by modifying the `pTempBuffer` and `dTempByteCount` variables in the `ADI_BOOT_DATA` structure.

Reset and Booting Registers

Two registers are used for reset and booting—the software reset register (`SWRST`) and the system reset configuration register (`SYSCR`).

Software Reset (`SWRST`) Register

A software reset can be initiated by setting bits [2:0] in the system software reset field in the software reset register (`SWRST`) shown in [Figure 26-23](#). Bit 3 can be used to generate reset upon core-double-fault. A core-double-fault resets both the core and the peripherals, but not the RTC block and most of the DPMC. Bit 15 indicates whether a software reset has occurred since the last time `SWRST` was read. Bit 14 indicates the

software watchdog timer has generated the software reset. Bit 13 indicates the core-double-fault has generated the software reset. Bits [15:13] are read-only and cleared when the register is read. Reading the SWRST also clears bits [15:13] in the SYSCR register. Bits [3:0] are read/write.

Only writing to bits[2:0], resets only the modules in the SCLK domain. It does not clear the core. The program executes normally at the instruction after the MMR write to SWRST. The system is kept in the reset state as long as the bits[2:0] are set to b#111. To release reset, write a zero again. An example is shown in Listing 26-3 on page 26-102. It is not recommended that this functionality be used directly. Rather, call the ROM function `bfrom_SysControl()` to perform a system reset.

Software Reset Register (SWRST)

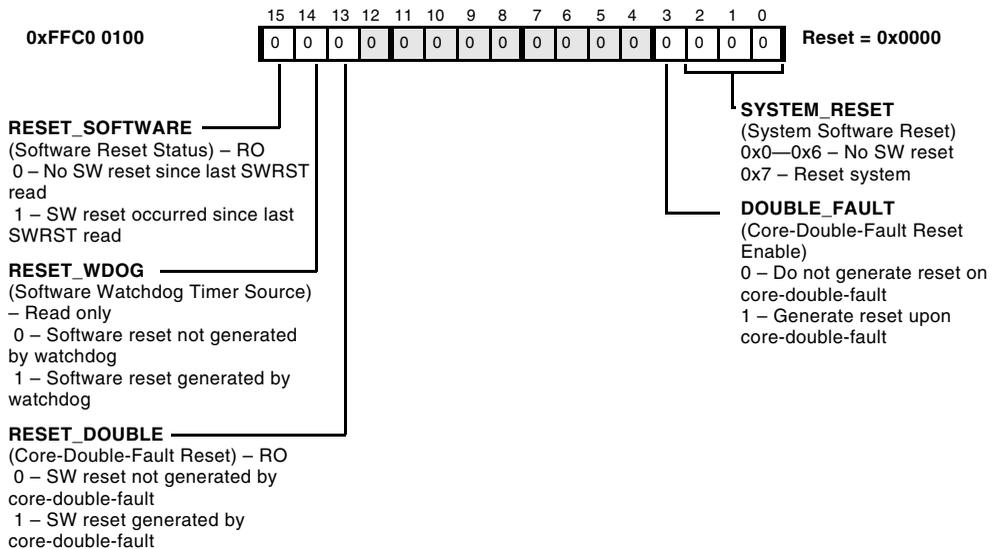


Figure 26-23. Software Reset Register

System Reset Configuration (SYSCR) Register

The values sensed from the `BMODE[2:0]` pins are mirrored into the system reset configuration register (`SYSCR`). The values are available for software access and modification after the hardware reset sequence. Software can modify only bits[7:4] in this register to customize boot processing upon a software reset.

The bits [15:13] are exact copies of the same bits in the `SWRST` register. Unlike the `SWRST` register, `SYSCR` can be read without clearing these bits. Reading `SWRST` also causes `SYSCR[15:13]` to clear.

The `WURESET` indicates whether there was a wake up from hibernate since the last hardware reset. The bit cannot be cleared by software.

The software reset configuration register (`SYSCR`) is shown in [Figure 26-24](#).

System Reset Configuration Register (SYSCR)

X – state is initialized from BMODE pins during hardware reset

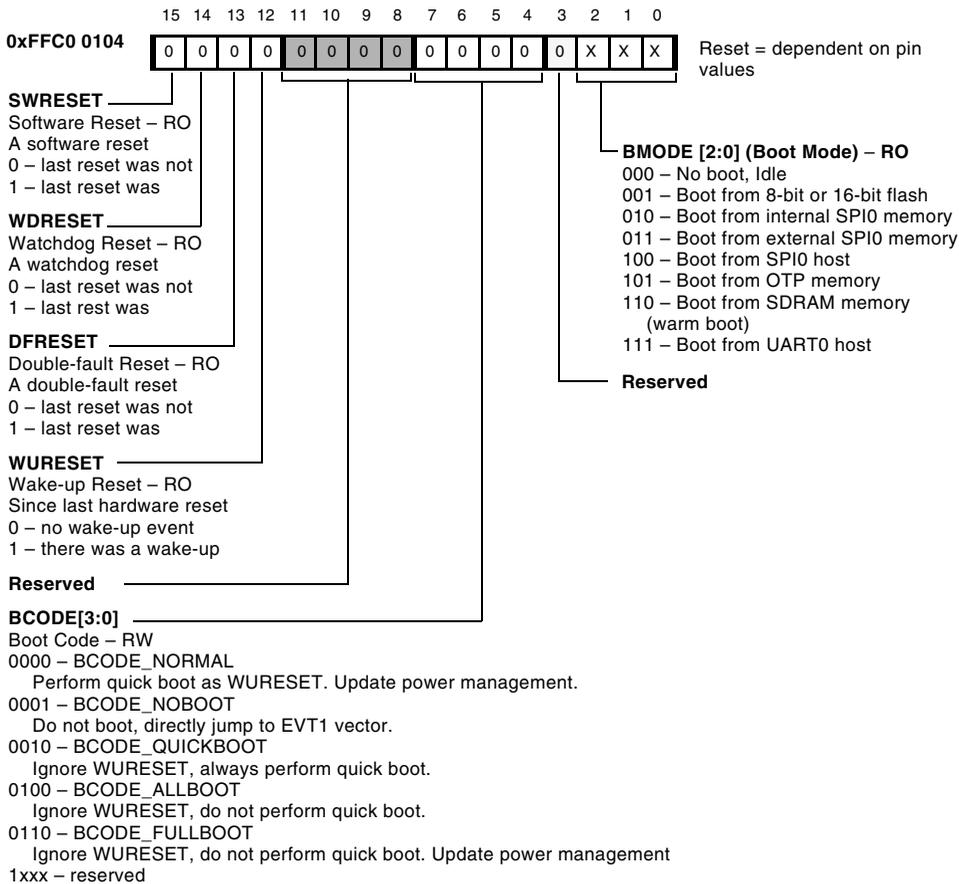
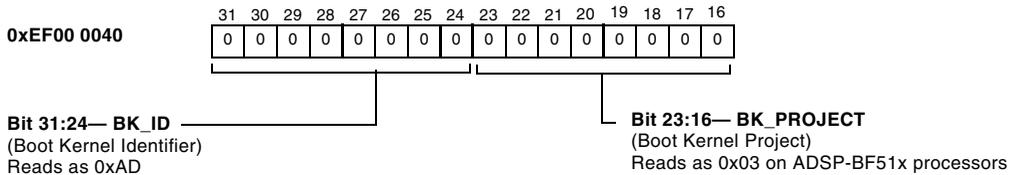


Figure 26-24. System Reset Configuration Register

Boot Code Revision Control (BK_REVISION)

The boot ROM reserves the 32-bits at address 0xEF00 0040 for a four byte version code as shown in [Figure 26-25](#).

Boot Code Revision BK_REVISION Word, 31–16



Boot Code Revision BK_REVISION Word, 15–0

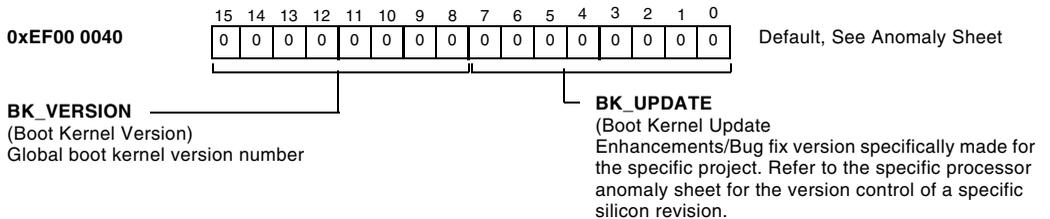
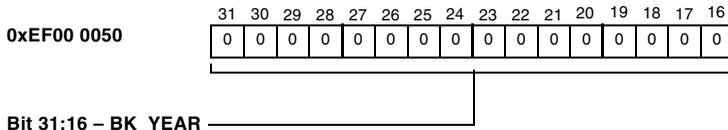


Figure 26-25. Boot Code Revision Code (BK_REVISION)

Boot Code Date Code (BK_DATECODE)

The boot ROM reserves the 32-bits at address 0xEF00 0050 for the build date as shown in [Figure 26-26](#).

Boot Code Date Code BK_DATECODE Word, 31–16



Boot Code Date Code BK_DATECODE Word, 15–0

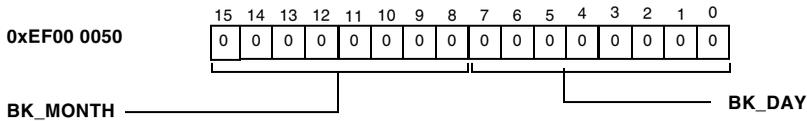
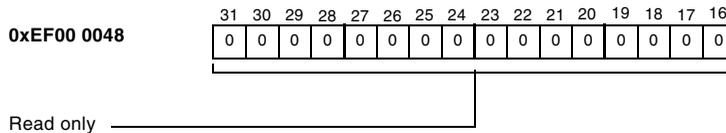


Figure 26-26. Boot Code Date Code (BK_DATECODE)

Zero Word (BK_ZEROS)

The boot ROM reserves the 32-bits at address 0xEF00 0048 which always reads as 0x0000 000 as shown in [Figure 26-27](#).

Zero Word BK_ZEROS, 31–16



Zero Word BK_ZEROS, 15–0

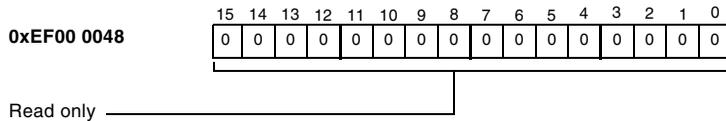
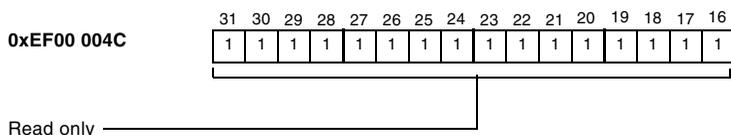


Figure 26-27. Zero Word (BK_ZEROS)

Ones Word (BK_ONES)

The boot ROM reserves the 32-bits at address 0xEF00 004C which always reads 0xFFFF FFFF as shown in [Figure 26-28](#).

Ones Word BK_ONES, 31–16



Ones Word BK_ONES, 15–0

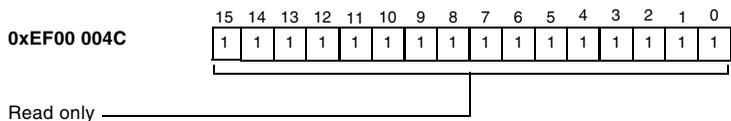


Figure 26-28. Ones Word (BK_ONES)

OTP Memory Pages for Booting

This section details the one-time programmable memory pages for booting.

Lower PBS00 Half Page

The 64-bit lower half of page 0x18 is always read by the preboot routine. These control bits customize the boot process and instruct the preboot routine whether to process further pages and whether the PLL settings have to be changed. Other bits customize the SPI .

Lower PBS00 Half Page (PBS00L, Bits 63–48)

One-Time Programmable

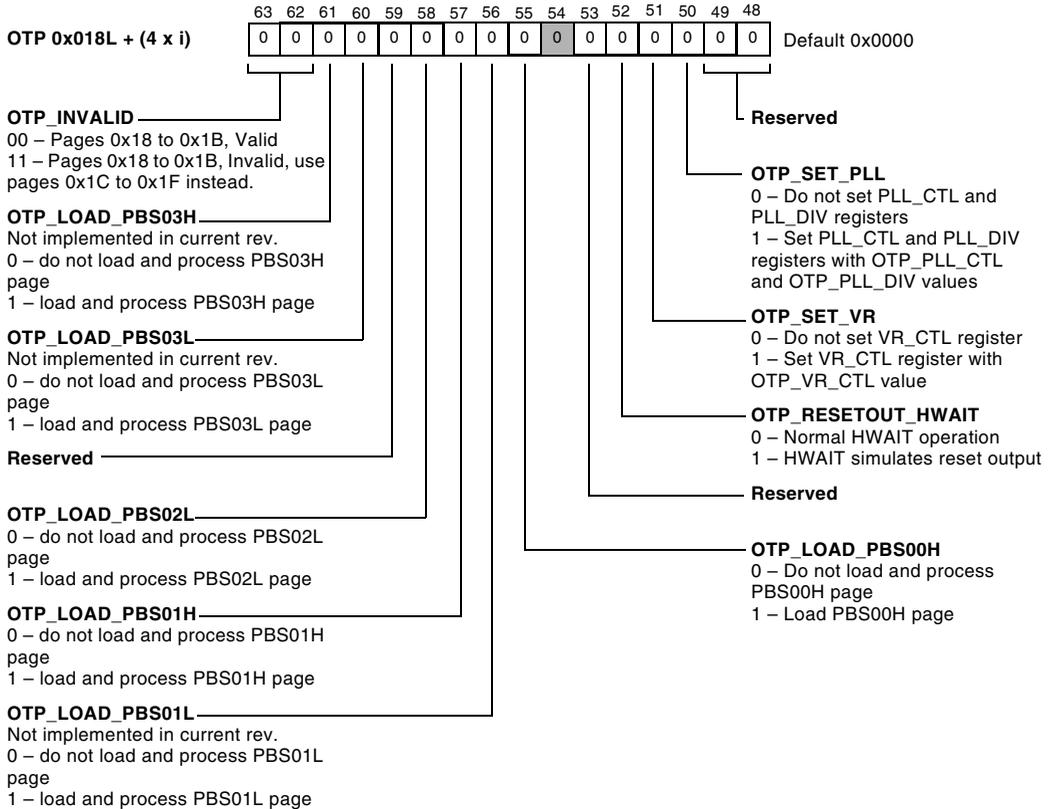


Figure 26-29. Lower PBS00 Half Page (PBS00L, Bits 63–48)

Lower PBS00 Half Page (PBS00L, Bits 47–32)

One-Time Programmable

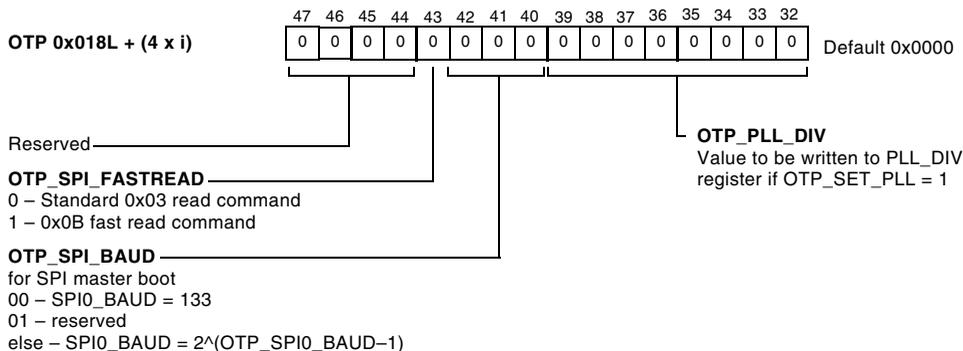
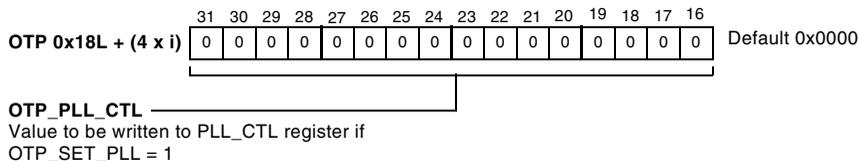


Figure 26-30. Lower PBS00 Half Page (PBS00L, Bits 47–32)

Lower PBS00 Half Page (PBS00L, Bits 31–16)

One-Time Programmable



Lower PBS00 Half Page (PBS00L, Bits 15–0)

One-Time Programmable

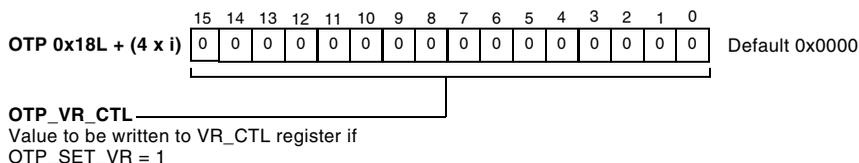


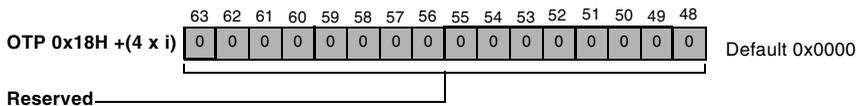
Figure 26-31. Lower PBS00 Half Page (PBS00L, Bits 31–0)

Upper PBS00 Half Page

The preboot routine loads the upper 64-bit half of page PBS00 only if the `OTP_LOAD_PBS00H` bit in the `PBS00L` page is set. Page `PBS00H` customizes the default setting of the asynchronous portion of the EBIU controller.

Upper PBS00 Half Page (PBS00H, Bits 63–48)

One-time Programmable



Upper PBS00 Half Page (PBS00H, Bits 47–32)

One-time Programmable

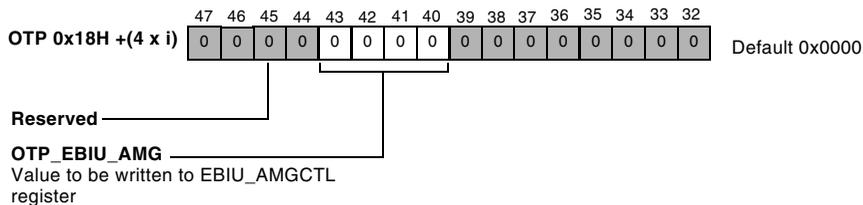
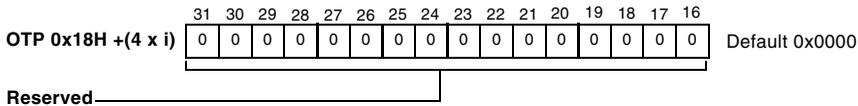


Figure 26-32. Upper PBS00 Half Page (PBS00H, Bits 63–32)

Upper PBS00 Half Page (PBS00H, Bits 31–16)

One-Time Programmable



Upper PBS00 Half Page (PBS00H, Bits 15–0)

One-Time Programmable

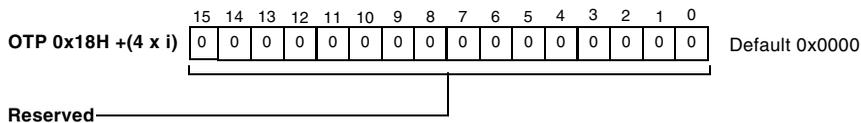


Figure 26-33. Upper PBS00 Half Page (PBS00H, Bits 31–0)

Lower PBS01 Half Page

The half page PBS01L is reserved and not used in the current silicon.



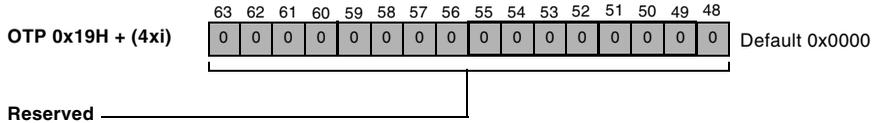
Do not use this page as it may be populated in future silicon revisions.

Upper PBS01 Half Page

The preboot routine loads the upper 64-bit half of page 0x19 only if the OTP_LOAD_PBS01H bit in the PBS00L page is set. This page allows the user to disable boot modes. If a disabled boot mode configuration is chosen by the BMODE[2:0] pins, the boot kernel goes into idle state. In OTP boot mode, this page determines where in OTP memory the boot stream resides.

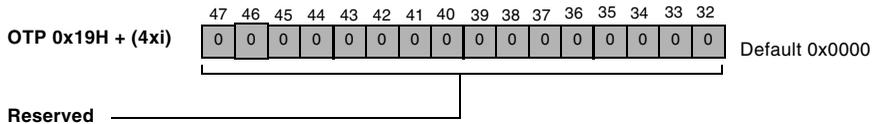
Upper PBS01 Half Page (PBS01H, Bits 63–48)

One-Time Programmable



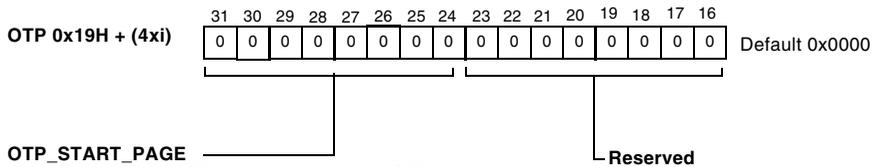
Upper PBS01 Half Page (PBS01H, Bits 47–32)

One-Time Programmable



Upper PBS01 Half Page (PBS01H, Bits 31–16)

One-Time Programmable



OTP start page for OTP boot mode. If 0x00, OTP boot starts at OTP page 0x40.

Figure 26-34. OTP Half Page (PBS01H, Bits 63–16)

Upper PBS01 Half Page (PBS01H, Bits 15–0)

One-Time Programmable

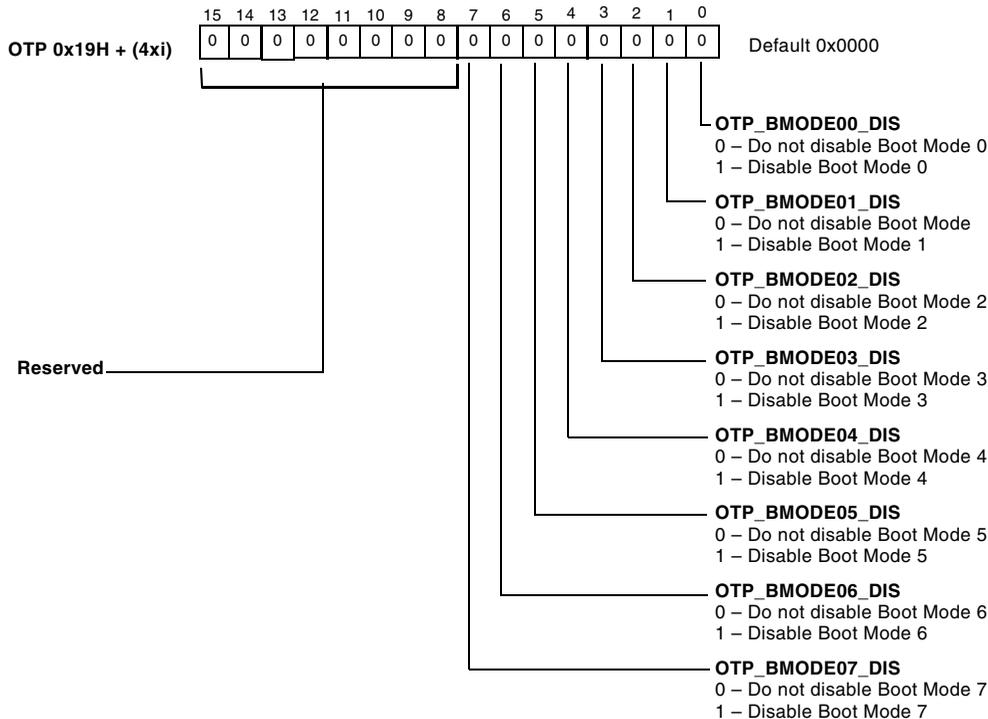


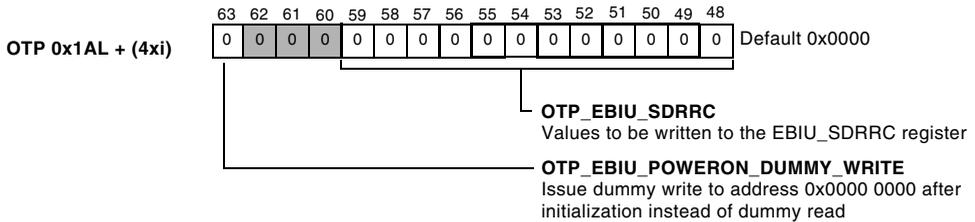
Figure 26-35. OTP Half Page PBS01H (PBS01H, Bits 15–0)

Lower PBS02 Half Page

The preboot routine loads the lower 64-bit half of page 0x1A only if the OTP_LOAD_PBS02L bit in half page PBS00L is set. Half pages PBS02L and PBS02H customize the SDRAM controller settings.

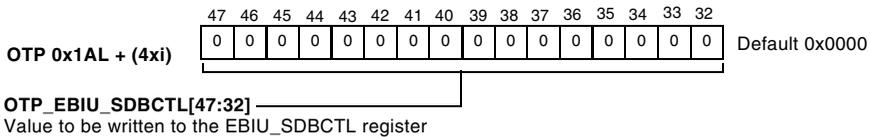
Lower PBS02 Half Page (PBS02L, Bits 63–48)

One-time Programmable



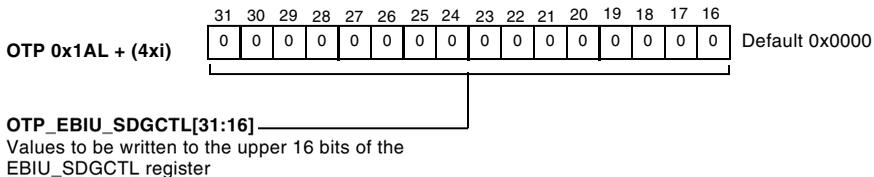
Lower PBS02 Half Page (PBS02L, Bits 47–32)

One-time Programmable



Lower PBS02 Half Page (PBS02L, Bits 31–16)

One-time Programmable



Lower PBS02 Half Page (PBS02L, Bits 15–0)

One-time Programmable

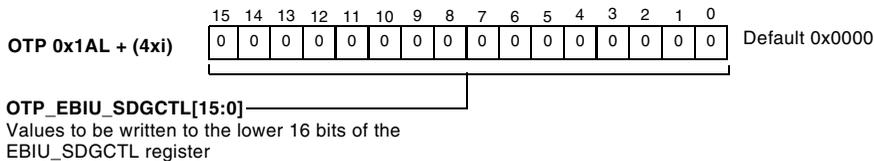


Figure 26-36. Lower PBS02 Half Page (PBS02L Bits 63–0)

Upper PBS02 Half Page

Page PBS02H is reserved. Do not use for any purpose.

Reserved Half Pages

The half pages PBS01L, PBS02H, PBS03L and PBS03H are reserved and not used in the current silicon.



Do not use these pages as they may be populated in future silicon revisions.

Data Structures

The boot kernel uses specific data structures for internal processing. Advanced users can customize the booting process by changing the content of the structure within the initcode routines. This section uses C language definitions for documentation purposes. VisualDSP++ users can use these structures directly in assembly programs by using the `.IMPORT` directive. The structures are supplied by the `bfrom.h` header file in your VisualDSP++ installation directory.

ADI_BOOT_HEADER

```
typedef struct {
    s32  dBlockCode;
    void* pTargetAddress;
    s32  dByteCount;
    s32  dArgument;
} ADI_BOOT_HEADER;
```

The structure `ADI_BOOT_HEADER` is used by the boot kernel to load and process a block header.

Every block header is loaded to L1 data memory location 0xFF80 7FF0–0xFF80 7FFF first or where `pHeader` points to. There it is analyzed by the boot kernel.

ADI_BOOT_BUFFER

```
typedef struct {
    void* pSource;
    s32   dByteCount;
} ADI_BOOT_BUFFER;
```

The structure `ADI_BOOT_BUFFER` is used for any kind of buffer. For the user, this structure is important when implementing advanced callback mechanisms.

ADI_BOOT_DATA

```
typedef struct {
    void* pSource;
    void* pDestination;
    s16* pControlRegister;
    s16* pDmaControlRegister;
    s32   dControlValue;
    s32   dByteCount;
    s32   dFlags;
    s16   uwDataWidth;
    s16   uwSrcModifyMult;
    s16   uwDstModifyMult;
    s16   uwHwait;
    s16   uwSsel;
    s16   uwUserShort;
    s32   dUserLong;
    s32   dReserved;
    ADI_BOOT_ERROR_FUNC* pErrorFunction;
    ADI_BOOT_LOAD_FUNC* pLoadFunction;
    ADI_BOOT_CALLBACK_FUNC* pCallbackFunction;
```

```

ADI_BOOT_HEADER* pHeader;
void* pTempBuffer;
void* pTempCurrent;
s32 dTempByteCount;
s32 dBlockCount;
s32 dClock;
void* pLogBuffer;
void* pLogCurrent;
s32 dLogByteCount;
} ADI_BOOT_DATA;

```

The structure `ADI_BOOT_DATA` is the main data structure. A pointer to a `ADI_BOOT_DATA` structure is passed to most complex subroutines, including load functions, `initcode`, and callback routines. The structure has two parts. While the first is closely related to internal memory load routines, the second provides access to global boot settings.

[Table 26-11](#) describes the data structures.

Table 26-11. Structure Variables, `ADI_BOOT_DATA`

Variable	Description
<code>pSource</code>	In the context of the boot kernel, the <code>pSource</code> pointer points either to the start address of the entire boot stream or to the header of the next boot block. In the context of memory load routines <code>pSource</code> points to the source address of the DMA work unit.
<code>pDestination</code>	The <code>pDestination</code> pointer is only used in memory load routines. It points to the destination address of the DMA work unit. It points to either <code>0xFF80_7FF0</code> when a header is loaded, or the target address when the payload data is loaded.
<code>pControlRegister</code>	This pointer holds the MMR address of the peripheral's main control register (for example <code>UARTx_LCR</code> or <code>SPIx_CTL</code>)
<code>pDmaControlRegister</code>	This pointer holds the MMR address of the <code>DMAX_CONFIG</code> register for the DMA channel in use.
<code>dControlValue</code>	The lower 16 bits of this value are written to the <code>pControlRegister</code> location each time a DMA work unit is started.

Table 26-11. Structure Variables, ADI_BOOT_DATA (Continued)

Variable	Description
dByteCount	Number of bytes to be transferred.
dFlags	The lower 16 bits of this variable hold the lower 16 bits of the current block code. The upper 16 bits hold global flags. See “dFlags Word” on page 26-91.
uwDataWidth	This instructs the memory load routine to use: 0 – 8-bit DMA 1 – 16-bit DMA 2 – 32-bit DMA
uwSrcModifyMult	This is the multiplication factor used by the DMA source. A value of 1 sets the source modifier to 1 for 8-bit DMA, 2 for 16-bit DMA, or 4 for 32-bit DMA.
uwDstModifyMult	This is the multiplication factor used by the DMA destination. A value of 1 sets the destination modifier to 1 for 8-bit DMA, 2 for 16-bit DMA, or 4 for 32-bit DMA.
uwHwait	This 16-bit value holds the GPIO used for HWAIT signaling. The value can change on the fly. The upper eight bits designate the port number (for example 01 for Port A, 02 for Port B). The lower four bits designate the GPIO in the port.
uwSsel	This 16-bit value holds the GPIO used for SPI slave select. The value can change on the fly. The upper eight bits designate the port number (for example 01 for Port A, 02 for Port B). The lower four bits designate the GPIO in the port.
uwUserShort	The programmer can use this 16-bit value for passing parameters between modules of a customized booting scheme.
dUserLong	The programmer can use this 32-bit value for passing parameters between modules of a customized booting scheme.
dReserved	This 32-bit value is reserved for future development.
pErrorFunction	This is the pointer to the error handler. See “Error Handler” on page 26-44.
pLoadFunction	This is the pointer to the function responsible for loading data. See “Load Functions” on page 26-44.

Table 26-11. Structure Variables, ADI_BOOT_DATA (Continued)

Variable	Description
pCallbackFunction;	This is the pointer to the callback function. See “Callback Routines” on page 26-41 .
pHeader	The pHeader pointer holds the address for intermediate storage of the block header. By default this value is set to 0xFF80 7FF0.
pTempBuffer	This pointer tells the boot kernel what memory to use for intermediate storage when the BFLAG_INDIRECT flag is set for a given block. The pointer defaults to 0xFF90 7E00. The value can be modified by the initcode routine, but there would be some impact to the VisualDSP++ tools.
pTempCurrent	Defaults to the pTempBuffer value. A load function can modify this value to manipulate subsequent callback and memory DMA routines.
dTempByteCount	This is the size of the intermediate storage buffer used when the BFLAG_INDIRECT flag is set for a given block. This value defaults to 256 and can be modified by an initcode routine. When increasing this value, the pTempBuffer must also be changed since by default the block is at the end of a physical data memory block.
dBlockCount	This 32-bit variable counts the boot blocks that are processed by the boot kernel. If the user sets this value to a negative value, the boot kernel exits when the variable increments to zero.
dClock	The dClock variable holds information about the clock divider used by individual (serial) boot modes.
pLogBuffer	Pointer to the circular log buffer. By default the log buffer resides in L1 scratch pad memory at address 0xFFB0 0400.
pLogCurrent	Pointer to the next free entry of the circular log buffer.
dLogByteCount	Size of the circular log buffer, default is 0x400 bytes.

dFlags Word

Figure 26-37 and Figure 26-38 describe the dFlags word. dFlags [15-0] is a copy of Block Code[15-0] of the block currently being processed.

dFlags Word, Bits 31–16

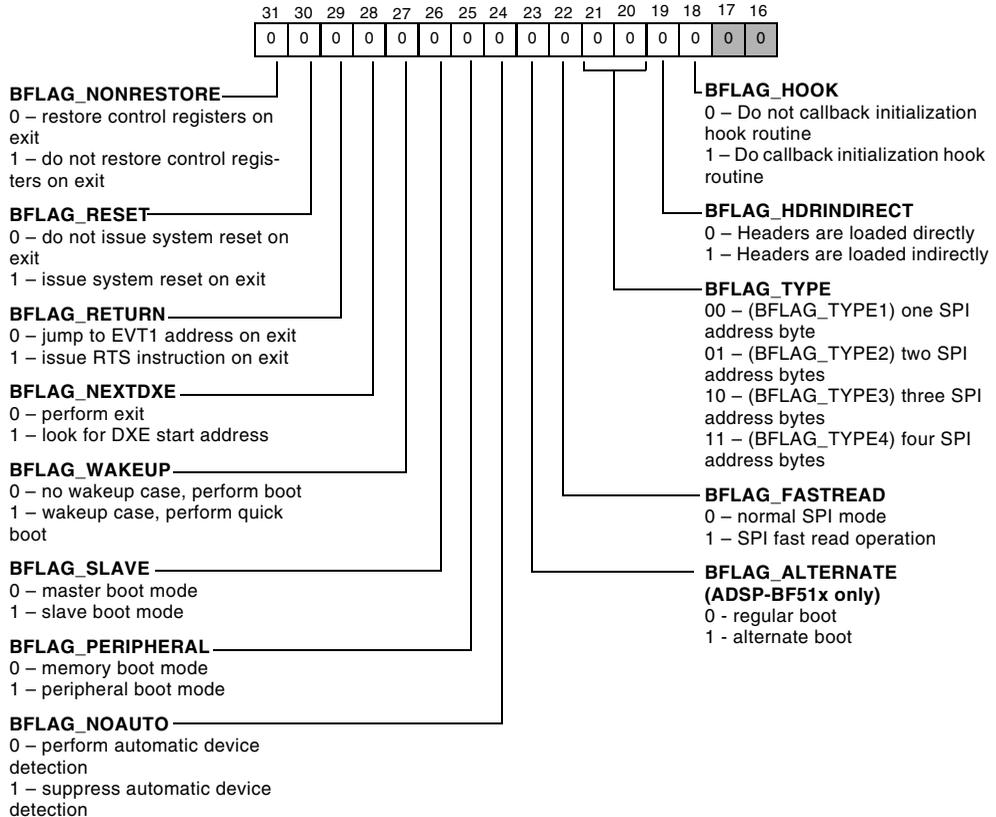


Figure 26-37. dFlags Word (Bits 31–16)

dFlags Word, Bits 15–0

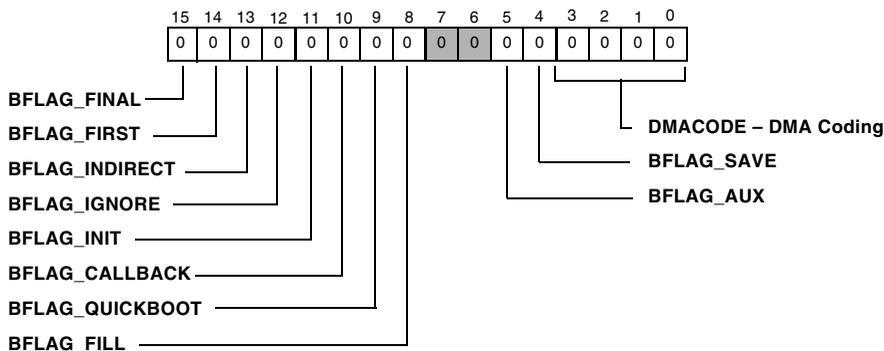


Figure 26-38. dFlags Word (Bits 15–0)

Callable ROM Functions for Booting

The following functions support boot management.

BFROM_FINALINIT

Entry address: 0xEF00 0002

Arguments: no arguments

C prototype: `void bfrom_FinalInit (void);`

The `bfrom_FinalInit` function never returns. It only executes a JUMP to the address stored in EVT1.

BFROM_PDMA

Entry address: 0xEF00 0004

Arguments: pointer to ADI_BOOT_DATA in R0

C prototype: `void bfrom_PDma (ADI_BOOT_DATA *p);`

This is the load function for peripherals such as SPI and UART that support DMA in their boot modes.

BFROM_MDMA

Entry address: 0xEF00 0006

Arguments: pointer to ADI_BOOT_DATA in R0

C prototype: `void bfrom_MDma (ADI_BOOT_DATA *p);`

This is the load function used for memory boot modes including the FIFO mode. This routine is also reused when the `BFLAG_FILL` or the `BFLAG_INDIRECT` flags are specified.

BFROM_MEMBOOT

Entry address: 0xEF00 0008

Arguments:

pointer to boot stream in R0

dFlags in R1

dBlockCount in R2

pCallHook passed over the stack in [FP+0x14]

updated block count returned in R0

C prototype:

```
s32 bfrom_MemBoot (void* pBootStream, s32 dFlags,  
s32 dBlockCount, ADI_BOOT_HOOK_FUNC* pCallHook);
```

This routine processes any boot stream that maps to the Blackfin memory starting from address `pBootStream`.

To boot a new application that may overwrite the calling application, the `dFlags` word is usually zero. When done, the routine does not return, but jumps to the `EVT1` vector address. If the `BFLAG_RETURN` flag is set, an `RTS` is executed instead and the routine returns to the parent function. In this way, fractions of an application can be loaded.

If the `dBlockCount` parameter is zero or a positive value, all boot blocks are processed until the `BFLAG_FINAL` flag is detected. If `dBlockCount` is a negative value, the negative number represents the number of blocks to be booted. For example, `-1` causes the kernel to return immediately, `-2` processes only one block.

The routine returns the updated source address `pSource` of the boot stream (for example, the first unused address after the processed boot stream).

The `BFLAG_NEXTDXE` flag suppresses boot loading. The boot kernel steps through the boot stream by analyzing the next-DXE pointers (in the `ARGUMENT` field of a `BFLAG_FIRST` block) and jumping to the next DXE. Assuming that the boot image is a chained list of boot streams, the boot kernel returns the absolute start address of the requested boot stream. In this example, the start address of the third boot stream (DXE) in a flash device is returned.

```
bfrom_MemBoot((void*)0x20000000,  
BFLAG_RETURN|BFLAG_NEXTDXE,-3, NULL);
```

In the above example, the routine would return `0x2000 0000` when `dBlockCount` was set to `-1`. If the parameter `dBlockCount` is zero or

positive when used along with the `BFLAG_NEXTDXE` command, the kernel returns when the `BFLAG_FIRST` flag on a header in the next-DXE chain is not set.

If the `BFLAG_HOOK` switch is set, the memboot routine call (`pCallHook` routine) after the `ADI_BOOT_DATA` structure is filled with default values. It then can overrule the default settings of the structure.

The `bfrom_MemBoot()` uses both MDMA channel pairs. Respective wake-up bits must be set in the `SIC_IWRx` registers.

BFROM_SPIBOOT

Entry address: `0xEF00 000A`

Arguments:

SPI address in `R0`

`dFlags` in `R1`

`dBlockCount` in `R2`

`pCallHook` passed over the stack in `[FP+0x14]`

updated block count returned in `R0`

C prototype:

```
s32 bfrom_SpiBoot (s32 dSpiAddress, s32 dFlags,  
s32 dBlockCount, ADI_BOOT_HOOK_FUNC* pCallHook);
```

This SPI master boot routine processes boot streams residing in SPI memories, using the SPI0 controller. The fourth argument `pCallHook` is passed over the stack. It provides a hook to call a callback routine after the `ADI_BOOT_DATA` structure is filled with default values. For example, the `pCallHook` routine may overwrite the default value of the `uwSsel` value in the `ADI_BOOT_DATA` structure. The coding follows the rules of

uwHWAIT (see [“Boot Host Wait \(HWAIT\) Feedback Strobe”](#) on [page 26-29](#)). A value of 0x070F represents GPIO_PG15 ($\overline{\text{SPIOSEL2}}$), 0x0804 represents PH4($\overline{\text{SPIOSEL3}}$) and so on. In this way, SPI memories which connect to a different pin than the $\overline{\text{SPIOSEL2}}$ output can be supported.

Additional bits in the dFlags word are relevant. The user should always set the BFLAG_PERIPHERAL flag but never the BFLAG_SLAVE bit. The BFLAG_NOAUTO flag instructs the system to skip the SPI device detection routine. The BFLAG_TYPE then tells the boot kernel what addressing mode is required for the SPI memory. (see [“SPI Device Detection Routine”](#) on [page 26-63](#)). The BFLAG_FASTREAD flag controls whether standard SPI read (0x3 command) or fast read (0xB) is performed. The three lower bits of the dFlags word are translated by the boot kernel into specific values to the SPIO_BAUD registers. This follows the truth table shown in [Table 26-10](#) on [page 26-62](#).

By default, the bfrom_SpiBoot() function attempts to boot from the on-chip SPI memory device connected to internal PH8 output. In this case, automatic device detection has to be avoided using the BFLAG_NOAUTO flag; and the BFLAG_TYPE field has to be set to b#01 to indicate 3-byte addressing mode. The use of the BFLAG_FASTREAD option is valid and optional.

When called with the BFLAG_ALTERNATE flag, the bfrom_SpiBoot() function attempts to boot from external SPI memory device instead. Unless the uwSsel variable in the ADI_BOOT_DATA structure is altered by a hook routine, the memory is expected to be connected to $\overline{\text{SPIOSEL2}}$. A pull-up resistor on this signal is required when automatic device detection is desired.

The bfrom_SpiBoot() routine does not deal with port muxing at all. When a part has been booted via SPI master mode after reset, the port muxing configuration is typically already ready for a runtime call to the bfrom_SpiBoot() routine. Otherwise ensure that the SPIO_MISO, SPIO_MOSI and SPIO_SCK signals are properly activated in the PORTx_FER and PORTx_MUX registers. The $\overline{\text{SPIOSEL2}}$ signal requires, however, that the respective

PORTX_FER bit be cleared, as the boot kernel toggles the signal in GPIO mode.

Similarly, the user shall clear the PH8 bit in the PORTH_FER register when booting from on-chip SPI memory and set PH8 when booting from an external device.

The `bfrom_SpiBoot()` routine uses the MDMA0 memory DMA channel pair and the DMA7 peripheral DMA. Respective wake-up bits must be set in the SIC_IWRx registers. If a different peripheral DMA channel has been assigned to the SPI0 controller, use the hook routine to store the MMR address of the respective DMAx_CONFIG register into the `pDmaControlRegister` variable in the ADI_BOOT_DATA structure. Similarly, when using a different SPI controller than SPI0, write the MMR address of the relevant SPIx_CTL register into the `pControlRegister` variable.

BFROM_OTPBOOT

Entry address: 0xEF00 000E

Arguments:

OTP byte address in R0

dFlags in R1

dBlockCount in R2

pCallHook passed over the stack in [FP+0x14]

Updated block count returned in R0

C prototype:

```
s32 bfrom_otpBoot (s32 dotpAddress, s32 dFlags,  
s32 dBlockCount, ADI_BOOT_HOOK_FUNC* pCallHook);
```

This OTP boot routine processes boot streams residing in the on-chip, serial OTP memory. Unlike the `bfrom_OtpRead()` function which uses the half-page addressing method, this one requires byte addressing. For example, set the `d0tpAddress` argument to `0x400` to process a boot stream starting from OTP page `0x40`. Remember that one OTP page spans 16 bytes.

The `bfrom_OtpBoot()` routine uses the MDMA0 memory DMA channel pair. The respective wake-up bit must be set in the `SIC_IWRx` registers.

BFROM_BOOTKERNEL

Entry address: `0xEF00 0020`

Arguments:

pointer to `ADI_BOOT_DATA` in R0

returns updated source address `pSource` in R0

C prototype:

```
s32 bfrom_BootKernel (ADI_BOOT_DATA *p);
```

This ROM entry provides access to the raw boot kernel routine. It is the user's responsibility to initialize the items passed in the `ADI_BOOT_DATA` structure. Pay particular attention that the function pointers (`pLoadFunction`, and `pErrorFunction`) point to functional routines.

BFROM_CRC32

Entry address: 0xEF00 0030

Arguments:

pointer to look-up table in R0

pointer to data in R1

dByteCount in R2

initial CRC value in R0

CRC value returned in R0

C prototype:

```
s32 bfrom_Crc32 (s32 *pLut, void *pData,  
                s32 dByteCount, s32 dInitial);
```

This routine calculates the CRC32 checksum for a given array of bytes. The look-up table is typically generated by the `BFROM_CRC32POLY` routine. During the boot process this routine is called by the `BFROM_CRC32CALLBACK` routine. The `dInitial` value is normally set to zero unless the CRC32 routine is called in multiple slices. Then, the `dInitial` parameter expects the result of the former run.

BFROM_CRC32POLY

Entry address: 0xEF00 0032

Arguments:

pointer to look-up table in R0

polynomial in R1

updated block count returned in R0

C prototype:

```
s32 bfrom_Crc32Poly (unsigned s32 *pLut, s32 dPolynomial);
```

This function generates a 1024-byte look-up table from a given CRC polynomial. During the boot process this routine is hidden by the `BFROM_CRC32INITCODE` routine.

BFROM_CRC32CALLBACK

Entry address: 0xEF00 0034

Arguments:

pointer to `ADI_BOOT_DATA` in R0

pointer to `ADI_BOOT_BUFFER` in R1* Callback Flags in R2

C prototype:

```
s32 bfrom_Crc32Callback (ADI_BOOT_DATA *pBS, ADI_BOOT_BUFFER *pCS, s32 dCbFlags);
```

This is a wrapper function that ensures the `BFROM_CRC32` subroutine fits into the boot process.

