

TMS320C62x DSP CPU and Instruction Set Reference Guide

Literature Number: SPRU731A
May 2010





Read This First

About This Manual

The TMS320C6000™ digital signal processor (DSP) platform is part of the TMS320™ DSP family. The TMS320C62x™ DSP generation and the TMS320C64x™ DSP generation comprise fixed-point devices in the C6000™ DSP platform, and the TMS320C67x™ DSP generation comprises floating-point devices in the C6000 DSP platform. The C62x™ and C64x™ DSPs are code-compatible.

This document describes the CPU architecture, pipeline, instruction set, and interrupts of the C62x DSP.

Notational Conventions

This document uses the following conventions.

- Hexadecimal numbers are shown with the suffix h. For example, the following number is 40 hexadecimal (decimal 64): 40h.

Related Documentation From Texas Instruments

The following documents describe the C6000™ devices and related support tools. Copies of these documents are available on the Internet at www.ti.com.
Tip: Enter the literature number in the search box provided at www.ti.com.

The current documentation that describes the C6000 devices, related peripherals, and other technical collateral, is available in the C6000 DSP product folder at: www.ti.com/c6000.

TMS320C6000 DSP Peripherals Overview Reference Guide (literature number SPRU190) describes the peripherals available on the TMS320C6000™ DSPs.

TMS320C6000 Technical Brief (literature number SPRU197) gives an introduction to the TMS320C62x™ and TMS320C67x™ DSPs, development tools, and third-party support.

TMS320C6000 Programmer's Guide (literature number SPRU198) describes ways to optimize C and assembly code for the TMS320C6000™ DSPs and includes application program examples.

TMS320C6000 Chip Support Library API Reference Guide (literature number SPRU401) describes a set of application programming interfaces (APIs) used to configure and control the on-chip peripherals.

Trademarks

Code Composer Studio, C6000, C62x, C64x, C67x, TMS320C2000, TMS320C5000, TMS320C6000, TMS320C62x, TMS320C64x, TMS320C67x, and VelociTI are trademarks of Texas Instruments.

Trademarks are the property of their respective owners.

Contents

1	Introduction	1-1
	<i>Provides features and options of the TMS320C62x DSP. An overview of the DSP architecture is also provided.</i>	
1.1	TMS320 DSP Family Overview	1-2
1.2	TMS320C6000 DSP Family Overview	1-2
1.3	TMS320C62x DSP Features and Options	1-4
1.4	TMS320C62x DSP Architecture	1-6
1.4.1	Central Processing Unit (CPU)	1-7
1.4.2	Internal Memory	1-7
1.4.3	Memory and Peripheral Options	1-7
2	CPU Data Paths and Control	2-1
	<i>Provides information about the data paths and control registers. The two register files and the data cross paths are described.</i>	
2.1	Introduction	2-2
2.2	General-Purpose Register Files	2-2
2.3	Functional Units	2-5
2.4	Register File Cross Paths	2-6
2.5	Memory, Load, and Store Paths	2-6
2.6	Data Address Paths	2-7
2.7	Control Register File	2-7
2.7.1	Register Addresses for Accessing the Control Registers	2-8
2.7.2	Pipeline/Timing of Control Register Accesses	2-9
2.7.3	Addressing Mode Register (AMR)	2-10
2.7.4	Control Status Register (CSR)	2-13
2.7.5	Interrupt Clear Register (ICR)	2-16
2.7.6	Interrupt Enable Register (IER)	2-17
2.7.7	Interrupt Flag Register (IFR)	2-18
2.7.8	Interrupt Return Pointer Register (IRP)	2-19
2.7.9	Interrupt Set Register (ISR)	2-20
2.7.10	Interrupt Service Table Pointer Register (ISTP)	2-21
2.7.11	Nonmaskable Interrupt (NMI) Return Pointer Register (NRP)	2-22
2.7.12	E1 Phase Program Counter (PCE1)	2-22