BFROM_CRC32INITCODE

Entry address: 0xEF00 0036

Arguments:

pointer to `ADI_BOOT_DATA` in R0

C prototype:

```
void bfrom_Crc32Initcode (ADI_BOOT_DATA *p);
```

This is an initcode residing in ROM with two jobs:

Register `BFROM_CRC32CALLBACK` as a callback routine to the `pCallback` pointer in `ADI_BOOT_DATA`.

Call `BFROM_CRC32POLY` to generate the look-up table.

This function is unlikely to be called by user code directly. This function is called as an initcode during the boot process when the CRC calculation is desired. See [“CRC Checksum Calculation”](#) on page 26-44 for details.

Programming Examples

This section contains the following programming examples:

- [“System Reset”](#)
- [“Exiting Reset to User Mode”](#) on page 26-102
- [“Exiting Reset to Supervisor Mode”](#) on page 26-103
- [“Initcode \(SDRAM Controller Setup\)”](#) on page 26-104
- [“Initcode \(Power Management Control\)”](#) on page 26-106
- [“Quickboot With Restore From SDRAM”](#) on page 26-109
- [“XOR Checksum”](#) on page 26-110
- [“Direct Code Execution”](#) on page 26-111
- [“Managing PBS Pages in OTP Memory”](#) on page 26-113

System Reset

To perform a system reset, use the code shown in [Listing 26-1](#) or [Listing 26-2](#).

Listing 26-1. System Reset in Assembly

```
#include <blackfin.h>
P0.L = LO(BFROM_SYSCONTROL);
P0.H = HI(BFROM_SYSCONTROL);
R0.L = LO(SYSCTRL_SYSRESET);
R0.H = HI(SYSCTRL_SYSRESET);
R1 = 0;
R2 = 0;
CALL (P0);
```

Listing 26-2. System Reset in C Language

```
bfrom_SysControl(SYSCTRL_SYSRESET, 0, NULL);
```

Exiting Reset to User Mode

To exit reset while remaining in user mode, use the code shown in [Listing 26-3](#).

Listing 26-3. Exiting Reset to User Mode

```
_reset:  P1.L = LO(_usercode); /* Point to start of user code */
        P1.H = HI(_usercode);
        RETI = P1;           /* Load address of _start into RETI */
        RTI;                /* Exit reset priority */
_reset.end:
_usercode:                /* Place user code here */
...

```

The reset handler most likely performs additional tasks not shown in the examples above. Stack pointers and EVT_x registers are initialized here.

Exiting Reset to Supervisor Mode

To exit reset while remaining in supervisor mode, use the code shown in [Listing 26-4](#).

Listing 26-4. Exiting Reset by Staying in Supervisor Mode

```
_reset:
    P0.L = L0(EVT15); /* Point to IVG15 in Event Vector Table */
    P0.H = HI(EVT15);
    P1.L = L0(_isr_IVG15); /* Point to start of IVG15 code */
    P1.H = HI(_isr_IVG15);
    [P0] = P1;          /* Initialize interrupt vector EVT15 */
    P0.L = L0(IMASK); /* read-modify-write IMASK register */
    R0 = [P0];         /* to enable IVG15 interrupts */
    R1 = EVT_IVG15 (Z);
    R0 = R0 | R1;      /* set IVG15 bit */
    [P0] = R0;         /* write back to IMASK */
    RAISE 15;         /* generate IVG15 interrupt request */
    /* IVG 15 is not served until reset handler returns */
    P0.L = L0(_usercode);
    P0.H = HI(_usercode);
    RETI = P0;        /* RETI loaded with return address */
    RTI;             /* Return from Reset Event */
_reset.end:
_usercode:          /* Wait in user mode till IVG15 */
    JUMP _usercode; /* interrupt is serviced */
_isr_IVG15:        /* IVG15 vectors here due to EVT15 */
    ...
```

Initcode (SDRAM Controller Setup)

Listing 26-5 shows an example of initcode to setup the SDRAM controller. The SDRAM controller must be initialized before data can be booted into it. Therefore, the SDRAM controller is typically initialized by an initcode or by the preboot functionality. The following initcode example assumes that the preboot did not do the job.

Listing 26-5. Example Initcode (SDRAM Controller Setup) in C Language

```
#include <ccb1kfn.h>
void init_SDRAM(void)
{
    while((*pEBIU_SDSTAT & SDCI) == 0){}
    /* clear SDRAM EAB sticky error status (W1C) */
    *pEBIU_SDSTAT |= SDEASE;
    /* SDRAM Refresh Rate Control Register */
    *pEBIU_SDRRC = 0x026B;
    /* SDRAM Memory Bank Control Register */
    *pEBIU_SDBCTL = (EBE|EBSZ_64|EBCAW_10);
    /* SDRAM Memory Global Control Register */
    *pEBIU_SDGCTL =
    (EMREN|SCTLE|PSS|TWR_2|TRCD_2|TRP_2|TRAS_4|PASR_ALL|CL_3);
    /* Finalize SDC initialization */
    pTmp = (u16*) 0x0;
    *pTmp = 0xBEEF;
    while((*pEBIU_SDSTAT & SDRS) == 1){}
}
```

Listing 26-6. Example Initcode (SDRAM Controller Setup) in Assembly

```
#include <blackfin.h>
/* Load Immediate 32-bit value into data or address register */
#define IMM32(reg,val) reg##.H=hi(val); reg##.L=lo(val)
.SECTION L1_code;
init_SDRAM:
link 0;
[--SP] = ASTAT;
[--SP] = (R7:7, P5:4);
IMM32(P5, EBIU_SDRRC);
PollSdcIdle:
R7 = w[P5 + EBIU_SDSTAT - EBIU_SDRRC] (z);
CC = bittst(R7,bitpos(SDCI));
if!CC jump PollSdcIdle;
/* clear SDRAM EAB sticky error status (W1C) */
R7 = SDEASE(z);
w[P5 + EBIU_SDSTAT - EBIU_SDRRC] = R7;
/* SDRAM Refresh Rate Control Register */
R7.L = 0x026B;
w[P5 + EBIU_SDRRC - EBIU_SDRRC] = R7;
/* SDRAM Memory Bank Control Register */
R7.L = (EBE|EBSZ_64|EBCAW_10);
w[P5 + EBIU_SDBCTL - EBIU_SDRRC] = R7;
/* SDRAM Memory Global Control Register */
IMM32(R7,(SCTLE|PSS|TWR_2|TRCD_2|TRP_2|TRAS_4|PASR_ALL|CL_3));
[P5 + EBIU_SDGCTL - EBIU_SDRRC] = R7;
/* Finalize SDC initialization */
/* a transfer is required to finalize SDC initialization! */
IMM32(P4,0x4);
nop;
R7 = [P4];
PollSdcPowerUpFinished:
R7 = w[P5 + EBIU_SDSTAT - EBIU_SDRRC] (z);
```

```
CC = bittst(R7,bitpos(SDRS));
if CC jump PollSdcPowerUpFinished;
(R7:7,P5:4) = [SP++];
ASTAT = [SP++];
unlink;
rts;
init_SDRAM.end;
```

Since this initcode needs to be executed only once, it can be volatile and can be overwritten by other boot blocks.

Initcode (Power Management Control)

The following examples show how to change PLL and the voltage regulator within an initcode. The examples assume that the preboot did not do the job already.

Because the ADSP-BF51x Blackfin is a low power processor, the maximum clock (~80MHz) of the SDRAM controller is lower than the maximum possible system clock (133MHz). See the current data sheets for the actual values if SDRAM is in use.

The ADSP-BF51x processors do not have an on-chip voltage regulator. Set the `bfrom_SysControl` option to `SYSCTRL_EXTVOLTAGE`.

Listing 26-7. Changing PLL and Voltage Regulator in C Language

```
#include <ccb1kfn.h>
#include <bfrom.h>

void init_DPM(ADI_BOOT_DATA* pBS)
{
    ADI_SYSCTRL_VALUES init_DPM;
    init_DPM.uwPllCtl = SET_MSEL(12);
    init_DPM.uwPllDiv = (SET_SSEL(4) | CSEL_DIV1);
    init_DPM.uwPllLockCnt = 0x0200;
    bfrom_SysControl(SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL |
        SYSCTRL_PLLDIV | SYSCTRL_LOCKCNT | SYSCTRL_WRITE, &init_DPM,
        NULL);
}
```

Listing 26-8. Changing PLL and Voltage Regulator in Assembly

```
#include <blackfin.h>
#include <bfrom.h>
.import "bfrom.h";

/* Load Immediate 32-bit value into data or address register */
#define IMM32(reg,val) reg##.H=hi(val); reg##.L=lo(val)

.SECTION L1_code;

init_DPM:

link sizeof(ADI_SYSCTRL_VALUES)+2;
[--SP] = (R7:0,P5:5);
SP += -12;
```

```

R0.L = SET_MSEL(12);
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offsetof(ADI_SYSCTRL_VALUES,uwPllCtl)] = R0;
R0.L = (SET_SSEL(4) | CSEL_DIV1);
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offsetof(ADI_SYSCTRL_VALUES,uwPllDiv)] = R0;
R0.L = 0x0200;
w[FP+-sizeof(ADI_SYSCTRL_VALUES)+offsetof(ADI_SYSCTRL_VALUES,uwPllLockCnt)] = R0;

R0 = (SYSCTRL_EXTVOLTAGE | SYSCTRL_PLLCTL | SYSCTRL_PLLDIV |
SYSCTRL_LOCKCNT | SYSCTRL_WRITE);
R1 = FP;
R1 += -sizeof(ADI_SYSCTRL_VALUES);
R2 = 0 (z);
IMM32(P5,BFROM_SYSCONTROL);
call(P5);

SP += 12;
(R7:0,P5:5) = [SP++];
unlink;
rts;

init_DPM.end:

```

Care must be taken that the reprogramming of the PLL does not break the communication with the booting host. For example, in the case of UART boot, the `UARTx_DLL` and `UARTx_DLH` registers must be updated to keep the old bit rate.

Quickboot With Restore From SDRAM

This example could be part of an advanced power saving concept. Assume the Blackfin is waking up from hibernate and processing any master boot mode. If the SDRAM has not been shut down, but was put in self-refresh mode, the content of the SDRAM will still be valid after wake up. The boot process would only have to initialize on-chip memories. Several boot blocks might be tagged by the `BFLAG_QUICKBOOT` flag.

Some applications might use a power-down handler that saves the contents of L1 memory to SDRAM before entering the hibernate state.

[Listing 26-9](#) assumes a suitable power-down handler was present that generated a partial boot stream in SDRAM at address `0x0001 0000` containing all the instructions required to restore the L1 memory contents.

Listing 26-9. Quickboot with Restore from SDRAM

```
void L1_recovery_initcode (ADI_BOOT_DATA *pBS)
{
    if (pBS->dFlags & BFLAG_WAKEUP) {
        bfrom_MemBoot((void*)0x00010000, BFLAG_RETURN, NULL);
    }
}
```

The boot stream generated at `0x0001 0000` will only be processed upon a wake-up condition. The `BFLAG_RETURN` ensures that the new instance of the boot kernel returns to the `initcode` rather than jumps to the `EVT1` vector.

XOR Checksum

[Listing 26-10](#) illustrates how an initcode can be used to register a callback routine. The routine is called after each boot block that has the `BFLAG_CALLBACK` flag set. The calculated XOR checksum is compared against the block header `ARGUMENT` field. When the checksum fails, this example goes into idle mode. Otherwise control is returned to the boot kernel.

Since this callback example accesses the data after it is loaded, it would fail if the target address were in L1 instruction space. Therefore the `BFLAG_INDIRECT` flag should also be set. The `xor_callback` routine could then perform the checksum calculation at an intermediate storage place. The boot kernel transfers the data from the temporary buffer to the final destination after the callback routine returns.

In general, the block size is bigger than the size of the temporary buffer. Therefore, the boot kernel may need to divide the processing of a single block into multiple steps. The callback routine may also need to be invoked multiple times—every time the temporary buffer is filled up and once for the remaining bytes. The boot kernel passes the `dFlags` parameter, so that the callback routines knows whether it is called the first time, the last time or neither. The `dUserLong` variable in the `ADI_BOOT_DATA` structure is used to store the intermediate results between function calls.

Listing 26-10. XOR Checksum

```
s32 xor_callback(ADI_BOOT_DATA* pBS, ADI_BOOT_BUFFER* pCS, s32
dFlags)
{
    s32 i;
    if ((pCS!= NULL) && (pBS->pHeader!= NULL)) {
        if (dFlags & CBFLAG_FIRST) {
            pBS->dUserLong = 0;
        }
    }
}
```

```

        for (i=0; i<pCS->dByteCount/sizeof(s32); i++)
    {
        pBS->dUserLong^= ((s32 *)pCS->pSource)[i];
    }
    if (dFlags & CBFLAG_FINAL) {
        if (pBS->dUserLong!= pBS->pHeader->dArgument) {
            idle ();
        }
    }
}
return 0;
}
void xor_initcode (ADI_BOOT_DATA *pBS)
{
    pBS->pCallbackFunction = xor_callback;
}

```

Note that the callback routine is not volatile. It should not be overwritten by subsequent boot blocks. It can, however, be overwritten after processing the last block with `BFLAG_CALLBACK` flag set.

The checksum algorithm must be booted first and cannot protect itself. Problems can be avoided by letting initcode and callback execute directly from off-chip flash memory. The ADSP-BF51x processors provide a CRC32 checksum algorithm in the on-chip L1 instruction ROM, that can be used for booting under this scenario. For more information see [“CRC Checksum Calculation” on page 26-44](#).

Direct Code Execution

This code example illustrates how to instruct the VisualDSP++ tools to generate a flash image that causes the boot kernel to start code execution at flash address 0x2000 0020 rather than performing a regular boot. See [“Direct Code Execution” on page 26-32](#).

First, a 32-byte data block is defined in an assembly file that contains the **initial block**.

```
.section bootblock;
.global _firstblock;
.var _firstblock[4] = 0xAD7BD006,
0x20000020, 0x00000010, 0x00000010;
```

Then, the linker is instructed to map the initial block to address **0x2000 0000** in the LDF file.

```
MEMORY
{
    MEM_ASYNC0
    {
        START(0x20000000)
        END(0x23FFFFFF)
        TYPE(ROM)
        WIDTH(8)
    }
}
PROCESSOR p0
{
    RESOLVE(_firstblock,0x20000000)
    RESOLVE(start,0x20000020)
    KEEP(start,_firstblock)
}
SECTIONS
{
    flash
    {
        INPUT_SECTION_ALIGN(4)
        INPUT_SECTIONS($OBJECTS(program)
$LIBRARIES(program))
        INPUT_SECTIONS($OBJECTS(bootblock))
    } >MEM_ASYNC0
```

```
}  
}
```

To invoke the `elfloader` utility, activate the `meminit` feature and use the command-line switches `-romsplitter` and `-maskaddr`. Refer to the application note *Running Programs from Flash on ADSP-BF533 Blackfin Processors (EE-239)* for further details.

Managing PBS Pages in OTP Memory

The following code snips illustrate how to read and write OTP memory, as it is required for the Preboot Settings (PBS). For detailed description of OTP API functions `bfrom_OtpCommand()`, `bfrom_OtpRead()` and `bfrom_OtpWrite()` used here, see [Chapter 3, “One-Time Programmable Memory”](#).

The first example reads PBS settings from OTP and stores them into an instance of the `ADI_PBS_BLOCK` structure. This is a union composite of the `ADI_PBS_HALFPAGES` or the `ADI_PBS_BITFIELDS` types. These structure types are defined in the `bfrom.h` header file. The `dPbsSet` variable describes the set of PBS pages of interest. A `0x00` value reads from OTP pages `0x18` to `0x1B`. A `0x01` value reads from OTP pages `0x1C` to `0x1F` and so on.

Listing 26-11. Reading a Set of PBS Pages from OTP Memory

```
#include <blackfin.h>  
#include <bfrom.h>  
ADI_PBS_BLOCK PBS;  
u32 dPbsSet = 0;  
bfrom_OtpCommand(OTP_INIT, OTP_INIT_VALUE);  
bfrom_OtpRead(PBS00+dPbsSet*4, OTP_LOWER_HALF,  
&(PBS.HalfPages.uqPbs00L));  
bfrom_OtpRead(PBS00+dPbsSet*4,  
OTP_UPPER_HALF, &(PBS.HalfPages.uqPbs00H));
```

```

bfrom_OtpRead(PBS01+dPbsSet*4,OTP_LOWER_HALF,
&(PBS.HalfPages.uqPbs01L));
bfrom_OtpRead(PBS01+dPbsSet*4,OTP_UPPER_HALF,
&(PBS.HalfPages.uqPbs01H));
bfrom_OtpRead(PBS02+dPbsSet*4,OTP_LOWER_HALF,
&(PBS.HalfPages.uqPbs02L));
bfrom_OtpCommand(OTP_CLOSE, 0);

```

The next example shows how PBS pages can be written.

Listing 26-12. Programming a Set of PBS Pages from OTP Memory

```

#include <blackfin.h>
#include <bfrom.h>
ADI_PBS_BLOCK PBS;
u32 dPbsSet = 0;
/* fill PBS with meaningful data */
bfrom_OtpCommand(OTP_INIT, OTP_INIT_VALUE);
bfrom_OtpWrite(PBS00+dPbsSet*4, OTP_LOWER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs00L));
bfrom_OtpWrite(PBS00+dPbsSet*4, OTP_UPPER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs00H));
bfrom_OtpWrite(PBS01+dPbsSet*4, OTP_LOWER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs01L));
bfrom_OtpWrite(PBS01+dPbsSet*4, OTP_UPPER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs01H));
bfrom_OtpWrite(PBS02+dPbsSet*4, OTP_LOWER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs02L));
bfrom_OtpWrite(PBS02+dPbsSet*4, OTP_UPPER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs02H));
bfrom_OtpWrite(PBS03+dPbsSet*4, OTP_LOWER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs03L));
bfrom_OtpWrite(PBS03+dPbsSet*4, OTP_UPPER_HALF |
OTP_CHECK_FOR_PREV_WRITE,&(PBS.HalfPages.uqPbs03H));
bfrom_OtpCommand(OTP_CLOSE, 0);

```

If a set of PBS pages has been written earlier, but need to be replaced by a new set, the old PBS pages have to be invalidated. Do not use the `OTP_CHECK_FOR_PREV_WRITE` option in this case.

Listing 26-13. Invalidating a Set of PBS Pages

```
#include <blackfin.h>
#include <bfrom.h>
u32 dPbsSet = 0;
u64 dlInvalidate = (u64)0xC000000000000000;
bfrom_OtpWrite(PBS00+dPbsSet*4,
bfrom_OtpCommand(OTP_INIT, OTP_INIT_VALUE);
OTP_LOWER_HALF | OTP_NO_ECC, &dlInvalidate);
bfrom_OtpCommand(OTP_CLOSE, 0);
dPbsSet++;
/* write next set as in Listing x-2 */
```

For production you may want to lock the PBS pages to protect them from being overwritten in the field. This can be performed by the following instructions.

Listing 26-14. Write-protecting a Set of PBS Pages

```
#include <blackfin.h>
#include <bfrom.h>
u32 dPbsSet = 0;
bfrom_OtpCommand(OTP_INIT, OTP_INIT_VALUE);
bfrom_OtpWrite(PBS00+dPbsSet*4, OTP_LOCK, NULL);
bfrom_OtpWrite(PBS01+dPbsSet*4, OTP_LOCK, NULL);
bfrom_OtpWrite(PBS02+dPbsSet*4, OTP_LOCK, NULL);
bfrom_OtpWrite(PBS03+dPbsSet*4, OTP_LOCK, NULL);
bfrom_OtpCommand(OTP_CLOSE, 0);
```

When locking PBS pages remember to duplicate the active set of PBS pages best reliability. In the above examples, if the `dPbsSet*4` contains the

final configuration, then program set 5 with the same data. For completeness, note that the above code example does not lock the ECC fields corresponding to the PBS pages. See [Chapter 3, “One-Time Programmable Memory”](#) for details.

27 SYSTEM DESIGN

This chapter provides hardware, software and system design information to aid users in developing systems based on the Blackfin 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 processor data sheet for pin information, including pin numbers.

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`).

Configuring and Servicing Interrupts

Managing Core and System Clocks

The processor produces a multiplication of the clock input provided on the CLKIN pin to generate the PLL VCO clock. This VCO clock is divided to produce the core clock (CCLK) and the system clock (SCLK). The core clock is based on a divider ratio that is programmed via the CSEL bit settings in the PLL_DIV register. The system clock is based on a divider ratio that is programmed via the SSEL bit settings in the PLL_DIV register. For detailed information about how to set and change CCLK and SCLK frequencies, see [Chapter 8, “Dynamic Power Management”](#).

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 [Chapter 5, “System Interrupts”](#).

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 the *Blackfin Processor Programming Reference*.

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` instruction immediately after semaphore release.

The `TESTSET` instruction can be used to implement binary semaphores or any other type of mutual exclusion method. The `TESTSET` instruction supports 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 27-1](#) provides an example of a query semaphore that checks the availability of a shared resource.

Listing 27-1. Query Semaphore

```
/* Query semaphore. Denotes "Busy" if its value is nonzero. Wait
until free (or reschedule thread-- see note below). P0 holds
address of semaphore. */
```

Data Delays, Latencies and Throughput

```
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 nonzero. 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 [Chapter 4, “Chip Bus Hierarchy”](#).

Bus Priorities

For an explanation of prioritization between the various internal buses, refer to [Chapter 4, “Chip Bus Hierarchy”](#).

External Memory Design Issues

This section describes design issues related to external memory.

Example Asynchronous Memory Interfaces

This section shows glueless connections to 16-bit wide SRAM. Note this interface does not require external assertion of \overline{ARDY} , since the internal wait state counter is sufficient for deterministic access times of memories.

Figure 27-1 shows the interface to 8-bit SRAM or flash.

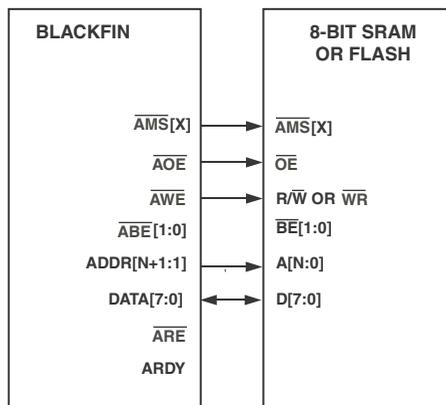


Figure 27-1. Interface to 8-Bit SRAM or Flash

External Memory Design Issues

Figure 27-2 shows the interface to 16-bit SRAM or flash.

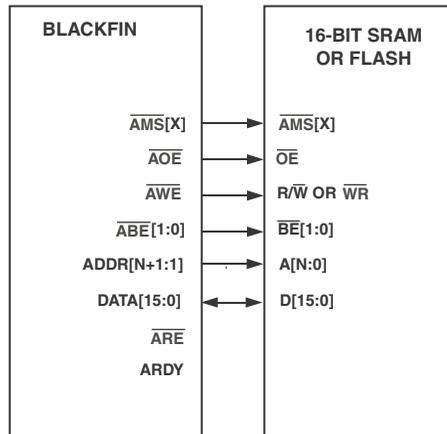


Figure 27-2. Interface to 16-Bit SRAM or Flash

Figure 27-3 shows the system interconnect required to support 16-bit memories. Note this application requires the 16-bit packing mode be enabled for this bank of memory. 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.

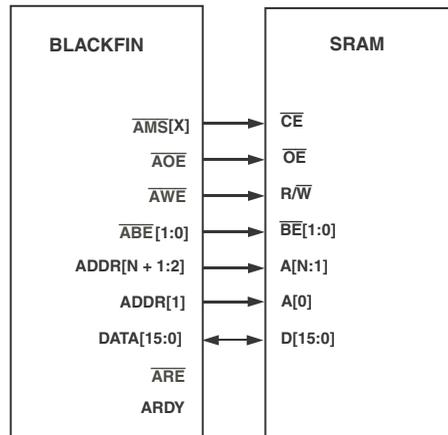


Figure 27-3. Interface to 16-Bit SRAM

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 can 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

High-Frequency Design Considerations

types of accesses on a bank-by-bank basis. Minimally, the external bus interface unit (EBIU) provides one cycle for the transition to occur.

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.

Signal Integrity

In addition to reducing signal length and capacitive loading, critical signals should be treated like transmission lines.

Capacitive loading and signal length of buses can be reduced by using a buffer for devices that operate with wait states (for example, SDRAMs). 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. Extra care should be taken with certain signals such as external memory, read, write, and acknowledge strobes.

Use simple signal integrity methods to prevent transmission line reflections that may cause extraneous extra clock and sync signals. Additionally, avoid overshoot and undershoot that can cause long term damage to input pins.

Some signals are especially critical for short trace length and usually require series termination. The `CLKIN` pin should have impedance matching series resistance at its driver. SPORT interface signals `TCLK`, `RCLK`, `RFS`, and `TFS` should use some termination. Although the serial ports may be operated at a slow rate, the output drivers still have fast edge rates and for longer distances the drivers often require resistive termination located at

the source. (Note also that `TFS` and `RFS` should not be shorted in multi-channel mode.) On the PPI interface, the `PPI_CLK` and `SYNC` signals also benefit from these standard signal integrity techniques. If these pins have multiple sources, it will be difficult to keep the traces short. Consider termination of SDRAM clocks, control, address, and data to improve signal quality and reduce unwanted EMI.

Adding termination to fix a problem on an existing board requires delays for new artwork and new boards. A transmission line simulator is recommended for critical signals. IBIS models are available from Analog Devices Inc. that will assist signal simulation software. Some signals can be corrected with a small zero or 22 ohm resistor located near the driver. The resistor value can be adjusted after measuring the signal at all endpoints.

For details, see the reference sources in [“Recommended Reading” on page 27-12](#) for suggestions on transmission line termination.

Other recommendations and suggestions to promote signal integrity:

- Use more than one ground plane on the Printed Circuit Board (PCB) to reduce crosstalk. Be sure to use lots of vias between the ground planes.
- Keep critical signals such as clocks, strobos, 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.
- 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

High-Frequency Design Considerations

package as shown in [Figure 27-4](#). 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. A ground plane should be located near the component side of the board to reduce the distance that ground current must travel through vias. The ground planes must not be densely perforated with vias or traces as their effectiveness is reduced.

VDDINT is the highest frequency and requires special attention. Two things help power filtering above 100 MHz. First, capacitors should be physically small to reduce the inductance. Surface mount capacitors of size 0402 give better results than larger sizes. Secondly, lower values of capacitance will raise the resonant frequency of the LC circuit. While a cluster of 0.1 μ F is acceptable below 50 MHz, a mix of 0.1 μ F, 0.01 μ F, 0.001 μ F and even 100 pF is preferred in the 500 MHz range.

Note that the instantaneous voltage on both internal and external power pins must at all times be within the recommended operating conditions as specified in the product data sheet. Local “bulk capacitance” (many microfarads) is also necessary. Although all capacitors should be kept close to the power consuming device, small capacitance values should be the closest and larger values may be placed further from the chip.

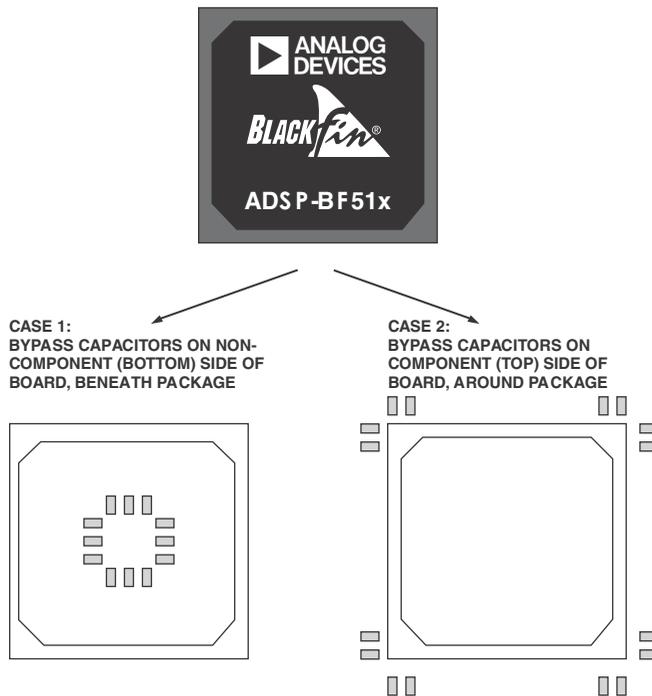


Figure 27-4. Bypass Capacitor Placement

5 Volt Tolerance

Outputs that connect to inputs on 5 V devices can float or be pulled up to 5 V. Most Blackfin pins are not 5 V tolerant. There are a few exceptions such as the TWI pins. Level shifters are required on all other Blackfin pins to keep the pin voltage at or below absolute maximum ratings.

High-Frequency Design Considerations

Test Point Access

The debug process is aided by test points on signals such as CLKOUT or SCLK, bank selects, PPICLK, and $\overline{\text{RESET}}$. If selection pins such as boot mode are connected directly to power or ground, they are inaccessible under a BGA chip. Use pull-up and pull-down resistors instead.

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 to be seen 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

For more information, refer to *High-Speed Digital Design: A Handbook of Black Magic*, Johnson & 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

Consult your CAD software tools vendor. Some companies offer demonstration versions of signal integrity software. Simply by using their free software, you can learn:

- Transmission lines are real
- Unterminated printed circuit board traces will ring and have overshoot and undershoot
- Simple termination will control signal integrity problems

Resetting the Processor

The reset pin requires a monotonic rise and fall. Therefore the pin should not be connected directly to an R/C time delay because such a circuit could be noise sensitive. 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 *Blackfin Processor Programming Reference*. The processor state after reset is also described in the programming reference.

Recommendations for Unused Pins

Most often, there is no need to terminate unused pins, but the handful that do require termination are listed at the end of the pin list description section of the product data sheet.

If the real-time clock is not used, `RTXI` should be pulled low.

Also note that unused peripherals may have separate power connections. These should be driven to the specified value.

Programmable Outputs

During power up, each GPIO pin is set to an input and any pins used in the system as an output should be connected to a pullup or pulldown resistor to maintain the desired state.

This would be particularly important in motor drive applications. It is also important for UART TX and RTS, SPI and serial TWI, or other communications interfaces. Some memory enable pull-ups may also be desired.

After the boot cycle, each GPIO pin may be set to input or output depending on ADSP-BF51x model number and the boot cycle chosen. The I/O / GPIO muxing of all pins may need to be reprogrammed to support the user's application. Care should be taken for compatibility of function and state, before boot, during boot, and application pin usage.

Voltage Regulation Interface

ADSP-BF51x processors must use an external voltage regulator to power the V_{DDINT} domain. The `EXT_WAKE` and \overline{PG} signals can facilitate communication with the external voltage regulator. `EXT_WAKE` is high-true for power-up and low only when the processor is in the hibernate state.

`EXT_WAKE` may be connected directly to the low-true shut down input of many common regulators.

The $\overline{\text{PG}}$ (power-good, low-true) signal that allows the processor to start only after the internal voltage has reached a chosen level. In this way, the startup time of the external regulator will be detected after hibernation.

If the processor never will enter the hibernate state, the $\overline{\text{PG}}$ signal can be grounded in this mode. This will always indicate “power good”, meaning that V_{DDINT} is at a safe operating level. Any delay required at initial power-on, to guarantee a safe operating level for V_{DDINT} , will be provided by the `RESET` signal.

If the external regulator for V_{DDINT} has a power-good signal output, it can be used to help the processor recover properly from its hibernate state. This signal may need to be inverted, as the processor’s input should be low-true in order to indicate a “power good” condition.

If the external regulator does not have a power-good output, the $\overline{\text{PG}}$ signal should be driven to a fixed level (just below the desired operating voltage) so that the $\overline{\text{PG}}$ pin voltage can be compared to V_{DDINT} by the internal startup logic. This can be accomplished with an external resistor divider from V_{DDEXT} or any other fixed stable voltage. A divider with impedance of 1M Ohm is sufficient to supply current to this $\overline{\text{PG}}$ input. To save even more current during hibernation, the `EXT_WAKE` signal may be used as the voltage source to the divider. `EXT_WAKE` is low during hibernation, but will go high before the V_{DDINT} voltage is applied by the external regulator. In all cases, care should be taken to account for the min and max values of V_{DDEXT} or V_{OH} for `EXT_WAKE`. The voltage applied to the $\overline{\text{PG}}$ pin is used as the threshold that is compared internally to the rising value of V_{DDINT} to signal the processor to start. The voltage at $\overline{\text{PG}}$ should be calculated such that the V_{DDINT} value has risen to the desired voltage range for the application.

Voltage Regulation Interface

A SYSTEM MMR ASSIGNMENTS

This appendix lists MMR addresses and register names for all system registers. [Table A-1](#) groups the registers by function/peripheral and indicates the section later in this chapter where individual registers for that group are listed. The tables in the later sections cross reference to individual register diagrams located in the chapter where that register is described. The diagrams show individual bit descriptions for each register.

Table A-1. Register Tables in this Chapter

Function/Peripheral
"System Reset and Interrupt Control Registers" on page A-4
"DMA/Memory DMA Control Registers" on page A-5
"External Bus Interface Unit Registers" on page A-9
"Ports Registers" on page A-9
"Timer Registers" on page A-13
"Core Timer Registers" on page A-3
"Watchdog Timer Registers" on page A-15
"GP Counter Registers" on page A-15
"Real-Time Clock Registers" on page A-16
"OTP and Security Registers" on page A-17
"Dynamic Power Management Registers" on page A-17
"Handshake MDMA Control Registers" on page A-7
"Processor-Specific Memory Registers" on page A-3
"Ethernet MAC Registers" on page A-18

Table A-1. Register Tables in this Chapter (Continued)

Function/Peripheral
“PPI Registers” on page A-25
“SPI Controller Registers” on page A-25
“SPORT Controller Registers” on page A-27
“UART Controller Registers” on page A-30
“Motor Control PWM Registers” on page A-31
“Removable Storage Interface (RSI) Registers” on page A-32
“IEEE 1588 PTP Registers” on page A-23
“TWI Registers” on page A-34

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.

Processor-Specific Memory Registers

Processor-specific memory registers (0xFFE0 0004 – 0xFFE0 0300) are listed in [Table A-2](#).

Table A-2. Processor-Specific Memory Registers

Memory-Mapped Address	Register Name	Description
0xFFE0 0004	DMEM_CONTROL	“L1 Data Memory Control Register” on page 2-6
0xFFE0 0300	DTEST_COMMAND	“Data Test Command Register” on page 2-7

Core Timer Registers

Core timer registers (0xFFE0 3000 – 0xFFE0 300C) are listed in [Table A-3](#).

Table A-3. Core Timer Registers

Memory-Mapped Address	Register Name	Description
0xFFE0 3000	TCNTL	“Core Timer Control Register” on page 11-5
0xFFE0 3004	TPERIOD	“Core Timer Period Register” on page 11-6
0xFFE0 3008	TSCALE	“Core Timer Scale Register” on page 11-7
0xFFE0 300C	TCOUNT	“Core Timer Count Register” on page 11-6

System Reset and Interrupt Control Registers

System reset and interrupt control registers (0xFFC0 0100 – 0xFFC0 01FF) are listed in [Table A-4](#).

Table A-4. System Reset and Interrupt Control Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0104	SYSCR	“System Reset Configuration (SYSCR) Register” on page 26-73
0xFFC0 010C	SIC_IMASK0	“System Interrupt Mask (SIC_IMASK) Register” on page 5-12
0xFFC0 014C	SIC_IMASK1	“System Interrupt Mask (SIC_IMASK) Register” on page 5-12
0xFFC0 0110	SIC_IAR0	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0114	SIC_IAR1	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0118	SIC_IAR2	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 011C	SIC_IAR3	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0150	SIC_IAR4	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0154	SIC_IAR5	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0158	SIC_IAR6	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 015C	SIC_IAR7	“System Interrupt Assignment (SIC_IAR) Register” on page 5-11
0xFFC0 0120	SIC_ISR0	“System Interrupt Status (SIC_ISR) Register” on page 5-12

Table A-4. System Reset and Interrupt Control Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0160	SIC_ISR1	“System Interrupt Status (SIC_ISR) Register” on page 5-12
0xFFC0 0124	SIC_IWR0	“System Interrupt Wakeup-Enable (SIC_IWR) Register” on page 5-12
0xFFC0 0164	SIC_IWR1	“System Interrupt Wakeup-Enable (SIC_IWR) Register” on page 5-12

DMA/Memory DMA Control Registers

DMA control registers (0xFFC0 0B00 – 0xFFC0 0FFF) are listed in [Table A-5](#).

Table A-5. DMA Traffic Control Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0B0C	DMA_TC_PER	“DMA Traffic Control Counter Period Register” on page 6-89
0xFFC0 0B10	DMA_TC_CNT	“DMA Traffic Control Counter Register” on page 6-89

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 A-6](#) and [Table A-7](#). [Table A-6](#) identifies the base address of each DMA channel, as well as the register prefix that identifies the channel. [Table A-7](#) then lists the register suffix and provides its offset from the Base Address.

As an example, the DMA channel 0 Y_MODIFY register is called DMA0_Y_MODIFY, and its address is 0xFFC0 0C1C. Likewise, the memory

DMA/Memory DMA Control Registers

DMA stream 0 source current address register is called MDMA_S0_CURR_ADDR, and its address is 0xFFC0 0E64.

Table A-6. DMA Channel Base Addresses

DMA Channel Identifier	MMR Base Address	Register Prefix
0	0xFFC0 0C00	DMA0_
1	0xFFC0 0C40	DMA1_
2	0xFFC0 0C80	DMA2_
3	0xFFC0 0CC0	DMA3_
4	0xFFC0 0D00	DMA4_
5	0xFFC0 0D40	DMA5_
6	0xFFC0 0D80	DMA6_
7	0xFFC0 0DC0	DMA7_
8	0xFFC0 0E00	DMA8_
9	0xFFC0 0E40	DMA9_
10	0xFFC0 0E80	DMA10_
11	0xFFC0 0EC0	DMA11_
MemDMA stream 0 destination	0xFFC0 0F00	MDMA_D0_
MemDMA stream 0 source	0xFFC0 0F40	MDMA_S0_
MemDMA stream 1 destination	0xFFC0 0F80	MDMA_D1_
MemDMA stream 1 source	0xFFC0 0FC0	MDMA_S1_

Table A-7. DMA Register Suffix and Offset

Register Suffix	Offset From Base	For individual bits, see this diagram:
NEXT_DESC_PTR	0x00	“DMA Next Descriptor Pointer Registers” on page 6-81
START_ADDR	0x04	“DMA Start Address Registers” on page 6-74
CONFIG	0x08	“DMA Configuration Registers” on page 6-67
X_COUNT	0x10	“DMA Inner Loop Count Registers” on page 6-75

Table A-7. DMA Register Suffix and Offset (Continued)

Register Suffix	Offset From Base	For individual bits, see this diagram:
X_MODIFY	0x14	“DMA Inner Loop Address Increment Registers” on page 6-77
Y_COUNT	0x18	“DMA Outer Loop Count Registers” on page 6-78
Y_MODIFY	0x1C	“DMA Outer Loop Address Increment Registers” on page 6-80
CURR_DESC_PTR	0x20	“DMA Current Descriptor Pointer Registers” on page 6-82
CURR_ADDR	0x24	“DMA Current Address Registers” on page 6-75
IRQ_STATUS	0x28	“DMA Interrupt Status Registers” on page 6-72
PERIPHERAL_MAP	0x2C	“DMA Peripheral Map Registers” on page 6-66
CURR_X_COUNT	0x30	“DMA Current Inner Loop Count Registers” on page 6-76
CURR_Y_COUNT	0x38	“DMA Current Outer Loop Count Registers” on page 6-79

Handshake MDMA Control Registers

HMDMA registers (0xFFC0 3300 – 0xFFC0 33FF) are listed in [Table A-8](#).

Table A-8. HMDMA Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 3300	HMDMA0_CONTROL	“Handshake MDMA Control Registers” on page 6-83
0xFFC0 3304	HMDMA0_ECINIT	“Handshake MDMA Initial Edge Count Registers” on page 6-87
0xFFC0 3308	HMDMA0_BCINIT	“Handshake MDMA Initial Block Count Registers” on page 6-84
0xFFC0 330C	HMDMA0_ECURGENT	“Handshake MDMA Edge Count Urgent Registers” on page 6-87
0xFFC0 3310	HMDMA0_ECOVERFLOW	“Handshake MDMA Edge Count Overflow Interrupt Registers” on page 6-88

Handshake MDMA Control Registers

Table A-8. HMDMA Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3314	HMDMA0_ECOUNTER	“Handshake MDMA Initial Edge Count Registers” on page 6-87
0xFFC0 3318	HMDMA0_BCOUNT	“Handshake MDMA Current Block Count Registers” on page 6-85
0xFFC0 3340	HMDMA1_CONTROL	“Handshake MDMA Control Registers” on page 6-83
0xFFC0 3344	HMDMA1_ECINIT	“Handshake MDMA Initial Edge Count Registers” on page 6-87
0xFFC0 3348	HMDMA1_BCINIT	“Handshake MDMA Initial Block Count Registers” on page 6-84
0xFFC0 334C	HMDMA1_ECURGENT	“Handshake MDMA Edge Count Urgent Registers” on page 6-87
0xFFC0 3350	HMDMA1_ECOVERFLOW	“Handshake MDMA Edge Count Overflow Interrupt Registers” on page 6-88
0xFFC0 3354	HMDMA1_ECOUNTER	“Handshake MDMA Initial Edge Count Registers” on page 6-87
0xFFC0 3358	HMDMA1_BCOUNT	“Handshake MDMA Current Block Count Registers” on page 6-85

External Bus Interface Unit Registers

External bus interface unit registers (0xFFC0 0A00 – 0xFFC0 0AFF) are listed in [Table A-9](#).

Table A-9. External Bus Interface Unit Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0A00	EBIU_AMGCTL	“Asynchronous Memory Global Control Register” on page 7-20
0xFFC0 0A04	EBIU_AMBCTL0	“Asynchronous Memory Bank Control 0 Register” on page 7-21
0xFFC0 0A08	EBIU_AMBCTL1	“Asynchronous Memory Bank Control 1 Register” on page 7-22
0xFFC0 0A10	EBIU_SDGCTL	“SDRAM Memory Global Control Register” on page 7-65
0xFFC0 0A14	EBIU_SDBCTL	“SDRAM Memory Bank Control Register” on page 7-61
0xFFC0 0A18	EBIU_SDRRC	“SDRAM Refresh Rate Control Register” on page 7-58
0xFFC0 0A1C	EBIU_SDSTAT	“SDRAM Control Status Register” on page 7-74

Ports Registers

Ports registers (port F: 0xFFC0 0700 – 0xFFC0 07FF, port G: 0xFFC0 1500 – 0xFFC0 15FF, port H: 0xFFC0 1700 – 0xFFC0 17FF, pin control: 0xFFC0 3200 – 0xFFC0 32FF) are listed in [Table A-10](#).

Table A-10. Ports Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0700	PORTFIO	“GPIO Data Registers” on page 9-34
0xFFC0 0704	PORTFIO_CLEAR	“GPIO Clear Registers” on page 9-35

Ports Registers

Table A-10. Ports Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0708	PORTFIO_SET	“GPIO Set Registers” on page 9-34
0xFFC0 070C	PORTFIO_TOGGLE	“GPIO Toggle Registers” on page 9-35
0xFFC0 0710	PORTFIO_MASKA	“GPIO Mask Interrupt A Registers” on page 9-37
0xFFC0 0714	PORTFIO_MASKA_CLEAR	“GPIO Mask Interrupt A Clear Registers” on page 9-41
0xFFC0 0718	PORTFIO_MASKA_SET	“GPIO Mask Interrupt A Set Registers” on page 9-39
0xFFC0 071C	PORTFIO_MASKA_TOGGLE	“GPIO Mask Interrupt A Toggle Registers” on page 9-43
0xFFC0 0720	PORTFIO_MASKB	“GPIO Mask Interrupt B Registers” on page 9-38
0xFFC0 0724	PORTFIO_MASKB_CLEAR	“GPIO Mask Interrupt B Clear Registers” on page 9-42
0xFFC0 0728	PORTFIO_MASKB_SET	“GPIO Mask Interrupt B Set Registers” on page 9-40
0xFFC0 072C	PORTFIO_MASKB_TOGGLE	“GPIO Mask Interrupt B Toggle Registers” on page 9-44
0xFFC0 0730	PORTFIO_DIR	“GPIO Direction Registers” on page 9-33
0xFFC0 0734	PORTFIO_POLAR	“GPIO Polarity Registers” on page 9-36
0xFFC0 0738	PORTFIO_EDGE	“Interrupt Sensitivity Registers” on page 9-36
0xFFC0 073C	PORTFIO_BOTH	“GPIO Set on Both Edges Registers” on page 9-37
0xFFC0 0740	PORTFIO_INEN	“GPIO Input Enable Registers” on page 9-33
0xFFC0 1500	PORTGIO	“GPIO Data Registers” on page 9-34
0xFFC0 1504	PORTGIO_CLEAR	“GPIO Clear Registers” on page 9-35
0xFFC0 1508	PORTGIO_SET	“GPIO Set Registers” on page 9-34

Table A-10. Ports Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 150C	PORTGIO_TOGGLE	“GPIO Toggle Registers” on page 9-35
0xFFC0 1510	PORTGIO_MASKA	“GPIO Mask Interrupt A Registers” on page 9-37
0xFFC0 1514	PORTGIO_MASKA_CLEAR	“GPIO Mask Interrupt A Clear Registers” on page 9-41
0xFFC0 1518	PORTGIO_MASKA_SET	“GPIO Mask Interrupt A Set Registers” on page 9-39
0xFFC0 151C	PORTGIO_MASKA_TOGGLE	“GPIO Mask Interrupt A Toggle Registers” on page 9-43
0xFFC0 1520	PORTGIO_MASKB	“GPIO Mask Interrupt B Registers” on page 9-38
0xFFC0 1524	PORTGIO_MASKB_CLEAR	“GPIO Mask Interrupt B Clear Registers” on page 9-42
0xFFC0 1528	PORTGIO_MASKB_SET	“GPIO Mask Interrupt B Set Registers” on page 9-40
0xFFC0 152C	PORTGIO_MASKB_TOGGLE	“GPIO Mask Interrupt B Toggle Registers” on page 9-44
0xFFC0 1530	PORTGIO_DIR	“GPIO Direction Registers” on page 9-33
0xFFC0 1534	PORTGIO_POLAR	“GPIO Polarity Registers” on page 9-36
0xFFC0 1538	PORTGIO_EDGE	“Interrupt Sensitivity Registers” on page 9-36
0xFFC0 153C	PORTGIO_BOTH	“GPIO Set on Both Edges Registers” on page 9-37
0xFFC0 1540	PORTGIO_INEN	“GPIO Input Enable Registers” on page 9-33
0xFFC0 1700	PORTHIO	“GPIO Data Registers” on page 9-34
0xFFC0 1704	PORTHIO_CLEAR	“GPIO Clear Registers” on page 9-35
0xFFC0 1708	PORTHIO_SET	“GPIO Set Registers” on page 9-34
0xFFC0 170C	PORTHIO_TOGGLE	“GPIO Toggle Registers” on page 9-35

Ports Registers

Table A-10. Ports Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 1710	PORTHIO_MASKA	“GPIO Mask Interrupt A Registers” on page 9-37
0xFFC0 1714	PORTHIO_MASKA_CLEAR	“GPIO Mask Interrupt A Clear Registers” on page 9-41
0xFFC0 1718	PORTHIO_MASKA_SET	“GPIO Mask Interrupt A Set Registers” on page 9-39
0xFFC0 171C	PORTHIO_MASKA_TOGGLE	“GPIO Mask Interrupt A Toggle Registers” on page 9-43
0xFFC0 1720	PORTHIO_MASKB	“GPIO Mask Interrupt B Registers” on page 9-38
0xFFC0 1724	PORTHIO_MASKB_CLEAR	“GPIO Mask Interrupt B Clear Registers” on page 9-42
0xFFC0 1728	PORTHIO_MASKB_SET	“GPIO Mask Interrupt B Set Registers” on page 9-40
0xFFC0 172C	PORTHIO_MASKB_TOGGLE	“GPIO Mask Interrupt B Toggle Registers” on page 9-44
0xFFC0 1730	PORTHIO_DIR	“GPIO Direction Registers” on page 9-33
0xFFC0 1734	PORTHIO_POLAR	“GPIO Polarity Registers” on page 9-36
0xFFC0 1738	PORTHIO_EDGE	“Interrupt Sensitivity Registers” on page 9-36
0xFFC0 173C	PORTHIO_BOTH	“GPIO Set on Both Edges Registers” on page 9-37
0xFFC0 1740	PORTHIO_INEN	“GPIO Input Enable Registers” on page 9-33
0xFFC0 3200	PORTF_FER	“Function Enable Registers” on page 9-32
0xFFC0 3204	PORTG_FER	“Function Enable Registers” on page 9-32
0xFFC0 3208	PORTH_FER	“Function Enable Registers” on page 9-32
0xFFC0 3210	PORTF_MUX	“Port F Multiplexer Control Register” on page 9-30

Table A-10. Ports Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3214	PORTG_MUX	“Port F Multiplexer Control Register” on page 9-30
0xFFC0 3218	PORTH_MUX	“Port F Multiplexer Control Register” on page 9-30
0xFFC0 3240	PORTF_HYSTERESIS	“Port F Hysteresis Register” on page 9-26
0xFFC0 3244	PORTG_HYSTERESIS	“Port G Hysteresis Register” on page 9-27
0xFFC0 3248	PORTH_HYSTERESIS	“Port H Hysteresis Register” on page 9-27
0xFFC0 3280	NONGPIO_DRIVE	“Non-GPIO Drive Strength Control Register” on page 9-28
0xFFC0 3288	NONGPIO_HYSTERESIS	“Non-GPIO Hysteresis (NONGPIO_HYSTERESIS) Register” on page 9-29

Timer Registers

Timer registers (0xFFC0 0600 – 0xFFC0 06FF) are listed in [Table A-11](#).

Table A-11. Timer Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0600	TIMER0_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0604	TIMER0_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0608	TIMER0_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 060C	TIMER0_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0610	TIMER1_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0614	TIMER1_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0618	TIMER1_PERIOD	“Timer Period Register” on page 10-43

Timer Registers

Table A-11. Timer Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 061C	TIMER1_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0620	TIMER2_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0624	TIMER2_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0628	TIMER2_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 062C	TIMER2_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0630	TIMER3_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0634	TIMER3_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0638	TIMER3_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 063C	TIMER3_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0640	TIMER4_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0644	TIMER4_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0648	TIMER4_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 064C	TIMER4_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0650	TIMER5_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0654	TIMER5_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0658	TIMER5_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 065C	TIMER5_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0660	TIMER6_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0664	TIMER6_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0668	TIMER6_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 066C	TIMER6_WIDTH	“Timer Width Register” on page 10-44
0xFFC0 0670	TIMER7_CONFIG	“Timer Configuration Register” on page 10-40
0xFFC0 0674	TIMER7_COUNTER	“Timer Counter Register” on page 10-41
0xFFC0 0678	TIMER7_PERIOD	“Timer Period Register” on page 10-43
0xFFC0 067C	TIMER7_WIDTH	“Timer Width Register” on page 10-44

Table A-11. Timer Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0680	TIMER_ENABLE	“Timer Enable Register” on page 10-36
0xFFC0 0684	TIMER_DISABLE	“Timer Disable Register” on page 10-37
0xFFC0 0688	TIMER_STATUS	“Timer Status Register” on page 10-39

Watchdog Timer Registers

Watchdog timer registers (0xFFC0 0200 – 0xFFC0 02FF) are listed in [Table A-12](#).

Table A-12. Watchdog Timer Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0200	WDOG_CTL	“Watchdog Control Register” on page 12-8
0xFFC0 0204	WDOG_CNT	“Watchdog Count Register” on page 12-6
0xFFC0 0208	WDOG_STAT	“Watchdog Status Register” on page 12-7

GP Counter Registers

GP Counter registers (0xFFC0 3500 – 0xFFC0 351C) are listed in [Table A-13](#).

Table A-13. GP Counter Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 3500	CNT_CONFIG	“Counter Configuration Register” on page 13-20
0xFFC0 3504	CNT_IMASK	“Counter Interrupt Mask Register” on page 13-21

Real-Time Clock Registers

Table A-13. GP Counter Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3508	CNT_STATUS	“Counter Status Register” on page 13-22
0xFFC0 350C	CNT_COMMAND	“Counter Command Register” on page 13-24
0xFFC0 3510	CNT_DEBOUNCE	“Counter Debounce Register” on page 13-25
0xFFC0 3514	CNT_COUNTER	“Counter Count Value Register” on page 13-26
0xFFC0 3518	CNT_MAX	“Counter Maximal Count Register” on page 13-27
0xFFC0 351C	CNT_MIN	“Counter Minimal Count Register” on page 13-27

Real-Time Clock Registers

Real-time clock registers (0xFFC0 0300 – 0xFFC0 03FF) are listed in [Table A-14](#).

Table A-14. Real-Time Clock Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0300	RTC_STAT	“RTC Status Register” on page 24-20
0xFFC0 0304	RTC_ICTL	“RTC Interrupt Control Register” on page 24-20
0xFFC0 0308	RTC_ISTAT	“RTC Status Register” on page 24-20
0xFFC0 030C	RTC_SWCNT	“RTC Stopwatch Count Register” on page 24-21
0xFFC0 0310	RTC_ALARM	“RTC Alarm Register” on page 24-22
0xFFC0 0314	RTC_PREN	“RTC Prescaler Enable Register” on page 24-22

OTP and Security Registers

OTP Controller and Security Controller registers (0xFFC0 3600 – 0xFFC0 368C) are listed in [Table A-15](#).

Table A-15. OTP Controller and Security Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0360C	OTP_TIMING	“OTP_TIMING Register” on page 3-14
0xFFC0 3620	SECURE_SYSSWT	“Secure System Switch Register, Bits 15:0” on page 25-47
0xFFC0 3624	SECURE_CONTROL	“Secure Control Register” on page 25-53
0xFFC0 3628	SECURE_STATUS	“Secure Status Register” on page 25-55
0xFFC0 3680	OTP_DATA0	“OTP_DATAx Registers” on page 3-3
0xFFC0 3684	OTP_DATA1	“OTP_DATAx Registers” on page 3-3
0xFFC0 3688	OTP_DATA2	“OTP_DATAx Registers” on page 3-3
0xFFC0 368C	OTP_DATA3	“OTP_DATAx Registers” on page 3-3

Dynamic Power Management Registers

Dynamic power management registers (0xFFC0 0000 – 0xFFC0 00FF) are listed in [Table A-16](#).

Table A-16. Dynamic Power Management Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0000	PLL_CTL	“PLL Control Register” on page 8-21
0xFFC0 0004	PLL_DIV	“PLL Divide Register” on page 8-21
0xFFC0 0008	VR_CTL	“Voltage Regulator Control Register” on page 8-23

Ethernet MAC Registers

Table A-16. Dynamic Power Management Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0_000C	PLL_STAT	“PLL Status Register” on page 8-22
0xFFC0_0010	PLL_LOCKCNT	“PLL Lock Count Register” on page 8-22

Ethernet MAC Registers

Ethernet MAC registers (0xFFC0_3000 – 0xFFC0_31FF) are listed in [Table A-17](#).

Table A-17. Ethernet MAC Registers

Memory-Mapped Address	Register Name	Description
0xFFC0_3000	EMAC_OPMODE	“EMAC_OPMODE Register” on page 22-60
0xFFC0_3004	EMAC_ADDRLO	“EMAC_ADDRLO Register” on page 22-67
0xFFC0_3008	EMAC_ADDRHI	“EMAC_ADDRHI Register” on page 22-68
0xFFC0_300C	EMAC_HASHLO	“EMAC_HASHLO Register” on page 22-69
0xFFC0_3010	EMAC_HASHHI	“EMAC_HASHHI Register” on page 22-70
0xFFC0_3014	EMAC_STAADD	“EMAC_STAADD Register” on page 22-72
0xFFC0_3018	EMAC_STADAT	“EMAC_STADAT Register” on page 22-74
0xFFC0_301C	EMAC_FLC	“EMAC_FLC Register” on page 22-75
0xFFC0_3020	EMAC_VLAN1	“EMAC_VLAN1 Register” on page 22-77
0xFFC0_3024	EMAC_VLAN2	“EMAC_VLAN2 Register” on page 22-77
0xFFC0_302C	EMAC_WKUP_CTL	“EMAC_WKUP_CTL Register” on page 22-78
0xFFC0_3030	EMAC_WKUP_FFMSK0	“EMAC_WKUP_FFMSK0 Register” on page 22-81
0xFFC0_3034	EMAC_WKUP_FFMSK1	“EMAC_WKUP_FFMSK1 Register” on page 22-82

Table A-17. Ethernet MAC Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3038	EMAC_WKUP_FFMSK2	“EMAC_WKUP_FFMSK2 Register” on page 22-83
0xFFC0 303C	EMAC_WKUP_FFMSK3	“EMAC_WKUP_FFMSK3 Register” on page 22-84
0xFFC0 3040	EMAC_WKUP_FFCMD	“EMAC_WKUP_FFCMD Register” on page 22-85
0xFFC0 3044	EMAC_WKUP_FFOFF	“EMAC_WKUP_FFOFF Register” on page 22-87
0xFFC0 3048	EMAC_WKUP_FFCRC0/1	“EMAC_WKUP_FFCRC0 Register” on page 22-88
0xFFC0 304C	EMAC_WKUP_FFCRC2/3	“EMAC_WKUP_FFCRC1 Register” on page 22-88
0xFFC0 3060	EMAC_SYSCTL	“EMAC_SYSCTL Register” on page 22-89
0xFFC0 3064	EMAC_SYSTAT	“EMAC_SYSTAT Register” on page 22-91
0xFFC0 3068	EMAC_RX_STAT	“EMAC_RX_STAT Register” on page 22-94
0xFFC0 306C	EMAC_RX_STKY	“EMAC_RX_STKY Register” on page 22-99
0xFFC0 3070	EMAC_RX_IRQE	“EMAC_RX_IRQE Register” on page 22-104
0xFFC0 3074	EMAC_TX_STAT	“EMAC_TX_STAT Register” on page 22-105
0xFFC0 3078	EMAC_TX_STKY	“EMAC_TX_STKY Register” on page 22-109
0xFFC0 307C	EMAC_TX_IRQE	“EMAC_TX_IRQE Register” on page 22-112
0xFFC0 3080	EMAC_MMC_CTL	“EMAC_MMC_CTL Register” on page 22-119
0xFFC0 3084	EMAC_MMC_RIRQS	“EMAC_MMC_RIRQS Register” on page 22-113
0xFFC0 3088	EMAC_MMC_RIRQE	“EMAC_MMC_RIRQE Register” on page 22-115
0xFFC0 308C	EMAC_MMC_TIRQS	“EMAC_MMC_TIRQS Register” on page 22-116

Ethernet MAC Registers

Table A-17. Ethernet MAC Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3090	EMAC_MMC_TIRQE	“EMAC_MMC_TIRQE Register” on page 22-118
0xFFC0 3100	EMAC_RXC_OK	“MAC Management Counter Registers” on page 22-53
0xFFC0 3104	EMAC_RXC_FCS	“MAC Management Counter Registers” on page 22-53
0xFFC0 3108	EMAC_RXC_ALIGN	“MAC Management Counter Registers” on page 22-53
0xFFC0 310C	EMAC_RXC_OCTET	“MAC Management Counter Registers” on page 22-53
0xFFC0 3110	EMAC_RXC_DMAOVF	“MAC Management Counter Registers” on page 22-53
0xFFC0 3114	EMAC_RXC_UNICST	“MAC Management Counter Registers” on page 22-53
0xFFC0 3118	EMAC_RXC_MULTI	“MAC Management Counter Registers” on page 22-53
0xFFC0 311C	EMAC_RXC_BROAD	“MAC Management Counter Registers” on page 22-53
0xFFC0 3120	EMAC_RXC_LNERRI	“MAC Management Counter Registers” on page 22-53
0xFFC0 3124	EMAC_RXC_LNERRO	“MAC Management Counter Registers” on page 22-53
0xFFC0 3128	EMAC_RXC_LONG	“MAC Management Counter Registers” on page 22-53
0xFFC0 312C	EMAC_RXC_MACCTL	“MAC Management Counter Registers” on page 22-53
0xFFC0 3130	EMAC_RXC_OPCODE	“MAC Management Counter Registers” on page 22-53
0xFFC0 3134	EMAC_RXC_PAUSE	“MAC Management Counter Registers” on page 22-53

Table A-17. Ethernet MAC Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3138	EMAC_RXC_ALLFRM	“MAC Management Counter Registers” on page 22-53
0xFFC0 313C	EMAC_RXC_ALLOCT	“MAC Management Counter Registers” on page 22-53
0xFFC0 3140	EMAC_RXC_TYPED	“MAC Management Counter Registers” on page 22-53
0xFFC0 3144	EMAC_RXC_SHORT	“MAC Management Counter Registers” on page 22-53
0xFFC0 3148	EMAC_RXC_EQ64	“MAC Management Counter Registers” on page 22-53
0xFFC0 314C	EMAC_RXC_LT128	“MAC Management Counter Registers” on page 22-53
0xFFC0 3150	EMAC_RXC_LT256	“MAC Management Counter Registers” on page 22-53
0xFFC0 3154	EMAC_RXC_LT512	“MAC Management Counter Registers” on page 22-53
0xFFC0 3158	EMAC_RXC_LT1024	“MAC Management Counter Registers” on page 22-53
0xFFC0 315C	EMAC_RXC_GE1024	“MAC Management Counter Registers” on page 22-53
0xFFC0 3180	EMAC_TXC_OK	“MAC Management Counter Registers” on page 22-53
0xFFC0 3184	EMAC_TXC_1COL	“MAC Management Counter Registers” on page 22-53
0xFFC0 3188	EMAC_TXC_GT1COL	“MAC Management Counter Registers” on page 22-53
0xFFC0 318C	EMAC_TXC_OCTET	“MAC Management Counter Registers” on page 22-53
0xFFC0 3190	EMAC_TXC_DEFER	“MAC Management Counter Registers” on page 22-53

Ethernet MAC Registers

Table A-17. Ethernet MAC Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3194	EMAC_TXC_LATECL	“MAC Management Counter Registers” on page 22-53
0xFFC0 3198	EMAC_TXC_XS_COL	“MAC Management Counter Registers” on page 22-53
0xFFC0 319C	EMAC_TXC_DMAUND	“MAC Management Counter Registers” on page 22-53
0xFFC0 31A0	EMAC_TXC_CRSERR	“MAC Management Counter Registers” on page 22-53
0xFFC0 31A4	EMAC_TXC_UNICST	“MAC Management Counter Registers” on page 22-53
0xFFC0 31A8	EMAC_TXC_MULTI	“MAC Management Counter Registers” on page 22-53
0xFFC0 31AC	EMAC_TXC_BROAD	“MAC Management Counter Registers” on page 22-53
0xFFC0 31B0	EMAC_TXC_ES_DFR	“MAC Management Counter Registers” on page 22-53
0xFFC0 31B4	EMAC_TXC_MACCTL	“MAC Management Counter Registers” on page 22-53
0xFFC0 31B8	EMAC_TXC_ALLFRM	“MAC Management Counter Registers” on page 22-53
0xFFC0 31BC	EMAC_TXC_ALLOCT	“MAC Management Counter Registers” on page 22-53
0xFFC0 31C0	EMAC_TXC_EQ64	“MAC Management Counter Registers” on page 22-53
0xFFC0 31C4	EMAC_TXC_LT128	“MAC Management Counter Registers” on page 22-53
0xFFC0 31C8	EMAC_TXC_LT254	“MAC Management Counter Registers” on page 22-53
0xFFC0 31CC	EMAC_TXC_LT512	“MAC Management Counter Registers” on page 22-53

Table A-17. Ethernet MAC Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 31D0	EMAC_TXC_LT1024	“MAC Management Counter Registers” on page 22-53
0xFFC0 31D4	EMAC_TXC_GE1024	“MAC Management Counter Registers” on page 22-53
0xFFC0 31D8	EMAC_TXC_ABORT	“MAC Management Counter Registers” on page 22-53

IEEE 1588 PTP Registers

IEEE 1588 PTP registers (0xFFC0 30A0 – 0xFFC0 30F8) are listed in [Table A-18](#).

Table A-18. IEEE 1588 PTP Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 30A0	EMAC_PTP_CTL	“Control Register” on page 23-13
0xFFC0 30A4	EMAC_PTP_IE	“Interrupt Enable Register” on page 23-15
0xFFC0 30A8	EMAC_PTP_ISTAT	“Interrupt Status Register” on page 23-17
0xFFC0 30AC	EMAC_PTP_FOFF	“Message Filter Offset Register” on page 23-19
0xFFC0 30B0	EMAC_PTP_FV1	“Message Filter Value Register 1” on page 23-20
0xFFC0 30B4	EMAC_PTP_FV2	“Message Filter Value Register 2” on page 23-21
0xFFC0 30B8	EMAC_PTP_FV3	“Message Filter Value Register 3” on page 23-22
0xFFC0 30BC	EMAC_PTP_ADDEND	“Addend Register” on page 23-23
0xFFC0 30C0	EMAC_PTP_ACCR	“Accumulator Register” on page 23-24

IEEE 1588 PTP Registers

Table A-18. IEEE 1588 PTP Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC030C4	EMAC_PTP_OFFSET	“Time Offset Register” on page 23-25
0xFFC030C8	EMAC_PTP_TIMELO	“Local Clock Time Low Register” on page 23-26
0xFFC030CC	EMAC_PTP_TIMEHI	“Local Clock Time High Register” on page 23-27
0xFFC030D0	EMAC_PTP_RXSNAPLO	“Receive Snapshot Low Register” on page 23-28
0xFFC030D4	EMAC_PTP_RXSNAPHI	“Receive Snapshot High Register” on page 23-29
0xFFC030D8	EMAC_PTP_TXSNAPLO	“Transmit Snapshot Low Register” on page 23-30
0xFFC030DC	EMAC_PTP_TXSNAPHI	“Transmit Snapshot High Register” on page 23-31
0xFFC030E0	EMAC_PTP_ALARMLO	“Target Alarm Time Low Register” on page 23-32
0xFFC030E4	EMAC_PTP_ALARMHI	“Target Alarm Time High Register” on page 23-33
0xFFC030E8	EMAC_PTP_ID_OFF	“Source ID Offset Register” on page 23-33
0xFFC030EC	EMAC_PTP_ID_SNAP	“Source ID Snapshot Register” on page 23-34
0xFFC030F0	EMAC_PTP_PPS_STARTLO	“PPS Start Low Register” on page 23-35
0xFFC030F4	EMAC_PTP_PPS_STARTHI	“PSS Start High Register” on page 23-36
0xFFC030F8	EMAC_PTP_PPS_PERIOD	“PSS Period Register” on page 23-37

PPI Registers

PPI registers (0xFFC0 1000 – 0xFFC0 10FF) are listed in [Table A-19](#).

Table A-19. PPI Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 1000	PPI_CONTROL	“PPI Control Register” on page 20-26
0xFFC0 1004	PPI_STATUS	“PPI Status Register” on page 20-30
0xFFC0 1008	PPI_COUNT	“PPI Transfer Count Register” on page 20-33
0xFFC0 100C	PPI_DELAY	“PPI Delay Count Register” on page 20-32
0xFFC0 1010	PPI_FRAME	“PPI Lines Per Frame Register” on page 20-34

SPI Controller Registers

SPI0 controller registers (0xFFC0 0500 – 0xFFC0 05FF) are listed in [Table A-20](#).

SPI1 controller registers (0xFFC0 3400 – 0xFFC0 34FF) are listed in [Table A-21](#).

Table A-20. SPI0 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0500	SPI0_CTL	“SPI Control Register” on page 17-36
0xFFC0 0504	SPI0_FLG	“SPI Flag Register (example with 7 slave selects)” on page 17-37
0xFFC0 0508	SPI0_STAT	“SPI Status Register” on page 17-39
0xFFC0 050C	SPI0_TDBR	“SPI Transmit Data Buffer Register” on page 17-42
0xFFC0 0510	SPI0_RDBR	“SPI Receive Data Buffer Register” on page 17-43

SPI Controller Registers

Table A-20. SPI0 Controller Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0514	SPI0_BAUD	“SPI Baud Rate Register” on page 17-34
0xFFC0 0518	SPI0_SHADOW	“SPI RDBR Shadow Register” on page 17-43