3	Instruction Set	3-1
	<i>Describes the assembly language instructions of the TMS320C62x DSP. Also described are parallel operations, conditional operations, resource constraints, and addressing modes.</i>	
3.1	Instruction Operation and Execution Notations	3-2
3.2	Instruction Syntax and Opcode Notations	3-5
3.3	Delay Slots	3-6
3.4	Parallel Operations	3-7
	3.4.1 Example Parallel Code	3-9
	3.4.2 Branching Into the Middle of an Execute Packet	3-9
3.5	Conditional Operations	3-10
3.6	Resource Constraints	3-11
	3.6.1 Constraints on Instructions Using the Same Functional Unit	3-11
	3.6.2 Constraints on Cross Paths (1X and 2X)	3-11
	3.6.3 Constraints on Loads and Stores	3-12
	3.6.4 Constraints on Long (40-Bit) Data	3-13
	3.6.5 Constraints on Register Reads	3-14
	3.6.6 Constraints on Register Writes	3-15
3.7	Addressing Modes	3-16
	3.7.1 Linear Addressing Mode	3-16
	3.7.2 Circular Addressing Mode	3-17
	3.7.3 Syntax for Load/Store Address Generation	3-18
3.8	Instruction Compatibility	3-20
3.9	Instruction Descriptions	3-20
	ABS (Absolute Value With Saturation)	3-24
	ADD (Add Two Signed Integers Without Saturation)	3-26
	ADDAB (Add Using Byte Addressing Mode)	3-30
	ADDAH (Add Using Halfword Addressing Mode)	3-32
	ADDAW (Add Using Word Addressing Mode)	3-34
	ADDK (Add Signed 16-Bit Constant to Register)	3-36
	ADDU (Add Two Unsigned Integers Without Saturation)	3-37
	ADD2 (Add Two 16-Bit Integers on Upper and Lower Register Halves)	3-39
	AND (Bitwise AND)	3-41
	B (Branch Using a Displacement)	3-43
	B (Branch Using a Register)	3-45
	B IRP (Branch Using an Interrupt Return Pointer)	3-47
	B NRP (Branch Using NMI Return Pointer)	3-49
	CLR (Clear a Bit Field)	3-51
	CMPEQ (Compare for Equality, Signed Integers)	3-54
	CMPGT (Compare for Greater Than, Signed Integers)	3-56
	CMPGTU (Compare for Greater Than, Unsigned Integers)	3-59
	CMPLT (Compare for Less Than, Signed Integers)	3-61
	CMPLTU (Compare for Less Than, Unsigned Integers)	3-64
	EXT (Extract and Sign-Extend a Bit Field)	3-66
	EXTU (Extract and Zero-Extend a Bit Field)	3-69

IDLE (Multicycle NOP With No Termination Until Interrupt)	3-72
LDB(U) (Load Byte From Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-73
LDB(U) (Load Byte From Memory With a 15-Bit Unsigned Constant Offset)	3-76
LDH(U) (Load Halfword From Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-78
LDH(U) (Load Halfword From Memory With a 15-Bit Unsigned Constant Offset) ...	3-81
LDW (Load Word From Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-83
LDW (Load Word From Memory With a 15-Bit Unsigned Constant Offset)	3-86
LMBD (Leftmost Bit Detection)	3-88
MPY (Multiply Signed 16 LSB by Signed 16 LSB)	3-90
MPYH (Multiply Signed 16 MSB by Signed 16 MSB)	3-92
MPYHL (Multiply Signed 16 MSB by Signed 16 LSB)	3-93
MPYHLU (Multiply Unsigned 16 MSB by Unsigned 16 LSB)	3-94
MPYHSLU (Multiply Signed 16 MSB by Unsigned 16 LSB)	3-95
MPYHSU (Multiply Signed 16 MSB by Unsigned 16 MSB)	3-96
MPYHU (Multiply Unsigned 16 MSB by Unsigned 16 MSB)	3-97
MPYHULS (Multiply Unsigned 16 MSB by Signed 16 LSB)	3-98
MPYHUS (Multiply Unsigned 16 MSB by Signed 16 MSB)	3-99
MPYLH (Multiply Signed 16 LSB by Signed 16 MSB)	3-100
MPYLHU (Multiply Unsigned 16 LSB by Unsigned 16 MSB)	3-101
MPYLSHU (Multiply Signed 16 LSB by Unsigned 16 MSB)	3-102
MPYLUHS (Multiply Unsigned 16 LSB by Signed 16 MSB)	3-103
MPYSU (Multiply Signed 16 LSB by Unsigned 16 LSB)	3-104
MPYU (Multiply Unsigned 16 LSB by Unsigned 16 LSB)	3-106
MPYUS (Multiply Unsigned 16 LSB by Signed 16 LSB)	3-107
MV (Move From Register to Register)	3-108
MVC (Move Between Control File and Register File)	3-110
MVK (Move Signed Constant Into Register and Sign Extend)	3-113
MVKH and MVKLH (Move 16-Bit Constant Into Upper Bits of Register)	3-115
MVKL (Move Signed Constant Into Register and Sign Extend—Used with MVKH)	3-117
NEG (Negate)	3-119
NOP (No Operation)	3-120
NORM (Normalize Integer)	3-122
NOT (Bitwise NOT)	3-124
OR (Bitwise OR)	3-125
SADD (Add Two Signed Integers With Saturation)	3-127
SAT (Saturate a 40-Bit Integer to a 32-Bit Integer)	3-130
SET (Set a Bit Field)	3-132
SHL (Arithmetic Shift Left)	3-135
SHR (Arithmetic Shift Right)	3-137
SHRU (Logical Shift Right)	3-139