Table A-21. SPI1 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 3400	SPI1_CTL	“SPI Control Register” on page 17-36
0xFFC0 3404	SPI1_FLG	“SPI Flag Register (example with 7 slave selects)” on page 17-37
0xFFC0 3408	SPI1_STAT	“SPI Status Register” on page 17-39
0xFFC0 340C	SPI1_TDBR	“SPI Transmit Data Buffer Register” on page 17-42
0xFFC0 3410	SPI1_RDBR	“SPI Receive Data Buffer Register” on page 17-43
0xFFC0 3414	SPI1_BAUD	“SPI Baud Rate Register” on page 17-34
0xFFC0 3418	SPI1_SHADOW	“SPI RDBR Shadow Register” on page 17-43

SPORT Controller Registers

SPORT0 controller registers (0xFFC0 0800 – 0xFFC0 08FF) are listed in [Table A-22](#). SPORT1 controller registers (0xFFC0 0900 – 0xFFC0 09FF) are listed in [Table A-23](#).

Table A-22. SPORT0 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0800	SPORT0_TCR1	“SPORT Transmit Configuration 1 Register” on page 19-47
0xFFC0 0804	SPORT0_TCR2	“SPORT Transmit Configuration 2 Register” on page 19-48
0xFFC0 0808	SPORT0_TCLKDIV	“SPORT Transmit Serial Clock Divider Register” on page 19-62
0xFFC0 080C	SPORT0_TFSDIV	“SPORT Transmit Frame Sync Divider Register” on page 19-63
0xFFC0 0810	SPORT0_TX	“SPORT Transmit Data Register” on page 19-57
0xFFC0 0818	SPORT0_RX	“SPORT Receive Data Register” on page 19-59
0xFFC0 0820	SPORT0_RCR1	“SPORT Receive Configuration 1 Register” on page 19-52
0xFFC0 0824	SPORT0_RCR2	“SPORT Receive Configuration 2 Register” on page 19-53
0xFFC0 0828	SPORT0_RCLKDIV	“SPORT Receive Serial Clock Divider Register” on page 19-62
0xFFC0 082C	SPORT0_RFSDIV	“SPORT Receive Frame Sync Divider Register” on page 19-63
0xFFC0 0830	SPORT0_STAT	“SPORT Status Register” on page 19-60
0xFFC0 0834	SPORT0_CHNL	“SPORT Current Channel Register” on page 19-65
0xFFC0 0838	SPORT0_MCMC1	“SPORT Multichannel Configuration Register 1” on page 19-64

SPORT Controller Registers

Table A-22. SPORT0 Controller Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 083C	SPORT0_MCMC2	“SPORT Multichannel Configuration Register 2” on page 19-64
0xFFC0 0840	SPORT0_MTCS0	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0844	SPORT0_MTCS1	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0848	SPORT0_MTCS2	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 084C	SPORT0_MTCS3	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0850	SPORT0_MRCS0	“SPORT Multichannel Receive Select Registers” on page 19-66
0xFFC0 0854	SPORT0_MRCS1	“SPORT Multichannel Receive Select Registers” on page 19-66
0xFFC0 0858	SPORT0_MRCS2	“SPORT Multichannel Receive Select Registers” on page 19-66
0xFFC0 085C	SPORT0_MRCS3	“SPORT Multichannel Receive Select Registers” on page 19-66

Table A-23. SPORT1 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0900	SPORT1_TCR1	“SPORT Transmit Configuration 1 Register” on page 19-47
0xFFC0 0904	SPORT1_TCR2	“SPORT Transmit Configuration 2 Register” on page 19-48
0xFFC0 0908	SPORT1_TCLKDIV	“SPORT Transmit Serial Clock Divider Register” on page 19-62
0xFFC0 090C	SPORT1_TFSDIV	“SPORT Transmit Frame Sync Divider Register” on page 19-63

Table A-23. SPORT1 Controller Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0910	SPORT1_TX	“SPORT Transmit Data Register” on page 19-57
0xFFC0 0918	SPORT1_RX	“SPORT Receive Data Register” on page 19-59
0xFFC0 0920	SPORT1_RCR1	“SPORT Receive Configuration 1 Register” on page 19-52
0xFFC0 0924	SPORT1_RCR2	“SPORT Receive Configuration 2 Register” on page 19-53
0xFFC0 0928	SPORT1_RCLKDIV	“SPORT Receive Serial Clock Divider Register” on page 19-62
0xFFC0 092C	SPORT1_RFSDIV	“SPORT Receive Frame Sync Divider Register” on page 19-63
0xFFC0 0930	SPORT1_STAT	“SPORT Status Register” on page 19-60
0xFFC0 0934	SPORT1_CHNL	“SPORT Current Channel Register” on page 19-65
0xFFC0 0938	SPORT1_MCMC1	“SPORT Multichannel Configuration Register 1” on page 19-64
0xFFC0 093C	SPORT1_MCMC2	“SPORT Multichannel Configuration Register 2” on page 19-64
0xFFC0 0940	SPORT1_MTCS0	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0944	SPORT1_MTCS1	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0948	SPORT1_MTCS2	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 094C	SPORT1_MTCS3	“SPORT Multichannel Transmit Select Registers” on page 19-67
0xFFC0 0950	SPORT1_MRCS0	“SPORT Multichannel Receive Select Registers” on page 19-66
0xFFC0 0954	SPORT1_MRCS1	“SPORT Multichannel Receive Select Registers” on page 19-66

UART Controller Registers

Table A-23. SPORT1 Controller Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 0958	SPORT1_MRCS2	“SPORT Multichannel Receive Select Registers” on page 19-66
0xFFC0 095C	SPORT1_MRCS3	“SPORT Multichannel Receive Select Registers” on page 19-66

UART Controller Registers

UART0 controller registers (0xFFC0 0400 – 0xFFC0 04FF) are listed in [Table A-24](#). UART1 controller registers (0xFFC0 2000 – 0xFFC0 20FF) are listed in [Table A-25](#).

Table A-24. UART0 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 0400	UART0_THR	“UART Transmit Holding Register” on page 15-26
0xFFC0 0400	UART0_RBR	“UART Receive Buffer Register” on page 15-26
0xFFC0 0400	UART0_DLL	“UART Divisor Latch Registers” on page 15-30
0xFFC0 0404	UART0_DLH	“UART Divisor Latch Registers” on page 15-30
0xFFC0 0404	UART0_IER	“UART Interrupt Enable Register” on page 15-27
0xFFC0 0408	UART0_IIR	“UART Interrupt Identification Register” on page 15-29
0xFFC0 040C	UART0_LCR	“UART Line Control Register” on page 15-21
0xFFC0 0410	UART0_MCR	“UART Modem Control Registers” on page 15-23
0xFFC0 0414	UART0_LSR	“UART Line Status Register” on page 15-24
0xFFC0 041C	UART0_SCR	“UART Scratch Register” on page 15-31
0xFFC0 0424	UART0_GCTL	“UART Global Control Register” on page 15-31

Table A-25. UART1 Controller Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 2000	UART1_THR	“UART Transmit Holding Register” on page 15-26
0xFFC0 2000	UART1_RBR	“UART Receive Buffer Register” on page 15-26
0xFFC0 2000	UART1_DLL	“UART Divisor Latch Registers” on page 15-30
0xFFC0 2004	UART1_DLH	“UART Divisor Latch Registers” on page 15-30
0xFFC0 2004	UART1_IER	“UART Interrupt Enable Register” on page 15-27
0xFFC0 2008	UART1_IIR	“UART Interrupt Identification Register” on page 15-29
0xFFC0 200C	UART1_LCR	“UART Line Control Register” on page 15-21
0xFFC0 2010	UART1_MCR	“UART Modem Control Registers” on page 15-23
0xFFC0 2014	UART1_LSR	“UART Line Status Register” on page 15-24
0xFFC0 201C	UART1_SCR	“UART Scratch Register” on page 15-31
0xFFC0 2024	UART1_GCTL	“UART Global Control Register” on page 15-31

Motor Control PWM Registers

Motor Control PWM registers (0xFFC0 3700 – 0xFFC0 37FF) are listed in [Table A-26](#).

Table A-26. Motor Control PWM Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 3700	PWM_CTRL	“PWM Control Register” on page 14-37
0xFFC0 3704	PWM_STAT	“PWM_STAT Register” on page 14-39
0xFFC0 3708	PWM_TM	“PWM_TM Register” on page 14-40
0xFFC0 370C	PWM_DT	“PWM Dead Time Register” on page 14-41
0xFFC0 3710	PWM_GATE	“PWM Chopping Control Register” on page 14-42

Removable Storage Interface (RSI) Registers

Table A-26. Motor Control PWM Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3714	PWM_CHA	“PWM_CHA Register” on page 14-43
0xFFC0 3718	PWM_CHB	“PWM_CHB Register” on page 14-43
0xFFC0 371C	PWM_CHC	“PWM_CHC Register” on page 14-43
0xFFC0 3720	PWM_SEG	“PWM_SEG Register” on page 14-44
0xFFC0 3724	PWM_SYNCWT	“PWM Sync Pulse Width Control Register” on page 14-46
0xFFC0 3728	PWM_CHAL	“PWM Channel AL Duty Control Register” on page 14-47
0xFFC0372C	PWM_CHBL	“PWM Channel BL Duty Control Register” on page 14-47
0xFFC0 3730	PWM_CHCL	“PWM_CHCL Register” on page 14-48
0xFFC0 3734	PWM_LSI	“PWM Low Side Invert Register” on page 14-48
0xFFC0 3738	PWM_STAT2	“PWM_STAT2 Register” on page 14-49

Removable Storage Interface (RSI) Registers

RSI registers (0xFFC0 3800 – 0xFFC0 3CFF) are listed in [Table A-27](#).

Table A-27. RSI Registers

Memory-Mapped Address	Register Name	Description
0xFFC0 3800	RSI_PWR_CONTROL	“RSI Power Control Register” on page 21-54
0xFFC0 3804	RSI_CLK_CONTROL	“RSI Clock Control Register” on page 21-55
0xFFC0 3808	RSI_ARGUMENT	“RSI Argument Register” on page 21-56
0xFFC0 380C	RSI_COMMAND	“RSI Command Register” on page 21-57

Table A-27. RSI Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 3810	RSI_RESP_CMD	“RSI Response Command Register” on page 21-59
0xFFC0 3814	RSI_RESPONSE0	“RSI Response Registers” on page 21-60
0xFFC0 3818	RSI_RESPONSE1	“RSI Response Registers” on page 21-60
0xFFC0 381C	RSI_RESPONSE2	“RSI Response Registers” on page 21-60
0xFFC0 3820	RSI_RESPONSE3	“RSI Response Registers” on page 21-60
0xFFC0 3824	RSI_DATA_TIMER	“RSI Data Timer Register” on page 21-61
0xFFC0 3828	RSI_DATA_LGTH	“RSI Data Length Register” on page 21-61
0xFFC0 382C	RSI_DATA_CONTROL	“RSI Data Control Register” on page 21-62
0xFFC0 3830	RSI_DATA_CNT	“RSI Data Counter Register” on page 21-64
0xFFC0 3834	RSI_STATUS	“RSI Status Register” on page 21-65
0xFFC0 3838	RSI_STATUSCL	“RSI Status Clear Register” on page 21-68
0xFFC0 383C	RSI_MASK0	“RSI Interrupt Mask Registers” on page 21-70
0xFFC0 3840	RSI_MASK1	“RSI Interrupt Mask Registers” on page 21-70
0xFFC0 3848	RSI_FIFO_CNT	“RSI FIFO Counter Register” on page 21-73
0xFFC0 384C	RSI_CEATA_CONTROL	“RSI CE_ATA Control Register” on page 21-74
0xFFC0 3880	RSI_FIFO	“RSI Data FIFO Register” on page 21-74
0xFFC0 38C0	RSI_ESTAT	“RSI Exception Status Register” on page 21-75
0xFFC0 38C4	RSI_EMASK	“RSI Exception Mask Register” on page 21-76
0xFFC0 38C8	RSI_CONFIG	“RSI Configuration Register” on page 21-78
0xFFC0 38CC	RSI_RD_WAIT_EN	“RSI Read Wait Enable Register” on page 21-79
0xFFC0 3FE0	RSI_PID0	“RSI Peripheral ID Registers” on page 21-80
0xFFC0 3FE4	RSI_PID1	“RSI Peripheral ID Registers” on page 21-80
0xFFC0 3FE8	RSI_PID2	“RSI Peripheral ID Registers” on page 21-80

TWI Registers

Table A-27. RSI Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0_3FEC	RSI_PID3	“RSI Peripheral ID Registers” on page 21-80
0xFFC0_3FF0	RSI_PID4	“RSI Peripheral ID Registers” on page 21-80
0xFFC0_3FF4	RSI_PID5	“RSI Peripheral ID Registers” on page 21-80
0xFFC0_3FF8	RSI_PID6	“RSI Peripheral ID Registers” on page 21-80
0xFFC0_3FFC	RSI_PID7	“RSI Peripheral ID Registers” on page 21-80

TWI Registers

Two Wire Interface (TWI) registers (0xFFC0_1400 – 0xFFC0_14FF) are listed in [Table A-28](#).

Table A-28. TWI Registers

Memory-Mapped Address	Register Name	Description
0xFFC0_1400	TWI_CLKDIV	“SCL Clock Divider Register” on page 16-26
0xFFC0_1404	TWI_CONTROL	“TWI Control Register” on page 16-25
0xFFC0_1408	TWI_SLAVE_CTL	“TWI Slave Mode Control Register” on page 16-26
0xFFC0_140C	TWI_SLAVE_STAT	“TWI Slave Mode Status Register” on page 16-28
0xFFC0_1410	TWI_SLAVE_ADDR	“TWI Slave Mode Address Register” on page 16-28
0xFFC0_1414	TWI_MASTER_CTL	“TWI Master Mode Control Register” on page 16-30
0xFFC0_1418	TWI_MASTER_STAT	“TWI Master Mode Status Register” on page 16-33
0xFFC0_141C	TWI_MASTER_ADDR	“TWI Master Mode Address Register” on page 16-32
0xFFC0_1420	TWI_INT_STAT	“TWI Interrupt Status Register” on page 16-40

Table A-28. TWI Registers (Continued)

Memory-Mapped Address	Register Name	Description
0xFFC0 1424	TWI_INT_MASK	“TWI Interrupt Mask Register” on page 16-40
0xFFC0 1428	TWI_FIFO_CTL	“TWI FIFO Control Register” on page 16-36
0xFFC0 142C	TWI_FIFO_STAT	“TWI FIFO Status Register” on page 16-38
0xFFC0 1480	TWI_XMT_DATA8	“TWI FIFO Transmit Data Single Byte Register” on page 16-43
0xFFC0 1484	TWI_XMT_DATA16	“TWI FIFO Transmit Data Double Byte Register” on page 16-44
0xFFC0 1488	TWI_RCV_DATA8	“TWI FIFO Receive Data Single Byte Register” on page 16-45
0xFFC0 148C	TWI_RCV_DATA16	“TWI FIFO Receive Data Double Byte Register” on page 16-46

TWI Registers

B TEST FEATURES

This appendix discusses the test features of the processor.

JTAG Standard

The processor is fully compatible with the IEEE 1149.1 standard, also known as the Joint Test Action Group (JTAG) standard.

The JTAG standard defines circuitry that may be built to assist in the test, maintenance, and support of assembled printed circuit boards. The circuitry includes a standard interface through which instructions and test data are communicated. A set of test features is defined, including a boundary-scan register, such that the component can respond to a minimum set of instructions designed to help test printed circuit boards.

The standard defines test logic that can be included in an integrated circuit to provide standardized approaches to:

- Testing the interconnections between integrated circuits once they have been assembled onto a printed circuit board
- Testing the integrated circuit itself
- Observing or modifying circuit activity during normal component operation

The test logic consists of a boundary-scan register and other building blocks. The test logic is accessed through a Test Access Port (TAP).

Boundary-Scan Architecture

Full details of the JTAG standard can be found in the document *IEEE Standard Test Access Port and Boundary-Scan Architecture*, ISBN 1-55937-350-4.

Boundary-Scan Architecture

The boundary-scan test logic consists of:

- A TAP comprised of five pins (see [Table B-1](#))
- A TAP controller that controls all sequencing of events through the test registers
- An instruction register (IR) that interprets 5-bit instruction codes to select the test mode that performs the desired test operation
- Several data registers defined by the JTAG standard

Table B-1. Test Access Port Pins

Pin Name	Input/Output	Description
TDI	Input	Test Data Input
TMS	Input	Test Mode Select
TCK	Input	Test Clock
$\overline{\text{TRST}}$	Input	Test Reset
TDO	Output	Test Data Out

The TAP controller is a synchronous, 16-state, finite-state machine controlled by the TCK and TMS pins. Transitions to the various states in the diagram occur on the rising edge of TCK and are defined by the state of the TMS pin, here denoted by either a logic 1 or logic 0 state. For full details of the operation, see the JTAG standard.

Figure B-1 shows the state diagram for the TAP controller.

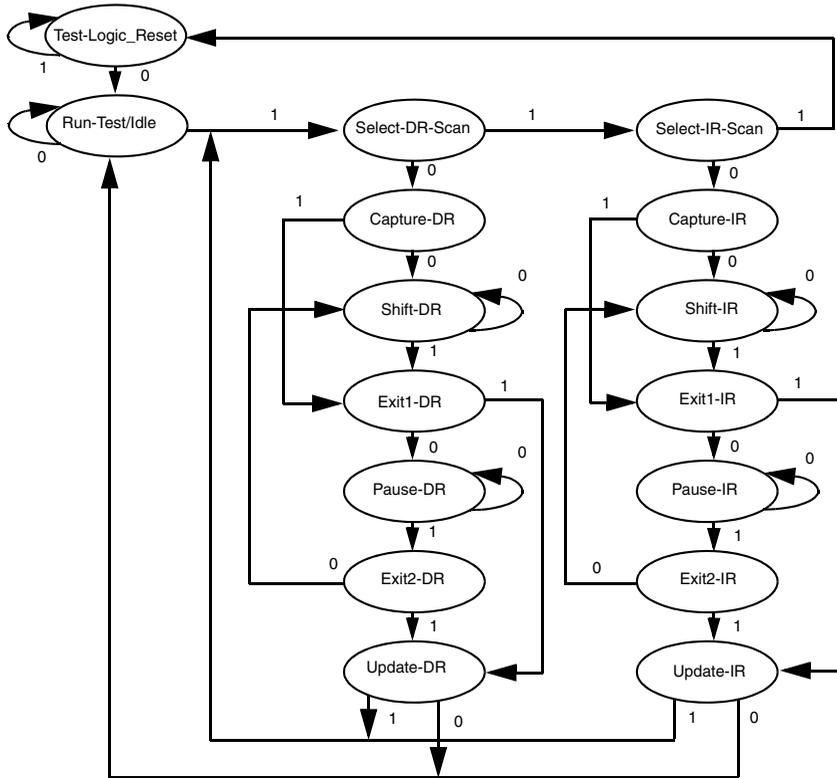


Figure B-1. TAP Controller State Diagram

Boundary-Scan Architecture

Note:

- The TAP controller enters the test-logic-reset state when TMS is held high after five TCK cycles.
- The TAP controller enters the test-logic-reset state when TRST is asynchronously asserted.
- An external system reset does not affect the state of the TAP controller, nor does the state of the TAP controller affect an external system reset.

Instruction Register

The instruction register is five bits wide and accommodates up to 32 boundary-scan instructions.

The instruction register holds both public and private instructions. The JTAG standard requires some of the public instructions; other public instructions are optional. Private instructions are reserved for the manufacturer's use.

The binary decode column of [Table B-2](#) lists the decode for the public instructions. The register column lists the serial scan paths.

Table B-2. Decode for Public JTAG-Scan Instructions

Instruction Name	Binary Decode 01234	Register
EXTEST	00000	Boundary-Scan
SAMPLE/PRELOAD	10000	Boundary-Scan
BYPASS	11111	Bypass

Figure B-2 shows the instruction bit scan ordering for the paths shown in Table B-2.

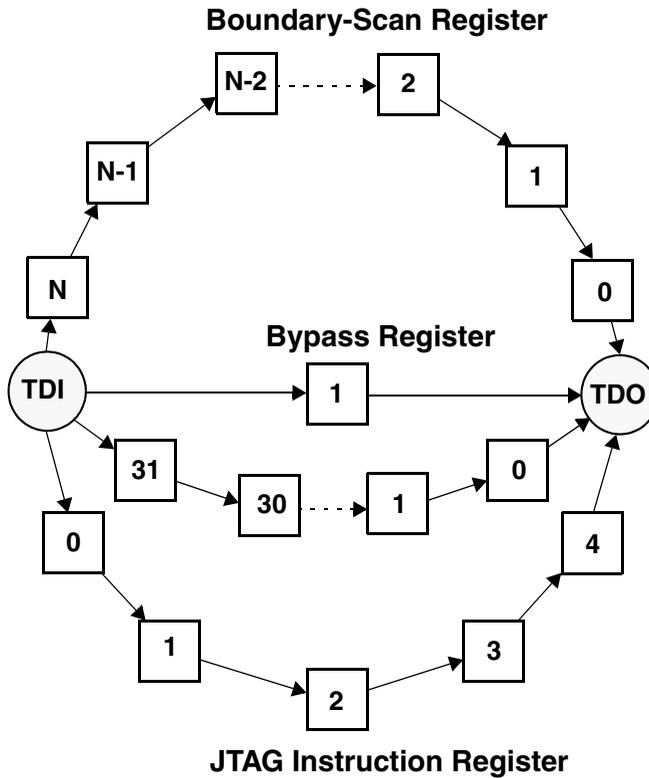


Figure B-2. Serial Scan Paths

Public Instructions

The following sections describe the public JTAG scan instructions.

EXTEST – Binary Code 00000

The EXTEST instruction selects the boundary-scan register to be connected between the TDI and TDO pins. This instruction allows testing of on-board circuitry external to the device.

The EXTEST instruction allows internal data to be driven to the boundary outputs and external data to be captured on the boundary inputs.

 To protect the internal logic when the boundary outputs are over-driven or signals are received on the boundary inputs, make sure that nothing else drives data on the processor's output pins.

SAMPLE/PRELOAD – Binary Code 10000

The SAMPLE/PRELOAD instruction performs two functions and selects the Boundary-Scan register to be connected between TDI and TDO. The instruction has no effect on internal logic.

The SAMPLE part of the instruction allows a snapshot of the inputs and outputs captured on the boundary-scan cells. Data is sampled on the rising edge of TCK.

The PRELOAD part of the instruction allows data to be loaded on the device pins and driven out on the board with the EXTEST instruction. Data is preloaded on the pins on the falling edge of TCK.

BYPASS – Binary Code 11111

The BYPASS instruction selects the BYPASS register to be connected to TDI and TDO. The instruction has no effect on the internal logic. No data inversion should occur between TDI and TDO.

Boundary-Scan Register

The boundary-scan register is selected by the `EXTEST` and `SAMPLE/PRELOAD` instructions. These instructions allow the pins of the processor to be controlled and sampled for board-level testing.

Boundary-Scan Architecture

I INDEX

Symbols

μ -law companding, [19-22](#), [19-27](#)

Numerics

16-bit flash interface, [27-6](#)
16-bit SRAM interface, [27-5](#), [27-6](#)
24 hours event flag bit, [24-21](#)
24 hours int en bit, [24-20](#)
2D DMA, [6-11](#)
2X input clock, [19-24](#)
5 volt tolerance, [27-11](#)
8-bit flash interface, [27-5](#)
8-bit SRAM interface, [27-5](#)

A

A[10] pin, [7-53](#)

ABE[1: 0] pins, [7-15](#)

aborted frames, MAC, [22-16](#)

aborts, DMA, [6-29](#)

accesses

 off-core, [4-4](#)

 to internal memory, [2-1](#)

access way/instruction address bit, [2-7](#)

ACCUMULATOR (accumulator value for local clock adjustment) bit, [23-24](#)

accumulator (EMAC_PTP_ACCR) register, [23-24](#)

accumulator value for local clock adjustment (ACCUMULATOR) bit, [23-24](#)

active descriptor queue, and DMA synchronization, [6-59](#)

active field select (FLD_SEL) bit, [20-26](#)

active low/high frame syncs, serial port, [19-32](#)

active mode, [1-29](#), [8-8](#)

ACTIVE_PLLDISABLED bit, [8-22](#)

ACTIVE_PLENABLED bit, [8-22](#)

active video only mode, PPI, [20-10](#)

ADCs, connecting to, [19-2](#)

ADDEND (addend value for local clock adjustment) bit, [23-23](#)

addend (EMAC_PTP_ADDEND) register, [23-23](#)

addend value for local clock adjustment (ADDEND) bit, [23-23](#)

address filter evaluation, MAC, [22-13](#)

address filter failed (RX_ADDR) bit, [22-96](#), [22-101](#)

address filter field (RX_ADDR) bit, [22-104](#)

address high[15: 0] field, [22-68](#)

address high (EMAC_ADDRHI) register, [22-68](#)

address low[15: 0] field, [22-67](#)

address low[31: 16] field, [22-67](#)

address low (EMAC_ADDRLO) register, [22-67](#)

address mapping, SDRAM, [7-26](#)

address not acknowledged (ANAK) bit, [16-33](#)

AFEXIT (authentication firmware exit) bit, [25-55](#)

AFVALID (authentication firmware valid) bit, [25-55](#)

Index

- AHAL_XOVR (channel A output crossover) bit, [14-44](#)
- AH_EN (AH output enable) bit, [14-44](#)
- AH output enable (AH_EN) bit, [14-44](#)
- alarm clock, RTC, [24-2](#)
- alarm event flag bit, [24-21](#)
- alarm int en (ALIE) bit, [23-15](#)
- alarm int en (RTC_AIE) bit, [24-20](#)
- alarm status (ALS) bit, [23-17](#)
- ALARM_TIME_HI (higher 32 bits of target alarm time) bit, [23-33](#)
- ALARM_TIME_LO (lower 32 bits of target alarm time) bit, [23-32](#)
- A-law companding, [19-22](#), [19-27](#)
- AL_EN (AL output enable) bit, [14-44](#)
- ALIE (alarm int en) bit, [23-15](#)
- alignment error (RX_ALIGN) bit, [22-97](#), [22-102](#), [22-104](#)
- alignment errors counter int en (RX_ALIGN_CNT) bit, [22-115](#)
- alignment errors counter int (RX_ALIGN_CNT) bit, [22-113](#)
- alignment errors (EMAC_RXC_ALIGN) register, [22-53](#)
- AL output enable (AL_EN) bit, [14-44](#)
- ALS (alarm status) bit, [23-17](#)
- alternate frame sync mode, [19-35](#)
- alternate timing, serial port, [19-34](#)
- AMBEN[2: 0] field, [7-17](#), [7-20](#)
- AMBEN (enable asynchronous memory banks) bit, [7-20](#)
- AMC, [1-9](#), [1-13](#)
 - bus contention, [7-10](#)
 - EBIU block diagram, [7-4](#)
 - features, [7-8](#)
 - signals, [7-9](#)
 - timing parameters, [7-18](#)
- AMCKEN bit, [7-17](#), [7-20](#)
- AMS, [7-8](#)
- ANAK (address not acknowledged) bit, [16-33](#), [16-35](#)
- application data, loading, [26-1](#)
- arbitration
 - DAB, [4-8](#), [4-9](#)
 - DCB, [4-8](#), [4-9](#)
 - DEB, [4-8](#), [4-9](#)
 - EAB, [4-11](#)
 - latency, [4-10](#)
 - TWI, [16-8](#)
- architecture, memory, [2-1](#)
- ARDY pin, [7-11](#), [7-15](#)
- array access bit, [2-7](#)
- ASEN (auxiliary snapshot control) bit, [23-13](#)
- ASIC/FPGA designs, [7-1](#)
- ASIE (auxiliary snapshot int en) bit, [23-15](#)
- ASL (auxiliary snapshot interrupt status) bit, [23-17](#)
- ASTP (enable automatic pad stripping) bit, [22-17](#), [22-60](#), [22-66](#)
- asynchronous
 - interfaces, supported, [7-1](#)
 - memory, [2-1](#), [7-2](#), [7-8](#)
 - memory bank address range (table), [7-8](#), [14-24](#)
 - memory controller, [1-13](#)
 - read, [7-12](#)
 - serial communications, [15-5](#)
 - write, [7-13](#)
- asynchronous memory bank control (EBIU_AMBCTLx) registers, [7-20](#), [7-21](#), [7-22](#)
- asynchronous memory banks enable (AMBEN) bit, [7-20](#)
- asynchronous memory controller. *See* AMC
- asynchronous memory global control (EBIU_AMGCTL) register, [7-17](#), [7-20](#)
- ASYNC memory banks, [7-3](#)

atomic operations, [27-3](#)
 authentication firmware exit (AFEXIT) bit,
[25-55](#)
 authentication firmware valid (AFVALID)
 bit, [25-55](#)
 autobaud, and general-purpose timers,
[10-31](#)
 autobaud detection, [15-13](#)
 autobuffer mode, [6-11](#), [6-28](#), [6-68](#)
 automatic pad stripping enable (ASTP) bit,
[22-66](#)
 auto-refresh
 command, [7-34](#), [7-51](#)
 timing, [7-58](#)
 auxiliary snapshot control (ASEN) bit,
[23-13](#)
 auxiliary snapshot int en (ASIE) bit, [23-15](#)
 auxiliary snapshot interrupt status (ASL)
 bit, [23-17](#)
 avoiding bus contention, [7-10](#)

B

BA[1: 0] pins, [7-32](#)
 backpressure enable (BKPRSEN) bit,
[22-75](#)
 bandwidth, and memory DMA operations,
[6-46](#)
 bank
 address, [7-62](#)
 size, [7-25](#), [7-62](#)
 size encodings (table), [7-27](#)
 width, [7-25](#)
 bank activate command, [7-33](#)
 bank activate command delay (TRAS) bit,
[7-65](#)
 bank precharge delay (TRP) bit, [7-65](#)
 baud rate
 SPI, [17-34](#)
 UART, [15-7](#), [15-13](#)
 baud rate[15: 0] field, [17-34](#)

BCINIT[15: 0] field, [6-84](#)
 BCODE[1
 0], [25-47](#), [25-48](#)
 BCODE[3
 0], [26-74](#)
 BCOUNT[15: 0] field, [6-85](#)
 BDI (block done interrupt generated) bit,
[6-83](#)
 BDIE (block done int en) bit, [6-40](#), [6-83](#)
 BFLAG_ALTERNATE bit, [26-91](#)
 BFLAG_AUX bit, [26-92](#)
 BFLAG_CALLBACK bit, [26-92](#)
 BFLAG_FASTREAD bit, [26-91](#)
 BFLAG_FILL bit, [26-92](#)
 BFLAG_FINAL bit, [26-92](#)
 BFLAG_FIRST bit, [26-92](#)
 BFLAG_HDRINDIRECT bit, [26-91](#)
 BFLAG_HOOK bit, [26-91](#)
 BFLAG_IGNORE bit, [26-92](#)
 BFLAG_INDIRECT bit, [26-92](#)
 BFLAG_INIT bit, [26-92](#)
 BFLAG_NEXTDXE bit, [26-91](#)
 BFLAG_NOAUTO bit, [26-91](#)
 BFLAG_NONRESTORE bit, [26-91](#)
 BFLAG_PERIPHERAL bit, [26-91](#)
 BFLAG_QUICKBOOT bit, [26-92](#)
 BFLAG_RESET bit, [26-91](#)
 BFLAG_RETURN bit, [26-91](#)
 BFLAG_SAVE bit, [26-92](#)
 BFLAG_SLAVE bit, [26-91](#)
 BFLAG_TYPE bit, [26-91](#)
 BFLAG_WAKEUP bit, [26-91](#)
 BHBL_XOVR (channel B output
 crossover) bit, [14-44](#)
 BH_EN (BH output enable) bit, [14-44](#)
 BH output enable (BH_EN) bit, [14-44](#)
 BI (break int) bit, [15-24](#), [15-25](#)
 binary decode, [B-4](#)
 bin x bit, [22-69](#), [22-70](#)
 bit order, selecting, [19-26](#)

Index

- bit order select (TLSBIT) bit, [19-47](#)
- bit rate generation, [15-12](#)
- BK_DATECODE (boot code date code) register, [26-76](#)
- BK_DAY bit, [26-76](#)
- BK_ID (boot kernel identifier) bit, [26-75](#)
- BK_MONTH bit, [26-76](#)
- BK_ONES (ones word) register, [26-78](#)
- BK_PROJECT (boot kernel project) bit, [26-75](#)
- BKPRSEN (enable backpressure) bit, [22-75](#)
- BK_REVISION (boot code revision control) register, [26-75](#)
- BK_UPDATE (boot kernel update) bit, [26-75](#)
- BK_VERSION (boot kernel version) bit, [26-75](#)
- BK_YEAR bit, [26-76](#)
- BK_ZEROS (zero word) register, [26-77](#)
- Blackfin processor family
 - I/O memory space, [1-9](#)
 - memory architecture, [1-7](#)
- BL_EN (BL output enable) bit, [14-44](#)
- block, DMA, [6-9](#)
- block code field, [26-22](#)
- block count, DMA, [6-37](#)
- block diagrams
 - bus hierarchy, [4-3](#)
 - core, [4-4](#)
 - core timer, [11-2](#)
 - DMA controller, [6-103](#)
 - EBIU, [7-4](#)
 - general-purpose timers, [10-57](#)
 - MAC, [22-3](#)
 - PLL, [8-3](#)
 - PPI, [20-3](#)
 - processor, [1-7](#)
 - RTC, [24-4](#)
 - SDRAM, [7-55](#), [7-56](#)
 - SPI, [17-3](#), [17-4](#)
 - SPORT, [19-6](#)
 - TWI, [16-3](#)
 - UART, [15-3](#)
 - watchdog timer, [12-3](#)
- block done int en (BDIE) bit, [6-83](#)
- block done interrupt, DMA, [6-40](#)
- block done interrupt generated (BDI) bit, [6-83](#)
- block flags, [26-25](#)
- block transfers, DMA, [6-37](#)
- BL output enable (BL_EN) bit, [14-44](#)
- BMODE[2:0], [26-74](#)
- BMODE[2:0] pins, [26-1](#), [26-5](#)
- BMODE[3:0], [25-48](#)
- BMODEx_DIS (boot mode disable) bits, [26-84](#)
- BNDMODE (boundary register mode) bits, [13-20](#)
- BOLMT[1:0] field, [22-60](#), [22-63](#)

- boot
 - call boot kernel at run time, [26-46](#)
 - load function, [26-45](#)
 - manager, [26-49](#)
 - quick, [26-39](#)
 - ROM functions, [26-50](#)
 - streams
 - multi-DXE, [26-50](#)
- boot code date code (BK_DATECODE) register, [26-76](#)
- boot code revision control (BK_REVISION) register, [26-75](#)
- boot host wait (HWAIT) pin, [26-29](#)
- booting
 - BFROM_MEMBOOT, [26-50](#)
 - BFROM_OTPBOOT, [26-50](#)
 - BFROM_SPIBOOT, [26-50](#)
 - boot stream, [26-19](#)
 - host boot scenarios, [26-19](#)
 - indirect, [26-40](#)
 - memory locations, [26-21](#)
 - SPI slave mode, [26-65](#)
- booting modes, [26-2](#)
- booting registers, [26-71](#)
- boot kernel, [26-1](#)
- boot kernel identifier (BK_ID) bit, [26-75](#)
- boot kernel project (BK_PROJECT) bit, [26-75](#)
- boot kernel update (BK_UPDATE) bit, [26-75](#)
- boot kernel version (BK_VERSION) bit, [26-75](#)
- boot management, [26-49](#)
- boot mode
 - flash boot, [26-57](#)
 - no-boot, [26-57](#)
 - SDRAM boot, [26-60](#)
 - SPI device detection, [26-63](#)
- boot ROM
 - internal, [26-1](#)
 - memory space, [2-4](#)
- boot stream, [26-1, 26-19](#)
- boot termination, [26-31](#)
- boundary register mode (BNDMODE) bits, [13-20](#)
- boundary-scan architecture, [B-2](#)
- boundary-scan register, [B-7](#)
- break int (BI) bit, [15-24](#)
- broadcast frames detected (RX_BROAD) bit, [22-95, 22-101, 22-104](#)
- broadcast frames received OK counter int en (RX_BROAD_CNT) bit, [22-115](#)
- broadcast frames received OK counter int (RX_BROAD_CNT) bit, [22-113](#)
- broadcast frames received ok (EMAC_RXC_BROAD) register, [22-54](#)
- broadcast frames transmitted OK counter int en (TX_BROAD_CNT) bit, [22-118](#)
- broadcast frames transmitted OK counter int (TX_BROAD_CNT) bit, [22-116](#)
- broadcast frames xmitted ok (EMAC_TXC_BROAD) register, [22-58](#)
- broadcast mode, [17-8, 17-14, 17-15](#)
- buffer read error (BUFRDERR) bit, [16-33](#)
- buffer registers, timers, [10-42](#)
- buffer write error (BUFWRERR) bit, [16-33](#)
- BUFRDERR (buffer read error) bit, [16-33, 16-35](#)
- BUFWRERR (buffer write error) bit, [16-33, 16-35](#)
- burst
 - length, [7-32](#)
 - type, [7-32](#)

Index

- bus agents
 - DAB, [4-9](#)
 - PAB, [4-6](#)
 - BUSBUSY (bus busy) bit, [16-33](#)
 - bus busy (BUSBUSY) bit, [16-33](#)
 - bus contention, avoiding, [7-10](#), [27-7](#)
 - bus cycles
 - asynchronous read, [7-12](#)
 - asynchronous write, [7-13](#)
 - bus error, EBIU, [7-7](#)
 - buses
 - See also* DAB, DCB, DEB, EAB, EPB, PAB
 - bandwidth, [1-6](#)
 - core, [4-4](#)
 - hierarchy, [4-3](#)
 - on-chip, [4-1](#)
 - PAB, [4-6](#)
 - peripheral, [4-6](#)
 - and peripherals, [1-6](#)
 - prioritization and DMA, [6-48](#)
 - BUS_MODE (data bus width) bit, [21-55](#)
 - bus standard, I²C, [1-14](#)
 - BxHT fields, [7-21](#), [7-22](#)
 - BxRAT fields, [7-21](#), [7-22](#)
 - BxRDYEN bits, [7-21](#), [7-22](#)
 - BxRDYPOL bits, [7-21](#), [7-22](#)
 - BxST fields, [7-21](#), [7-22](#)
 - BxTT fields, [7-21](#), [7-22](#)
 - BxWAT fields, [7-21](#), [7-22](#)
 - bypass, capacitor placement, [27-11](#)
 - BYPASS bit, [8-21](#)
 - bypass clock divisor (CLKDIV_BYPASS) bit, [21-55](#)
 - BYPASS instruction, [B-6](#)
 - BYPASS register, [B-6](#)
 - byte
 - address, [7-62](#)
 - enables, [7-15](#)
 - byte enable x bit, [22-81](#), [22-82](#), [22-83](#), [22-84](#)
- ## C
- callback routines, [26-41](#)
 - capacitors, [27-9](#)
 - capture mode. *See* WDTH_CAP mode
 - capture wakeup frames (CAPWKFRM) bit, [22-79](#)
 - CAPWKFRM (capture wakeup frames) bit, [22-79](#)
 - card detect int (CEATA_INT_DET) bit, [21-75](#)
 - card detect int en (SD_CARD_DET_MASK) bit, [21-76](#)
 - card detect int (SD_CARD_DET) bit, [21-75](#)
 - carrier sense errors counter int en (TX_CRS_CNT) bit, [22-118](#)
 - carrier sense errors counter int (TX_CRS_CNT) bit, [22-116](#)
 - carrier sense errors (EMAC_TXC_CRSEERR) register, [22-57](#)
 - CAS latency, [7-32](#), [7-35](#)
 - CAW (column address width), [7-62](#)
 - CCIR-656. *See* ITU-R 656
 - CCITT G.711 specification, [19-27](#)
 - CCLK (core clock), [4-4](#), [8-5](#), [8-7](#)
 - CCOR (counter clear-on-read mode) bit, [22-42](#), [22-119](#)
 - CDGINV (CDG pin polarity invert) bit, [13-20](#)
 - CDG pin polarity invert (CDGINV) bit, [13-20](#)
 - CDPRIO (core priority) bit, [4-8](#), [7-18](#), [7-20](#), [7-43](#)
 - CEATA_CCS_EN (command completion signal enable) bit, [21-62](#)

- CEATA_EN (CE-ATA mode enable) bit, [21-62](#)
- CEATA_INT_DET (card detect int) bit, [21-75](#)
- CEATA_INT_DET_MASK (command completion signal detect enable) bit, [21-76](#)
- CE-ATA mode enable (CEATA_EN) bit, [21-62](#)
- CEATA_TX_CCSD (transmit command completion signal disable) bit, [21-74](#)
- channel A duty two's complement (PWMCHA) bit, [14-42](#)
- channel A duty two's complement (PWM_CHAL) bit, [14-47](#)
- channel A output crossover (AHAL_XOVR) bit, [14-44](#)
- channel B duty two's complement (PWMCHB) bit, [14-43](#)
- channel B duty two's complement (PWM_CHBL) bit, [14-47](#)
- channel B output crossover (BHBL_XOVR) bit, [14-44](#)
- channel C duty two's complement (PWMCHC) bit, [14-43](#)
- channel C duty two's complement (PWM_CHCL) bit, [14-47](#)
- channel C output crossover (CHCL_XOVR) bit, [14-44](#)
- channel running (DMA_RUN) bit, [6-72](#)
- channels
 - defined, serial, [19-21](#)
 - serial port TDM, [19-21](#)
 - serial select offset, [19-21](#)
- CHCL_XOVR (channel C output crossover) bit, [14-44](#)
- CH_EN (CH output enable) bit, [14-44](#)
- CHNL[9: 0] field, [19-64](#), [19-65](#)
- CH output enable (CH_EN) bit, [14-44](#)
- circuit board testing, [B-1](#), [B-6](#)
- circular addressing, [6-56](#)
- CKDIV (divider for selected PTP_CLK output) bit, [23-13](#)
- CKOEN (clock output control) bit, [23-13](#)
- CKS (PTP) bit, [23-13](#)
- CKS (PTP clock source) bit, [23-13](#)
- CL[1: 0] field, [7-65](#), [7-66](#)
- CL (CAS latency), [7-35](#)
- clear command response CRC fail (CMD_CRC_FAIL_STAT) bit, [21-68](#)
- clear command response received (CMD_RESP_END_STAT) bit, [21-68](#)
- clear command sent (CMD_SENT_STAT) bit, [21-68](#)
- clear command timeout (CMD_TIMEOUT_STAT) bit, [21-68](#)
- clear data block end (DAT_BLK_END_STAT) bit, [21-68](#)
- clear data CRC failure (DAT_CRC_FAIL_STAT) bit, [21-68](#)
- clear data timeout (DAT_TIMEOUT_STAT) bit, [21-68](#)
- clear end of data (DAT_END_STAT) bit, [21-68](#)
- clearing interrupt requests, [5-13](#)
- clear Pxn bit, [9-35](#)
- clear Pxn interrupt A enable bit, [9-41](#)
- clear Pxn interrupt B enable bit, [9-42](#)
- clear receive FIFO overrun error (RX_OVERRUN_STAT) bit, [21-68](#)
- clear start bit error (START_BIT_ERR_STAT) bit, [21-68](#)

Index

- clear transmit FIFO underrun error (TX_UNDERRUN_STAT) bit, 21-68
- CL_EN (CL output enable) bit, 14-44
- CLKBUFOE (CLKIN buffer output enable) bit, 8-23
- CLKBUF pin, 22-4, 22-46
- CLKDIV_BYPASS (bypass clock divisor) bit, 21-55
- CLKDIV (clock divisor) bit, 21-55
- CLK_EN (RSI_CLOCK enable) bit, 21-55
- CLKHI[7: 0] field, 16-26
- CLKHI bit, 16-26
- CLKIN buffer output enable (CLKBUFOE) bit, 8-23
- CLKIN (input clock), 1-28, 4-4, 8-1, 8-2
- CLKLOW[7: 0] field, 16-26
- CLKLOW bit, 16-26
- CLKOUT pin, 7-7, 7-66, 7-71
- CLK_SEL (timer clock select) bit, 10-12, 10-19, 10-40, 10-45
- clock
 - clock signals, 1-28
 - control, 8-1
 - EBIU, 7-1
 - external, 1-28
 - frequency for SPORT, 19-61
 - internal, 4-4
 - MAC, 22-4
 - managing, 27-1
 - peripheral, 8-7
 - RTC, 24-3
 - source for general-purpose timers, 10-3
 - SPI clock signal, 17-4
 - system, 1-28
 - system (SCLK), 27-2
 - types, 27-1
- clock divide modulus registers, 19-61
- clock divisor (CLKDIV) bit, 21-55
- clock domain synchronization, PPI, 20-15
- clock drive/sample edge select (TCKFE) bit, 19-47
- clock falling edge select (RCKFE) bit, 19-52
- clock input (CLKIN) pin, 27-1
- clock output control (CKOEN) bit, 23-13
- clock phase, SPI, 17-11, 17-13
- clock phase (CPHA) bit, 17-36
- clock polarity, SPI, 17-11
- clock polarity (CPOL) bit, 17-36
- clock rate
 - core timer, 11-1
 - SPORT, 19-2
- clock ratio, changing, 8-6
- clocks, overview, 1-28
- clock signals, 1-28
- CL output enable (CL_EN) bit, 14-44
- CMD_ACT (command active) bit, 21-65
- CMD_ACT_MASK (command active) bit, 21-70
- CMD_CRC_FAIL (command response CRC fail) bit, 21-65
- CMD_CRC_FAIL_MASK (command response CRC fail) bit, 21-70
- CMD_CRC_FAIL_STAT (clear command response CRC fail) bit, 21-68
- CMD_EN (command enable) bit, 21-57
- CMD_IDX (command index) bit, 21-57
- CMD_INT_EN (command int en) bit, 21-57
- CMD_LRSP_EN (long response enable) bit, 21-57
- CMD_PEND_EN (pend enable) bit, 21-57
- CMD_RESP_END (command response received) bit, 21-65
- CMD_RESP_END_MASK (command response received) bit, 21-70

- CMD_RESP_END_STAT (clear command response received) bit, [21-68](#)
- CMD_RSP_EN (wait for response) bit, [21-57](#)
- CMD_SENT (command sent) bit, [21-65](#)
- CMD_SENT_MASK (command sent) bit, [21-70](#)
- CMD_SENT_STAT (clear command sent) bit, [21-68](#)
- CMD_TIMEOUT (command timeout) bit, [21-65](#)
- CMD_TIMEOUT_MASK (command timeout) bit, [21-70](#)
- CMD_TIMEOUT_STAT (clear command timeout) bit, [21-68](#)
- CNT_COMMAND (counter command) register, [13-19](#), [13-22](#), [13-24](#)
- CNT_CONFIG (counter configuration) register, [13-19](#), [13-20](#)
- CNT_COUNTER (counter) register, [13-19](#)
- CNT_DEBOUNCE (counter debounce) register, [13-19](#), [13-25](#)
- CNTE (counter enable) bit, [13-20](#)
- CNT_IMASK (counter interrupt mask) register, [13-21](#)
- CNT_IMASK (interrupt mask) register, [13-19](#)
- CNT_MAX (maximal count) register, [13-19](#), [13-26](#)
- CNT_MIN (minimal count) register, [13-19](#), [13-26](#)
- CNTMODE (counter operating mode) bits, [13-20](#)
- CNT_STATUS (counter status) register, [13-19](#), [13-22](#)
- codecs, connecting to, [19-2](#)
- collision error detected, excessive (TX_ECOLL) bit, [22-105](#)
- collision error int en, excessive (TX_ECOLL) bit, [22-112](#)
- collision errors detected, excessive (TX_ECOLL) bit, [22-109](#)
- column address, [7-62](#)
- strobe latency, [7-32](#)
- column read/write, SDRAM, [7-31](#)
- command active (CMD_ACT) bit, [21-65](#)
- command active (CMD_ACT_MASK) bit, [21-70](#)
- command (CNT_COMMAND) register, [13-19](#), [13-22](#)
- command completion signal detect enable (CEATA_INT_DET_MASK) bit, [21-76](#)
- command completion signal enable (CEATA_CCS_EN) bit, [21-62](#)
- command enable (CMD_EN) bit, [21-57](#)
- command index (CMD_IDX) bit, [21-57](#)
- command index of last received response (RESP_CMD) bit, [21-59](#)
- command inhibit command, [7-53](#)
- command int en (CMD_INT_EN) bit, [21-57](#)
- command response CRC fail (CMD_CRC_FAIL) bit, [21-65](#)
- command response CRC fail (CMD_CRC_FAIL_MASK) bit, [21-70](#)
- command response received (CMD_RESP_END) bit, [21-65](#)
- command response received (CMD_RESP_END_MASK) bit, [21-70](#)

Index

- commands
 - auto-refresh, [7-34](#), [7-51](#), [7-58](#)
 - bank activate, [7-33](#)
 - command inhibit, [7-53](#)
 - DMA control, [6-31](#), [6-32](#)
 - EMRS, [7-48](#)
 - ERMS, [7-33](#)
 - MRS, [7-33](#), [7-47](#)
 - no operation (NOP), [7-53](#)
 - precharge, [7-34](#)
 - precharge all, [7-34](#), [7-50](#)
 - read, [7-33](#)
 - read/write, [7-49](#)
 - SDC, [7-45](#)
 - self-refresh, [7-51](#)
 - transfer initiate, [17-17](#), [17-18](#)
 - write, [7-33](#)
 - write with data mask, [7-49](#)
- command sent, [21-70](#)
- command sent (CMD_SENT) bit, [21-65](#)
- command sent (CMD_SENT_MASK) bit, [21-70](#)
- command timeout (CMD_TIMEOUT) bit, [21-65](#)
- command timeout (CMD_TIMEOUT_MASK) bit, [21-70](#)
- companding, [19-15](#), [19-22](#)
 - A-law, [19-22](#), [19-27](#)
 - defined, [19-27](#)
 - lengths supported, [19-28](#)
 - multichannel operations, [19-22](#)
 - μ -law, [19-22](#), [19-27](#)
- compare mask of frame type field (EFTM) bit, [23-13](#)
- compare mask of IP frame type field (IPTM) bit, [23-13](#)
- compare mask of IP version field (IPVM) bit, [23-13](#)
- compare mask of PTP control field (PTPCM) bit, [23-13](#)
- compare mask of UDP event port field (UDPEM) bit, [23-13](#)
- completion interrupt status (DMA_DONE) bit, [6-72](#)
- configuration
 - SDC, [7-54](#)
 - SDRAM, [7-24](#)
 - SPORT, [19-10](#)
- configuration (CNT_CONFIG) register, [13-19](#)
- congestion, on DMA channels, [6-45](#)
- contention, bus, avoiding, [7-10](#)
- continuous polling, and MMC register, [22-43](#)
- continuous transition, DMA, [6-27](#)
- control bit summary, general-purpose timers, [10-44](#)
- control byte sequences, PPI, [20-8](#)
- control (EMAC_PTP_CTL) register, [23-13](#)
- control frame int en (RX_CTL) bit, [22-100](#), [22-104](#)
- control frames, MAC, [22-16](#)
- control frames transmitted counter int en (TX_MACCTL_CNT) bit, [22-118](#)
- control frames transmitted counter int (TX_MACCTL_CNT) bit, [22-116](#)
- control register
 - data memory, [2-6](#)
 - EBIU, [7-5](#)
- core
 - block diagram, [4-4](#)
 - core bus, [4-4](#)
 - core clock (CCLK), [8-5](#), [27-2](#)
 - core clock/system clock ratio control, [8-5](#)
 - timer, [5-5](#)
 - waking from idle state, [5-6](#)

- core and system reset, code example, [26-102](#)
- core clock (CCLK), [11-2](#)
- core clock. *See* CCLK
- core double-fault reset, [26-5](#)
- core-double-fault reset enable (DOUBLE_FAULT) bit, [26-72](#)
- core-double-fault reset (RESET_DOUBLE) bit, [26-72](#)
- core event controller (CEC), [5-2](#)
- core-only software reset, [26-5](#)
- core priority (CDPRIO) bit, [7-20](#)
- core select (CSEL) bit, [8-21](#)
- core select (CSEL) bits, [8-21](#)
- core time control (TCNTL) register, [11-5](#)
- core timer, [11-1](#) to [11-8](#)
 - block diagram, [11-2](#)
 - clock rate, [11-1](#)
 - features, [11-2](#)
 - initialization, [11-3](#)
 - internal interfaces, [11-3](#)
 - low power mode, [11-3](#)
 - operation, [11-3](#)
 - registers, [11-4](#)
 - scaling, [11-7](#)
- core timer control (TCNTL) register, [11-3](#), [11-5](#)
- core timer count (TCOUNT) register, [11-3](#), [11-5](#), [11-6](#)
- core timer period (TPERIOD) register, [11-6](#)
- core timer scale (TSCALE) register, [11-3](#), [11-7](#)
- counter, RTC, [24-2](#)
- counter clear-on-read mode (CCOR) bit, [22-119](#)
- counter (CNT_COUNTER) register, [13-19](#)
- counter command (CNT_COMMAND) register, [13-24](#)
- counter configuration (CNT_CONFIG) register, [13-20](#)
- counter count value (CNT_COUNTER) register, [13-26](#)
- counter debounce (CNT_DEBOUNCE) register, [13-25](#)
- counter enable (CNTE) bit, [13-20](#)
- counter interrupt mask (CNT_IMASK) register, [13-21](#)
- counter max (CNT_MAX) bit, [13-27](#)
- counter maximal count (CNT_MAX) register, [13-27](#)
- counter min (CNT_MIN) bit, [13-27](#)
- counter minimal count (CNT_MIN) register, [13-27](#)
- counter module registers, [13-19](#)
- counter operating mode (CNTMODE) bits, [13-20](#)
- counter overflow int en (COVxIE) bits, [13-21](#)
- counter overflow interrupt identifier (COVxII) bits, [13-22](#)
- counter rollover enable (CROLL) bit, [22-119](#)
- counter status (CNT_STATUS) register, [13-22](#)
- counter zeroed by zero marker (CZMZIE) bit, [13-22](#)
- counter zeroed int en (CZMZIE) bit, [13-21](#)
- count to zero int en (CZEROIE) bit, [13-21](#)
- count to zero interrupt identifier (CZEROII) bit, [13-22](#)
- count value[15: 0] field, [11-6](#)
- count value[31: 16] field, [11-6](#)
- COVxIE (counter overflow int en) bits, [13-21](#)
- COVxII (counter overflow interrupt identifier) bits, [13-22](#)
- CPHA (clock phase) bit, [17-36](#)

Index

- CPOL (clock polarity) bit, [17-36](#)
- CRC-16 hash value calculation, [22-37](#)
- CRC-32 calculation, MAC, [22-71](#)
- CRC32 checksum generation, [26-44](#)
- CRC state, MAC, [22-36](#)
- CROLL (counter rollover enable) bit, [22-43](#), [22-119](#)
- CROSSCORE software, [1-31](#)
- crosstalk, [27-9](#)
- crystal, external, [1-28](#)
- CSEL[1: 0] field, [8-5](#), [8-21](#), [27-2](#)
- CSEL (core select) bit, [8-21](#)
- CTYPE (DMA channel type) bit, [6-66](#)
- CUD and CDZ input disable (INPDIS) bit, [13-20](#)
- CUDINV (CUD pin polarity invert) bit, [13-20](#)
- CUD pin polarity invert (CUDINV) bit, [13-20](#)
- current address (DMA_x_CURR_ADDR) registers, [6-74](#)
- current address field, [6-75](#)
- current address (MDMA_{yy}_CURR_ADDR) registers, [6-74](#)
- current descriptor pointer (DMA_x_CURR_DESC_PTR) registers, [6-81](#)
- current descriptor pointer (MDMA_{yy}_CURR_DESC_PTR) registers, [6-81](#)
- current frame status (EMAC_RX_STAT) register, [22-93](#)
- current frame status (EMAC_TX_STAT) register, [22-105](#)
- current inner loop count (DMA_x_CURR_X_COUNT) registers, [6-76](#)
- current inner loop count (MDMA_{yy}_CURR_X_COUNT) registers, [6-76](#)
- current outer loop count (DMA_x_CURR_Y_COUNT) registers, [6-78](#)
- current outer loop count (MDMA_{yy}_CURR_Y_COUNT) registers, [6-78](#)
- current time, [24-12](#)
- CURR_X_COUNT[15: 0] field, [6-76](#)
- CURR_Y_COUNT[15: 0] field, [6-79](#)
- customer support, [-lxii](#)
- CZEROIE (count to zero int en) bit, [13-21](#)
- CZEROII (count to zero interrupt identifier) bit, [13-22](#)
- CZMEIE (zero marker error int en) bit, [13-21](#)
- CZMEII (CZM error interrupt identifier) bit, [13-22](#)
- CZM error interrupt identifier (CZMEII) bit, [13-22](#)
- CZMIE (CZM pin int en) bit, [13-21](#)
- CZMII (CZM pin interrupt identifier) bit, [13-22](#)
- CZMINV (CZM pin polarity invert) bit, [13-20](#)
- CZM pin int en (CZMIE) bit, [13-21](#)
- CZM pin interrupt identifier (CZMII) bit, [13-22](#)
- CZM pin polarity invert (CZMINV) bit, [13-20](#)
- CZM zeroes counter enable (ZMZC) bit, [13-20](#)
- CZMZIE (counter zeroed by zero marker) bit, [13-22](#)
- CZMZIE (counter zeroed int en) bit, [13-21](#)

D

DAB, [4-8](#), [6-5](#), [6-41](#), [6-90](#)

arbitration, [4-8](#), [4-9](#)

bus agents (masters), [4-9](#)

clocking, [8-2](#)

latencies, [4-10](#)

performance, [4-10](#)

throughput, [4-10](#)

DAB_TRAFFIC_COUNT[2: 0] field, [6-90](#)

DAB_TRAFFIC_PERIOD bit, [6-89](#)

DAG0 port preference (PORT_PREF0) bit, [2-6](#)

DAG1 port preference (PORT_PREF1) bit, [2-6](#)

data

I/O mask function, [7-33](#)

masks, [7-40](#)

sampling, serial, [19-32](#)

data bank access bit, [2-7](#)

DATA_BLK_LGTH (data block length) bit, [21-62](#)

data block end (DAT_BLK_END) bit, [21-65](#)

data block end (DAT_BLK_END_MASK) bit, [21-70](#)

data block length (DATA_BLK_LGTH) bit, [21-62](#)

data bus width (BUS_MODE) bit, [21-55](#)

data cacheability protection lookaside

buffer enable (ENDCPLB) bit, [2-6](#)

data cache select/address bit, [2-7](#)

data corruption, avoiding with SPI, [17-14](#)

DATA_COUNT (number of bytes still to be transferred) bit, [21-64](#)

data CRC failure (DAT_CRC_FAIL) bit, [21-65](#)

data CRC failure

(DAT_CRC_FAIL_MASK) bit, [21-70](#)

DATA_DIR (data transfer direction) bit, [21-62](#)

data DMA enable (DATA_DMA_EN) bit, [21-62](#)

DATA_DMA_EN (data DMA enable) bit, [21-62](#)

data-driven interrupts, [6-72](#)

data enable (DATA_EN) bit, [21-62](#)

DATA_EN (data enable) bit, [21-62](#)

data formats, SPORT, [19-27](#)

data-independent transmit frame sync select (DITFS) bit, [19-47](#)

data input modes for PPI, [20-14](#) to [20-16](#)

data/instruction access bit, [2-7](#)

data int en (DI_EN) bit, [6-67](#)

data interrupt timing select (DI_SEL) bit, [6-67](#)

data length (DLEN) bit, [20-26](#)

DATA_LENGTH (number of bytes to be transferred) bit, [21-61](#)

data memory control

(DMEM_CONTROL) register, [2-6](#)

data memory control register

(DMEM_CONTROL), [2-6](#)

DATA_MODE (data transfer mode) bit, [21-62](#)

data move, serial port operations, [19-37](#)

data not acknowledged (DPAK) bit, [16-33](#)

data output modes for PPI, [20-17](#) to [20-18](#)

data ready (DR) bit, [15-24](#)

data receive active (RX_ACT) bit, [21-65](#)

data receive active (RX_ACT_MASK) bit, [21-70](#)

data structures, [26-86](#)

boot_struct, [26-88](#)

buffer_struct, [26-87](#)

header_struct, [26-86](#)

data test command

(DTEST_COMMAND) register, [2-7](#)

Index

- data test command register
 - (DTEST_COMMAND), [2-6](#)
- data timeout (DAT_TIMEOUT) bit,
 - [21-65](#)
- data timeout (DAT_TIMEOUT_MASK) bit,
 - [21-70](#)
- data transfer count (DCNT) bit, [16-30](#)
- data transfer (DATA_MODE) bit, [21-62](#)
- data transfer direction (DATA_DIR) bit,
 - [21-62](#)
- data transfers
 - DMA, [4-10](#), [6-2](#)
 - SPI, [17-14](#)
- data transmit active (TX_ACT) bit, [21-65](#)
- data transmit active (TX_ACT_MASK) bit,
 - [21-70](#)
- data word, serial data formats, [19-55](#)
- DAT_BLK_END (data block end) bit,
 - [21-65](#)
- DAT_BLK_END_MASK (data block end) bit,
 - [21-70](#)
- DAT_BLK_END_STAT (clear data block end) bit, [21-68](#)
- DAT_CRC_FAIL (data CRC failure) bit,
 - [21-65](#)
- DAT_CRC_FAIL_MASK (data CRC failure) bit, [21-70](#)
- DAT_CRC_FAIL_STAT (clear data CRC failure) bit, [21-68](#)
- DAT_END (end of data) bit, [21-65](#)
- DAT_END_MASK (end of data) bit,
 - [21-70](#)
- DAT_END_STAT (clear end of data) bit,
 - [21-68](#)
- DAT_TIMEOUT (data timeout) bit,
 - [21-65](#)
- DAT_TIMEOUT_MASK (data timeout) bit,
 - [21-70](#)
- DAT_TIMEOUT_STAT (clear data timeout) bit, [21-68](#)
- day[14: 0] field, [24-22](#)
- day alarm event flag bit, [24-21](#)
- day alarm int en bit, [24-20](#)
- day counter[14: 0] field, [24-20](#)
- DBF (disable broadcast frame reception) bit,
 - [22-60](#), [22-65](#)
- DCB, [4-8](#), [6-5](#), [6-41](#), [6-90](#)
 - arbitration, [4-8](#), [4-9](#)
- DCBS (L1 data cache bank select) bit, [2-6](#)
- DCB_TRAFFIC_COUNT field, [6-90](#)
- DCB_TRAFFIC_PERIOD field, [6-90](#)
- DC (deferral check) bit, [22-60](#), [22-63](#)
- DCIE (down count int en) bit, [13-21](#)
- DCII (down count interrupt identifier) bit,
 - [13-22](#)
- DCNT[7: 0] field, [16-30](#), [16-31](#)
- DCNT (data transfer count) bit, [16-30](#)
- DEB, [4-8](#), [6-5](#), [6-41](#), [6-90](#)
 - arbitration, [4-8](#), [4-9](#)
 - and EBIU, [7-4](#)
 - frequency, [4-11](#)
 - performance, [4-11](#)
- DEBE (debounce enable) bit, [13-20](#)
- debounce (CNT_DEBOUNCE) register,
 - [13-19](#)
- debounce delay (DPRESCALE) bit, [13-25](#)
- debounce enable (DEBE) bit, [13-20](#)
- DEB_TRAFFIC_COUNT field, [6-89](#), [6-90](#)
- DEB_TRAFFIC_PERIOD field, [6-90](#)
- debugging, [1-32](#)
 - test point access, [27-12](#)
- deep sleep, and RTC, [24-9](#)
- deep sleep mode, [1-29](#), [7-45](#), [8-9](#)
- deferral check (DC) bit, [22-63](#)
- deferral int en, excessive (TX_EDEFER) bit, [22-112](#)
- deferrals detected, excessive (TX_EDEFER) bit, [22-105](#), [22-109](#)
- delay count (PPI_DELAY) register, [20-32](#)

- descriptor
 - array mode, DMA, [6-15](#), [6-68](#)
 - chains, DMA, [6-26](#)
 - list mode, DMA, [6-14](#), [6-68](#), [6-69](#)
 - pairs, DMA, [22-12](#)
- descriptor-based DMA, [6-13](#)
- descriptor fetch (DFETCH) bit, [6-72](#)
- descriptor queue, [6-57](#)
 - management, [6-56](#)
 - synchronization, [6-57](#)
- descriptor structures
 - alternative, [22-25](#)
 - DMA, [6-55](#)
 - MDMA, [6-62](#)
- destination channels, memory DMA, [6-7](#)
- development tools, [1-31](#)
- DF (divide frequency) bit, [8-3](#), [8-4](#), [8-21](#)
- DFETCH (descriptor fetch) bit, [6-14](#), [6-21](#), [6-72](#)
- dFlags variable, [26-91](#)
- dFlags Word register, [26-91](#)
- DFRESET (double-fault reset) bit, [25-47](#), [25-48](#), [26-74](#)
- DI_EN (data int en) bit, [6-14](#), [6-67](#), [6-69](#)
- direct code execution, [26-33](#)
 - initial header, [26-32](#), [26-34](#)
- direction errors, MAC, [22-29](#)
- direction (PORT_DIR) bit, [20-26](#)
- direct memory access. *See* DMA
- disable automatic Tx CRC generation (DTXCRC) bit, [22-64](#)
- disable automatic Tx padding (DTXPAD) bit, [22-64](#)
- disable broadcast frame reception (DBF) bit, [22-65](#)
- disable/enable PTP_TSYNC module (EN) bit, [23-13](#)
- disable for trip input (PWMTRIP_DSBL) bit, [14-37](#)
- disable preamble generation (STADISPRES) bit, [22-72](#)
- disable receive own frames (DRO) bit, [22-61](#)
- disable Tx retry on collision (DRTYP) bit, [22-62](#)
- disabling PLL, [8-11](#)
- discarded frames, MAC, [22-15](#)
- DI_SEL (data interrupt timing select) bit, [6-67](#), [6-69](#)
- DITFS (data-independent transmit frame sync select) bit, [19-36](#), [19-47](#), [19-50](#), [19-61](#)
- divide frequency (DF) bit, [8-21](#)
- divider for selected PTP_CLK output (CKDIV) bit, [23-13](#)
- divisor latch access (DLAB) bit, [15-21](#)
- divisor latch high byte[15: 8] field, [15-30](#)
- divisor latch low byte[7: 0] field, [15-30](#)
- divisor reset, UART, [15-30](#)
- DLAB (divisor latch access) bit, [15-21](#), [15-26](#), [15-27](#)
- DLEN (data length) field, [20-25](#), [20-26](#)
- DMA, [6-1](#) to [6-102](#)
 - autobuffer mode, [6-11](#), [6-28](#), [6-68](#)
 - bandwidth, [6-45](#)
 - block count, [6-37](#)
 - block diagram, [6-103](#)
 - block done interrupt, [6-40](#)
 - block transfers, [6-9](#), [6-37](#)
 - buses, [4-8](#)
 - channels, [6-41](#), [6-50](#), [6-63](#), [6-64](#), [6-65](#)
 - congestion, [6-45](#)
 - connecting asynchronous FIFO, [6-38](#)
 - control commands, [6-31](#), [6-32](#), [6-34](#)
 - controller, [1-11](#), [6-2](#)

Index

DMA, *(continued)*

- descriptor structure, [6-96](#)
- direction, [6-70](#)
- DMA error interrupt, [6-73](#)
- double buffer scheme, [6-54](#)
- EBIU port DMA, [6-4](#)
- errors, [6-29](#), [6-30](#)
- ethernet MAC port DMA, [22-23](#),
[22-121](#), [22-122](#)
- external interfaces, [6-4](#)
- features, [6-2](#)
- finish control command, [6-33](#), [6-34](#)
- first data memory access, [6-21](#)
- flow chart, [6-18](#), [6-19](#)
- FLOW mode, [6-16](#)
- FLOW value, [6-20](#)
- handshake DMA, [1-12](#)
- handshake operation, [6-36](#)
- initializing, [6-17](#)
- internal interfaces, [6-4](#)
and L1 memory, [6-5](#)
- large model mode, [6-69](#)
- latency, [6-24](#)
- linked list, [22-122](#)
- memory conflict, [6-48](#)
- memory DMA, [1-12](#), [6-5](#), [6-6](#), [6-7](#)
- memory read, [6-25](#)
- operation flow, [6-17](#)
- orphan access, [6-28](#)
- overflow interrupt, [6-40](#)
- overview, [1-11](#)
- performance, [6-40](#), [6-42](#)
- pipelining requests, [6-37](#)
- polling registers, [6-51](#)
- PPI port DMA, [20-13](#), [20-22](#), [20-35](#)
- prioritization and traffic control, [6-44](#) to
[6-50](#)
- registers, [6-9](#), [6-65](#)
- remapping peripheral assignment, [6-6](#)

DMA, *(continued)*

- request data control command, [6-34](#)
- request data urgent control command,
[6-34](#)
- restart control command, [6-32](#), [6-33](#)
- small model mode, [6-68](#)
- software management, [6-50](#)
- SPI port DMA, [17-10](#), [17-23](#), [17-26](#),
[17-42](#)
- SPORT (serial port) DMA, [19-4](#), [19-23](#),
[19-37](#)
- stop mode, [6-11](#), [6-68](#)
- synchronization, [6-27](#), [6-50](#) to [6-60](#),
[20-13](#)
- throughput, [6-40](#), [6-42](#), [6-45](#)
- traffic control, [6-45](#), [6-48](#)
- types supported, [1-11](#)
- UART port DMA, [15-17](#), [15-27](#)
- word size, changing, [6-27](#), [6-28](#)
- work units, [6-14](#), [6-22](#), [6-24](#)
- DMA0 memory access override
(DMA0OVR) bit, [25-47](#)
- DMA0OVR (DMA0 memory access
override) bit, [25-47](#)
- DMA2D (DMA mode) bit, [6-67](#), [6-70](#)
- DMA bus. *See* DAB
- DMACFG field, [6-20](#), [6-62](#)
- DMA channel enable (DMAEN) bit, [6-67](#)
- DMA channel type (CTYPE) bit, [6-66](#)
- DMACODE bit, [26-92](#)
- DMA code (DMACODE) field, [26-22](#)
- DMA configuration (DMAx_CONFIG)
registers, [6-67](#)
- DMA configuration
(MDMA_yy_CONFIG) registers,
[6-67](#)
- DMA core bus. *See* DCB

- DMA descriptor, [6-13](#), [6-96](#)
 - array mode, [6-15](#), [6-22](#), [6-68](#)
 - chains, [6-26](#)
 - descriptor pairs, [22-23](#)
 - element offsets, [6-16](#)
 - list mode, [6-14](#), [6-22](#), [6-68](#), [6-69](#)
 - queue, [6-56](#), [6-57](#)
 - recommended size, [6-16](#)
 - structures, [6-55](#)
 - transfers, comparison, [6-3](#)
 - transfers, initializing, [6-94](#)
 - variable size descriptors, [6-15](#)
- DMA direction (WNR) bit, [6-67](#), [6-70](#)
- DMA_DONE (completion interrupt status) bit, [6-10](#), [6-72](#)
- DMA_DONE interrupt, [6-71](#)
- DMAEN (DMA channel enable) bit, [6-17](#), [6-61](#), [6-67](#), [6-70](#)
- DMA_ERR (error interrupt status) bit, [6-10](#), [6-72](#)
- DMA_ERROR interrupt, [6-29](#)
- DMA error interrupts, [6-72](#)
- DMA examples
 - 1-D interrupt-driven, [6-53](#)
 - 1-D unsynchronized FIFO, [6-55](#)
 - 2-D, polled, [6-54](#)
 - 2-D array, [6-91](#)
 - 2-D interrupt-driven, [6-54](#)
 - connection, receive, [6-39](#)
 - connection, transmit, [6-38](#)
 - header file to define descriptor structure, [6-95](#)
 - HMDMA1 block enable, [6-100](#)
 - HMDMA with delayed processing, [6-101](#)
 - polling DMA status, [6-94](#)
 - programming, [6-91](#) to [6-102](#)
 - register-based 2D memory DMA, [6-92](#)
- DMA examples *(continued)*
 - software-triggered descriptor fetch, [6-97](#)
 - two descriptors in small list flow mode, [6-95](#)
 - two-dimensional memory DMA setup, [6-93](#)
- DMA external bus. *See* DEB
- DMA mode (DMA2D) bit, [6-67](#)
- DMA operations, [1-11](#), [4-11](#), [6-2](#)
 - continuous transfers, [6-27](#), [6-53](#)
 - receive transfers, [6-26](#)
 - refresh transfers, [6-22](#)
 - round robin operations, [6-47](#)
 - serial port block transfers, [19-37](#)
 - single-buffer transfers, [6-53](#)
 - transfer startup, [6-17](#)
 - transfer stopping, [6-28](#)
 - transfer termination without abort, [6-28](#)
 - transfer triggering, [6-60](#)
 - transmit, [6-25](#), [6-34](#), [6-35](#), [22-23](#)
 - transmit finish, [6-34](#), [6-35](#)
 - transmit restart, [6-34](#), [6-35](#)
 - two-dimensional transfers, [6-11](#)
 - urgent transfers, [6-45](#)
- DMA overrun detected (RX_DMAO0) bit, [22-104](#)
- DMA overrun detected (RX_DMAO) bit, [22-96](#), [22-101](#)
- DMA peripherals, [5-6](#), [6-4](#), [6-5](#), [6-6](#), [6-104](#), [15-17](#), [15-27](#), [17-10](#), [17-23](#), [17-26](#), [17-42](#), [19-4](#), [20-13](#), [20-22](#), [20-35](#), [22-121](#)
 - peripherals supported, [1-6](#)
 - switching peripherals, [6-73](#)
- DMA queue completion interrupt, [6-59](#)
- DMA registers, [6-62](#), [6-63](#)
- DMA_RUN (channel running) bit, [6-10](#), [6-21](#), [6-58](#), [6-61](#), [6-71](#), [6-72](#)
- DMARx pin, [6-37](#)
- DMA start address field, [6-74](#)

Index

- DMA_TC_CNT (DMA traffic control counter) register, [6-89](#)
- DMA_TC_PER (DMA traffic control counter period) register, [6-47](#), [6-89](#)
- DMA traffic control counter (DMA_TC_CNT) register, [6-89](#)
- DMA traffic control counter period (DMA_TC_PER) register, [6-89](#)
- DMA traffic control registers, [6-88](#)
- DMA_TRAFFIC_PERIOD field, [6-90](#)
- DMA underrun (TX_DMAU) bit, [22-105](#)
- DMAx_CONFIG (DMA configuration) registers, [6-8](#), [6-17](#), [6-25](#), [6-67](#)
- DMAx_CURR_ADDR (current address) registers, [6-74](#)
- DMAx_CURR_DESC_PTR (current descriptor pointer) registers, [6-81](#)
- DMAx_CURR_X_COUNT (current inner loop count) registers, [6-76](#)
- DMAx_CURR_Y_COUNT (current outer loop count) registers, [6-78](#)
- DMAx_IRQ_STATUS (interrupt status) registers, [6-71](#), [6-72](#)
- DMAx_NEXT_DESC_PTR (next descriptor pointer) registers, [6-17](#), [6-80](#)
- DMAx_PERIPHERAL_MAP (peripheral map) registers, [5-6](#), [6-66](#)
- DMAx_START_ADDR (start address) registers, [6-17](#), [6-74](#)
- DMAx_X_COUNT (inner loop count) registers, [6-75](#)
- DMAx_X_MODIFY (inner loop address increment) registers, [6-17](#), [6-77](#)
- DMAx_Y_COUNT (outer loop count) registers, [6-78](#)
- DMAx_Y_MODIFY (outer loop address increment) registers, [6-17](#), [6-79](#)
- DMC (L1 data memory configure) field, [2-6](#)
- DMEM_CONTROL (data memory control) register, [2-4](#), [2-6](#)
- DNAK (data not acknowledged) bit, [16-33](#), [16-35](#)
- DOUBLE_FAULT (core-double-fault reset enable) bit, [25-39](#), [26-72](#)
- double-fault reset (DFRESET) bit, [26-74](#)
- DOUBLE_RESET, [25-39](#)
- double-update mode (PWM_DBL) bit, [14-37](#)
- double word index field, [2-7](#)
- down count int en (DCIE) bit, [13-21](#)
- down count interrupt identifier (DCII) bit, [13-22](#)
- DPMC, [8-2](#), [8-7](#) to [8-19](#)
- DPRESCALE (debounce delay) bit, [13-25](#)
- DR (data ready) bit, [15-12](#), [15-16](#), [15-24](#)
- drive/tolerate (TWL_DT) pin, [9-28](#)
- DRO (disable receive own frames) bit, [22-60](#), [22-61](#)
- DRQ[1: 0] field, [6-46](#), [6-82](#), [6-83](#)
- DRTY (disable Tx retry on collision) bit, [22-60](#), [22-62](#)
- DRxPRI signal, [19-5](#)
- DRxSEC signal, [19-5](#)
- DSP libraries, [1-33](#)
- DTEST_COMMAND (data test command) register, [2-6](#), [2-7](#)
- DTXCRC (disable automatic Tx CRC generation) bit, [22-60](#), [22-64](#)
- DTXPAD (disable automatic Tx padding) bit, [22-60](#), [22-64](#)
- DTxPRI signal, [19-5](#)
- DTxSEC signal, [19-5](#)
- dynamic power management, [1-28](#), [8-1](#), [8-2](#)

E

EAB

- arbitration, [4-11](#)
- clocking, [8-2](#)
- and EBIU, [7-4](#)
- frequency, [4-11](#)
- performance, [4-11](#)

early frame sync, [19-34](#)

EAV signal, [20-5](#)

EBCAW (external bank column address width) field, [7-61](#)

EBC (external bus controller), [7-5](#)

EBE (external bank enable) bit, [7-54](#), [7-60](#), [7-61](#)

EBIU, [1-12](#), [7-1](#) to [7-79](#)

- asynchronous interfaces supported, [7-1](#)

- block diagram, [7-4](#)

- bus errors, [7-7](#)

- byte enables, [7-15](#)

- clock, [7-1](#), [8-2](#)

- control registers, [7-5](#)

- core transfers to SRAM, [7-23](#)

- and DMA, [6-4](#)

- error detection, [7-7](#)

- external access extension, [7-15](#)

- overview, [7-1](#)

- read access period, [7-15](#)

- request priority, [7-1](#)

- shared pins, [7-6](#)

- as slave, [7-4](#)

- status register, [7-5](#)

- system clock, [7-7](#)

- timing characteristics, programmable, [7-11](#)

- wait states, [7-15](#)

EBIU_AMBCTLx (asynchronous memory bank control) registers, [7-20](#), [7-21](#), [7-22](#)

EBIU_AMGCTL (asynchronous memory global control) register, [7-20](#)

EBIU_SDBCTL (SDRAM memory global control) register, [7-60](#), [7-61](#)

EBIU_SDGCTL (SDRAM memory global control) register, [7-55](#), [7-64](#), [7-65](#)

EBIU_SDRRC (SDRAM refresh rate control) register, [7-58](#)

EBIU_SDSTAT (SDRAM control status) register, [7-74](#)

EBSZ (external bank size) field, [7-26](#), [7-60](#), [7-61](#)

EBUFE (external buffering enabled) bit, [7-57](#), [7-65](#), [7-72](#)

ECINIT[15: 0] field, [6-87](#)

ECOUNT[15: 0] field, [6-86](#)

edge detection, GPIO, [9-18](#)

EFT (ethernet frame type matching value) bit, [23-20](#)

EFTM (compare mask of frame type field) bit, [23-13](#)

EFTOF (offset of frame type EFT field) bit, [23-19](#)

elfloader.exe utility, [26-19](#)

ELSI (enable RX status int) bit, [15-27](#)

ELSI (enable Rx status int) bit, [15-7](#), [15-28](#)

EMAC_ADDRHI (address high) register, [22-51](#), [22-68](#)

EMAC_ADDRLO (address low) register, [22-51](#), [22-67](#)

EMAC_FLC (flow control) register, [22-51](#), [22-74](#)

EMAC_HASHHI (multicast hash table high) register, [22-51](#), [22-68](#)

EMAC_HASHLO (multicast hash table low) register, [22-51](#), [22-68](#)

EMAC_MMC_CTL (management counters control) register, [22-52](#), [22-119](#)

EMAC_MMC_RIRQE (MMC RX int en) register, [22-52](#), [22-114](#)

Index

- EMAC_MMC_RIRQS (MMC RX interrupt status) register, [22-52](#), [22-112](#)
- EMAC_MMC_TIRQE (MMC TX int en) register, [22-52](#), [22-117](#)
- EMAC_MMC_TIRQS (MMC TX interrupt status) register, [22-52](#), [22-114](#)
- EMAC_OPMODE (operating mode) register, [22-51](#), [22-60](#)
- EMAC_PTP_ACCR (accumulator) register, [23-24](#)
- EMAC_PTP_ADDEND (addend) register, [23-23](#)
- EMAC_PTP_ALARMHI (target alarm time high) register, [23-33](#)
- EMAC_PTP_ALARMLO (target alarm time low) register, [23-32](#)
- EMAC_PTP_CTL (control) register, [23-13](#)
- EMAC_PTP_FOFF (message filter offset) register, [23-19](#)
- EMAC_PTP_FV1 (message filter value 1) register, [23-20](#)
- EMAC_PTP_FV2 (message filter value 2) register, [23-21](#)
- EMAC_PTP_FV3 (message filter value 3) register, [23-22](#)
- EMAC_PTP_ID_OFF (source ID offset) register, [23-33](#)
- EMAC_PTP_ID_SNAP (source ID snapshot) register, [23-34](#)
- EMAC_PTP_IE (PTP int en) register, [23-15](#)
- EMAC_PTP_ISTAT (interrupt status) register, [23-17](#)
- EMAC_PTP_OFFSET (time offset) register, [23-25](#)
- EMAC_PTP_PPS_PERIOD (PPS period) register, [23-37](#)
- EMAC_PTP_PPS_STARTHI (PPS start high) register, [23-36](#)
- EMAC_PTP_PPS_STARTLO (PPS start low) register, [23-35](#)
- EMAC_PTP_RXSNAPHI (receive snapshot high) register, [23-29](#)
- EMAC_PTP_RXSNAPLO (receive snapshot low) register, [23-28](#)
- EMAC_PTP_TIMEHI (local clock time high) register, [23-27](#)
- EMAC_PTP_TIMELO (local clock time low) register, [23-26](#)
- EMAC_PTP_TXSNAPHI (transmit snapshot high) register, [23-31](#)
- EMAC_PTP_TXSNAPLO (transmit snapshot low) register, [23-30](#)
- EMAC_RXC_ALIGN (alignment errors) register, [22-53](#)
- EMAC_RXC_ALLFRM (frames received all) register, [22-55](#)
- EMAC_RXC_ALLOCT (octets received all) register, [22-55](#)
- EMAC_RXC_BROAD (broadcast frames received ok) register, [22-54](#)
- EMAC_RXC_DMAOVF (frames lost due to int MAC rcv error) register, [22-53](#)
- EMAC_RXC_EQ64 (frames len eq 64 received) register, [22-56](#)
- EMAC_RXC_FCS (frame check sequence errors) register, [22-53](#)
- EMAC_RXC_GE1024 (frames len 1024-max received) register, [22-56](#)
- EMAC_RXC_LNERRI (in range length errors) register, [22-54](#)
- EMAC_RXC_LNERRO (out of range length field) register, [22-54](#)
- EMAC_RXC_LONG (frame too long errors) register, [22-54](#)
- EMAC_RXC_LT1024 (frames len 512-1023 received) register, [22-56](#)

- EMAC_RXC_LT128 (frames len 65-127 received) register, [22-56](#)
- EMAC_RXC_LT256 (frames len 128-255 received) register, [22-56](#)
- EMAC_RXC_LT512 (frames len 256-511 received) register, [22-56](#)
- EMAC_RXC_MACCTL (MAC control frames received) register, [22-55](#)
- EMAC_RXC_MULTI (multi cast frames received ok) register, [22-54](#)
- EMAC_RXC_OCTET (octets received ok) register, [22-53](#)
- EMAC_RXC_OK (frames received ok) register, [22-53](#)
- EMAC_RXC_OPCODE (unsupported opcodes received) register, [22-55](#)
- EMAC_RXC_PAUSE (pause MAC ctrl frames received) register, [22-55](#)
- EMAC_RXC_SHORT (frames len lt 64 received) register, [22-56](#)
- EMAC_RXC_TYPED (typed frames received) register, [22-55](#)
- EMAC_RXC_UNICST (unicast frames received ok) register, [22-53](#)
- EMAC_RX_IRQE (frame status int en) register, [22-52](#), [22-104](#)
- EMAC_RX_STAT (current frame status) register, [22-52](#), [22-93](#)
- EMAC_RX_STKY (sticky frame status) register, [22-52](#), [22-99](#)
- EMAC. *See* MAC or EMAC_ registers
- EMAC_STAADD (station management address) register, [22-10](#), [22-51](#), [22-72](#)
- EMAC_STADAT (station management data) register, [22-10](#), [22-51](#), [22-74](#)
- EMAC_SYSCTL (system control) register, [22-10](#), [22-52](#), [22-89](#)
- EMAC_SYSTAT (system status) register, [22-52](#), [22-90](#)
- EMAC_TXC_1COL (single collision frames) register, [22-56](#)
- EMAC_TXC_ALLFRM (frames transmitted all) register, [22-58](#)
- EMAC_TXC_ALLOCT (octets transmitted all) register, [22-58](#)
- EMAC_TXC_BROAD (broadcast frames xmitted ok) register, [22-58](#)
- EMAC_TXC_CRSEERR (carrier sense errors) register, [22-57](#)
- EMAC_TXC_DEFER (frames with deferred xmissions) register, [22-57](#)
- EMAC_TXC_DMAUND (frames lost due to int MAC xmit error) register, [22-57](#)
- EMAC_TXC_EQ64 (frames len eq 64 transmitted) register, [22-58](#)
- EMAC_TXC_GT1COL (multiple collision frames) register, [22-57](#)
- EMAC_TXC_LATECL (late collisions) register, [22-57](#)
- EMAC_TXC_LT128 (frames len 65-127 transmitted) register, [22-58](#)
- EMAC_TXC_MACCTL (MAC control frames transmitted) register, [22-58](#)
- EMAC_TXC_MULTI (multicast frames xmitted ok) register, [22-58](#)
- EMAC_TXC_OCTET (octets transmitted ok) register, [22-57](#)
- EMAC_TXC_OK (frames transmitted ok) register, [22-56](#)
- EMAC_TXC_UNICST (unicast frames xmitted ok) register, [22-58](#)
- EMAC_TXC_XS_COL (frames aborted due to xs colls) register, [22-57](#)
- EMAC_TXC_XS_DFR (frames with excessive deferral) register, [22-58](#)
- EMAC_TX_IRQE (frame status int en) register, [22-52](#), [22-112](#)
- EMAC_TX_STAT (current frame status) register, [22-52](#), [22-105](#)

Index

- EMAC_TX_STKY (sticky frame status) register, [22-52](#), [22-109](#)
- EMAC_VLANx (VLANx tag) registers, [22-51](#), [22-76](#)
- EMAC_WKUP_CTL (wakeup frame control and status) register, [22-51](#), [22-77](#)
- EMAC_WKUP_FFCMD (wakeup frame filter commands) register, [22-51](#), [22-85](#)
- EMAC_WKUP_FFCRCx (wakeup frame filter CRCx) registers, [22-51](#), [22-87](#)
- EMAC_WKUP_FFMSKx (wakeup frameX byte mask) registers, [22-51](#), [22-80](#), [22-81](#), [22-82](#), [22-83](#), [22-84](#)
- EMAC_WKUP_FFOFF (wakeup frame filter offsets) register, [22-51](#), [22-87](#)
- EMISO (enable MISO) bit, [17-35](#), [17-36](#)
- EMREN (extended mode register enabled) bit, [7-48](#), [7-65](#), [7-73](#)
- EMRS command, [7-48](#)
- EMUDABL (emulator disable) bit, [25-47](#)
- emulation, and timer counter, [10-41](#)
- emulation behavior select (EMU_RUN) bit, [10-40](#)
- emulation override (EMUOVR) bit, [25-47](#)
- emulator disable (EMUDABL) bit, [25-47](#)
- EMUOVR (emulation override) bit, [25-47](#)
- EMU_RUN (emulation behavior select) bit, [10-40](#), [10-41](#), [10-45](#)
- enabling interrupts, [5-5](#)
- ENDCPLB (enable cacheability lookaside buffer) bit, [2-6](#)
- end of data (DAT_END) bit, [21-65](#)
- end of data (DAT_END_MASK) bit, [21-70](#)
- EN (enable/disable PTP_TSYNC module) bit, [23-13](#)
- entire field mode, PPI, [20-9](#)
- EPROM, [1-9](#)
- EPS (even parity select) bit, [15-21](#)
- ERBFI (enable receive buffer full int) bit, [15-7](#), [15-12](#), [15-26](#), [15-27](#)
- ERMS command, [7-33](#)
- ERR_DET (error detected) bit, [20-29](#), [20-30](#)
- ERR_NCOR (error not corrected) bit, [20-29](#), [20-30](#)
- error code. *See* OTP error code
- error detected (ERR_DET) bit, [20-30](#)
- error interrupt status (DMA_ERR) bit, [6-72](#)
- error not corrected (ERR_NCOR) bit, [20-30](#)
- errors, DMA, [6-29](#)
 - not detected by hardware, [6-30](#)
- errors, timer, [10-8](#)
- error signals, SPI, [17-39](#) to [17-41](#)
- error type (ERR_TYP) bit, [10-40](#)
- ERR_TYP (error type) bits, [10-7](#), [10-28](#), [10-38](#), [10-40](#), [10-45](#)
- ETBEI (enable transmit buffer empty int) bit, [15-6](#), [15-11](#), [15-17](#), [15-26](#), [15-27](#)
- ethernet controller architecture, [22-3](#)
- ethernet frame header, [22-49](#)
- ethernet frame type matching value (EFT) bit, [23-20](#)
- ethernet MAC. *See* MAC or EMAC_ registers
- even parity select (EPS) bit, [15-21](#)
- events
 - controller, [5-2](#)
 - definition, [5-3](#)
 - handling, [5-2](#)
 - system, [5-3](#)
 - types of, [5-2](#)
 - vector table (EVT), [5-2](#)
- EVT1 register, [26-7](#)
- EXT_CLK mode, [10-31](#) to [10-33](#), [10-42](#), [10-44](#)

- extended mode register enabled (EMREN) bit, 7-65
 - extended mode register set command, 7-48
 - external access bus. *See* EAB
 - external bank column address width (EBCAW) bit, 7-61
 - external bank enable (EBE) bit, 7-61
 - external bank size (EBSZ) bit, 7-61
 - external buffering enabled (EBUFE) bit, 7-65
 - external bus interface unit. *See* EBIU
 - external crystal, 1-28
 - external emulator/debugger, 10-41
 - external memory, 1-9, 2-5
 - design issues, 27-5
 - interfaces, 7-6
 - map, 2-1, 7-3
 - external PHY, 22-4
 - external SDRAM memory, 7-26, 7-32
 - external sync (PWM_EXTSYNC) bit, 14-37
 - external sync select (PWM_SYNCSEL) bit, 14-37
 - EXTEST instruction, B-6
 - EZ-KIT Lite card, 1-33
- F**
- fast back-to-back read to write (FBBRW) bit, 7-65
 - FAST (fast mode) bit, 16-30, 16-31
 - fast mode, TWI, 16-10
 - fast mode (FAST) bit, 16-30
 - FBBRW (fast back-to-back read to write) bit, 7-65, 7-73
 - FDMODE (full duplex mode) bit, 22-60, 22-61
 - FE (framing error) bit, 15-24, 15-25
 - FFE (force framing error on transmit) bit, 15-31, 15-32
 - field indicator (FLD) bit, 20-30
 - FIFO, asynchronous connection, 6-38
 - FIFO_COUNT (number of 32-bit words remaining) bit, 21-73
 - FIFO underrun (UNDR) bit, 20-30
 - filtering, MAC, 22-13
 - finish control command, DMA, 6-33, 6-34
 - flash memory, 1-9, 7-1, 27-5
 - FLCBUSY (flow control busy status) bit, 22-75
 - FLCE (flow control enable) bit, 22-16, 22-75
 - FLCPAUSE[15: 0] field, 22-75
 - FLD (field indicator) bit, 20-30
 - FLD_SEL (active field select) bit, 20-4, 20-26, 20-28
 - flex descriptors, 6-3
 - FLGx (slave select value) bit, 17-37, 17-38
 - FLOW[2: 0] field, 6-22, 6-23, 6-55, 6-67, 6-68
 - flow charts
 - DMA operations, 6-18, 6-19
 - GP I/O operations, 9-21, 9-24
 - GP timer operations, 10-32
 - GP timers operations, 10-6
 - PPI operations, 20-24
 - SPI operations, 17-29, 17-30
 - timer PWM_OUT operations, 10-11
 - timer WIDTH_CAP operations, 10-23
 - TWI operations, 16-22, 16-23
 - flow control busy status (FLCBUSY) bit, 22-75
 - flow control (EMAC_FLC) register, 22-74
 - flow control enable (FLCE) bit, 22-75
 - FLOW mode, DMA, 6-16, 6-20
 - FLOW (next operation) bit, 6-14, 6-15
 - FLSx (slave select enable) bit, 17-8, 17-37
 - force framing error on transmit (FFE) bit, 15-31
 - force parity error on transmit (FPE) bit, 15-31

Index

- force reload (RBC) bit, [6-83](#)
- FPE (force parity error on transmit) bit, [15-31](#), [15-32](#)
- frame buffer, ethernet, [22-122](#)
- frame check sequence errors counter int en (RX_FCS_CNT) bit, [22-115](#)
- frame check sequence errors counter int (RX_FCS_CNT) bit, [22-113](#)
- frame check sequence errors (EMAC_RXC_FCS) register, [22-53](#)
- frame CRC error (RX_CRC) bit, [22-97](#), [22-102](#), [22-104](#)
- frame deferral int en (TX_DEFER) bit, [22-105](#), [22-110](#), [22-112](#)
- framed serial transfers, characteristics, [19-30](#)
- framed/unframed data, [19-29](#)
- frame filter evaluation, MAC, [22-14](#)
- frame fragment detected (RX_FRAG) bit, [22-96](#), [22-102](#), [22-104](#)
- frame length (TX_FRLLEN) bit, [22-105](#)
- frames aborted due to excess collisions counter interrupt (TX_ABORTC_CNT) bit, [22-118](#)
- frames aborted due to excess collisions counter int (TX_ABORTC_CNT) bit, [22-116](#)
- frames aborted due to xs colls (EMAC_TXC_XS_COL) register, [22-57](#)
- frames len 1024-max received counter int en (RX_GE1024_CNT) bit, [22-115](#)
- frames len 1024-max received counter int (RX_GE1024_CNT) bit, [22-113](#)
- frames len 1024-max received (EMAC_RXC_GE1024) register, [22-56](#)
- frames len 1024-max transmitted counter int en (TX_GE1024_CNT) bit, [22-118](#)
- frames len 1024-max transmitted counter int (TX_GE1024_CNT) bit, [22-116](#)
- frames len 128-255 received counter int en (RX_LT256_CNT) bit, [22-115](#)
- frames len 128-255 received counter int (RX_LT256_CNT) bit, [22-113](#)
- frames len 128-255 received (EMAC_RXC_LT256) register, [22-56](#)
- frames len 128-255 transmitted counter int en (TX_LT256_CNT) bit, [22-118](#)
- frames len 128-255 transmitted counter int (TX_LT256_CNT) bit, [22-116](#)
- frames len 256-511 received counter int en (RX_LT512_CNT) bit, [22-115](#)
- frames len 256-511 received counter int (RX_LT512_CNT) bit, [22-113](#)
- frames len 256-511 received (EMAC_RXC_LT512) register, [22-56](#)
- frames len 256-511 transmitted counter int en (TX_LT512_CNT) bit, [22-118](#)
- frames len 256-511 transmitted counter int (TX_LT512_CNT) bit, [22-116](#)
- frames len 512-1023 received counter int en (RX_LT1024_CNT) bit, [22-115](#)
- frames len 512-1023 received counter int (RX_LT1024_CNT) bit, [22-113](#)
- frames len 512-1023 received (EMAC_RXC_LT128) register, [22-56](#)
- frames len 512-1023 transmitted counter int en (TX_LT1024_CNT) bit, [22-118](#)
- frames len 512-1023 transmitted counter int (TX_LT1024_CNT) bit, [22-116](#)
- frames len 65-127 received counter int en (RX_LT128_CNT) bit, [22-115](#)
- frames len 65-127 received counter int (RX_LT128_CNT) bit, [22-113](#)

- frames len 65-127 received
(EMAC_RXC_LT128) register,
[22-56](#)
- frames len 65-127 transmitted counter int
en (TX_LT128_CNT) bit, [22-118](#)
- frames len 65-127 transmitted counter int
(TX_LT128_CNT) bit, [22-116](#)
- frames len 65-127 transmitted
(EMAC_TXC_LT128) register,
[22-58](#)
- frames len eq 64 received
(EMAC_RXC_EQ64) register, [22-56](#)
- frames len eq 64 transmitted
(EMAC_TXC_EQ64) register, [22-58](#)
- frames len eq to 64 received counter int en
(RX_EQ64_CNT) bit, [22-115](#)
- frames len eq to 64 received counter int
(RX_EQ64_CNT) bit, [22-113](#)
- frames len eq to 64 transmitted counter int
en (TX_EQ64_CNT) bit, [22-118](#)
- frames len eq to 64 transmitted counter int
(TX_EQ64_CNT) bit, [22-116](#)
- frames len lt 64 received counter int en
(RX_SHORT_CNT) bit, [22-115](#)
- frames len lt 64 received counter int
(RX_SHORT_CNT) bit, [22-113](#)
- frames len lt 64 received
(EMAC_RXC_SHORT) register,
[22-56](#)
- frames lost counter int en
(RX_LOST_CNT) bit, [22-115](#)
- frames lost counter int (RX_LOST_CNT)
bit, [22-113](#)
- frames lost due to internal MAC transmit
error counter int en
(TX_LOST_CNT) bit, [22-118](#)
- frames lost due to internal MAC transmit
error counter int (TX_LOST_CNT)
bit, [22-116](#)
- frames lost due to int mac rcv error
(EMAC_RXC_DMAOVF) register,
[22-53](#)
- frames lost due to int mac xmit error
(EMAC_TXC_DMAUND) register,
[22-57](#)
- frames received all counter int en
(RX_ALLF_CNT) bit, [22-115](#)
- frames received all counter int
(RX_ALLF_CNT) bit, [22-113](#)
- frames received all
(EMAC_RXC_ALLFRM) register,
[22-55](#)
- frames received OK counter int en
(RX_OK_CNT) bit, [22-115](#)
- frames received OK counter int
(RX_OK_CNT) bit, [22-113](#)
- frames received ok (EMAC_RXC_OK)
register, [22-53](#)
- frame start detect, PPI, [20-34](#)
- frame status int en (EMAC_RX_IRQE)
register, [22-104](#)
- frame status int en (EMAC_TX_IRQE)
register, [22-112](#)
- frames transmitted all counter int en
(TX_ALLF_CNT) bit, [22-118](#)
- frames transmitted all counter int
(TX_ALLF_CNT) bit, [22-116](#)
- frames transmitted all
(EMAC_TXC_ALLFRM) register,
[22-58](#)
- frames transmitted OK counter int en
(TX_OK_CNT) bit, [22-118](#)
- frames transmitted OK counter int
(TX_OK_CNT) bit, [22-116](#)
- frames transmitted ok (EMAC_TXC_OK)
register, [22-56](#)
- frames transmitted OK (TX_OK) bit,
[22-109](#)

Index

- frames with deferred transmission counter int en (TX_DEFER_CNT) bit, [22-118](#)
 - frames with deferred transmission counter int (TX_DEFER_CNT) bit, [22-116](#)
 - frames with deferred xmissions (EMAC_TXC_DEFER) register, [22-57](#)
 - frames with excessive deferral counter int en (TX_EXDEF_CNT) bit, [22-118](#)
 - frames with excessive deferral counter int (TX_EXDEF_CNT) bit, [22-116](#)
 - frames with excessive deferral (EMAC_TXC_XS_DFR) register, [22-58](#)
 - frame sync
 - active high/low, [19-32](#)
 - early, [19-34](#)
 - external, [19-31](#)
 - internal, [19-25](#), [19-31](#), [19-62](#)
 - late, [19-34](#)
 - multichannel mode, [19-18](#)
 - sampling edge, [19-32](#)
 - SPORT options, [19-29](#)
 - frame sync divider[15: 0] field, [19-63](#)
 - frame synchronization, [19-3](#), [20-19](#)
 - frame sync polarity, PPI and timer, [20-20](#)
 - frame sync signal, [19-49](#), [19-50](#), [19-54](#)
 - frame sync to data relationship (FSDR) bit, [19-64](#)
 - frame too long errors counter int en (RX_LONG_CNT) bit, [22-115](#)
 - frame too long errors counter int (RX_LONG_CNT) bit, [22-113](#)
 - frame too long errors (EMAC_RXC_LONG) register, [22-54](#)
 - frame too long (RX_LONG) bit, [22-97](#), [22-102](#), [22-104](#)
 - frame track error (FT_ERR) bit, [20-30](#), [20-33](#)
 - frame transmission complete (TX_COMP) bit, [22-105](#), [22-109](#), [22-112](#)
 - frame transmit OK int en (TX_OK) bit, [22-105](#), [22-112](#)
 - framing error (FE) bit, [15-24](#)
 - FREQ bit, [8-23](#)
 - frequencies, clock and frame sync, [19-24](#)
 - frequency, DEB, [4-11](#)
 - frequency, EAB, [4-11](#)
 - FSDR (frame sync to data relationship) bit, [19-21](#), [19-64](#)
 - F signal, [20-30](#)
 - FT_ERR (frame track error) bit, [20-30](#), [20-33](#)
 - full duplex, [17-2](#), [19-4](#), [19-5](#)
 - full duplex mode (FDMODE) bit, [22-61](#)
 - FULL_ON bit, [8-22](#)
 - full-on mode, [1-28](#), [8-8](#)
 - function enable (PORTx_FER) registers, [9-10](#), [9-32](#)
- ## G
- GCALL (general call) bit, [16-28](#), [16-29](#)
 - general call address, TWI, [16-10](#)
 - general call enable (GEN) bit, [16-26](#)
 - general call (GCALL) bit, [16-28](#)
 - general-purpose interrupts, [5-2](#), [5-3](#)
 - general-purpose I/O
 - overview, [1-13](#)
 - general-purpose I/O. *See* GPIO
 - general-purpose ports, [9-1](#) to [9-45](#)
 - interrupt channels, [9-19](#)
 - interrupt generation flow, [9-19](#)
 - I/O pins, [9-3](#), [9-18](#)
 - latency, [9-13](#)
 - throughput, [9-13](#)
 - general-purpose ports. *See* GPIO

- general-purpose timers, 10-1 to 10-57
 - aborting immediately, 10-22
 - autobaud mode, 10-31
 - block diagram, 10-57
 - buffer registers, 10-42
 - capture mode, 10-5
 - clock source, 10-3
 - code examples, 10-47
 - control bit summary, 10-44
 - counter, 10-4
 - disable timing, 10-22
 - enabling, 10-5, 10-33
 - error detection, 10-7
 - EXT_CLK mode, 10-42
 - external interface, 10-3
 - features, 10-2
 - generating maximum frequency, 10-15
 - illegal states, 10-7, 10-9
 - internal interface, 10-3, 10-4
 - interrupts, 10-4, 10-5, 10-6, 10-14, 10-28, 10-49, 10-50
 - measurement report, 10-24, 10-26, 10-27
 - non-overlapping clock pulses, 10-52
 - output pad disable, 10-12
 - overflow, 10-4
 - PPI related operations, 10-56
 - programming model, 10-33
 - PULSE_HI toggle mode, 10-15
 - PWM mode, 10-5
 - PWM_OUT mode, 10-10 to 10-22, 10-42, 10-43
 - register accesses, 10-35
 - registers, 10-34
 - setup, 10-47
 - signal generation, 10-48
 - single pulse generation, 10-12
 - startup errors, 10-8
 - three timers with same period, 10-17
 - two timers with non-overlapping clocks, 10-17
 - waveform generation, 10-13
 - WDTH_CAP mode, 10-23, 10-42, 10-54
- GEN (general call enable) bit, 16-26, 16-27
- get more data (GM) bit, 17-36
- glitch filtering, UART, 15-9
- global unicast wakeup enable (GUWKE) bit, 22-79
- glueless connection, 27-5
- GM (get more data) bit, 17-20, 17-36
- good received frame int en (RX_OK) bit, 22-98, 22-103, 22-104
- GPIO, 1-13, 9-1 to 9-45
 - assigned to same interrupt channel, 9-22
 - both edges control, 9-37
 - clearing interrupt conditions, 9-20
 - clear registers, 9-17
 - code examples, 9-44
 - configuration, 9-14
 - data registers, 9-15, 9-16
 - direction control, 9-33
 - direction registers, 9-14, 9-15, 9-19
 - edge detection, 9-18
 - edge-sensitive, 9-16
 - flow chart, 9-24
 - function enable registers, 9-14, 9-18
 - input buffers, 9-15
 - input driver, 9-15
 - input drivers, 9-19
 - input enable control, 9-33
 - input enable registers, 9-15, 9-18
 - interrupt enable control, 9-37, 9-38
 - interrupt flow chart, 9-21
 - interrupts, 5-14, 9-19, 9-22, 9-23
 - interrupt sensitivity registers, 9-18

Index

GPIO, *(continued)*
mask data registers, [9-20](#)
mask interrupt clear registers, [9-21](#)
mask interrupt set registers, [9-20](#)
mask interrupt toggle registers, [9-22](#)
mask registers, [9-19](#)
pins, [9-14](#)
polarity control, [9-36](#)
polarity registers, [9-18](#)
registers, [9-26](#)
sensitivity control, [9-36](#)
set registers, [9-16](#)
toggle registers, [9-17](#)
using as input, [9-15](#)
write operations, [9-15](#)
writes to registers, [9-17](#)
GPIO clear (PORTxIO_CLEAR) registers, [9-35](#)
GPIO data (PORTxIO) registers, [9-34](#)
GPIO direction (PORTxIO_DIR) registers, [9-33](#), [14-40](#), [14-41](#), [14-42](#), [14-43](#), [14-46](#), [14-47](#), [14-48](#), [14-49](#)
GPIO input enable (PORTxIO_INEN) registers, [9-33](#)
GPIO mask interrupt set/clear/toggle (PORTxIO_MASKA/B_X) registers, [9-37](#), [9-39](#), [9-40](#), [9-41](#), [9-42](#), [9-43](#), [9-44](#)
GPIO mask int (PORTxIO_MASKA/B) registers, [9-37](#), [9-38](#)
GPIO pins, [9-14](#)
GPIO polarity (PORTxIO_POLAR) registers, [9-36](#)
GPIO set on both edges (PORTxIO_BOTH) registers, [9-37](#)
GPIO set (PORTxIO_SET) registers, [9-34](#)
GPIO toggle (PORTxIO_TOGGLE) registers, [9-35](#)
GP modes, PPI, [20-14](#)
ground plane, [27-9](#), [27-10](#)

GUWKE (global unicast wakeup enable) bit, [22-79](#)

H

H.100 standard protocol, [19-21](#), [19-23](#)
handshake MDMA, [1-12](#), [6-8](#), [6-40](#)
handshake MDMA configuration (HMDMAx_BCINIT) registers, [6-37](#), [6-84](#)
handshake MDMA control (HMDMAx_CONTROL) registers, [6-82](#)
handshake MDMA control registers, [6-83](#)
handshake MDMA current block count (HMDMAx_BCOUNT) registers, [6-37](#), [6-85](#)
handshake MDMA current edge count (HMDMAx_ECOUNTER) registers, [6-37](#), [6-38](#), [6-86](#)
handshake MDMA edge count overflow int (HMDMAx_ECOVERFLOW) registers, [6-88](#)
handshake MDMA edge count urgent (HMDMAx_ECURGENT) registers, [6-87](#)
handshake MDMA enable (HMDMAEN) bit, [6-83](#)
handshake MDMA initial block count (HMDMAx_BCINIT) registers, [6-84](#)
handshake MDMA initial edge count (HMDMAx_ECINIT) registers, [6-38](#), [6-87](#)
handshaking memory DMA (HMDMA), [6-2](#), [6-4](#)
hardware reset, [26-4](#), [26-5](#), [26-7](#)
hash filter multicast addresses (HM) bit, [22-66](#)
hash filter unicast addresses (HU) bit, [22-66](#)
header checksum (HDRCHK) field, [26-26](#)

- header signature (HDRSGN), [26-27](#)
 - hibernate state, [1-30](#), [7-45](#), [8-10](#)
 - higher 32 bits of local time clock
(LOCAL_TIME_HI) bit, [23-27](#)
 - higher 32 bits of PPS start time
(PPS_STARTHI) bit, [23-36](#)
 - higher 32 bits of receive snapshot of local
clock time (RX_SNAP_HI) bit, [23-29](#)
 - higher 32 bits of target alarm time
(ALARM_TIME_HI) bit, [23-33](#)
 - higher 32 bits of transmit snapshot of local
clock time (TX_SNAP_HI) bit, [23-31](#)
 - high-frequency design considerations, [27-8](#)
 - HMDMA, [6-2](#), [6-8](#)
 - HMDMAEN (handshake MDMA enable)
bit, [6-36](#), [6-38](#), [6-83](#)
 - HMDMA_x_BCINIT (handshake MDMA
configuration) registers, [6-37](#), [6-84](#)
 - HMDMA_x_BCOUNT (handshake
MDMA current block count)
registers, [6-37](#), [6-85](#)
 - HMDMA_x_CONTROL (handshake
MDMA control) registers, [6-82](#)
 - HMDMA_x_ECINIT (handshake MDMA
initial edge count) registers, [6-38](#),
[6-87](#)
 - HMDMA_x_ECOUNT (handshake
MDMA current edge count) registers,
[6-37](#), [6-38](#), [6-86](#)
 - HMDMA_x_ECOVERFLOW (handshake
MDMA edge count overflow int)
registers, [6-88](#)
 - HMDMA_x_ECURGENT (handshake
MDMA edge count urgent) registers,
[6-87](#)
 - HM (hash filter multicast addresses) bit,
[22-60](#), [22-66](#)
 - HMVIP compatibility mode, [19-24](#)
 - hold, for EBIU asynchronous memory
controller, [7-19](#)
 - horizontal blanking, [20-6](#)
 - horizontal tracking, PPI, [20-31](#)
 - horizontal tracking overflow error
(LT_ERR_OVR) bit, [20-30](#)
 - horizontal tracking underflow error
(LT_ERR_UNDR) bit, [20-30](#)
 - hours[3: 0] field, [24-20](#), [24-22](#)
 - hours[4] bit, [24-20](#), [24-22](#)
 - hours event flag bit, [24-21](#)
 - hours int en bit, [24-20](#)
 - HU (hash filter unicast addresses) bit,
[22-60](#), [22-66](#)
- ## I
- I²C bus standard, [1-14](#), [16-2](#)
 - I2C. *See* TWI
 - I²S, [1-21](#)
 - format, [19-10](#)
 - serial devices, [19-3](#)
 - ICIE (illegal gray/binary code int en) bit,
[13-21](#)
 - ICII (illegal gray/binary code interrupt
identifier) bit, [13-22](#)
 - idle state
 - waking from, [5-6](#)
 - IEEE 1149.1 standard. *See* JTAG standard
 - IEEE 802.3, [1-15](#), [22-4](#), [22-8](#), [22-42](#)
 - IFE (inverse filtering) bit, [22-60](#), [22-65](#)
 - illegal gray/binary code interrupt identifier
(ICII) bit, [13-22](#)
 - IMASK (interrupt mask) register
 - initialization, [5-8](#)
 - IN_DELAY bit, [8-21](#)
 - INIT bit, [26-36](#)
 - initcall address/symbol command, [26-37](#)
 - initcode routines, [26-35](#)
 - initialization
 - IMASK register, [5-8](#)
 - interrupt, [5-8](#)
 - SDRAM, [7-70](#)

Index

- initializing
 - DMA, [6-17](#)
- init initcode.dxe command, [26-37](#)
- inner loop address increment registers
 - (DMAx_X_MODIFY), [6-77](#)
 - (MDMA_yy_X_MODIFY), [6-77](#)
- inner loop count registers
 - (DMAx_X_COUNT), [6-75](#)
 - (MDMA_yy_X_COUNT), [6-75](#)
- INPDIS (CUD and CDZ input disable)
 - bit, [13-20](#)
- input buffers, GPIO, [9-15](#)
- input clock. *See* CLKIN
- input delay bit, [8-21](#)
- input driver, GPIO, [9-15](#)
- in-range length errors counter int
 - (RX_IRL_CNT) bit, [22-113](#), [22-115](#)
- in range length errors
 - (EMAC_RXC_LNERRI) register, [22-54](#)
- instruction bit scan ordering, [B-5](#)
- instruction register (IR), [B-2](#), [B-4](#)
- instructions, [1-30](#)
 - private, [B-4](#)
 - public, [B-4](#)
 - See also* instructions by name
- interfaces
 - external memory, [7-6](#)
 - internal memory, [7-5](#)
 - on-chip, [4-2](#)
 - overview, [4-3](#)
 - RTC, [24-3](#)
 - system, [4-2](#)
- inter IC bus, [16-2](#)
- interlaced video, [20-6](#)
- interleaving SPORT data, [19-5](#), [19-56](#)
- internal address mapping (table), [7-61](#)
- internal bank, [7-31](#)
- internal boot ROM, [26-1](#)
- internal clocks, [4-4](#)
- internal/external frame syncs. *See* frame sync
- internal loopback enable (LB) bit, [22-61](#)
- internal MAC error int en (TX_MACE)
 - bit, [22-112](#)
- internal MAC errors detected (TX_MACE)
 - bit, [22-109](#)
- internal memory, [1-9](#)
 - accesses, [2-1](#)
 - interfaces, [7-5](#)
- internal receive clock select (IRCLK) bit, [19-52](#)
- internal receive frame sync select (IRFS) bit, [19-52](#)
- internal SDRAM banks, [7-27](#)
- internal transmit clock select (ITCLK) bit, [19-47](#)
- internal transmit frame sync select (ITFS)
 - bit, [19-47](#)
- internal TSR register, UART, [15-6](#)
- internal voltage regulator, [1-27](#)
- interrupt mask (CNT_IMASK) register, [13-19](#), [13-21](#)
- interrupt request enable (IRQ_ENA) bit, [10-40](#)
- interrupts, [5-1](#) to [5-15](#)
 - channel assignment, [5-3](#), [5-4](#), [9-19](#)
 - configuration, [27-2](#)
 - core wake up, [5-6](#)
 - definition, [5-3](#)
 - DMA usage, [5-6](#), [6-29](#), [6-40](#), [6-57](#), [6-59](#), [6-73](#)
 - enabling, [5-5](#)
 - ethernet event, [22-38](#)
 - general-purpose usage, [5-2](#), [5-3](#)
 - GPIO event, [9-22](#)
 - GPIO usage, [9-18](#), [9-19](#), [9-22](#), [9-23](#)
 - handshake MDMA, [6-40](#)

- interrupts, *(continued)*
- initialization, 5-8
 - inputs and outputs, 5-4
 - interrupt handler and DMA
 - synchronization, 6-58
 - MAC configuration, 22-46
 - masking function, 5-7
 - peripheral usage, 5-1, 5-2, 5-3, 5-4 to 5-7, 5-8, 5-15
 - prioritization, 5-4
 - processing, 5-1, 5-8
 - programming examples, 5-13 to 5-15
 - remote wakeup, 22-35
 - request clearing, 5-13
 - reset, 26-7
 - routing, 5-16, 5-17
 - RTC usage, 24-3, 24-6, 24-14
 - service, 27-2
 - shared, 5-4
 - software, 5-3
 - source identification, 5-5, 5-9
 - SPI usage, 17-16, 17-46
 - SPORT usage, 19-37, 19-57, 19-59
 - system, 5-1
 - system control, 5-2
 - timer usage, 10-4, 10-5, 10-14, 10-28
 - UART interrupt priority, 15-12
 - UART usage, 15-11, 15-28
 - wake up core, 5-6
 - interrupt sensitivity (PORTxIO_EDGE)
 - registers, 9-36
 - interrupt service routine, determining
 - source of interrupt, 5-5
 - interrupt status, DMA
 - (DMAx_IRQ_STATUS) registers, 6-71, 6-72
 - interrupt status, EMAC
 - (EMAC_PTP_ISTAT) register, 23-17
 - interrupt status, MDMA
 - (MDMA_yy_IRQ_STATUS) registers, 6-71, 6-72
 - inverse filtering (IFE) bit, 22-65
 - I/O interface to peripheral serial device, 19-4
 - I/O memory space, 1-9
 - I/O pins, general-purpose, 9-14
 - IP checksum, MAC, 22-19
 - IPT (IP type matching value) bit, 23-20
 - IPTM (compare mask of IP frame type field) bit, 23-13
 - IPTOF (offset of frame type IP field) bit, 23-19
 - IP type matching value (IPT) bit, 23-20
 - IP version matching value (IPV) bit, 23-20
 - IPV (IP version matching value) bit, 23-20
 - IPVM (compare mask of IP version field) bit, 23-13
 - IRCLK (internal receive clock select) bit, 19-52, 19-54
 - IrDA mode enable (IREN) bit, 15-31
 - IrDA mode operation, 15-8, 15-9, 15-31
 - IrDA RX polarity change (RPOLC) bit, 15-31
 - IrDA SIR physical layer, 15-5
 - IrDA TX polarity change (TPOLC) bit, 15-31
 - IREN (enable IrDA mode) bit, 15-31
 - IRFS (internal receive frame sync select) bit, 19-31, 19-52, 19-54
 - IR instruction register, B-2, B-4
 - IRPOL bit, 15-10
 - IRQ bit, 10-46
 - IRQ_ENA (interrupt request enable) bit, 10-40, 10-44, 10-46
 - ISR
 - supporting multiple interrupt sources, 5-7, 5-18
 - ISR and multiple interrupt sources, 5-9

Index

ISR for the MMC interrupt, structure, [22-43](#)
issue stop condition (STOP) bit, [16-30](#)
ITCLK (internal transmit clock select) bit, [19-47](#), [19-49](#)
ITFS (internal transmit frame sync select) bit, [19-18](#), [19-31](#), [19-47](#), [19-50](#)
ITHR[15: 0] field, [6-88](#)
ITU-R 601 recommendation, [20-15](#)
ITU-R 656 modes, [1-19](#), [1-20](#), [20-5](#), [20-9](#), [20-11](#), [20-28](#), [20-29](#)
 active video only submode, [20-9](#), [20-10](#)
 DLEN field usage, [20-25](#)
 entire field submode, [20-9](#)
 frame start detect, [20-34](#)
 frame synchronization, [20-11](#)
 SAV codes, [20-31](#)
 vertical blanking interval only submode, [20-9](#), [20-10](#)

J

JTAG port, [1-33](#), [B-1](#), [B-2](#), [B-4](#)
JTAG_SE bit, [9-29](#)

L

L1

 data cache, [2-4](#)
 data memory, [1-9](#)
 data memory subbanks, [2-3](#)
 data SRAM, [2-3](#)
 instruction memory, [1-9](#), [2-2](#)
 memory and core, [4-4](#)
 memory and DMA controller, [6-5](#)
 scratchpad RAM, [1-9](#)
L1DADABL (L1 data bank A memory disable) bit, [25-47](#)
L1 data bank A memory disable (L1DADABL) bit, [25-47](#)

L1 data bank B memory disable (L1DBDABL) bit, [25-47](#)
L1data cache bank select (DCBS) bit, [2-6](#)
L1 data memory configure (DMC) bit, [2-6](#)
L1DBDABL (L1 data bank B memory disable) bit, [25-47](#)
L1IDABL (L1 instruction memory disable) bit, [25-47](#)
L1 instruction memory
 address alignment, [2-2](#)
 subbanks, [2-3](#)
L1 instruction memory disable (L1IDABL) bit, [25-47](#)
LARFS (late receive frame sync) bit, [19-34](#), [19-52](#), [19-55](#)
large descriptor mode, DMA, [6-14](#)
large model mode, DMA, [6-69](#)
late collision detected (RX_LATE) bit, [22-96](#), [22-101](#), [22-104](#)
late collision detected (TX_RETRY) bit, [22-105](#)
late collision error (TX_LATE) bit, [22-105](#), [22-109](#), [22-112](#)
late collisions counter int (TX_LATE_CNT) bit, [22-116](#), [22-118](#)
late collisions detected (TX_RETRY) bit, [22-109](#), [22-112](#)
late collisions (EMAC_TXC_LATECL) register, [22-57](#)
late frame sync, [19-17](#), [19-34](#)
latency
 DAB, [4-10](#)
 DMA, [6-24](#)
 general-purpose ports, [9-13](#)
 powerup, [7-54](#)
 SDC, [7-41](#)
 SDRAM, [7-70](#)
late receive frame sync (LARFS) bit, [19-52](#)

- late transmit frame sync (LATFS) bit, [19-47](#)
- LATFS (late transmit frame sync) bit, [19-34](#), [19-47](#), [19-51](#)
- LB (internal loopback enable) bit, [22-60](#), [22-61](#)
- LCRTE (enable Tx retry on late collision) bit, [22-27](#), [22-60](#), [22-62](#)
- left/right order (R/TRFST) bits, [19-48](#), [19-53](#)
- length error detected (RX_LEN) bit, [22-97](#), [22-102](#), [22-104](#)
- level shifters, [27-11](#)
- lines per frame (PPI_FRAME) register, [20-33](#), [20-34](#)
- line terminations, SPORT, [19-8](#)
- linked list, DMA, [22-122](#)
- little endian byte order, [16-43](#)
- loader, [26-19](#)
- local clock time high (EMAC_PTP_TIMEHI) register, [23-27](#)
- local clock time low (EMAC_PTP_TIMELO) register, [23-26](#)
- LOCAL_TIME_HI (higher 32 bits of local time clock) bit, [23-27](#)
- LOCAL_TIME_LO (lower 32 bits of local time clock) bit, [23-26](#)
- LOCKCNT field, [8-22](#)
- locked transfers, DMA, [4-10](#)
- long response enable (CMD_LRSP_EN) bit, [21-57](#)
- loopback mode enable (LOOP_ENA) bit, [15-23](#)
- loopback operation
 - ethernet MAC port, [22-61](#)
 - general-purpose timers, [10-54](#)
 - PPI port, [20-10](#)
 - UART port, [15-23](#)
- LOOP_ENA (loopback mode enable) bit, [15-23](#)
- losses of carrier detected (TX_LOSS) bit, [22-105](#), [22-110](#)
- loss of carrier int en (TX_LOSS) bit, [22-112](#)
- lost arbitration (LOSTARB) bit, [16-33](#)
- LOSTARB (lost arbitration) bit, [16-33](#), [16-36](#)
- lower 32 bits of local time clock (LOCAL_TIME_LO) bit, [23-26](#)
- lower 32 bits of PPS start time (PPS_STARTLO) bit, [23-35](#)
- lower 32 bits of receive snapshot of local clock time (RX_SNAP_LO) bit, [23-28](#)
- lower 32 bits of target alarm time (ALARM_TIME_LO) bit, [23-32](#)
- lower 32 bits of transmit snapshot of local clock time (TX_SNAP_LO) bit, [23-30](#)
- lower PBS00 half page (PBS00L) register, [26-79](#), [26-80](#)
- lower PBS02 half page (PBS02L) register, [26-85](#)
- low receive frame sync select (LRFS) bit, [19-52](#)
- low transmit frame sync select (LTFS) bit, [19-47](#)
- LRFS (low receive frame sync select) bit, [19-12](#), [19-30](#), [19-32](#), [19-52](#), [19-55](#)
- LSB first (LSBF) bit, [17-36](#)
- LSBF (LSB first) bit, [17-36](#)
- LT_ERR_OVR (horizontal tracking overflow error) bit, [20-30](#), [20-31](#)
- LT_ERR_UNDR (horizontal tracking underflow error) bit, [20-30](#), [20-31](#)
- LTFS (low transmit frame sync select) bit, [19-18](#), [19-30](#), [19-32](#), [19-47](#), [19-50](#)

Index

M

MAC, 1-15, 22-1 to 22-126
addresses, 22-13, 22-47, 22-124
block diagram, 22-3
clocking, 22-4
code examples, 22-120
configuration, 22-45, 22-46, 22-47, 22-48
continuous polling, 22-43
control frames, 22-16
CRC usage, 22-36, 22-37, 22-71
CSMA/CD protocol, 22-3
ethernet event interrupts, 22-38
filters, 22-13
internal interface, 22-7
interrupts, 22-46
IP checksum, 22-19
late collisions, 22-27
magic packet detection, 22-33
management counters, 22-42
MII management interface, 22-8
MMC interrupt, 22-39, 22-43
operation in sleep state, 22-32
peripheral, 22-10
PHYINT interrupt, 22-39
PHY usage, 22-48, 22-124, 22-126
pins, multiplexed, 22-5, 22-6, 22-45
power management states, 22-7
protocol compliance, 22-8
receiving data, 22-49
registers, 22-47, 22-50
RX automatic pad stripping, 22-17
RX receive status priority, 22-20
speculative read, 22-42
station management, 22-9, 22-40, 22-47, 22-48
transmitting data, 22-50
TX transmit status priority, 22-29
type definitions, 22-121
wake up (from hibernate/sleep), 22-30,

22-31
MAC control frames received
(EMAC_RXC_MACCTL) register, 22-55
MAC control frames transmitted
(EMAC_TXC_MACCTL) register, 22-58
MAC DMA
configuration, 22-121
data transfer, 22-49
descriptors, 22-12, 22-23, 22-25, 22-26, 22-122
linked lists, 22-122
receive operation, 22-11
RX buffer structure, 22-18
RX direction error detected, 22-40
RX direction errors, 22-22
transmit operation, 22-23
TX data alignment, 22-27
TX direction error detected, 22-40
TX direction errors, 22-29
MAC frames
aborted frames, 22-16
discarded frames, 22-15
ethernet frame buffer, 22-122
ethernet frame header, 22-49
filter evaluation, 22-14
filtering and reception, 22-13
remote wakeup filters, 22-34
RX frame status, 22-19, 22-39, 22-41
RX IP frame checksum, 22-21
transfer frame protocol, 22-9
TX frame status, 22-28, 22-39, 22-41, 22-42
wakeup frame detected, 22-40
MADDR[6: 0] field, 16-32
MADDR (master mode address) bit, 16-32
magic packet, 22-33
magic packet received status (MPKS) bit, 22-78

- magic packet wakeup enable (MPKE) bit, [22-79](#)
- management counters control (EMAC_MMC_CTL) register, [22-119](#)
- mask block done int (MBDI) bit, [6-83](#)
- mask for PTP control field of event messages (RX_EM) bit, [23-22](#)
- mask for PTP control field of general messages (RX_GN) bit, [23-22](#)
- master mode address (MADDR) bit, [16-32](#)
- master mode enable (MEN) bit, [16-30](#)
- master (MSTR) bit, [17-36](#)
- masters
 - DAB, [4-9](#)
 - PAB, [4-6](#)
- master transfer complete (MCOMP) bit, [16-40](#)
- master transfer direction (MDIR) bit, [16-30](#)
- master transfer error interrupt mask (MERRM) bit, [16-40](#)
- master transfer error (MERR) bit, [16-40](#)
- master transfer in progress (MPROG) bit, [16-33](#)
- MAXCIE (max count int en) bit, [13-21](#)
- MAXCII (max count interrupt identifier) bit, [13-22](#)
- max count int en (MAXCIE) bit, [13-21](#)
- max count interrupt identifier (MAXCII) bit, [13-22](#)
- maximal count (CNT_MAX) register, [13-19, 13-26](#)
- MBDI bit, [6-40](#)
- MBDI (mask block done int) bit, [6-83](#)
- MCCRM[1: 0] field, [19-64](#)
- MCDRXPE (multichannel DMA receive packing) bit, [19-64](#)
- MCDTXPE (multichannel DMA transmit packing) bit, [19-64](#)
- MCMEN (multichannel frame mode enable) bit, [19-16, 19-64](#)
- MCOMP (master transfer complete) bit, [16-40, 16-42](#)
- MCOMP (master transfer complete interrupt mask) bit, [16-40](#)
- MCOMP (master transfer complete interrupt mask (MCOMP)) bit, [16-40](#)
- MDCDIV[5: 0] field, [22-89](#)
- MDIO station management interface, [22-8](#)
- MDIR (master transfer direction) bit, [16-30, 16-31](#)
- MDMA channels, [6-6](#)
- MDMA controllers, [6-6](#)
- MDMA_ROUND_ROBIN_COUNT field, [6-47, 6-89](#)
- MDMA_ROUND_ROBIN_PERIOD field, [6-47, 6-89, 6-90](#)
- MDMA_yy_CONFIG (DMA configuration) registers, [6-67](#)
- MDMA_yy_CURR_ADDR (current address) registers, [6-74](#)
- MDMA_yy_CURR_DESC_PTR (current descriptor pointer) registers, [6-81](#)
- MDMA_yy_CURR_X_COUNT (current inner loop count) registers, [6-76](#)
- MDMA_yy_CURR_Y_COUNT (current outer loop count) registers, [6-78](#)
- MDMA_yy_IRQ_STATUS (interrupt status) registers, [6-71, 6-72](#)
- MDMA_yy_NEXT_DESC_PTR (next descriptor pointer) registers, [6-80](#)
- MDMA_yy_PERIPHERAL_MAP (peripheral map) registers, [6-66](#)
- MDMA_yy_START_ADDR (start address) registers, [6-74](#)
- MDMA_yy_X_COUNT (inner loop count) registers, [6-75](#)

Index

- MDMA_yy_X_MODIFY (inner loop address increment) registers, 6-77
- MDMA_yy_Y_COUNT (outer loop count) registers, 6-78
- MDMA_yy_Y_MODIFY (outer loop address increment) registers, 6-79
- measurement report, general-purpose timers, 10-24, 10-26, 10-27
- memory, 2-1 to 2-7
 - accesses to internal, 2-1
 - architecture, 1-7, 2-1
 - asynchronous, 2-1, 7-2, 27-5
 - boot ROM, 2-4
 - configurations, 1-8
 - external, 1-9, 2-1, 2-5, 7-6
 - flash, 1-9
 - internal, 1-9, 7-31
 - I/O space, 1-9
 - L1 data, 1-9, 2-3, 4-4
 - L1 data cache, 2-4, 4-4
 - L1 instruction, 1-9, 2-2, 2-3, 4-4
 - L1 scratchpad, 1-9, 4-4
 - off-chip, 1-8, 1-9
 - on-chip, 1-7, 1-9
 - OTP, 1-10
 - SDRAM, 2-1, 7-26
 - SPORT transfers, 19-37
 - unpopulated, 7-9
- memory conflict, DMA, 6-48
- memory DMA, 1-12, 6-6
 - bandwidth, 6-43
 - buffers, 6-8
 - channels, 6-7
 - descriptor structures, 6-62
 - handshake operation, 6-8
 - priority, 6-46
 - scheduling, 6-46
 - timing, 6-44
 - transfers, 4-11, 4-12, 6-2, 6-5, 6-8
 - word size, 6-7
- memory map, 2-2, 7-3
- memory-mapped GPIO registers, 9-26
- memory-mapped registers. *See* MMRs
- memory-to-memory transfers, 6-7
- MEN (master mode enable) bit, 16-30, 16-32
- MERR (master transfer error) bit, 16-40, 16-41
- MERRM (master transfer error interrupt mask) bit, 16-40
- message filter offset (EMAC_PTP_FOFF) register, 23-19
- message filter value 1 (EMAC_PTP_FV1) register, 23-20
- message filter value 2 (EMAC_PTP_FV2) register, 23-21
- message filter value 3 (EMAC_PTP_FV3) register, 23-22
- MFD[3: 0] field, 19-20, 19-64
- MII protocol, 22-5, 22-6, 22-8, 22-47
- MINCIE (minimum count int en) bit, 13-21
- MINCII (min count interrupt identifier) bit, 13-22
- min count interrupt identifier (MINCII) bit, 13-22
- minimal count (CNT_MIN) register, 13-19, 13-26
- minimum count int en (MINCIE) bit, 13-21
- minutes[5: 0] field, 24-20, 24-22
- minutes event flag bit, 24-21
- minutes int en bit, 24-20
- MISO enable (EMISO) bit, 17-36
- MISO pin, 17-5, 17-11, 17-14, 17-15, 17-20
- MMC, continuous polling, 22-43
- MMC counter enable (MMCE) bit, 22-119

- MMC counter interrupt status (MMCINT) bit, [22-91](#)
- MMCE (MMC counter enable) bit, [22-42](#), [22-119](#)
- MMC interrupt, [22-39](#), [22-43](#)
- MMCINT (MMC counter interrupt status) bit, [22-91](#)
- MMC RX int en (EMAC_MMC_RIRQE) register, [22-114](#)
- MMC RX interrupt status (EMAC_MMC_RIRQS) register, [22-112](#)
- MMC TX interrupt status (EMAC_MMC_TIRQS) register, [22-114](#)
- MMRs, [1-9](#)
 - address range, [A-2](#)
 - for PPI, [20-25](#)
 - memory-related, [2-5](#)
 - width, [A-2](#)
- mode fault error, [17-16](#), [17-40](#)
- mode fault error (MODF) bit, [17-39](#)
- modes
 - broadcast, [17-8](#), [17-14](#), [17-15](#)
 - multichannel, [19-14](#)
 - SPI port, [17-14](#), [17-15](#), [17-17](#), [17-19](#)
 - SPORT (serial port), [19-10](#)
 - UART port, [15-15](#), [15-17](#)
- MODF (mode fault error) bit, [17-39](#), [17-40](#)
- module (PTP_TSYNC) registers, [23-11](#)
- module registers (PTP_TSYNC), [23-11](#)
- MOSI pin, [17-5](#), [17-11](#), [17-14](#), [17-15](#), [17-20](#)
- moving data, serial port, [19-37](#)
- MPKE (magic packet wakeup enable) bit, [22-33](#), [22-79](#)
- MPKS (magic packet received status) bit, [22-33](#), [22-78](#)
- MPROG (master transfer in progress) bit, [16-33](#), [16-36](#)
- MRS command, [7-33](#), [7-47](#)
- MSEL (multiple select) field, [8-3](#), [8-4](#), [8-21](#)
- MSTR (master) bit, [17-35](#), [17-36](#)
- μ-law companding, [19-22](#), [19-27](#)
- multibank operation, [7-42](#)
- multicast frame detected (TX_MULTI) bit, [22-109](#)
- multicast frame int en (RX_MULTI) bit, [22-95](#), [22-101](#), [22-104](#)
- multicast frames received OK counter int en (RX_MULTI_CNT) bit, [22-115](#)
- multicast frames received OK counter int (RX_MULTI_CNT) bit, [22-113](#)
- multicast frames received ok (EMAC_RXC_MULTI) register, [22-54](#)
- multicast frames transmitted OK counter int en (TX_MULTI_CNT) bit, [22-118](#)
- multicast frames transmitted OK counter int (TX_MULTI_CNT) bit, [22-116](#)
- multicast frames xmitted ok (EMAC_TXC_MULTI) register, [22-58](#)
- multicast hash table high (EMAC_HASHHI) register, [22-68](#)
- multicast hash table low (EMAC_HASHLO) register, [22-68](#)
- multichannel DMA receive packing (MCDRXPE) bit, [19-64](#)
- multichannel DMA transmit packing (MSDTXPE) bit, [19-64](#)
- multichannel frame, [19-19](#)
- multichannel frame delay field, [19-20](#)
- multichannel frame mode enable (MCMEM) bit, [19-64](#)
- multichannel mode, SPORT, [19-14](#), [19-16](#), [19-18](#)

Index

multichannel operation, SPORT, [19-14](#) to [19-23](#)
multiple collision frames counter int en (TX_MCOLL_CNT) bit, [22-118](#)
multiple collision frames counter int (TX_MCOLL_CNT) bit, [22-116](#)
multiple collision frames (EMAC_TXC_GTICOL) register, [22-57](#)
multiple interrupt sources, [5-9](#)
multiple select (MSEL) bit, [8-21](#)
multiple slave SPI systems, [17-8](#)
multiplexed SDRAM addressing scheme figure, [7-27](#)
multiplexing, MII and RMII, [22-5](#)
multiplexing, pins, [9-1](#)
MVIP-90 compatibility mode, [19-24](#)
MW_EN (SDIO interrupt moving window enable) bit, [21-78](#)

N

NAK (not acknowledge) bit, [16-26](#), [16-27](#)
NDPH bit, [6-20](#)
NDPL bit, [6-20](#)
NDSIZE[3: 0] field, [6-15](#), [6-67](#), [6-69](#)
legal values, [6-31](#)
next descriptor pointer (NDP) field, [6-81](#), [6-82](#)
next descriptor pointer registers (DMAx_NEXT_DESC_PTR), [6-80](#)
(MDMA_yy_NEXT_DESC_PTR), [6-80](#)
NINT (pending int) bit, [15-28](#), [15-29](#)
NMI (reflects detection of NMI) bit, [25-55](#)
NMI_RST_BMODE_SE bit, [9-29](#)
no carrier int en (TX_CRS) bit, [22-112](#)
no carrier (TX_CRS) bit, [22-105](#), [22-110](#)
NONGPIO_DRIVE (non-GPIO ports drive strength control) register, [9-28](#)

NONGPIO_HYSTERESIS (non-GPIO hysteresis) register, [9-29](#)
non-GPIO hysteresis (NONGPIO_HYSTERESIS) register, [9-29](#)
non-GPIO ports drive strength control (NONGPIO_DRIVE) register, [9-28](#)
NOP (no operation) command, [7-53](#)
NOPREBOOT, [25-47](#), [25-48](#)
normal frame sync mode, [19-34](#)
normal timing, serial port, [19-34](#)
not acknowledge (NAK) bit, [16-26](#)
NTSC systems, [20-6](#)
number of 32-bit words remaining (FIFO_COUNT) bit, [21-73](#)
number of bytes still to be transferred (DATA_COUNT) bit, [21-64](#)
number of bytes to be transferred (DATA_LENGTH) bit, [21-61](#)

O

octets received all counter int en (RX_ALLO_CNT) bit, [22-115](#)
octets received all counter int (RX_ALLO_CNT) bit, [22-113](#)
octets received all (EMAC_RXC_ALLOCT) register, [22-55](#)
octets received OK counter int en (RX_OCTET_CNT) bit, [22-115](#)
octets received OK counter int (RX_OCTET_CNT) bit, [22-113](#)
octets received ok (EMAC_RXC_OCTET) register, [22-53](#)
octets transmitted all counter int en (TX_ALLO_CNT) bit, [22-118](#)
octets transmitted all counter int (TX_ALLO_CNT) bit, [22-116](#)

- octets transmitted all
 - (EMAC_TXC_ALLOC) register, [22-58](#)
- octets transmitted OK counter int en
 - (TX_OCTET_CNT) bit, [22-118](#)
- octets transmitted OK counter int
 - (TX_OCTET_CNT) bit, [22-116](#)
- octets transmitted ok
 - (EMAC_TXC_OCTET) register, [22-57](#)
- OE (overrun error) bit, [15-24](#)
- off-chip
 - bus connections, [4-8](#)
 - infrared driver, [15-8](#)
 - line drivers, [15-6](#)
 - memory, [1-8](#), [1-9](#), [4-10](#)
 - peripherals, [6-2](#)
 - signals, [9-18](#)
- off-core accesses, [4-4](#)
- offset of frame type EFT field (EFTOF) bit, [23-19](#)
- offset of frame type IP field (IPTOF) bit, [23-19](#)
- OFFSET (offset to local time) bit, [23-25](#)
- offset of PTP control field (PTPCOF) bit, [23-19](#)
- offset of UDP event port field (UEVOF) bit, [23-19](#)
- offsets, DMA descriptor elements, [6-16](#)
- offset to local time (OFFSET) bit, [23-25](#)
- OIE (overflow int en) bit, [6-83](#)
- OI (overflow interrupt generated) bit, [6-83](#)
- on-chip
 - buses, [4-8](#)
 - internal voltage regulator, [1-27](#)
 - I/O devices, [1-9](#)
 - memory, [1-7](#), [1-9](#)
 - peripherals, [1-10](#), [6-2](#)
 - PLL, [1-28](#)
 - SDRAM controller, [7-45](#)
- ones word (BK_ONES) register, [26-78](#)
- one-time-programmable (OTP) memory, [1-10](#)
- open drain drivers, [17-2](#)
- open drain outputs, [17-14](#)
- open memory page, [7-33](#)
- operating mode (EMAC_OPMODE) register, [22-60](#)
- operating modes, [8-7](#)
 - active, [1-29](#), [8-8](#)
 - deep sleep, [1-29](#), [8-9](#)
 - full-on, [1-28](#), [8-8](#)
 - hibernate state, [1-30](#), [8-10](#)
 - PPI port, [20-4](#)
 - sleep, [1-29](#), [8-9](#)
 - transition, [8-10](#), [8-12](#)
- OPT_ACC_VIO_ERROR error code, [3-22](#)
- OPT error codes
 - OPT_ACC_VIO_ERROR, [3-22](#)
 - OTP_DATA_MULT_ERROR, [3-22](#)
 - OTP_DATA_SB_WARN, [3-22](#)
 - OTP_ECC_MULT_ERROR, [3-22](#)
 - OTP_ECC_SB_WARN, [3-22](#)
 - OTP_MASTER_ERROR, [3-22](#)
 - OTP_PREV_WR_ERROR, [3-22](#)
 - OTP_READ_ERROR, [3-22](#)
 - OTP_SB_DEFECT_ERROR, [3-22](#)
 - OTP_SUCCESS, [3-22](#)
 - OTP_WRITE_ERROR, [3-22](#)
- optimization, of DMA performance, [6-40](#)
- oscilloscope probes, [27-12](#)
- OTP_BMODEx_DIS bits, [26-84](#)
- OTP_DATA_MULT_ERROR error code, [3-22](#)
- OTP_DATA_SB_WARN error code, [3-22](#)
- OTP_DATAx registers, [3-3](#)
- OTP_EBIU_AMG bit, [26-81](#)
- OTP_EBIU_POWERON_DUMMY_WRITE bit, [26-85](#)

Index

- OTP_EBIU_SDBCTL bit, [26-85](#)
- OTP_EBIU_SDGCTL bit, [26-85](#)
- OTP_EBIU_SDRRC bit, [26-85](#)
- OTP_ECC_MULT_ERROR error code, [3-22](#)
- OTP_ECC_SB_WARN error code, [3-22](#)
- OTP_INVALID bit, [26-79](#)
- OTP_LOAD_PBSxH/L bits, [26-79](#)
- OTP_MASTER_ERROR error code, [3-22](#)
- OTP memory
 - bfrom_OtpCommand(), [3-10](#)
 - BFROM_OTP_READ, [3-17](#)
 - bfrom_OtpRead(), [3-10](#)
 - BFROM_OTP_WRITE, [3-18](#)
 - bfrom_OtpWrite(), [3-10](#)
 - BFROM_TOP_COMMAND, [3-15](#)
 - error correction, [3-7](#)
 - map, [3-3](#)
 - overview, [3-1](#)
- OTP_PLL_CTL bit, [26-80](#)
- OTP_PLL_DIV bit, [26-80](#)
- OTP_PREV_WR_ERROR error code, [3-22](#)
- OTP_READ_ERROR error code, [3-22](#)
- OTP_RESETOUT_HWAIT bit, [26-79](#)
- OTP_SB_DEFECT_ERROR error code, [3-22](#)
- OTP secrets enable (OTPSEN) bit, [25-47](#)
- OTPSEN (OTP secrets enable) bit, [25-47](#)
- OTP_SET_FCTL bit, [26-81](#)
- OTP_SET_MODE bit, [26-81](#)
- OTP_SET_PLL bit, [26-79](#)
- OTP_SET_VR bit, [26-79](#)
- OTP_SPI_BAUD bit, [26-80](#)
- OTP_SPI_FASTREAD bit, [26-80](#)
- OTP_START_PAGE bit, [26-83](#)
- OTP_SUCCESS error code, [3-22](#)
- OTP_TIMING (OTP timing) register, [3-14](#)
- OTP timing (OTP_TIMING) register, [3-14](#)
- OTP_TP_x bits, [3-14](#)
- OTP_USB_CALIB, [26-83](#)
- OTP_VR_CTL bit, [26-80](#)
- OTP_WRITE_ERROR error code, [3-22](#)
- OUT_DELAY bit, [8-21](#)
- OUT_DIS bit, [10-38](#), [10-45](#), [10-57](#)
- OUT_DIS (output pad disable) bit, [10-40](#)
- outer loop address increment registers
 - (DMA_x_Y_MODIFY), [6-79](#)
 - (MDMA_{yy}_Y_MODIFY), [6-79](#)
- outer loop count (DMA_x_Y_COUNT) registers, [6-78](#)
- outer loop count (MDMA_{yy}_Y_COUNT) registers, [6-78](#)
- out-of-range length field counter int en (RX_ORL_CNT) bit, [22-115](#)
- out-of-range length field counter int (RX_ORL_CNT) bit, [22-113](#)
- out-of-range length field detected (RX_RANGE) bit, [22-95](#), [22-101](#), [22-104](#)
- out of range length field (EMAC_RXC_LNERRO) register, [22-54](#)
- output delay bit, [8-21](#)
- output pad disable, timer, [10-12](#)
- output pad disable (OUT_DIS) bit, [10-40](#)
- overflow int en (OIE) bit, [6-83](#)
- overflow interrupt, DMA, [6-40](#)
- overflow interrupt generated (OI) bit, [6-83](#)
- overrun error (OE) bit, [15-24](#)

P

- PAB, 4-6
 - arbitration, 4-6
 - bus agents (masters, slaves), 4-6
 - clocking, 8-1
 - and EBIU, 7-5
 - performance, 4-7
- PACK_EN (packing mode enable) bit, 20-26, 20-27
- packing, serial port, 19-23
- packing mode enable (PACK_EN) bit, 20-26
- page hit/miss, 7-41, 7-50
- page size, 7-62
- PAL systems, 20-6
- PAM (pass all multicast mode) bit, 22-60, 22-65
- parallel peripheral interface. *See* PPI
- parity enable (PEN) bit, 15-21
- parity error (PE) bit, 15-24
- partial array self-refresh (PASR) bits, 7-65
- PASR feature, 7-27
- PASR (partial array self-refresh) bits, 7-65, 7-67
- pass all multicast mode (PAM) bit, 22-65
- pass bad frames (PBF) bit, 22-65
- pass control frames (PCF) bit, 22-75
- pass short frames (PSF) bit, 22-64
- PAUSE MAC control frames received counter int en (RX_PAUSE_CNT) bit, 22-115
- PAUSE MAC control frames received counter interrupt (RX_PAUSE_CNT) bit, 22-113
- PAUSE MAC ctrl frames received (EMAC_RXC_PAUSE) register, 22-55
- PBF (pass bad frames) bit, 22-60, 22-65
- PBSxH (upper PBSx half page) registers, 26-81, 26-82, 26-83, 26-84
- PBSxL (lower PBSx half page) registers, 26-79, 26-80, 26-85
- PC133 SDRAM controller, 1-12
- PCF (pass control frames) bit, 22-15, 22-75
- PD_DAT3 (RSI_DATA3 pull-down enable) bit, 21-78
- PDWN (power down) bit, 8-21
- pend enable (CMD_PEND_EN) bit, 21-57
- pending int (NINT) bit, 15-29
- PEN (parity enable) bit, 15-21
- PE (parity error) bit, 15-24
- performance
 - DAB, 4-10
 - DCB, 4-10
 - DEB, 4-10, 4-11
 - DMA, 6-42
 - EAB, 4-11
 - general-purpose ports, 9-13
 - memory DMA, 6-43
 - memory DMA transfers, 4-11, 4-12
 - optimization, DMA, 6-40
 - PAB, 4-7
 - SDRAM, 7-29
- PERIOD_CNT bit, 10-12, 10-20, 10-24, 10-25, 10-44
- PERIOD_CNT (period count) bit, 10-40
- period count (PERIOD_CNT) bit, 10-40
- period value[15: 0] field, 11-6
- period value[31: 16] field, 11-6
- peripheral
 - bus errors generated by SPORT, 19-38
 - DMA, 6-5
 - DMA channels, 6-41
 - DMA transfers, 6-2
 - error interrupts, 6-72
 - interrupt request lines, 5-15
 - supporting interrupts, 5-1
- peripheral access bus. *See* PAB

Index

- peripheral DMA start address registers, [6-74](#)
- peripheral ID (RSI_PID) bit, [21-80](#)
- peripheral interrupts, [5-2](#), [5-3](#), [5-4](#) to [5-7](#)
- peripheral map registers
 - (DMAx_PERIPHERAL_MAP), [6-66](#)
 - (MDMA_yy_PERIPHERAL_MAP), [6-66](#)
- peripheral pins, default configuration, [9-14](#)
- peripherals, [1-5](#)
 - and buses, [1-6](#)
 - compatible with SPI, [17-3](#)
 - and DMA controller, [6-31](#)
 - DMA support, [1-6](#)
 - enabling, [9-3](#)
 - interrupt generated by, [5-8](#)
 - interrupts, clearing, [5-13](#)
 - level-sensitivity of interrupts, [5-15](#)
 - mapping to DMA, [6-6](#), [6-104](#)
 - multiplexing, [9-1](#)
 - switching from DMA to non-DMA, [6-73](#)
 - timing, [4-4](#)
 - used to wake from idle, [5-6](#)
- PfX pins, [9-16](#), [17-7](#)
- phase locked loop. *See* [PLL](#)
- PHY, [22-4](#)
 - configuring, [22-48](#)
 - control routines, [22-124](#)
 - initialization, minimum requirements, [22-126](#)
 - read access, [22-126](#)
- PHYAD[4: 0] field, [22-72](#)
- PHYCLKOE bit, [22-46](#)
- PHY error int en (RX_PHY) bit, [22-96](#), [22-101](#), [22-104](#)
- PHYIE (PHYINT int en) bit, [22-89](#)
- PHYINT int en (PHYIE) bit, [22-89](#)
- PHYINT interrupt, [22-39](#)
- PHYINT interrupt status (PHYINT) bit, [22-91](#)
- PHYINT (PHYINT interrupt status) bit, [22-91](#)
- PHY/PF15 wake enable (PHYWE) bit, [8-23](#)
- PHYWE (PHY/PF15 wake enable) bit, [8-23](#)
- pin information, [27-1](#)
- pins, [27-1](#)
 - GPIO, [9-14](#)
 - MAC, [22-6](#)
 - multiplexing, [9-1](#)
 - SDRAM, [7-53](#)
 - terminations for SPORT, [19-8](#)
 - unused, [27-14](#)
- pin state during SDC commands (table), [7-46](#)
- pin status (PS) bit, [6-83](#)
- pipeline, lengths of, [6-52](#)
- pipelining
 - DMA requests, [6-37](#)
 - SDC supported, [7-72](#)
- PJSE bit, [9-30](#), [9-31](#), [9-32](#)
- PLL, [8-1](#) to [8-41](#)
 - active mode, [8-8](#)
 - architecture, [8-2](#), [8-3](#)
 - BYPASS bit, [8-9](#)
 - clock derivation, [4-4](#)
 - clock derivation, CCLK, [8-3](#)
 - clock derivation, SCLK, [8-1](#), [8-3](#)
 - clock derivation, SDRAM, [8-10](#)
 - clock ratios and dividers, [8-3](#), [8-6](#)
 - configuration and control, [8-1](#), [8-3](#), [8-10](#), [8-11](#)
 - deep sleep mode, [8-9](#)
 - DMA access, [8-8](#), [8-9](#)
 - DPMC (dynamic power management controller), [8-2](#), [8-7](#)
 - enabling/disabling the PLL, [8-11](#)

- PLL, *(continued)*
- hibernate state, 8-10
 - multiplier select (MSEL) field, 8-3
 - operating modes, 8-7, 8-8, 8-9, 8-10, 8-13
 - PDWN bit, 8-11
 - performance mode, 8-8
 - PLL_OFF bit, 8-11
 - PLL status, 8-7
 - power domains, 8-15
 - power savings, 8-7
 - power to, applying/removing, 8-11
 - registers, 8-19
 - RTC interrupt, 8-9
 - sleep mode, 8-9
 - STOPCK bit, 8-11
 - voltage control, 8-7
 - PLL control (PLL_CTL) register, 8-3, 8-4, 8-20, 8-21
 - PLL_CTL (PLL control) register, 8-3, 8-4, 8-20, 8-21
 - PLL divide (PLL_DIV) register, 8-21
 - PLL divide register, 4-4
 - PLL_DIV (PLL divide) register, 8-5, 8-20, 8-21
 - PLL_LOCKCNT (PLL lock count) register, 8-20, 8-22
 - PLL lock count (PLL_LOCKCNT) register, 8-22
 - PLL_LOCKED bit, 8-22
 - PLL_OFF bit, 8-21
 - PLL registers, 8-19
 - PLL_STAT (PLL status) register, 8-20, 8-22
 - PLL status (PLL_STAT) register, 8-22
 - PLL VCO frequency, changing, 7-44
 - PMAP[3: 0] field, 6-5, 6-44, 6-66
 - PMAP bit, 6-66
 - polarity, GPIO, 9-18
 - polarity change (POLC) bit, 20-26
 - POLC (polarity change) bit, 20-4, 20-25, 20-26
 - polling DMA registers, 6-51
 - POLS bit, 20-4, 20-25, 20-26
 - PORT_CFG (port configuration) bits, 20-4, 20-26, 20-28
 - port configuration (PORT_CFG) bit, 20-26
 - port connection, SPORT, 19-7
 - PORT_DIR (direction) bit, 20-4, 20-26, 20-28
 - port enable (PORT_EN) bit, 20-26
 - PORT_EN (enable) bit, 20-26, 20-29
 - PORTF_FER (function enable) register, 9-10
 - port F/G/H
 - GPIO, 9-14
 - interrupt A channel, 9-22
 - MAC operations, 22-45
 - peripherals, 9-2, 9-3, 9-5
 - pins, 9-3, 17-38
 - pins, test access, B-2
 - structure, 9-4, 9-5, 9-6
 - port F/G/H hysteresis control (PORTx_HYSTERESIS) registers, 9-26
 - port F/G/H hysteresis (PORTx_HYSTERESIS) registers, 9-26, 9-27
 - port FG/H hysteresis (PORTx_HYSTERESIS) registers, 9-27
 - port F/G/H multiplexer control (PORTx_MUX) registers, 9-30, 9-31, 9-32
 - PORT_PREF0 bit, 2-6
 - PORT_PREF0 (DAG0 port preference) bit, 2-6
 - PORT_PREF1 bit, 2-6

Index

- PORT_PREF1 (DAG1 port preference)
 - bit, 2-6
- port width, PPI, 20-27
- PORTx_FER (function enable) registers,
 - 9-3, 9-4, 9-10, 9-14, 9-18, 9-32
- PORTx_HYSTERESIS (port F/G/H hysteresis) registers, 9-26, 9-27
- PORTxIO_BOTH (GPIO set on both edges) registers, 9-37
- PORTxIO_CLEAR (GPIO clear) registers, 9-35
- PORTxIO_DIR (GPIO direction)
 - registers, 9-33, 14-40, 14-41, 14-42, 14-43, 14-46, 14-47, 14-48, 14-49
- PORTxIO_EDGE (interrupt sensitivity) registers, 9-36
- PORTxIO (GPIO data) registers, 9-34
- PORTxIO_INEN (GPIO input enable) registers, 9-18, 9-33
- PORTxIO_MASKA/B_CLEAR (GPIO mask interrupt clear) registers, 9-21, 9-41, 9-42
- PORTxIO_MASKA/B (GPIO mask int) registers, 9-37, 9-38
- PORTxIO_MASKA/B_SET (GPIO mask interrupt set) registers, 9-39, 9-40
- PORTxIO_MASKA/B_TOGGLE (GPIO mask interrupt toggle) registers, 9-43, 9-44
- PORTxIO_POLAR (GPIO polarity) registers, 9-36
- PORTxIO_SET (GPIO set) registers, 9-34
- PORTxIO_TOGGLE (GPIO toggle) registers, 9-35
- PORTx_MUX (port F/G/H multiplexer control) registers, 9-30, 9-31, 9-32
- PORTx_MUX (port multiplexer control) registers, 9-3, 9-30, 9-31, 9-32
- PORTx_MUX (port multiplexer control) registers, 9-4, 9-10
- power
 - dissipation, 8-15
 - domains, 8-15
 - plane, 27-10
 - power down (PDWN) bit, 8-21
 - power management, 1-28, 8-1 to 8-30
 - power on (PWR_ON) bit, 21-54
 - power save enable (PWR_SV_EN) bit, 21-55
 - power-up
 - SDRAM, 7-47, 7-70
 - sequence mode, 7-69
 - start delay, 7-69
 - start enable, 7-70
 - power-up latency, SDC, 7-54
 - power-up sequence mode (PSM) bit, 7-65
 - power-up sequence start enable (PSSE) bit, 7-65
 - power-up start delay (PUPSD) bit, 7-65
- PPI, 1-19, 20-2 to 20-37
 - active video only mode, 20-10
 - clock input, 20-3
 - control byte sequences, 20-8
 - control signal polarities, 20-25
 - data input modes, 20-14 to 20-16
 - data movement, 20-9
 - data output modes, 20-17 to 20-18
 - data width, 20-25
 - delay before starting, 20-32
 - DMA operations, 20-13, 20-22, 20-35, 20-36, 20-37
 - edge-sensitive inputs, 20-20
 - enabling, 20-29, 20-36
 - entire field mode, 20-9

- PPI, *(continued)*
- FIFO, [20-31](#)
 - frame, number of lines per, [20-33](#)
 - frame start detect, [20-34](#)
 - frame sync, internal/external, [20-15](#), [20-16](#), [20-17](#), [20-18](#), [20-19](#), [20-21](#)
 - frame sync modes, [20-15](#), [20-16](#), [20-18](#), [20-19](#)
 - frame track error, [20-30](#), [20-33](#)
 - general-purpose mode, [13-3](#), [13-5](#), [13-8](#), [20-12](#), [20-14](#), [20-19](#)
 - general-purpose timer operations, [10-56](#), [20-20](#)
 - hardware signalling, [20-15](#)
 - horizontal tracking, [20-31](#)
 - ITU-R 601 recommendation, [20-15](#)
 - ITU-R 656 mode, [20-5](#), [20-9](#), [20-11](#)
 - loopback feature, [20-10](#)
 - no frame syncs modes, [20-15](#), [20-17](#)
 - operating modes, [20-4](#), [20-25](#)
 - port width, [20-27](#)
 - preamble, [20-7](#)
 - programming model, [20-21](#)
 - registers, [20-25](#), [20-36](#)
 - samples, number of, [20-32](#)
 - synchronization, [20-11](#), [20-13](#), [20-20](#), [20-32](#)
 - timer pins, [20-20](#)
 - transfers, [20-18](#), [20-29](#)
 - valid data detection, [20-15](#)
 - vertical blanking interval only mode, [20-10](#)
 - video, interlaced, [20-6](#)
 - video, progressive, [20-6](#)
 - video frame partitioning, [20-7](#)
 - video processing/streams, [20-5](#), [20-8](#)
- PPI_CLK cycle count, [20-32](#)
- PPI_CLK pin, [20-3](#)
- PPI_CLK signal, [20-25](#)
- PPI_CONTROL (PPI control) register, [20-25](#), [20-26](#)
- PPI control (PPI_CONTROL) register, [20-25](#), [20-26](#)
- PPI_COUNT field, [20-33](#)
- PPI_COUNT (PPI transfer count) register, [20-32](#), [20-33](#)
- PPI delay count (PPI_DELAY) register, [20-32](#)
- PPI_DELAY (delay count) register, [20-32](#)
- PPI_DELAY field, [20-32](#)
- PPI_DELAY (PPI delay count) register, [20-32](#)
- PPI_FRAME field, [20-34](#)
- PPI_FRAME (PPI lines per frame) register, [20-33](#), [20-34](#)
- PPI_FSx signals, [20-25](#), [20-31](#)
- PPI lines per frame (PPI_FRAME) register, [20-34](#)
- PPI_STATUS (PPI status) register, [20-29](#), [20-30](#)
- PPI status (PPI_STATUS) register, [20-30](#)
- PPI transfer count (PPI_COUNT) register, [20-33](#)
- PPSEN (pulse-per-second control) bit, [23-13](#)
- PPS period (EMAC_PTP_PPS_PERIOD) register, [23-37](#)
- PPS_PERIOD (PPS signal period) bit, [23-37](#)
- PPS signal period (PPS_PERIOD) bit, [23-37](#)
- PPS start high (EMAC_PTP_PPS_STARTHI) register, [23-36](#)
- PPS_STARTHI (higher 32 bits of PPS start time) bit, [23-36](#)
- PPS_STARTLO (lower 32 bits of PPS start time) bit, [23-35](#)

Index

- PPS start low
 - (EMAC_PTP_PPS_STARTLO) register, [23-35](#)
- preamble, PPI, [20-7](#)
- preboot, controlled by OTP programming, [26-3](#)
- preboot routine, [26-9](#)
- precharge command, [7-34](#), [7-50](#)
- PREN bit, [24-22](#)
- PRESCALE field, [16-25](#)
- prescaler, RTC, [24-2](#)
- prescaler enable (RTC_PREN) register, [24-4](#), [24-5](#), [24-19](#)
- PRESCALE value, [16-4](#)
- prioritization, DMA, [6-44](#) to [6-50](#)
- prioritization, interrupt, [5-4](#)
- priority, DMA operation, [6-46](#)
- private instructions, [B-4](#)
- private OTP memory map, [3-5](#)
- probes, oscilloscope, [27-12](#)
- processor
 - block diagram, [1-7](#)
 - dynamic power management, [8-1](#)
 - test features, [B-1](#)
- programmable timing characteristics, EBIU, [7-11](#)
- program Pxn bit, [9-34](#)
- progressive video, [20-6](#)
- promiscuous mode (PR) bit, [22-65](#)
- PR (promiscuous mode) bit, [22-60](#), [22-65](#)
- PSF (pass short frames) bit, [22-60](#), [22-64](#)
- PSM (power-up sequence mode) bit, [7-47](#), [7-55](#), [7-65](#), [7-69](#)
- PS (pin status) bit, [6-83](#)
- PSSE (power-up sequence start enable) bit, [7-44](#), [7-47](#), [7-55](#), [7-65](#), [7-70](#)
- PSSE (slave select enable) bit, [17-35](#), [17-36](#)
- PTP clock source (CKS) bit, [23-13](#)
- PTPCM (compare mask of PTP control field) bit, [23-13](#)
- PTPCOF (offset of PTP control field) bit, [23-19](#)
- PTP int en (EMAC_PTP_IE) register, [23-15](#)
- PTP_TSYNC module enable (EN) bit, [23-13](#)
- PTP_TSYNC module registers, [23-11](#)
- public instructions, [B-4](#)
- public JTAG scan instructions, [B-6](#)
- public OTP memory map, [3-4](#)
- PU_DAT3 (RSI_DATA3 pull-up enable) bit, [21-78](#)
- PU_DAT (pull-up enable) bit, [21-78](#)
- pull-up enable (PU_DAT) bit, [21-78](#)
- PULSE_HI bit, [10-14](#), [10-16](#), [10-24](#), [10-40](#), [10-44](#)
- PULSE_HI toggle mode, [10-15](#)
- pulse-per-second control (PPSEN) bit, [23-13](#)
- pulse width count and capture mode. *See* WDTM_CAP mode
- pulse width modulation mode. *See* PWM_OUT mode
- pulse width modulation mode. *See* PWM_OUT mode
- pulse width modulator, [1-23](#)
- PUPSD (power-up start delay) bit, [7-65](#), [7-69](#)
- PWM_AH output signal for S/W observation (PWM_AH) bit, [14-49](#)
- PWM_AH (PWM_AH output signal for S/W observation) bit, [14-49](#)
- PWM_AL output signal for S/W observation (PWM_AL) bit, [14-49](#)
- PWM_AL (PWM_AL output signal for S/W observation) bit, [14-49](#)
- PWM_BH output signal for S/W observation (PWM_BH) bit, [14-49](#)
- PWM_BH (PWM_BH output signal for S/W observation) bit, [14-49](#)

- PWM_BL output signal for S/W observation (PWM_BL) bit, [14-49](#)
- PWM_BL (PWM_BL output signal for S/W observation) bit, [14-49](#)
- PWMCHA (channel A duty two's complement) bit, [14-42](#)
- PWM_CHAL (channel A duty two's complement) bit, [14-37](#), [14-47](#)
- PWM channel AL duty control (PWM_CHAL) register, [14-47](#)
- PWM channel BL duty control (PWM_CHBL) register, [14-47](#)
- PWM channel CL duty control (PWM_CHCL) register, [14-47](#)
- PWM_CHA (PWM duty control) register, [14-36](#), [14-42](#)
- PWMCHB (channel B duty two's complement) bit, [14-43](#)
- PWM_CHBL (channel B duty two's complement) bit, [14-37](#), [14-47](#)
- PWM_CHB (PWM duty control) register, [14-36](#), [14-43](#)
- PWMCHC (channel C duty two's complement) bit, [14-43](#)
- PWM_CHCL (channel C duty two's complement) bit, [14-37](#), [14-47](#)
- PWM_CHC (PWM duty control) register, [14-37](#), [14-42](#), [14-43](#)
- PWM_CH output signal for S/W observation (PWM_CH) bit, [14-49](#)
- PWM_CH (PWM_CH output signal for S/W observation) bit, [14-49](#)
- PWM_CLK signal, [10-20](#)
- PWM_CL output signal for S/W observation (PWM_CL) bit, [14-49](#)
- PWM_CL (PWM_CL output signal for S/W observation) bit, [14-49](#)
- PWM control (PWM_CTRL) register, [14-37](#)
- PWM crossover and output enable (PWM_SEG) register, [14-44](#)
- PWM_CTRL (PWM control) register, [14-36](#), [14-37](#)
- PWM_DBL (double-update mode) bit, [14-37](#)
- PWM dead time (PWM_DT) register, [14-41](#)
- PWM dead time unsigned (PWM_DT) bit, [14-41](#)
- PWM_DT (PWM dead time) register, [14-41](#)
- PWM_DT (PWM dead time unsigned) bit, [14-36](#), [14-41](#)
- PWM duty control (PWM_CHA) register, [14-42](#)
- PWM duty control (PWM_CHB) register, [14-43](#)
- PWM duty control (PWM_CHC) register, [14-42](#), [14-43](#)
- PWM enable (PWM_EN) bit, [14-37](#)
- PWM_EN (PWM enable) bit, [14-37](#)
- PWM_EXTSYNC (external sync) bit, [14-37](#)
- PWM_GATE (chopping control) register, [14-36](#)
- PWM low side invert (PWM_LSI) register, [14-48](#)
- PWM_LSI (PWM low side invert) register, [14-37](#), [14-48](#)
- PWM_OUT mode, [10-10](#) to [10-22](#), [10-42](#) control bit and register usage, [10-44](#) error prevention, [10-43](#) externally clocked, [10-19](#) PULSE_HI toggle mode, [10-15](#) stopping the timer, [10-21](#)
- PWM output polarity (PWM_POLARITY) bit, [14-37](#)
- PWM period (PWM_TM) register, [14-40](#)

Index

- PWM period unsigned (PWM_TM) bit, [14-40](#)
 - PWM peripheral, [14-1](#) to [14-49](#)
 - PWM_PHASE (PWM phase) bit, [14-39](#)
 - PWM phase (PWM_PHASE) bit, [14-39](#)
 - PWM_POLARITY (PWM output polarity) bit, [14-37](#)
 - PWM polarity (PWM_POL) bit, [14-39](#)
 - PWM_POL (PWM polarity) bit, [14-39](#)
 - PWM registers, [14-36](#)
 - PWM_SEG (PWM crossover and output enable) register, [14-37](#), [14-44](#)
 - PWM simulation status (PWM_STAT2) register, [14-49](#)
 - PWM_SR_LSI_A/B/C (PWM SR mode low side invert channel) bits, [14-48](#)
 - PWM SR mode low side invert channel A/B/C (PWM_SR_LSI_x) bits, [14-48](#)
 - PWM SR mode (PWM_SR) bit, [14-39](#)
 - PWM_SRMODE (PWM SR mode) bit, [14-37](#)
 - PWM SR mode (PWM_SRMODE) bit, [14-37](#)
 - PWM_SR (PWM SR mode) bit, [14-39](#)
 - PWM_STAT2 (PWM simulation status) register, [14-37](#), [14-49](#)
 - PWM_STAT (PWM status) register, [14-36](#), [14-39](#)
 - PWM status (PWM_STAT) register, [14-39](#)
 - PWM sync enable (PWM_SYNC_EN) bit, [14-37](#)
 - PWM_SYNC_EN (PWM sync enable) bit, [14-37](#)
 - PWM SYNC int en for sync (PWMSYNCINT_EN) bit, [14-37](#)
 - PWMSYNCINT_EN (int en for sync) bit, [14-37](#)
 - PWM_SYNCINT (PWM sync int) bit, [14-39](#)
 - PWM sync int (PWM_SYNCINT) bit, [14-39](#)
 - PWM sync pulse width control (PWM_SYNCWT) register, [14-46](#)
 - PWM sync pulse width unsigned (PWMSYNCWT) bit, [14-46](#)
 - PWM_SYNCSEL (external sync select) bit, [14-37](#)
 - PWM_SYNCWT (PWM sync pulse width control) register, [14-37](#), [14-46](#)
 - PWMSYNCWT (PWM sync pulse width unsigned) bit, [14-46](#)
 - PWM_TM (PWM period) register, [14-40](#)
 - PWM_TM (PWM period unsigned) bit, [14-36](#), [14-40](#)
 - PWMTRIP_DSBL (disable for trip input) bit, [14-37](#)
 - PWM TRIP int en for trip (PWMTRIPINT_EN) bit, [14-37](#)
 - PWMTRIPINT_EN (int en for trip) bit, [14-37](#)
 - PWM_TRIPINT (PWM trip int) bit, [14-39](#)
 - PWM trip int (PWM_TRIPINT) bit, [14-39](#)
 - PWM_TRIP (PWM trip) bit, [14-39](#)
 - PWM trip (PWM_TRIP) bit, [14-39](#)
 - PWR_ON (power on) bit, [21-54](#)
 - PWR_SV_EN (power save enable) bit, [21-55](#)
 - Pxn (pin control) bits, [9-32](#), [9-33](#), [9-36](#), [9-37](#), [9-38](#)
-
- ## Q
- query semaphore, [27-3](#)
 - quick boot, [26-39](#)
-
- ## R
- RAF (receive all frames) bit, [22-60](#), [22-64](#)

- RAS to CAS delay (TRCD) bit, 7-65
- RBC (force reload) bit, 6-37, 6-83
- RBSY (receive error) bit, 17-39, 17-41
- RCKFE (clock falling edge select) bit, 19-32, 19-52, 19-55
- RCVDATA16 (receive FIFO 16-bit data) bits, 16-46
- RCVDATA8 (receive FIFO 8-bit data) bits, 16-45
- RCVFLUSH (receive buffer flush) bit, 16-36, 16-37
- RCVINTLEN (receive buffer interrupt length) bit, 16-36, 16-37
- RCVSERVM (receive FIFO service interrupt mask) bit, 16-40
- RCVSERV (receive FIFO service) bit, 16-40, 16-41
- RCVSTAT (receive FIFO status) bits, 16-38
- RDIV bits, 7-39, 7-54, 7-58
- RDTYPE bits, 19-27, 19-52, 19-54
- read, asynchronous, 7-12
- read access, EBIU, 7-18
- read access, SDRAM, 7-49
- read command, 7-33
- read/write access bit, 2-7
- read/write command, 7-49
- real-time clock. *See* RTC
- real-time clock. *See* RTC
- real time clock wake enable (WAKE) bit, 8-23
- receive all frames (RAF) bit, 22-64
- receive bit order (RLSBIT) bit, 19-52
- receive buffer bits, 15-26
- receive buffer flush (RCVFLUSH) bit, 16-36
- receive buffer full int en (ERBFI) bit, 15-27
- receive buffer interrupt length (RCVINTLEN) bit, 16-36
- receive complete (RX_COMP) bit, 22-98, 22-103, 22-104
- receive configuration (SPORT_x_RCR1, SPORT_x_RCR2) registers, 19-51
- receive data available (RX_DAT_RDY) bit, 21-65
- receive data available (RX_DAT_RDY_MASK) bit, 21-70
- receive data buffer bits, 17-43
- receive data field, 19-59
- received filtered frame int en (RX_ACCEPT) bit, 22-93, 22-100, 22-104
- receive enable (RE) bit, 22-66
- receive enable (RSPEN) bit, 19-51, 19-52
- receive error (RBSY) bit, 17-39
- receive event int en (RXEIE) bit, 23-15
- receive event interrupt status (RXEL) bit, 23-17
- receive FIFO, SPORT, 19-58
- receive FIFO 16-bit data (RCVDATA16) bit, 16-46
- receive FIFO 8-bit data (RCVDATA8) bit, 16-45
- receive FIFO empty (RX_FIFO_ZERO) bit, 21-65
- receive FIFO empty (RX_FIFO_ZER/W_MASK) bit, 21-70
- receive FIFO full (RX_FIFO_FULL) bit, 21-65
- receive FIFO full (RX_FIFO_FULL_MASK) bit, 21-70
- receive FIFO not empty status (RXNE) bit, 19-60
- receive FIFO overrun error (RX_OVERRUN) bit, 21-65
- receive FIFO overrun error (RX_OVERRUN_MASK) bit, 21-70

Index

- receive FIFO service interrupt mask (RCVSRVM) bit, [16-40](#)
- receive FIFO service (RCVSERV) bit, [16-40](#)
- receive FIFO status (RCVSTAT) bit, [16-38](#)
- receive FIFO watermark (RX_FIFO_STAT) bit, [21-65](#)
- receive FIFO watermark (RX_FIFO_STAT_MASK) bit, [21-70](#)
- receive frame DMA word alignment (RXDWA) bit, [22-89](#)
- receive frame sync required select (RFSR) bit, [19-52](#)
- receive frame TCP/UDP checksum enable (RXCKS) bit, [22-89](#)
- receive general int en (RXGIE) bit, [23-15](#)
- receive general interrupt status (RXGL) bit, [23-17](#)
- receive overrun error int en (RXOVE) bit, [23-15](#)
- receive snapshot high (EMAC_PTP_RXSNAPHI) register, [23-29](#)
- receive snapshot low (EMAC_PTP_RXSNAPLO) register, [23-28](#)
- receive snapshot overrun status (RXOV) bit, [23-17](#)
- receive stereo frame sync enable (RSFSE) bit, [19-53](#)
- reception error, SPI, [17-41](#)
- reflects detection of NMI (NMI) bit, [25-55](#)
- refresh, SDRAM, [7-31](#)
- REGAD field, [22-72](#)
- register-based DMA, [6-9](#)
- registers
 - See also* registers by name
 - registers, system, [A-2](#)
 - remote wakeup frame enable (RWKE) bit, [22-79](#)
 - remote wakeup frame filters, MAC, [22-34](#)
 - repeat start (RSTART) bit, [16-30](#)
 - REP (request polarity) bit, [6-38](#), [6-83](#)
 - request data control command, DMA, [6-34](#)
 - request data urgent control command, DMA, [6-34](#)
 - request polarity (REP) bit, [6-83](#)
 - RE (receive enable) bit, [22-41](#), [22-49](#), [22-60](#), [22-66](#)
 - reserved SDRAM, [7-2](#)
 - reset all counters (RSTC) bit, [22-119](#)
 - reset disable (RSTDABL) bit, [25-47](#)
 - RESET_DOUBLE (core-double-fault reset) bit, [25-39](#), [26-72](#)
 - RESET pin, [26-5](#)
 - reset registers, [26-71](#)
 - resets
 - core and system, [26-102](#)
 - core double-fault, [26-5](#)
 - core-only software, [26-5](#)
 - hardware, [26-4](#), [26-7](#)
 - interrupts, [26-7](#)
 - software, [26-6](#)
 - SPI, [17-15](#)
 - system software, [26-4](#)
 - watchdog timer, [26-4](#), [26-6](#)
 - RESET_SOFTWARE (software reset status) bit, [25-39](#), [26-72](#)
 - reset vector, [26-1](#)
 - RESET_WDOG (software watchdog timer source) bit, [12-5](#), [25-39](#), [25-42](#), [26-72](#)
 - resource sharing, with semaphores, [27-2](#)
 - RESP_CMD (command index of last received response) bit, [21-59](#)
 - response (RSI_RESPONSEx) bits, [21-60](#)
 - restart control command, DMA, [6-32](#), [6-33](#)

- restart or finish control command,
 - transmit, [6-34](#), [6-35](#)
- restrictions, DMA, [6-24](#), [6-34](#)
- RETI register, [26-7](#)
- RFSR (receive frame sync required select)
 - bit, [19-29](#), [19-30](#), [19-52](#), [19-54](#)
- RFSx signals, [19-5](#), [19-18](#), [19-29](#)
- RLSBIT (receive bit order) bit, [19-52](#), [19-54](#)
- RMII_IO (RMII port speed selector) bit, [22-60](#), [22-61](#)
- RMII mode (RMII) bit, [22-62](#)
- RMII port speed selector (RMII_IO) bit, [22-61](#)
- RMII protocol
 - multiplexing, [22-5](#)
 - pins, [22-6](#)
- RMII (RMII mode) bit, [22-60](#), [22-62](#)
- ROM, [1-9](#), [7-1](#)
- rotary counter registers, [13-19](#)
- round robin operation, MDMA, [6-47](#)
- routing of interrupts, [5-16](#), [5-17](#)
- ROVF (sticky receive overflow status) bit, [19-60](#), [19-61](#)
- row activation, SDRAM, [7-31](#)
- row address, [7-62](#)
- row precharge, SDRAM, [7-31](#)
- RPOLC (IrDA RX polarity change) bit, [15-31](#)
- RPOLC (IrDA Rx polarity change) bit, [15-32](#)
- RRFST (left/right order) bit, [19-12](#), [19-53](#), [19-55](#)
- RSCLKx signals, [19-5](#), [19-28](#)
- RSFSE (receive stereo frame sync enable)
 - bit, [19-10](#), [19-53](#), [19-55](#)
- RSI_ARGUMENT (RSI argument)
 - register, [21-56](#)
- RSI argument (RSI_ARGUMENT)
 - register, [21-56](#)
- RSI_CEATA_CONTROL (RSI CE-ATA control) register, [21-74](#)
- RSI CE-ATA control
 - (RSI_CEATA_CONTROL) register, [21-74](#)
- RSI_CLK_EN (RSI clocks enable) bit, [21-78](#)
- RSI_CLOCK enable (CLK_EN) bit, [21-55](#)
- RSI clocks enable (RSI_CLK_EN) bit, [21-78](#)
- RSI_CMD_OD (RSI command open drain) bit, [21-54](#)
- RSI command open drain
 - (RSI_CMD_OD) bit, [21-54](#)
- RSI_COMMAND (RSI command)
 - register, [21-57](#)
- RSI command (RSI_COMMAND)
 - register, [21-57](#)
- RSI_CONFIG (RSI configuration)
 - register, [21-78](#)
- RSI configuration (RSI_CONFIG)
 - register, [21-78](#)
- RSI_DATA3 pull-down enable
 - (PD_DAT3) bit, [21-78](#)
- RSI_DATA3 pull-up enable (PU_DAT3)
 - bit, [21-78](#)
- RSI_DATA_CNT (RSI data counter)
 - register, [21-64](#)
- RSI_DATA_CONTROL (RSI data control) register, [21-62](#)
- RSI data control
 - (RSI_DATA_CONTROL) register, [21-62](#)
- RSI data counter (RSI_DATA_CNT)
 - register, [21-64](#)
- RSI data FIFO (RSI_FIFO) register, [21-74](#)
- RSI data length (RSI_DATA_LGTH)
 - register, [21-61](#)

Index

- RSI_DATA_LGTH (RSI data length) register, [21-61](#)
- RSI_DATA_TIMER (RSI data timer) register, [21-61](#)
- RSI data timer (RSI_DATA_TIMER) register, [21-61](#)
- RSI_EMASK (RSI exception mask) register, [21-76](#)
- RSI_ESTAT (RSI exception status) register, [21-75](#)
- RSI exception mask (RSI_EMASK) register, [21-76](#)
- RSI exception status (RSI_ESTAT) register, [21-75](#)
- RSI_FIFO_CNT (RSI FIFO counter) register, [21-73](#)
- RSI FIFO counter (RSI_FIFO_CNT) register, [21-73](#)
- RSI_FIFO (RSI data FIFO) register, [21-74](#)
- RSI interrupt mask (RSI_MASKx) registers, [21-70](#)
- RSI_MASKx (RSI interrupt mask) registers, [21-70](#)
- RSI peripheral ID (RSI_PIDx) registers, [21-80](#)
- RSI_PID (peripheral ID) bit, [21-80](#)
- RSI_PIDx (RSI peripheral ID) registers, [21-80](#)
- RSI power control (RSI_PWR_CONTROL) register, [21-54](#), [21-55](#)
- RSI_PWR_CONTROL (RSI power control) register, [21-54](#), [21-55](#)
- RSI_RD_WAIT_EN (RSI read wait enable) register, [21-79](#)
- RSI read wait enable (RSI_RD_WAIT_EN) register, [21-79](#)
- RSI read wait request enable (SDIO_RWR) bit, [21-79](#)
- RSI registers, [21-52](#)
- RSI reset (RSI_RST) bit, [21-78](#)
- RSI_RESP_CMD (RSI response command) register, [21-59](#)
- RSI response command (RSI_RESP_CMD) register, [21-59](#)
- RSI response (RSI_RESPONSEx) registers, [21-60](#)
- RSI_RESPONSEx (response) bits, [21-60](#)
- RSI_RESPONSEx (RSI response) registers, [21-60](#)
- RSI_RST (RSI reset) bit, [21-78](#)
- RSI status clear (RSI_STATUSCL) register, [21-68](#)
- RSI_STATUSCL (RSI status clear) register, [21-68](#)
- RSI_STATUS (RSI status) register, [21-65](#)
- RSI status (RSI_STATUS) register, [21-65](#)
- RSPEN (receive enable) bit, [19-9](#), [19-51](#), [19-52](#), [19-53](#)
- RSTART (repeat start) bit, [16-30](#), [16-31](#)
- RSTC (reset all counters) bit, [22-42](#), [22-119](#)
- RSTDABL (reset disable) bit, [25-47](#)

- RTC, [1-26](#), [24-1](#) to [24-27](#)
 - alarm clock features, [24-2](#), [24-26](#)
 - block diagram, [24-3](#)
 - clock operation, [24-3](#), [24-5](#)
 - counters, [24-2](#)
 - deep sleep, [24-9](#)
 - digital watch features, [1-4](#), [24-1](#)
 - disabling prescaler, [24-5](#)
 - enabling prescaler, [24-4](#), [24-23](#)
 - event flags, [24-10](#)
 - initializing, [24-5](#)
 - interfaces, [24-3](#)
 - interrupts, [8-9](#), [24-6](#), [24-14](#), [24-15](#)
 - latency, [24-8](#)
 - prescaler, [24-2](#)
 - programming, [24-6](#), [24-23](#)
 - reads, [24-9](#)
 - registers, [24-19](#)
 - setting time of day, [24-12](#)
 - state transitions, [24-16](#)
 - stopwatch, [24-2](#), [24-13](#), [24-24](#)
 - synchronization, [24-5](#)
 - system state transition events, [24-17](#)
 - test mode, [24-5](#)
 - writes, [24-8](#)
- RTC_ALARM (RTC alarm) register, [24-2](#), [24-19](#), [24-22](#)
- RTC alarm (RTC_ALARM) register, [24-22](#)
- RTC_ICTL (RTC interrupt control) register, [24-19](#), [24-20](#)
- RTC interrupt control (RTC_ICTL) register, [24-20](#)
- RTC_ISTAT (RTC interrupt status) register, [24-19](#), [24-21](#)
- RTC_PREN (prescaler enable) register, [24-4](#), [24-5](#), [24-19](#)
- RTC_PREN (RTC prescaler enable) register, [24-22](#)
- RTC prescaler enable (RTC_PREN) register, [24-22](#)
- RTC_STAT (RTC status) register, [24-12](#), [24-19](#), [24-20](#)
- RTC status (RTC_STAT) register, [24-12](#), [24-19](#), [24-20](#)
- RTC stopwatch count (RTC_SWCNT) register, [24-21](#)
- RTC_SWCNT (RTC stopwatch count) register, [24-2](#), [24-13](#), [24-19](#), [24-21](#)
- RUVF (sticky receive underflow status) bit, [19-59](#), [19-60](#), [19-61](#)
- RWKE (remote wakeup frame enable) bit, [22-34](#), [22-79](#)
- RWKS (wakeup frame received status) bits, [22-78](#)
- RX_ACCEPT (received filtered frame int en) bit, [22-93](#), [22-100](#), [22-104](#)
- RX_ACT (data receive active) bit, [21-65](#)
- RX_ACT_MASK (data receive active) bit, [21-70](#)
- RX_ADDR (address filter failed) bit, [22-96](#), [22-101](#), [22-104](#)
- RX_ALIGN (alignment error) bit, [22-97](#), [22-102](#), [22-104](#)
- RX_ALIGN_CNT (alignment errors counter int) bit, [22-113](#)
- RX_ALIGN_CNT (alignment errors counter int en) bit, [22-115](#)
- RX_ALLF_CNT (frames received all counter int) bit, [22-113](#)
- RX_ALLF_CNT (frames received all counter int en) bit, [22-115](#)
- RX_ALLO_CNT (octets received all counter int) bit, [22-113](#)
- RX_ALLO_CNT (octets received all counter int en) bit, [22-115](#)
- RX_BROAD (broadcast frame detected) bit, [22-101](#), [22-104](#)

Index

- RX_BROAD (broadcast frames detected) bit, [22-95](#)
- RX_BROAD_CNT (broadcast frames received OK counter int) bit, [22-113](#)
- RX_BROAD_CNT (broadcast frames received OK counter int en) bit, [22-115](#)
- RXCKS (enable receive frame TCP/UDP checksum computation) bit, [22-89](#)
- RX_COMP (receive complete) bit, [22-98](#), [22-103](#), [22-104](#)
- RX_CRC (frame CRC error) bit, [22-97](#), [22-102](#), [22-104](#)
- RX_CTL (control frame int en) bit, [22-95](#), [22-100](#), [22-104](#)
- RX data buffer status (RXS) bit, [17-39](#)
- RX_DAT_RDY_MASK (receive data available) bit, [21-70](#)
- RX_DAT_RDY (receive data available) bit, [21-65](#)
- RX DMA direction error detected, [22-40](#)
- RX DMA direction error status (RXDMAERR) bit, [22-91](#)
- RXDMAERR (RX DMA direction error status) bit, [22-91](#)
- RX_DMAO (DMA overrun detected) bit, [22-96](#), [22-101](#), [22-104](#)
- RXDWA (receive frame DMA word alignment) bit, [22-17](#), [22-49](#), [22-89](#)
- RXEIE (receive event int en) bit, [23-15](#)
- RXEL (receive event interrupt status) bit, [23-17](#)
- RX_EM (mask for PTP control field of event messages) bit, [23-22](#)
- RX_EQ64_CNT (frames length equal to 64 received counter int) bit, [22-113](#)
- RX_EQ64_CNT (frames length equal to 64 received counter int en) bit, [22-115](#)
- RX_FCS_CNT (frame check sequence errors counter int) bit, [22-113](#)
- RX_FCS_CNT (frame check sequence errors counter int en) bit, [22-115](#)
- RX_FIFO_FULL_MASK (receive FIFO full) bit, [21-70](#)
- RX_FIFO_FULL (receive FIFO full) bit, [21-65](#)
- RX_FIFO_STAT_MASK (receive FIFO watermark) bit, [21-70](#)
- RX_FIFO_STAT (receive FIFO watermark) bit, [21-65](#)
- RX_FIFO_ZERO (receive FIFO empty) bit, [21-65](#)
- RX_FIFO_ZERO/W_MASK (receive FIFO empty) bit, [21-70](#)
- RX_FRAG (frame fragment detected) bit, [22-96](#), [22-102](#), [22-104](#)
- RX frame status interrupt, [22-39](#)
- RX frame-status interrupt status (RXFSINT) bit, [22-91](#)
- RX_FRLEN[10:0] field, [22-98](#)
- RXFSINT (RX frame-status interrupt status) bit, [22-91](#)
- RX_GE1024_CNT (frames length 1024-max received counter int) bit, [22-113](#)
- RX_GE1024_CNT (frames length 1024-max received counter int en) bit, [22-115](#)
- RXGIE (receive general int en) bit, [23-15](#)
- RXGL (receive general interrupt status) bit, [23-17](#)
- RX_GN (mask for PTP control field of general messages) bit, [23-22](#)
- RX hold register, [19-59](#)
- RX_IRL_CNT (in-range length errors counter int) bit, [22-113](#)
- RX_IRL_CNT (in-range length errors counter int en) bit, [22-115](#)
- RX_LATE (late collision detected) bit, [22-96](#), [22-101](#), [22-104](#)

- RX_LEN (length error detected) bit, [22-97](#), [22-102](#), [22-104](#)
- RX_LONG_CNT (frame too long errors counter int) bit, [22-113](#)
- RX_LONG_CNT (frame too long errors counter int en) bit, [22-115](#)
- RX_LONG (frame too long) bit, [22-97](#), [22-102](#), [22-104](#)
- RX_LOST_CNT (frames lost counter int) bit, [22-113](#)
- RX_LOST_CNT (frames lost counter int en) bit, [22-115](#)
- RX_LT1024_CNT (frames length 512-1023 received counter int) bit, [22-113](#)
- RX_LT1024_CNT (frames length 512-1023 received counter int en) bit, [22-115](#)
- RX_LT128_CNT (frames length 65-127 received counter int) bit, [22-113](#)
- RX_LT128_CNT (frames length 65-127 received counter int en) bit, [22-115](#)
- RX_LT256_CNT (frames length 128-255 received counter int) bit, [22-113](#)
- RX_LT256_CNT (frames length 128-255 received counter int en) bit, [22-115](#)
- RX_LT512_CNT (frames length 256-511 received counter int) bit, [22-113](#)
- RX_LT512_CNT (frames length 256-511 received counter int en) bit, [22-115](#)
- RX modes with external frame syncs, [20-20](#)
- RX_MULTI_CNT (multicast frames received OK counter int) bit, [22-113](#)
- RX_MULTI_CNT (multicast frames received OK counter int en) bit, [22-115](#)
- RX_MULTI (multicast frame int en) bit, [22-95](#), [22-101](#), [22-104](#)
- RXNE (receive FIFO not empty status) bit, [19-60](#)
- RX_OCTET_CNT (octets received OK counter int) bit, [22-113](#)
- RX_OCTET_CNT (octets received OK counter int en) bit, [22-115](#)
- RX_OK_CNT (frames received OK counter int) bit, [22-113](#)
- RX_OK_CNT (frames received OK counter int en) bit, [22-115](#)
- RX_OK (good received frame int en) bit, [22-98](#), [22-103](#), [22-104](#)
- RX_OPCODE_CNT (unsupported opcodes received counter int) bit, [22-113](#)
- RX_OPCODE_CNT (unsupported opcodes received counter int en) bit, [22-115](#)
- RX_ORL_CNT (out-of-range length field counter int) bit, [22-113](#)
- RX_ORL_CNT (out-of-range length field counter int en) bit, [22-115](#)
- RXOVE (receive overrun error int en) bit, [23-15](#)
- RX_OVERRUN_MASK (receive FIFO overrun error) bit, [21-70](#)
- RX_OVERRUN (receive FIFO overrun error) bit, [21-65](#)
- RX_OVERRUN_STAT (clear receive FIFO overrun error) bit, [21-68](#)
- RXOV (receive snapshot overrun status) bit, [23-17](#)
- RX_PAUSE_CNT (PAUSE MAC control frames received counter int) bit, [22-113](#)
- RX_PAUSE_CNT (PAUSE MAC control frames received counter int en) bit, [22-115](#)
- RX_PHY (PHY error int en) bit, [22-96](#), [22-101](#), [22-104](#)
- RX_RANGE (out-of-range length field detected) bit, [22-101](#), [22-104](#)

Index

RX_RANGE (out-of-range length fields detected) bit, [22-95](#)
RXREQ signal, [15-7](#)
RxSEC enable (RXSE) bit, [19-53](#)
RXSE (RxSEC enable) bit, [19-53](#), [19-55](#)
RX_SHORT_CNT (frames length less than 64 received counter int) bit, [22-113](#)
RX_SHORT_CNT (frames length less than 64 received counter int en) bit, [22-115](#)
RX_SNAP_HI (higher 32 bits of receive snapshot of local clock time) bit, [23-29](#)
RX_SNAP_LO (lower 32 bits of receive snapshot of local clock time) bit, [23-28](#)
RXS (RX data buffer status) bit, [17-22](#), [17-39](#)
RX status int en (ELSI) bit, [15-27](#)
RX_TYPED_CNT (typed frames received counter int) bit, [22-113](#)
RX_TYPED_CNT (typed frames received counter int en) bit, [22-115](#)
RX_TYPE (typed frame int en) bit, [22-94](#), [22-100](#), [22-104](#)
RX_UCTL (unsupported control frame) bit, [22-95](#), [22-100](#), [22-104](#)
RX_UNI_CNT (unicast frames received OK counter int) bit, [22-113](#)
RX_UNI_CNT (unicast frames received OK counter int en) bit, [22-115](#)
RX_VLAN1 (VLAN1 frame int en) bit, [22-94](#), [22-100](#), [22-104](#)
RX_VLAN2 (VLAN2 frame detected) bit, [22-100](#), [22-104](#)
RX_VLAN2 (VLAN2 frames detected) bit, [22-93](#)

S

SA10 pin, [7-53](#)
SADDR (slave mode address) bits, [16-28](#)
SAMPLE/PRELOAD instruction, [B-6](#)
sampling clock period, UART, [15-8](#)
sampling edge, SPORT, [19-32](#)
SAV codes, [20-31](#)
SAV signal, [20-5](#)
SB (set break) bit, [15-21](#)
scale value[7: 0] field, [11-6](#)
scaling, of core timer, [11-7](#)
scan paths, [B-5](#)
SCCB compatibility (SCCB) bit, [16-25](#)
SCCB (SCCB compatibility) bit, [16-25](#)
scheduling, memory DMA, [6-46](#)
SCKELOW, [8-23](#)
SCKE low reset (SCKELOW) bit, [8-23](#)
SCKELOW (SCKE low reset) bit, [8-23](#)
SCKE pin, [7-52](#)
SCK signal, [17-4](#), [17-11](#), [17-14](#), [17-15](#)
SCL clock divider (TWI_CLKDIV) register, [16-25](#)
SCLK, [1-28](#), [4-4](#), [8-5](#)
 changing frequency, [7-44](#)
 derivation, [8-1](#)
 EBIU, [7-1](#)
 operating mode, [8-7](#)
 status, [8-7](#)
SCLOVR (serial clock override) bit, [16-30](#)
SCL pin, [16-5](#)
SCLSEN (serial clock sense) bit, [16-33](#), [16-34](#)
SCL serial clock, [16-25](#)
SCL (serial clock) signal, [16-4](#)
SCOMPM (slave transfer complete interrupt mask) bit, [16-40](#)
SCOMP (slave transfer complete) bit, [16-40](#), [16-42](#)
SCRATCH bits, [15-31](#)
scratchpad memory and booting, [26-21](#)

- SCTLE bit, [7-52](#), [7-55](#), [7-64](#), [7-65](#), [7-66](#),
[7-71](#)
- SDAOVR (serial data override) bit, [16-30](#)
- SDA pin, [16-5](#)
- SDASEN (serial data sense) bit, [16-33](#),
[16-34](#)
- SDA (serial data) signal, [16-4](#)
- SDC, [1-12](#), [7-4](#), [7-24](#)
 - address mapping, [7-26](#)
 - address muxing, [7-41](#)
 - architecture, [7-40](#)
 - code examples, [7-75](#)
 - commands, [7-45](#)
 - component configurations, [7-25](#)
 - core and DMA arbitration, [7-43](#)
 - core transfers to SDRAM, [7-75](#)
 - disabled CLKOUT, [7-76](#)
 - initialization, [7-78](#)
 - no burst mode, [7-41](#)
 - operation, [7-39](#)
 - pin state during commands, [7-46](#)
 - power-up latency, [7-54](#)
 - registers, [7-57](#)
 - SA10 pin, [7-53](#)
 - self-refresh mode, [7-76](#)
 - transfers to SDRAM using byte mask,
[7-76](#)
- SD_CARD_DET (card detect int) bit,
[21-75](#)
- SD_CARD_DET_MASK (card detect int
en) bit, [21-76](#)
- SDC EAB sticky error status (SDEASE)
bit, [7-74](#)
- SDC idle (SDCI) bit, [7-74](#)
- SDCI (SDC idle) bit, [7-74](#)
- SDC powerup active (SDPUA) bit, [7-74](#)
- SDC powerup delay (SDRS) bit, [7-74](#)
- SDC self-refresh active (SDSRA) bit, [7-74](#)
- SDEASE (SDC EAB sticky error status)
bit, [7-74](#), [7-75](#)
- SDIO 4-bit enable (SDIO4_EN) bit,
[21-78](#)
- SDIO4_EN (SDIO 4-bit enable) bit,
[21-78](#)
- SDIO_INT_DET_MASK (SDIO int en)
bit, [21-76](#)
- SDIO_INT_DET (SDIO interrupt detect)
bit, [21-75](#)
- SDIO int en (SDIO_INT_DET_MASK)
bit, [21-76](#)
- SDIO interrupt detect (SDIO_INT_DET)
bit, [21-75](#)
- SDIO interrupt moving window enable
(MW_EN) bit, [21-78](#)
- SDIO_RWR (RSI read wait request
enable) bit, [21-79](#)
- SDIR (slave transfer direction) bit, [16-28](#),
[16-29](#)
- SDPUA (SDC powerup active) bit, [7-74](#),
[7-75](#)
- SDQM[1:0] encodings during writes, [7-50](#)

Index

- SDRAM, 1-9, 2-1
 - A[10] pin, 7-53
 - address connections, 7-41
 - address mapping, 7-26
 - banks, 2-5, 7-2, 7-3, 7-27, 7-31
 - bursts, 7-32
 - clock enables, setting, 7-64
 - column read/write, 7-31
 - configuration, 7-24, 7-60
 - exiting self-refresh mode, 7-52
 - external SDRAM, 7-26, 7-32, 7-69
 - initialization, 7-47, 7-70
 - interface, 7-28, 7-45
 - latency, 7-32, 7-70
 - memory size, 7-32
 - memory space, 7-2
 - multibank operation, 7-42
 - parallel connection, 7-27
 - performance, 7-29
 - power-up sequence, 7-44, 7-47, 7-70
 - read transfers, 7-49
 - refresh, 7-31
 - reserved, 7-2
 - row activation, 7-31
 - row precharge, 7-31
 - self-refresh mode, 7-34, 7-51
 - sharing, 7-69
 - sizes, 2-5, 7-24, 7-60
 - smaller than 16M byte, 7-62
 - specification of system, 7-27
 - stall cycles, 7-30
 - start address, 7-2
 - system block diagram, 7-56, 7-57
 - timing, 7-34
- SDRAM commands, 7-33
 - auto-refresh command, 7-34, 7-51
 - bank activate command, 7-33
 - ERMS command, 7-33
 - interface controller commands, 7-45
 - MRS command, 7-33
 - NOP command, 7-53
 - precharge all command, 7-34
 - precharge command, 7-34
 - read command, 7-33
 - read/write command, 7-49
 - write command, 7-33
- SDRAM controller. *See* SDC
- SDRAM control status (EBIU_SDSTAT)
 - register, 7-74
- SDRAM control status register (EBIU_SDSTAT), 7-74
- SDRAM memory bank control (EBIU_SDBCTL) register, 7-60, 7-61
- SDRAM memory global control (EBIU_SDGCTL) register, 7-64, 7-65
- SDRAM refresh rate control (EBIU_SDRRC) register, 7-58
- SDRS (SDC powerup delay) bit, 7-54, 7-74, 7-75
- SDSRA (SDC self-refresh active) bit, 7-74
- SECMODE (secure mode control state) bit, 25-55
- seconds (1 Hz) event flag bits, 24-20, 24-21, 24-22
- SECSTAT (secure status) bit, 25-55
- SECURE_CONTROL (secure control) register, 25-53
- secure control (SECURE_CONTROL) register, 25-53
- secure mode control state (SECMODE) bit, 25-55
- secure status (SECSTAT) bit, 25-55

- SECURE_STATUS (secure status)
 - register, [25-55](#)
- secure status (SECURE_STATUS)
 - register, [25-55](#)
- SECURE_SYSSWT (secure system switch)
 - register, [25-47](#)
- secure system switch (SECURE_SYSSWT)
 - register, [25-47](#)
- SECURE x (SECUREx) bits, [25-53](#)
- security registers, [25-45](#)
- self-refresh command, [7-51](#)
- self-refresh mode, [7-34](#), [7-51](#), [7-71](#), [7-76](#)
 - entering, [7-51](#)
 - exiting, [7-52](#)
- self-refresh setting (SRFS) bit, [7-65](#)
- semaphores, [27-2](#)
 - coherency, [27-3](#)
 - example code, [27-3](#)
 - query, [27-3](#)
- send zero (SZ) bit, [17-36](#)
- SEN (slave enable) bit, [16-26](#), [16-27](#)
- serial
 - clock frequency, [17-34](#)
 - communications, [15-5](#)
 - data transfer, [19-4](#)
 - scan paths, [B-4](#)
- serial clock divide modulus [15: 0] field, [19-62](#)
- serial clock override (SCLOVR) bit, [16-30](#)
- serial clock sense (SCLSEN) bit, [16-33](#)
- serial data override (SDAOVR) bit, [16-30](#)
- serial data sense (SDASEN) bit, [16-33](#)
- serial peripheral interface. *See* SPI
- serial port. *See* SPORT
- serial scan paths, [B-5](#)
- SERRM (slave transfer error interrupt mask) bit, [16-40](#)
- SERR (slave transfer error) bit, [16-40](#), [16-42](#)
- set break (SB) bit, [15-21](#)
- set index bits, [2-7](#)
- SET_PHYAD macro, [22-125](#)
- set Pxn bit, [9-34](#)
- set Pxn interrupt A enable bit, [9-39](#)
- set Pxn interrupt B enable bit, [9-40](#)
- SET_REGAD macro, [22-125](#)
- setup EBIU asynchronous memory controller, [7-18](#)
- setup SDRAM clock enables, [7-64](#)
- shared interrupts, [5-4](#)
- SIC_IAR (system interrupt assignment)
 - register, [5-11](#)
- SIC_IMASK (system interrupt mask)
 - register, [5-5](#)
- SIC registers, [5-10](#)
- SIC. *See* system interrupt controller
- signal integrity, [27-8](#)
- signalling, via semaphores., [27-2](#)
- sine wave input, [1-28](#)
- single collision frames counter int en (TX_SCOLL_CNT) bit, [22-118](#)
- single collision frames counter int (TX_SCOLL_CNT) bit, [22-116](#)
- single collision frames (EMAC_TXC_1COL) register, [22-56](#)
- single pulse generation, timer, [10-12](#)
- SINITM (slave transfer initiated interrupt mask) bit, [16-40](#)
- SINIT (slave transfer initiated) bit, [16-40](#), [16-42](#)
- size of accesses, timer registers, [10-35](#)
- size of words (SIZE) bit, [17-36](#)
- SIZE (size of words) bit, [17-35](#), [17-36](#)
- skip enable (SKIP_EN) bit, [20-26](#)
- SKIP_EN (skip enable) bit, [20-25](#), [20-26](#)
- SKIP_EO (skip even odd) bit, [20-26](#), [20-27](#)
- skip even odd (SKIP_EO) bit, [20-26](#)
- slave enable (SEN) bit, [16-26](#)

Index

- slave mode address (SADDR) bit, 16-28
- slave mode control (TWI_SLAVE_CTL) register, 16-26
- slave mode setup, in TWI, 16-11, 16-51
- slave overflow interrupt mask (SOVFM) bit, 16-40
- slave overflow (SOVF) bit, 16-40
- slaves, EBIU, 7-4
- slaves, PAB, 4-6
- slave select enable (PSSE) bit, 17-36
- slave select value (FLGx) bit, 17-37
- slave SPI device, 17-5, 17-38
- slave transfer complete interrupt mask (SCOMPM) bit, 16-40
- slave transfer complete (SCOMP) bit, 16-40
- slave transfer direction (SDIR) bit, 16-28
- slave transfer error interrupt mask (SERRM) bit, 16-40
- slave transfer error (SERR) bit, 16-40
- slave transfer initiated interrupt mask (SINITM) bit, 16-40
- slave transfer initiated (SINIT) bit, 16-40
- slave transmit data valid (STDVAL) bit, 16-26
- sleep mode, 1-29, 8-9, 22-32
- SLEN[4: 0] field, 19-48, 19-49, 19-53, 19-54
 - restrictions, 19-26
 - word length formula, 19-26
- small descriptor mode, DMA, 6-14
- small model mode, DMA, 6-68
- software
 - interrupts, 5-3
 - management of DMA, 6-50
 - watchdog timer, 1-27, 12-1
- software reset, 26-6, 26-71
- software reset status
 - (RESET_SOFTWARE) bit, 26-72
- software reset (SWRESET) bit, 26-74
- software reset (SWRST) register, 25-39, 26-72
- software watchdog timer source
 - (RESET_WDOG) bit, 26-72
- source channels, memory DMA, 6-7
- source ID offset (EMAC_PTP_ID_OFF) register, 23-33
- source ID offset (SRC_ID_OFF) bit, 23-33
- source ID snapshot
 - (EMAC_PTP_ID_SNAP) register, 23-34
- source ID (SRC_ID) bit, 23-34
- SOVFM (slave overflow interrupt mask) bit, 16-40
- SOVF (slave overflow) bit, 16-40, 16-42
- speculative read and MAC, 22-42
- SPE (SPI enable) bit, 17-35, 17-36
- SPI, 1-23, 17-2 to 17-52
 - boot mode, 26-65
 - clock polarity, 17-11
 - data corruption, avoiding, 17-14
 - DMA operations, 17-10, 17-23, 17-26, 17-47
 - errors, 17-16, 17-39 to 17-41
 - examples, code, 17-43
 - full-duplex synchronous serial interface, 17-2
 - initialization, 17-44
 - interfaces, 17-10
 - interrupts, 17-16, 17-46
 - master mode, 17-14, 17-17, 17-23
 - mode fault error, 17-40
 - multiple slave systems, 17-8
 - operations, 17-14 to 17-21
 - peripherals, compatible, 17-3
 - reception error, 17-41
 - registers, 17-33
 - reset effects, 17-15

- SPI, *(continued)*
 - slave mode, 17-5, 17-15, 17-19, 17-21, 17-26, 17-37, 26-65
 - slave select, 17-38
 - SPI_FLG mapping to port pins, 17-38
 - system enabling, 17-34
 - timing, 17-6
 - transmit collisions, 17-41
 - word length, 17-35
- SPI baud rate (SPI_BAUD) register, 17-34
- SPI_BAUD (SPI baud rate) register, 17-33, 17-34
- SPI clock
 - phase, 17-11, 17-13, 17-16
 - polarity, 17-15
 - SCK signal, 17-4
 - signal, 17-3, 17-15
- SPI control (SPI_CTL) register, 17-36
- SPI_CTL (SPI control) register, 17-5, 17-33, 17-34, 17-36
- SPI data transfers, 17-14
 - completion detection, 17-39
 - formats, 17-11 to 17-13
 - initiate command, 17-17, 17-18
 - modes, 17-18
 - preparation, 17-21
 - protocol, 17-13
 - starting, 17-21, 17-45, 17-50
 - stopping, 17-21, 17-47, 17-50
 - transmit/receive, switching, 17-22
- SPI enable (SPE) bit, 17-36
- SPI finished (SPIF) bit, 17-39
- SPI flag (SPI_FLG) register, 17-37
- SPI_FLG (SPI flag) register, 17-7, 17-8, 17-33, 17-37
- SPIF (SPI finished) bit, 17-9, 17-22, 17-39
- SPI_RDBR (shadow) bits, 17-43
- SPI RDBR shadow (SPI_SHADOW register), 17-33
- SPI RDBR shadow (SPI_SHADOW) register, 17-43
- SPI_RDBR (SPI receive data buffer) register, 17-33, 17-42, 17-43
- SPI receive data buffer (SPI_RDBR) register, 17-43
- SPI_SHADOW (SPI RDBR shadow) register, 17-33, 17-43
- SPIISO_IEDSBL bit, 9-29
- SPIISO_PUPEN bit, 9-29
- SPISS signal, 17-6, 17-8, 17-11
- SPI_STAT (SPI status) register, 17-33, 17-39
- SPI status (SPI_STAT) register, 17-39
- SPI_TDBR data buffer status (TXS) bit, 17-39
- SPI_TDBR (SPI transmit data buffer) register, 17-33, 17-41, 17-42
- SPI transmit data buffer (SPI_TDBR) register, 17-42
- SPORT, 1-21, 19-1 to 19-74
 - 2X clock recovery control, 19-24
 - active low vs. active high frame syncs, 19-32
 - active pulses, shortened, 19-10
 - bit order selection, 19-26
 - channels, 19-14
 - companding, 19-28
 - configuration, 19-10
 - connections, 19-7, 19-8
 - disabling, 19-10
 - DMA data packing, 19-23

Index

- SPORT, *(continued)*
- enable/disable, [19-9](#)
 - enabling multichannel mode, [19-16](#)
 - errors, [19-37](#), [19-38](#)
 - H.100 standard protocol, [19-23](#)
 - initialization code, [19-54](#)
 - internal memory access, [19-37](#)
 - interrupts, [19-37](#), [19-57](#), [19-59](#)
 - modes, [19-10](#)
 - multichannel operation, [19-14](#) to [19-23](#)
 - operations, [19-9](#)
 - packing data, multichannel DMA, [19-23](#)
 - peripheral bus access, [19-38](#)
 - receive FIFO, [19-58](#)
 - receive operations, [19-4](#), [19-28](#), [19-59](#)
 - registers, [19-44](#), [19-45](#), [19-56](#), [19-59](#)
 - sampling edge, [19-32](#)
 - signals, [19-5](#)
 - stereo serial operations, [19-8](#), [19-10](#),
[19-17](#)
 - terminations, [19-8](#)
 - throughput, [19-5](#)
 - timing, [19-38](#)
 - transmit operations, [19-4](#)
 - transmitter FIFO, [19-56](#)
 - unpacking data, multichannel DMA,
[19-23](#)
 - window offset, [19-20](#)
- SPORT_CHNL (SPORT current channel)
register, [19-64](#), [19-65](#)
- SPORT clock, [19-28](#)
- frequency, [19-24](#), [19-61](#)
 - rate, [19-2](#), [19-26](#)
 - single clock, receive and transmit with,
[19-28](#)
 - transmit clock, [19-28](#)
- SPORT current channel (SPORT_CHNL)
register, [19-65](#)
- SPORT data transfers
- formats, [19-27](#), [19-55](#)
 - framed serial transfers, [19-30](#)
 - memory moves, [19-37](#)
 - multichannel timing, [19-16](#)
 - protocols, communication, [19-2](#), [19-23](#)
 - single word, [19-37](#)
 - unframed data flow, [19-30](#)
 - word length, [19-26](#), [19-56](#), [19-59](#)
- SPORT frame sync, [19-31](#), [19-34](#)
- framed vs. unframed, [19-29](#)
 - framing signals, [19-29](#)
 - frequencies, [19-24](#)
 - internal vs. external, [19-31](#)
 - late frame sync, [19-17](#)
 - modes, [19-17](#)
 - multichannel frame, [19-19](#)
- SPORT_MCMC_x (SPORT multichannel
configuration) registers, [19-63](#), [19-64](#)
- SPORT_MRCS_x (SPORT multichannel
receive select) registers, [19-21](#), [19-22](#),
[19-65](#), [19-66](#)
- SPORT_MTCS_x (SPORT multichannel
transmit select) registers, [19-21](#),
[19-22](#), [19-66](#), [19-67](#)
- SPORT multichannel configuration
(SPORT_MCMC_x) registers, [19-64](#)
- SPORT multichannel receive select
(SPORT_MRCS_x) registers, [19-66](#)
- SPORT multichannel transmit select
(SPORT_MTCS_x) registers, [19-67](#)
- SPORT_RCLKDIV (SPORT receive serial
clock divider) registers, [19-61](#), [19-62](#)
- SPORT_RCR_x (SPORT receive
configuration) registers, [19-51](#), [19-52](#),
[19-53](#)
- SPORT receive configuration
(SPORT_RCR_x) registers, [19-52](#),
[19-53](#)

- SPORT receive data (SPORT_RX) register, [19-59](#)
- SPORT receive frame sync divider (SPORT_RFSDIV) register, [19-63](#)
- SPORT receive serial clock divider (SPORT_RCLKDIV) register, [19-62](#)
- SPORT_RFSDIV (SPORT receive frame sync divider) registers, [19-62](#), [19-63](#)
- SPORT_RX (SPORT receive data) registers, [19-18](#), [19-58](#), [19-59](#)
- SPORT_STAT (SPORT status) registers, [19-60](#)
- SPORT status (SPORT_STAT) register, [19-60](#)
- SPORT_TCLKDIV (SPORT transmit serial clock divider) registers, [19-61](#), [19-62](#)
- SPORT_TCRx (SPORT transmit configuration) registers, [19-47](#), [19-48](#)
- SPORT_TCRx (transmit configuration) registers, [19-46](#)
- SPORT_TFSDIV (SPORT transmit frame sync divider) registers, [19-62](#), [19-63](#)
- SPORT transmit configuration (SPORT_TCRx) registers, [19-47](#), [19-48](#)
- SPORT transmit data (SPORT_TX) register, [19-57](#)
- SPORT transmit frame sync divider (SPORT_TFSDIV) register, [19-63](#)
- SPORT transmit serial clock divider (SPORT_TCLKDIV) register, [19-62](#)
- SPORT_TX (SPORT transmit data) registers, [19-18](#), [19-36](#), [19-56](#), [19-57](#)
- SRAM ADDR[13: 12] field, [2-7](#)
- SRAM operation, [1-9](#), [7-1](#), [27-5](#)
- SRC_ID_OFF (source ID offset) bit, [23-33](#)
- SRC_ID (source ID) bit, [23-34](#)
- SRFS (self-refresh setting) bit, [7-51](#), [7-65](#), [7-66](#), [7-71](#)
- SSEL (system select) bit, [8-21](#)
- SSEL (system select) bits, [4-4](#), [8-5](#), [8-21](#), [27-2](#)
- STABUSY (STA busy status) bit, [22-72](#), [22-73](#)
- STA busy status (STABUSY) bit, [22-72](#)
- STADATA[15: 0] field, [22-74](#)
- STADISPRE (disable preamble generation) bit, [22-72](#), [22-73](#)
- STAIE (station management transfer done int en) bit, [22-72](#), [22-73](#)
- stall cycles, SDRAM, [7-30](#)
- STAOP (station management operation code) bit, [22-72](#), [22-73](#)
- start address registers
 - (DMAx_START_ADDR), [6-74](#)
 - (MDMA_yy_START_ADDR), [6-74](#)
- START_BIT_ERR_MASK (start bit error) bit, [21-70](#)
- start bit error (START_BIT_ERR) bit, [21-65](#)
- start bit error (START_BIT_ERR_MASK) bit, [21-70](#)
- START_BIT_ERR (start bit error) bit, [21-65](#)
- START_BIT_ERR_STAT (clear start bit error) bit, [21-68](#)
- state transitions, RTC, [24-16](#)
- station management
 - read transfer, MAC, [22-48](#)
 - transfer done, [22-40](#)
 - write transfer, MAC, [22-47](#)
- station management address (EMAC_STAADD) register, [22-72](#)
- station management data (EMAC_STADAT) register, [22-74](#)
- station management operation code (STAOP) bit, [22-72](#)

Index

- station management transfer done int en (STAIE) bit, 22-72
- station management transfer done interrupt status (STMDONE) bit, 22-91
- status (CNT_STATUS) register, 13-19, 13-22
- STATUS (status) bits, 15-28, 15-29
- status (STATUS) bits, 15-29
- STB (stop bits) bit, 15-21
- STDVAL (slave transmit data valid) bit, 16-26, 16-27
- stereo serial
 - data, 19-3
 - device, SPORT connection, 19-8
 - frame sync modes, 19-17
 - operation, SPORT, 19-10
- stick parity (SP) bit, 15-21
- sticky frame status (EMAC_RX_STKY) register, 22-99
- sticky frame status (EMAC_TX_STKY) register, 22-109
- sticky receive overflow status (ROVF) bit, 19-60
- sticky receive underflow status (RUVF) bit, 19-60
- STMDONE (station management transfer done interrupt status) bit, 22-90, 22-91
- stop bits (STB) bit, 15-21
- STOPCK (stop clock) bit, 8-21
- stop clock (STOPCK) bit, 8-21
- STOP (issue stop condition) bit, 16-30, 16-31
- stop mode, DMA, 6-11, 6-68
- stopping DMA transfers, 6-28
- stopwatch count[15: 0] field, 24-21
- stopwatch event flag bit, 24-21
- stopwatch interrupt enable bit, 24-20
- stopwatch operation, 24-2, 24-13
- STP (stick parity) bit, 15-21
- streams, memory DMA, 6-7
- subbank access bits, 2-7
- subbanks, L1 memory, 2-3
- supervisor mode, 26-7
- surface-mount capacitors, 27-10
- SWRESET (software reset) bit, 25-47, 25-48, 26-74
- SWRST (software reset) register, 25-39, 26-71, 26-72
- synchronization, descriptor queue, 6-57
- synchronization, DMA, 6-50 to 6-60
- synchronization, interrupt-based methods, 6-51
- synchronized transition, DMA, 6-27
- synchronous serial data transfer, 19-4
- synchronous serial ports. *See* SPORT
- SYNC (work unit transitions) bit, 6-25, 6-26, 6-61, 6-67, 6-69, 15-17
- SYSCR (system reset configuration) register, 25-47, 25-48, 25-53, 25-55, 26-73, 26-74
- system
 - events, 5-3
 - interrupts, 5-1, 5-2, 5-8, 10-7
 - peripherals, 1-5
 - system clock (SCLK), 1-28, 7-43, 8-1, 27-2
- system control (EMAC_SYSCTL) register, 22-89
- system design, 27-1 to 27-13
 - high frequency considerations, 27-8
 - recommendations, 27-9, 27-12
- system interrupt assignment (SIC_IAR) register, 5-11

- system interrupt controller (SIC), [5-2](#)
 - controlling interrupts, [5-4](#)
 - enabling flexible interrupt handling, [10-7](#)
 - enabling individual peripheral interrupts, [5-4](#)
 - main functions of, [5-4](#)
 - peripheral interrupt events, [5-18](#)
 - registers, [5-10](#)
 - system interrupt mask (SIC_IMASK)
 - register, [5-5](#)
 - system peripheral clock. *See* SCLK
 - system reset configuration (SYSCR)
 - register, [25-47](#), [25-48](#), [25-53](#), [25-55](#), [26-73](#), [26-74](#)
 - SYSTEM_RESET (system software reset)
 - bits, [25-39](#), [25-42](#), [26-72](#)
 - system select (SSEL) bit, [8-21](#)
 - system software reset, [26-4](#)
 - system software reset (SYSTEM_RESET)
 - bit, [26-72](#)
 - system status (EMAC_SYSTAT) register, [22-90](#)
 - SZ (send zero) bit, [17-20](#), [17-36](#)
- T**
- TAP registers
 - boundary-scan, [B-7](#)
 - BYPASS, [B-6](#)
 - instruction, [B-2](#), [B-4](#)
 - TAP (test access port), [B-1](#), [B-2](#)
 - target address, [26-27](#)
 - target alarm time high
 - (EMAC_PTP_ALARMHI) register, [23-33](#)
 - target alarm time low
 - (EMAC_PTP_ALARMLO) register, [23-32](#)
 - TAUTORLD bit, [11-3](#), [11-5](#)
 - TCKFE (clock drive/sample edge select)
 - bit, [19-32](#), [19-47](#), [19-51](#)
 - TCNTL (core timer control) register, [11-3](#), [11-5](#)
 - TCOUNT (core timer count) register, [11-3](#), [11-5](#), [11-6](#)
 - TCP/IP-style checksums, MAC, [22-21](#)
 - TCSR (temperature compensated self-refresh) bit, [7-65](#), [7-73](#)
 - TDM interfaces, [19-4](#)
 - TDTYPE[1: 0] field, [19-27](#), [19-47](#), [19-49](#)
 - temperature compensated self-refresh (TCSR) bit, [7-65](#)
 - TEMT (TSR and UART_THR empty)
 - bit, [15-24](#)
 - TEMT (TSR and UARTx_THR empty)
 - bit, [15-6](#), [15-25](#)
 - termination, DMA, [6-28](#)
 - terminations, SPORT pin/line, [19-8](#)
 - test access port (TAP), [B-1](#), [B-2](#)
 - test clock (TCK), [B-6](#)
 - test features, [B-1](#) to [B-7](#)
 - testing circuit boards, [B-1](#), [B-6](#)
 - test-logic-reset state, [B-4](#)
 - test point access, [27-12](#)
 - TESTSET instruction, [4-10](#), [27-3](#)
 - TE (transmitter enable) bit, [22-42](#), [22-50](#), [22-60](#), [22-64](#)
 - TFSR (transmit frame sync required select)
 - bit, [19-29](#), [19-30](#), [19-47](#), [19-50](#)
 - TFSx pins, [19-5](#), [19-18](#), [19-29](#), [19-36](#)
 - THR empty (THRE) bit, [15-24](#)
 - THRE (THR empty) bit, [15-6](#), [15-11](#), [15-16](#), [15-24](#), [15-25](#)
 - throughput
 - DAB, [4-10](#)
 - DMA, [6-40](#), [6-42](#)
 - general-purpose ports, [9-13](#)
 - SPORT, [19-5](#)

Index

- TIMDISx (timer disable) bits, [10-36](#),
[10-37](#)
- time-division-multiplexed (TDM) mode,
[19-14](#)
- See also* SPORT, multichannel operation
- TIMENx (timer enable) bits, [10-35](#), [10-36](#)
- time offset (EMAC_PTP_OFFSET)
register, [23-25](#)
- timer clock select (CLK_SEL) bit, [10-40](#)
- TIMER_CONFIG (timer configuration)
registers, [10-5](#), [10-38](#), [10-40](#)
- timer configuration (TIMER_CONFIG)
registers, [10-5](#), [10-38](#), [10-40](#)
- timer counter[15: 0] field, [10-41](#)
- timer counter[31: 16] field, [10-41](#)
- timer counter overflow (TOVF_ERRx)
bits, [10-39](#)
- TIMER_COUNTER (timer counter)
register, [10-41](#)
- TIMER_COUNTER (timer counter)
registers, [10-4](#), [10-40](#), [10-41](#)
- timer counter (TIMER_COUNTER)
registers, [10-4](#), [10-40](#), [10-41](#)
- timer disable (TIMDISx) bits, [10-37](#)
- TIMER_DISABLE (timer disable) register,
[10-5](#), [10-36](#), [10-37](#), [10-44](#)
- timer disable (TIMER_DISABLE) register,
[10-5](#), [10-36](#), [10-37](#)
- TIMER_ENABLE bit, [10-44](#)
- timer enable (TIMENx) bits, [10-35](#), [10-36](#)
- TIMER_ENABLE (timer enable) register,
[10-5](#), [10-35](#), [10-36](#), [20-23](#)
- timer enable (TIMER_ENABLE) register,
[10-5](#), [10-35](#), [10-36](#)
- timer errors, [10-8](#)
- timer input select (TIN_SEL) bit, [10-40](#),
[10-45](#)
- timer int (TIMILx) bits, [10-4](#), [10-39](#)
- timer mode (TMODE) bit, [10-40](#)
- timer period[15: 0] field, [10-43](#)
- timer period[31: 16] field, [10-43](#)
- timer period bits, [10-43](#)
- TIMER_PERIOD (timer period) register,
[10-42](#), [10-43](#)
- TIMER_PERIOD (timer period) registers,
[10-4](#), [10-42](#), [10-43](#)
- timer period (TIMER_PERIOD) registers,
[10-4](#), [10-42](#), [10-43](#)
- timers, [10-1](#) to [10-56](#)
 - core, [11-1](#) to [11-8](#)
 - EXT_CLK mode, [10-31](#) to [10-33](#)
 - overview, [1-23](#)
 - watchdog, [1-27](#), [12-1](#) to [12-10](#)
 - WIDTH_CAP mode, [15-15](#)
- timer slave enable status (TRUNx) bits,
[10-39](#)
- TIMER_STATUS (timer status) register,
[10-5](#), [10-37](#), [10-39](#)
- timer status (TIMER_STATUS) register,
[10-5](#), [10-37](#), [10-39](#)
- timer width[15: 0] field, [10-44](#)
- timer width[31: 16] field, [10-44](#)
- timer width bits, [10-44](#)
- TIMER_WIDTH (timer width) register,
[10-44](#)
- TIMER_WIDTH (timer width) registers,
[10-42](#), [10-44](#)
- timer width (TIMER_WIDTH) registers,
[10-42](#), [10-44](#)
- timestamp lock control (TL) bit, [23-13](#)
- TIMILx (timer int) bits, [10-4](#), [10-39](#)
- timing
 - auto-refresh, [7-58](#)
 - memory DMA, [6-44](#)
 - multichannel transfer, [19-16](#)
 - peripherals, [4-4](#)
 - SDRAM, [7-34](#)
 - SPI, [17-6](#)
 - SPORT, [19-38](#)

- TIMOD (transfer initiation mode) bits,
 - 17-16, 17-18, 17-35, 17-36
- TIN_SEL (timer input select) bit, 10-40, 10-45
- TINT bit, 11-3, 11-5
- TLSBIT (bit order select) bit, 19-47, 19-49
- TL (timestamp lock control) bit, 23-13
- TMODE (timer mode) bits, 10-11, 10-40, 10-44
- TMPWR bit, 11-3, 11-5
- TMRCLK input, 10-56
- TMREN bit, 11-3, 11-5
- TMR pin, 10-45
- TMRx pins, 10-3, 10-15
- TOGGLE_HI bit, 10-40, 10-45
- TOGGLE_HI mode, 10-16
- toggle Pxn bit, 9-35
- toggle Pxn interrupt A enable bit, 9-43
- toggle Pxn interrupt B enable bit, 9-44
- tools, development, 1-31
- TOVF_ERRx (timer counter overflow)
 - bits, 10-4, 10-7, 10-15, 10-24, 10-28, 10-38, 10-39, 10-46
- TOVF (transmit overflow status) bit, 19-57, 19-60, 19-61
- TPERIOD (core timer period) register, 11-6
- TPOLC (IrDA TX polarity change) bit, 15-31, 15-32
- traffic control, DMA, 6-44 to 6-50
- transfer count (PPI_COUNT) register, 20-32, 20-33
- transfer frame protocol, MAC, 22-9
- transfer initiate command, 17-17, 17-18
- transfer initiation from SPI master, 17-18
- transfer initiation mode (TIMOD) bit, 17-36
- transfer rate
 - memory DMA channels, 6-42
 - peripheral DMA channels, 6-42
- transfers, memory-to-memory, 6-7
- transfer type (XFR_TYPE) bit, 20-26
- transitions
 - continuous DMA, 6-24
 - DMA work unit, 6-24
 - operating mode, 8-10, 8-12
 - synchronized DMA, 6-24
- transmission aborted frames counter int en (TX_ABORT_CNT) bit, 22-118
- transmission aborted frames counter int (TX_ABORT_CNT) bit, 22-116
- transmission error, SPI, 17-41
- transmission error (TXE) bit, 17-39, 19-60
- transmit buffer empty int en (ETBEI) bit, 15-27
- transmit buffer flush (XMTFLUSH) bit, 16-36
- transmit buffer interrupt length (XMTINTLEN) bit, 16-36
- transmit clock, serial (TSCLKx) pins, 19-28
- transmit collision error, SPI, 17-41
- transmit collision error (TXCOL) bit, 17-39
- transmit command completion signal disable (CEATA_TX_CCSD) bit, 21-74
- transmit complete (TX_COMP) bit, 22-105
- transmit configuration (SPORT_TCRx) registers, 19-46
- transmit data[15: 0] field, 19-57
- transmit data[31: 16] field, 19-57
- transmit data available (TX_DAT_RDY) bit, 21-65
- transmit data available (TX_DAT_RDY_MASK) bit, 21-70
- transmit data buffer bits, 17-42
- transmit enable (TSPEN) bit, 19-47

Index

- transmit FIFO 16-bit data
(XMTDATA16) bit, [16-44](#)
- transmit FIFO 8-bit data (XMTDATA)
bit, [16-43](#)
- transmit FIFO empty (TX_FIFO_ZERO)
bit, [21-65](#)
- transmit FIFO empty
(TX_FIFO_ZER/W_MASK) bit,
[21-70](#)
- transmit FIFO full (TX_FIFO_FULL) bit,
[21-65](#)
- transmit FIFO full
(TX_FIFO_FULL_MASK) bit,
[21-70](#)
- transmit FIFO service interrupt mask
(XMTSERVM) bit, [16-40](#)
- transmit FIFO service (XMTSERV) bit,
[16-40](#)
- transmit FIFO status (XMTSTAT) bit,
[16-38](#)
- transmit FIFO underrun error
(TX_UNDERRUN) bit, [21-65](#)
- transmit FIFO underrun error
(TX_UNDERRUN_MASK) bit,
[21-70](#)
- transmit FIFO watermark
(TX_FIFO_STAT) bit, [21-65](#)
- transmit FIFO watermark
(TX_FIFO_STAT_MASK) bit,
[21-70](#)
- transmit frame sync required select (TFSR)
bit, [19-47](#)
- transmit hold bits, [15-26](#)
- transmit hold register empty (TXHRE) bit,
[19-60](#)
- transmit int en (TXIE) bit, [23-15](#)
- transmit overflow status (TOVF) bit, [19-60](#)
- transmit overrun error int en (TXOVE) bit,
[23-15](#)
- transmit snapshot high
(EMAC_PTP_TXSNAPHI) register,
[23-31](#)
- transmit snapshot low
(EMAC_PTP_TXSNAPLO) register,
[23-30](#)
- transmit snapshot overrun status (TXOV)
bit, [23-17](#)
- transmit snapshot status (TXTL) bit, [23-17](#)
- transmit stereo frame sync enable (TSFSE)
bit, [19-48](#)
- transmitter enable (TE) bit, [22-64](#)
- transmit underflow status (TUVF) bit,
[19-60](#)
- TRAS (bank activate command delay) bits,
[7-65](#), [7-67](#)
- t_{RAS} operation, [7-35](#)
- t_{RCD} operation, [7-36](#)
- TRCD (RAS to CAS delay) bits, [7-65](#), [7-68](#)
- t_{RC} operation, [7-37](#)
- t_{REFI} operation, [7-38](#)
- t_{REF} operation, [7-38](#)
- t_{RFC} operation, [7-37](#)
- TRFST (left/right order) bit, [19-48](#), [19-51](#)
- triggering DMA transfers, [6-60](#)
- TRP (bank precharge delay) bits, [7-65](#),
[7-68](#)
- t_{RP} operation, [7-37](#)
- t_{RRD} operation, [7-36](#)
- TRUN_x (timer slave enable status) bits,
[10-21](#), [10-37](#), [10-39](#), [10-46](#)
- TSCALE (core timer scale) register, [11-3](#),
[11-7](#)
- TSCLK_x signal, [19-5](#)
- TSFSE (transmit stereo frame sync enable)
bit, [19-9](#), [19-10](#), [19-48](#), [19-51](#)
- TSPEN (transmit enable) bit, [19-46](#),
[19-47](#), [19-48](#)
- TSR and UART_THR empty (TEMT)
bit, [15-24](#)

- TUVF (transmit underflow status) bit,
 - [19-36, 19-57, 19-60, 19-61](#)
- TWI, [1-14, 16-2 to 16-58](#)
 - block diagram, [16-3](#)
 - bus arbitration, [16-8](#)
 - clock generation, [16-7](#)
 - clock setup, [16-12](#)
 - controller, [16-2](#)
 - electrical specifications, [16-58](#)
 - fast mode, [16-10](#)
 - features, [16-2](#)
 - general call address, [16-9](#)
 - general setup, [16-10](#)
 - I²C compatibility, [1-14](#)
 - master mode, [16-12, 16-14](#)
 - peripheral interface, [16-5](#)
 - pins, [16-5](#)
 - slave mode, [16-11](#)
 - start and stop conditions, [16-8](#)
 - synchronization, [16-7](#)
 - transfer protocol, [16-6](#)
- TWI_CLKDIV (SCL clock divider)
 - register, [16-25](#)
- TWI_CONTROL (TWI control) register,
 - [16-4, 16-25](#)
- TWI control (TWI_CONTROL) register,
 - [16-25](#)
- TWI_DT (drive/tolerate) pin, [9-28](#)
- TWI enable (TWI_ENA) bit, [16-25](#)
- TWI_ENA (TWI enable) bit, [16-25](#)
- TWI FIFO control (TWI_FIFO_CTL)
 - register, [16-36](#)
- TWI_FIFO_CTL (TWI FIFO control)
 - register, [16-36](#)
- TWI FIFO receive data double byte
 - (TWI_RCV_DATA16) register,
 - [16-46](#)
- TWI FIFO receive data single byte
 - (TWI_RCV_DATA8) register, [16-45](#)
- TWI_FIFO_STAT (TWI FIFO status)
 - register, [16-38](#)
- TWI FIFO status (TWI_FIFO_STAT)
 - register, [16-38](#)
- TWI FIFO transmit data double byte
 - (TWI_XMT_DATA16) register,
 - [16-44](#)
- TWI FIFO transmit data single byte
 - (TWI_XMT_DATA8) register,
 - [16-43](#)
- TWI interrupt mask (TWI_INT_MASK)
 - register, [16-40](#)
- TWI interrupt status (TWI_INT_STAT)
 - register, [16-40](#)
- TWI_INT_MASK (TWI interrupt mask)
 - register, [16-40](#)
- TWI_INT_STAT (TWI interrupt status)
 - register, [16-40](#)
- TWI_MASTER_ADDR (TWI master mode address) register, [16-32](#)
- TWI_MASTER_CTL (TWI master mode control) register, [16-29](#)
- TWI master mode address
 - (TWI_MASTER_ADDR) register,
 - [16-32](#)
- TWI master mode control
 - (TWI_MASTER_CTL) register,
 - [16-29](#)
- TWI master mode status
 - (TWI_MASTER_STAT) register,
 - [16-33](#)
- TWI_MASTER_STAT (TWI master mode status) register, [16-33](#)
- TWI_RCV_DATA16 (TWI FIFO receive data double byte) register, [16-46](#)
- TWI_RCV_DATA8 (TWI FIFO receive data single byte) register, [16-45](#)
- TWI_SLAVE_ADDR (TWI slave mode address) register, [16-28](#)

Index

- TWI_SLAVE_CTL (slave mode control) register, [16-26](#)
- TWI slave mode address (TWI_SLAVE_ADDR) register, [16-28](#)
- TWI slave mode status (TWI_SLAVE_STAT) register, [16-29](#)
- TWI_SLAVE_STAT (TWI slave mode status) register, [16-29](#)
- TWI_XMT_DATA16 (TWI FIFO transmit data double byte) register, [16-44](#)
- TWI_XMT_DATA8 (TWI FIFO transmit data single byte) register, [16-43](#)
- two-dimensional DMA, [6-11](#)
- two-wire interface. *See* TWI
- †WR operation, [7-36](#)
- TWR (write to precharge delay) bits, [7-65](#), [7-69](#)
- TX_ABORTC_CNT (frames aborted due to excess collisions counter int) bit, [22-116](#)
- TX_ABORTC_CNT (frames aborted due to excess collisions counter int en) bit, [22-118](#)
- TX_ABORT_CNT (transmission aborted frames counter int) bit, [22-116](#)
- TX_ABORT_CNT (transmission aborted frames counter int en) bit, [22-118](#)
- TX_ACT (data transmit active) bit, [21-65](#)
- TX_ACT_MASK (data transmit active) bit, [21-70](#)
- TX_ALLF_CNT (frames transmitted all counter int) bit, [22-116](#)
- TX_ALLF_CNT (frames transmitted all counter int en) bit, [22-118](#)
- TX_ALLO_CNT (octets transmitted all counter int) bit, [22-116](#)
- TX_ALLO_CNT (octets transmitted all counter int en) bit, [22-118](#)
- TX broadcast frame detected (TX_BROAD) bit, [22-109](#), [22-112](#)
- TX_BROAD_CNT (broadcast frames transmitted OK counter int) bit, [22-116](#)
- TX_BROAD_CNT (broadcast frames transmitted OK counter int en) bit, [22-118](#)
- TX_BROAD (TX broadcast frame detected) bit, [22-109](#)
- TX_BROAD (TX broadcast frame int en) bit, [22-112](#)
- TX_CCNT[3: 0] field, [22-105](#), [22-107](#)
- TXCOL flag, [17-41](#)
- TXCOL (transmit collision error) bit, [17-39](#)
- TX_COMP (frame transmission complete) bit, [22-105](#)
- TX_COMP (frame transmissions complete) bit, [22-109](#)
- TX_COMP (frame transmit complete int en) bit, [22-112](#)
- TX_CRS_CNT (carrier sense errors counter int) bit, [22-116](#)
- TX_CRS_CNT (carrier sense errors counter int en) bit, [22-118](#)
- TX_CRS (no carrier) bit, [22-105](#), [22-109](#), [22-110](#)
- TX_CRS (no carrier int en) bit, [22-112](#)
- TX_DAT_RDY_MASK (transmit data available) bit, [21-70](#)
- TX_DAT_RDY (transmit data available) bit, [21-65](#)
- TX_DEFER_CNT (frames with deferred transmission counter int) bit, [22-116](#)
- TX_DEFER_CNT (frames with deferred transmission counter int en) bit, [22-110](#), [22-118](#)

- TX_DEFER (frame deferral int en) bit, [22-105](#), [22-109](#), [22-112](#)
- TX DMA direction error detected, [22-40](#)
- TX DMA direction error status (TXDMAERR) bit, [22-91](#)
- TXDMAERR (TX DMA direction error status) bit, [22-29](#), [22-91](#)
- TX_DMAU (DMA underrun) bit, [22-105](#)
- TXDWA (transmit frame DMA word alignment) bit, [22-27](#)
- TX_ECOLL (excessive collision errors detected) bit, [22-105](#), [22-109](#)
- TX_ECOLL (excessive collision errors int en) bit, [22-112](#)
- TX_EDEFER (excessive deferrals detected) bit, [22-105](#), [22-109](#)
- TX_EDEFER (excessive deferrals int en) bit, [22-112](#)
- TX_EQ64_CNT (frames length equal to 64 transmitted counter int) bit, [22-116](#)
- TX_EQ64_CNT (frames length equal to 64 transmitted counter int en) bit, [22-118](#)
- TX_ER pin, [22-5](#)
- TXE (transmission error) bit, [17-39](#), [17-41](#), [19-57](#), [19-60](#)
- TX_EXDEF_CNT (frames with excessive deferral counter int) bit, [22-116](#)
- TX_EXDEF_CNT (frames with excessive deferral counter int en) bit, [22-118](#)
- TX_FIFO_FULL_MASK (transmit FIFO full) bit, [21-70](#)
- TX_FIFO_FULL (transmit FIFO full) bit, [21-65](#)
- TX_FIFO_STAT_MASK (transmit FIFO watermark) bit, [21-70](#)
- TX_FIFO_STAT (transmit FIFO watermark) bit, [21-65](#)
- TX_FIFO_ZERO (transmit FIFO empty) bit, [21-65](#)
- TX_FIFO_ZER/W_MASK (transmit FIFO empty) bit, [21-70](#)
- TX frame status interrupt, [22-39](#)
- TX frame-status interrupt status (TXFSINT) bit, [22-91](#)
- TX_FRLLEN (frame length) bits, [22-105](#)
- TXFSINT (TX frame-status interrupt status) bit, [22-91](#)
- TXF (transmit FIFO full status) bit, [19-60](#)
- TX_GE1024_CNT (frames length 1024-max transmitted counter int) bit, [22-116](#)
- TX_GE1024_CNT (frames length 1024-max transmitted counter int en) bit, [22-118](#)
- TX hold register, [19-56](#)
- TXHRE (transmit hold register empty) bit, [19-60](#)
- TXIE (transmit int en) bit, [23-15](#)
- TX_LATE_CNT (late collisions counter int) bit, [22-116](#)
- TX_LATE_CNT (late collisions counter int en) bit, [22-118](#)
- TX_LATE (late collision error) bit, [22-105](#), [22-109](#)
- TX_LATE (late collision error int en) bit, [22-112](#)
- TX_LOSS (losses of carrier detected) bit, [22-105](#), [22-109](#), [22-110](#)
- TX_LOSS (loss of carrier int en) bit, [22-112](#)
- TX_LOST_CNT (frames lost due to internal MAC transmit error counter int) bit, [22-116](#)
- TX_LOST_CNT (frames lost due to internal MAC transmit error counter int en) bit, [22-118](#)

Index

- TX_LT1024_CNT (frames length 512-1023 transmitted counter int) bit, [22-116](#)
- TX_LT1024_CNT (frames length 512-1023 transmitted counter int en) bit, [22-118](#)
- TX_LT128_CNT (frames length 65-127 transmitted counter int) bit, [22-116](#)
- TX_LT128_CNT (frames length 65-127 transmitted counter int en) bit, [22-118](#)
- TX_LT256_CNT (frames length 128-255 transmitted counter int) bit, [22-116](#)
- TX_LT256_CNT (frames length 128-255 transmitted counter int en) bit, [22-118](#)
- TX_LT512_CNT (frames length 256-511 transmitted counter int) bit, [22-116](#)
- TX_LT512_CNT (frames length 256-511 transmitted counter int en) bit, [22-118](#)
- TX_MACCTL_CNT (MAC control frames transmitted counter int) bit, [22-116](#)
- TX_MACCTL_CNT (MAC control frames transmitted counter int en) bit, [22-118](#)
- TX_MACE (internal MAC error int en) bit, [22-112](#)
- TX_MACE (internal MAC errors detected) bit, [22-109](#)
- TX_MCOLL_CNT (multiple collision frames counter int) bit, [22-116](#)
- TX_MCOLL_CNT (multiple collision frames counter int en) bit, [22-118](#)
- TX multicast, TX broadcast[1: 0] field, [22-105](#)
- TX multicast frame int en (TX_MULTI) bit, [22-112](#)
- TX_MULTI_CNT (multicast frames transmitted OK counter int) bit, [22-116](#)
- TX_MULTI_CNT (multicast frames transmitted OK counter int en) bit, [22-118](#)
- TX_MULTI (multicast frame detected) bit, [22-109](#)
- TX_MULTI (TX multicast frame int en) bit, [22-112](#)
- TX_OCTET_CNT (octets transmitted OK counter int) bit, [22-116](#)
- TX_OCTET_CNT (octets transmitted OK counter int en) bit, [22-118](#)
- TX_OK_CNT (frames transmitted OK counter int) bit, [22-116](#)
- TX_OK_CNT (frames transmitted OK counter int en) bit, [22-118](#)
- TX_OK (frames transmitted OK) bit, [22-105](#), [22-109](#)
- TX_OK (frame transmit OK int en) bit, [22-112](#)
- TXOVE (transmit overrun error int en) bit, [23-15](#)
- TXOV (transmit snapshot overrun status) bit, [23-17](#)
- TXREQ signal, [15-6](#)
- TX_RETRY (late collisions detected) bit, [22-105](#), [22-109](#)
- TX_RETRY (late collisions int en) bit, [22-112](#)
- TX retry on late collision en (LCRTE) bit, [22-62](#)
- TX_SCOLL_CNT (single collision frames counter int) bit, [22-116](#)
- TX_SCOLL_CNT (single collision frames counter int en) bit, [22-118](#)
- TxSEC enable (TXSE) bit, [19-48](#)
- TXSE (TxSEC enable) bit, [19-48](#), [19-51](#)

TX_SNAP_HI (higher 32 bits of transmit snapshot of local clock time) bit, 23-31

TX_SNAP_LO (lower 32 bits of transmit snapshot of local clock time) bit, 23-30

TXSR operation, 7-38

TXS (SPI_TDBR data buffer status) bit, 17-22, 17-39

TXTL (transmit snapshot status) bit, 23-17

TX_UNDERRUN_MASK (transmit FIFO underrun error) bit, 21-70

TX_UNDERRUN_STAT (clear transmit FIFO underrun error) bit, 21-68

TX_UNDERRUN (transmit FIFO underrun error) bit, 21-65

TX_UNI_CNT (unicast frames transmitted OK counter int) bit, 22-116

TX_UNI_CNT (unicast frames transmitted OK counter int en) bit, 22-118

type definitions, MAC, 22-121

typed frame int en (RX_TYPE) bit, 22-94, 22-100, 22-104

typed frames received counter int en (RX_TYPED_CNT) bit, 22-115

typed frames received counter int (RX_TYPED_CNT) bit, 22-113

typed frames received (EMAC_RXC_TYPED) register, 22-55

U

UART, 1-24, 15-2 to 15-42

- autobaud detection, 15-13
- baud rate, 15-7, 15-13
- bit rate, 15-12, 15-13
- bitstream, 15-5
- booting, 15-14
- character transmission, 15-36
- clock, 4-4, 15-8, 15-12
- connected to peripheral access bus, 15-4
- data communication via infrared signals, 15-5
- data words, 15-5
- diagram, 15-3
- divisor reset, 15-30
- DMA operation, 15-17, 15-18, 15-19, 15-27
- errors during reception, 15-7
- examples, 15-32
- external interfaces, 15-3
- features, 15-2
- glitch filtering, 15-9
- initialization, 15-32
- interrupts, 15-11, 15-12, 15-28, 15-38
- IrDA mode, 15-2, 15-8, 15-9, 15-10
- loopback mode, 15-23
- mixing modes, 15-18
- modem status, 15-4
- non-DMA operation, 15-15, 15-18, 15-19, 15-38
- overview, 1-24
- polling, 15-37
- receive operation, 15-7, 15-10
- registers, 15-6, 15-19
- sampling, 15-8, 15-10
- standard compliance, 15-2
- strings, 15-36
- SYNC bit usage, 15-40
- transmissions, 15-6, 15-36, 15-40

Index

- UART controller enable (UCEN) bit, [15-31](#)
- UART divisor latch high byte (UART_DLH) register, [15-30](#)
- UART divisor latch low byte (UART_DLL) register, [15-30](#)
- UART_DLH (UART divisor latch high byte) register, [15-30](#)
- UART_DLH (UART divisor latch high byte) registers, [15-20](#)
- UART_DLL (UART divisor latch low byte) register, [15-30](#)
- UART_DLL (UART divisor latch low byte) registers, [15-19](#), [15-30](#)
- UART_GCTL (UART global control) register, [15-31](#)
- UART_GCTL (UART global control) registers, [15-20](#), [15-31](#)
- UART global control (UART_GCTL) register, [15-31](#)
- UART_IER (UART int en) register, [15-27](#)
- UART_IER (UART int en) registers, [15-20](#), [15-26](#)
- UART_IIR (UART interrupt identification) register, [15-29](#)
- UART_IIR (UART interrupt identification) registers, [15-12](#), [15-20](#)
- UART int en (UART_IER) register, [15-27](#)
- UART interrupt identification (UART_IIR) register, [15-29](#)
- UART_LCR (UART line control) register, [15-21](#)
- UART_LCR (UART line control) registers, [15-6](#), [15-20](#), [15-21](#)
- UART line control (UART_LCR) register, [15-21](#)
- UART line status (UART_LSR) register, [15-24](#)
- UART_LSR (UART line status) register, [15-24](#)
- UART_LSR (UART line status) registers, [15-20](#), [15-24](#)
- UART_MCR (UART modem control) register, [15-23](#)
- UART_MCR (UART modem control) registers, [15-20](#)
- UART modem control (UART_MCR) register, [15-23](#)
- UART_RBR (UART receive buffer) register, [15-26](#)
- UART_RBR (UART receive buffer) registers, [15-7](#), [15-19](#)
- UART receive buffer registers (UARTx_RBR), [15-7](#)
- UART receive buffer (UART_RBR) register, [15-26](#)
- UART scratch (UART_SCR) register, [15-31](#)
- UART_SCR (UART scratch) register, [15-31](#)
- UART_SCR (UART scratch) registers, [15-20](#), [15-31](#)
- UART_THR (UART transmit holding) register, [15-26](#)
- UART_THR (UART transmit holding) registers, [15-6](#), [15-19](#)
- UART transmit holding (UART_THR) register, [15-26](#)
- UCEN (enable UART clocks) bit, [15-6](#), [15-12](#), [15-30](#), [15-32](#)
- UCEN (UART controller enable) bit, [15-31](#)
- UCIE (up count int en) bit, [13-21](#)
- UCII (up count interrupt identifier) bit, [13-22](#)
- UDPEM (compare mask of UDP event port field) bit, [23-13](#)
- UDP event port matching value (UDP_EVP) bit, [23-21](#)

- UDP_EVP (UDP event port matching value) bit, [23-21](#)
 - UDP general port matching value (UDP_GNP) bit, [23-21](#)
 - UDP_GNP (UDP general port matching value) bit, [23-21](#)
 - UEVOF (offset of UDP event port field) bit, [23-19](#)
 - UNDR (FIFO underrun) bit, [20-30](#), [20-31](#)
 - unframed/framed, serial data, [19-29](#)
 - unicast frames received OK counter int en (RX_UNI_CNT) bit, [22-115](#)
 - unicast frames received OK counter int (RX_UNI_CNT) bit, [22-113](#)
 - unicast frames received ok (EMAC_RXC_UNICST) register, [22-53](#)
 - unicast frames transmitted OK counter int en (TX_UNI_CNT) bit, [22-118](#)
 - unicast frames transmitted OK counter int (TX_UNI_CNT) bit, [22-116](#)
 - unicast frames xmitted ok (EMAC_TXC_UNICST) register, [22-58](#)
 - universal asynchronous receiver/transmitter. *See* UART
 - unpopulated memory, [7-9](#)
 - unsupported control frame (RX_UCTL) bit, [22-95](#), [22-100](#), [22-104](#)
 - unsupported opcodes received counter int en (RX_OPCODE_CNT) bit, [22-115](#)
 - unsupported opcodes received counter interrupt (RX_OPCODE_CNT) bit, [22-113](#)
 - unsupported opcodes received (EMAC_RXC_OPCODE) register, [22-55](#)
 - unused pins, [27-14](#)
 - up count int en (UCIE) bit, [13-21](#)
 - up count interrupt identifier (UCII) bit, [13-22](#)
 - Upper PBS00 Half Page (PBS00H, Bits 31–0), [26-82](#), [26-83](#), [26-84](#), [26-85](#), [26-91](#), [26-92](#)
 - upper PBS00 half page (PBS00H, bits 31–0) register, [26-82](#)
 - Upper PBS00 Half Page (PBS00H, Bits 63–32), [26-81](#)
 - upper PBS00 half page (PBS00H, bits 63–32) register, [26-81](#)
 - upper PBS01 half page (PBS01H, bits 15–0) register, [26-84](#)
 - upper PBS01 half page (PBS01H, bits 63–16) register, [26-83](#)
 - urgency threshold enable (UTE) bit, [6-39](#), [6-83](#)
 - user mode, [26-7](#)
 - UTE (urgency threshold enable) bit, [6-39](#), [6-83](#)
 - UTHE[15: 0] field, [6-87](#)
- ## V
- VCO, multiplication factors, [8-4](#)
 - VCO signal, [8-1](#)
 - VDDEXT pins, [27-9](#)
 - VDDINT pins, [27-9](#)
 - VDK, [1-32](#)
 - vertical blanking, [20-6](#)
 - vertical blanking interval only submode, [20-10](#)
 - video frame partitioning, [20-7](#)
 - video streams
 - CIF, [20-8](#)
 - NTSC, [20-5](#)
 - PAL, [20-5](#)
 - QCIF, [20-8](#)
 - VisualDSP++, [26-19](#)
 - debugger, [1-32](#)
 - development environment, [1-31](#)

Index

VLAN1 frame int en (RX_VLAN1) bit,
22-94, 22-100, 22-104
VLAN1TAG[15: 0] field, 22-77
VLAN1 tag (EMAC_VLAN1) register,
22-76
VLAN2 frame detected (RX_VLAN2) bit,
22-100, 22-104
VLAN2 frames detected (RX_VLAN2) bit,
22-93
VLAN2TAG[15: 0] field, 22-77
VLAN2 tag (EMAC_VLAN2) register,
22-76
voltage, 8-15
 control, 8-7
 dynamic control, 8-15
voltage controlled oscillator (VCO), 8-3
voltage regulator control (VR_CTL)
 register, 8-20, 8-23
VR_CTL (voltage regulator control)
 register, 8-20, 8-23
VR registers, 8-19
VSTAT bit, 8-22

W

W1C operations, 6-10
W1LCNT_MAX bit, 13-24
W1LCNT_MIN bit, 13-24
W1LCNT_ZERO bit, 13-24
W1LMAX_CNT bit, 13-24
W1LMAX_MIN bit, 13-24
W1LMAX_ZERO bit, 13-24
W1LMIN_CNT bit, 13-24
W1LMIN_ZERO bit, 13-24
W1ZMONCE bit, 13-24
wait for response (CMD_RSP_EN) bit,
21-57
wait states, additional, 7-15
WAKEDET (wake-up detected status) bit,
22-34

WAKEDET (wakeup detected status) bit,
22-91
wake from hibernate, MAC, 22-30
wake from sleep, MAC, 22-31
WAKE (real time clock wake enable) bit,
8-23
wakeup detected status (WAKEDET) bit,
22-91
wakeup filter 0 address type bit, 22-85
wakeup filter 0 pattern CRC[15: 0] field,
22-88
wakeup filter 0 pattern offset[7: 0] field,
22-87
wakeup filter 1 address type bit, 22-85
wakeup filter 1 pattern CRC[15: 0] field,
22-88
wakeup filter 1 pattern offset[7: 0] field,
22-87
wakeup filter 2 address type bit, 22-85
wakeup filter 2 pattern CRC[15: 0] field,
22-88
wakeup filter 2 pattern offset[7: 0] field,
22-87
wakeup filter 3 address type bit, 22-85
wakeup filter 3 pattern CRC[15: 0] field,
22-88
wakeup filter 3 pattern offset[7: 0] field,
22-87
wakeup filter enable bits, 22-85
wakeup frame0 byte mask
 (EMAC_WKUP_FFMSK0) register,
22-81
wakeup frame1 byte mask
 (EMAC_WKUP_FFMSK1) register,
22-82
wakeup frame2 byte mask
 (EMAC_WKUP_FFMSK2) register,
22-83

- wakeup frame3 byte mask
(EMAC_WKUP_FFMSK3) register,
22-84
- wakeup frame control and status
(EMAC_WKUP_CTL) register,
22-77
- wakeup frame detected, 22-40
- wakeup frame filter commands
(EMAC_WKUP_FFCMD MAC)
register, 22-85
- wakeup frame filter CRC0/1
(EMAC_WKUP_FFCRC0) register,
22-87
- wakeup frame filter CRC2/3
(EMAC_WKUP_FFCRC1) register,
22-87
- wakeup frame filter offsets
(EMAC_WKUP_FFOFF) register,
22-87
- wakeup frame received status (RWKS) bit,
22-78
- wakeup frameX byte mask
(EMAC_WKUP_FFMSKx) registers,
22-80
- wakeup function, 5-7
- wake-up reset (WURESET) bit, 26-74
- watchdog control (WDOG_CTL) register,
12-7, 12-8
- watchdog count[15: 0] field, 12-6
- watchdog count[31: 16] field, 12-6
- watchdog count (WDOG_CNT) register,
12-5, 12-6
- watchdog reset (WDRESET) bit, 26-74
- watchdog status[15: 0] field, 12-7
- watchdog status[31: 16] field, 12-7
- watchdog status (WDOG_STAT) register,
12-3, 12-4, 12-6, 12-7
- watchdog timer, 1-27, 12-1 to 12-10
 - block diagram, 12-3
 - disabling, 12-5
 - and emulation mode, 12-2
 - enabling with zero value, 12-5
 - features, 12-2
 - internal interface, 12-3
 - overview, 1-27
 - registers, 12-5
 - reset, 12-5, 26-4, 26-6
 - starting, 12-4
- watchdog timer registers, 12-5
- waveform generation, pulse width
modulation, 10-13
- WDEN[7: 0] field, 12-7
- WDEN bit, 12-8
- WDEV[1: 0] field, 12-4, 12-7
- WDEV bit, 12-8
- WDOG_CNT (watchdog count) register,
12-5, 12-6
- WDOG_CTL (watchdog control) register,
12-7, 12-8
- WDOG_STAT (watchdog status) register,
12-3, 12-4, 12-6, 12-7
- WDRESET (watchdog reset) bit, 25-47,
25-48, 26-74
- WDSIZE[1: 0] field, 6-67, 6-70
- WDTH_CAP mode, 10-23, 10-42, 10-44
- WLS (word length select) bits, 15-21
- WNR (DMA direction) bit, 6-67, 6-70
- WOFF[9: 0] field, 19-20, 19-64
- WOM (write open drain master) bit,
17-14, 17-36
- word length
 - SPI, 17-35
 - SPORT, 19-26
 - SPORT receive data, 19-59
 - SPORT transmit data, 19-56
- word length select (WLS) bit, 15-21

Index

work unit
 completion, [6-22](#)
 DMA, [6-14](#)
 interrupt timing, [6-25](#)
 restrictions, [6-24](#)
 transitions, [6-24](#)
work unit transition (SYNC) bit, [6-67](#)
WRDO-W1C bit, [12-8](#)
write
 asynchronous, [7-13](#)
 command, [7-33](#)
 with data mask command, [7-49](#)
write access for EBIU asynchronous
 memory controller, [7-19](#)
write complete bit, [24-21](#)
write complete int en bit, [24-20](#)
write-one-to-clear (W1C) operations, [6-10](#)
write open drain master (WOM) bit, [17-36](#)
write operation, GPIO, [9-15](#)
write pending status bit, [24-21](#)
write to precharge delay (TWR) bit, [7-65](#)
WSIZE[3: 0] field, [19-20](#), [19-64](#)
WURESET, [25-47](#), [25-48](#)
WURESET (wake-up reset) bit, [26-74](#)

X

X_COUNT[15: 0] field, [6-75](#)
XFR_TYPE[1: 0] field, [20-4](#), [20-26](#),
 [20-28](#), [20-29](#)
XFR_TYPE (transfer type) bit, [20-26](#)
X_MODIFY[15: 0] field, [6-77](#)

XMTDATA16[15: 0] field, [16-44](#)
XMTDATA16 (transmit FIFO 16-bit
 data) bit, [16-44](#)
XMTDATA8[7: 0] field, [16-43](#)
XMTDATA (transmit FIFO 8-bit data)
 bit, [16-43](#)
XMTFLUSH (transmit buffer flush) bit,
 [16-36](#), [16-38](#)
XMTINTLEN (transmit buffer interrupt
 length) bit, [16-36](#), [16-37](#)
XMTSERVM (transmit FIFO service
 interrupt mask) bit, [16-40](#)
XMTSERV (transmit FIFO service) bit,
 [16-40](#), [16-41](#)
XMTSTAT[1: 0] field, [16-38](#), [16-39](#)
XMTSTAT (transmit FIFO status) bit,
 [16-38](#)

Y

YCbCr format, [20-27](#)
Y_COUNT[15: 0] field, [6-78](#)
Y_MODIFY[15: 0] field, [6-80](#)

Z

zero marker error int en (CZMEIE) bit,
 [13-21](#)
zero word (BK_ZEROS) register, [26-77](#)
ZMZC (CZM zeroes counter enable) bit,
 [13-20](#)