SMPY (Multiply Signed 16 LSB by Signed 16 LSB With Left Shift and Saturation)	3-141
SMPYH (Multiply Signed 16 MSB by Signed 16 MSB With Left Shift and Saturation)	3-143
SMPYHL (Multiply Signed 16 MSB by Signed 16 LSB With Left Shift and Saturation)	3-144
SMPYLH (Multiply Signed 16 LSB by Signed 16 MSB With Left Shift and Saturation)	3-146
SSHL (Shift Left With Saturation)	3-148
SSUB (Subtract Two Signed Integers With Saturation)	3-150
STB (Store Byte to Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-152
STB (Store Byte to Memory With a 15-Bit Unsigned Constant Offset)	3-154
STH (Store Halfword to Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-156
STH (Store Halfword to Memory With a 15-Bit Unsigned Constant Offset)	3-159
STW (Store Word to Memory With a 5-Bit Unsigned Constant Offset or Register Offset)	3-161
STW (Store Word to Memory With a 15-Bit Unsigned Constant Offset)	3-163
SUB (Subtract Two Signed Integers Without Saturation)	3-165
SUBAB (Subtract Using Byte Addressing Mode)	3-168
SUBAH (Subtract Using Halfword Addressing Mode)	3-170
SUBAW (Subtract Using Word Addressing Mode)	3-171
SUBC (Subtract Conditionally and Shift—Used for Division)	3-173
SUBU (Subtract Two Unsigned Integers Without Saturation)	3-175
SUB2 (Subtract Two 16-Bit Integers on Upper and Lower Register Halves)	3-177
XOR (Bitwise Exclusive OR)	3-179
ZERO (Zero a Register)	3-181
4 Pipeline	4-1
<i>Describes phases, operation, and discontinuities for the TMS320C62x CPU pipeline.</i>	
4.1 Pipeline Operation Overview	4-2
4.1.1 Fetch	4-2
4.1.2 Decode	4-3
4.1.3 Execute	4-5
4.1.4 Pipeline Operation Summary	4-6
4.2 Pipeline Execution of Instruction Types	4-11
4.2.1 Single-Cycle Instructions	4-12
4.2.2 Two-Cycle Instructions	4-13
4.2.3 Store Instructions	4-13
4.2.4 Load Instructions	4-15
4.2.5 Branch Instructions	4-17
4.3 Performance Considerations	4-18
4.3.1 Pipeline Operation With Multiple Execute Packets in a Fetch Packet	4-18
4.3.2 Multicycle NOPs	4-20
4.3.3 Memory Considerations	4-22

5	Interrupts	5-1
	<i>Describes CPU interrupts, including reset and the nonmaskable interrupt (NMI). It details the related CPU control registers and their functions in controlling interrupts.</i>	
5.1	Overview	5-2
5.1.1	Types of Interrupts and Signals Used	5-2
5.1.2	Interrupt Service Table (IST)	5-5
5.1.3	Summary of Interrupt Control Registers	5-9
5.2	Globally Enabling and Disabling Interrupts	5-10
5.3	Individual Interrupt Control	5-12
5.3.1	Enabling and Disabling Interrupts	5-12
5.3.2	Status of Interrupts	5-13
5.3.3	Setting and Clearing Interrupts	5-13
5.3.4	Returning From Interrupt Servicing	5-14
5.4	Interrupt Detection and Processing	5-15
5.4.1	Setting the Nonreset Interrupt Flag	5-15
5.4.2	Conditions for Processing a Nonreset Interrupt	5-15
5.4.3	Actions Taken During Nonreset Interrupt Processing	5-17
5.4.4	Setting the $\overline{\text{RESET}}$ Interrupt Flag	5-18
5.4.5	Actions Taken During $\overline{\text{RESET}}$ Interrupt Processing	5-19
5.5	Performance Considerations	5-20
5.5.1	General Performance	5-20
5.5.2	Pipeline Interaction	5-20
5.6	Programming Considerations	5-21
5.6.1	Single Assignment Programming	5-21
5.6.2	Nested Interrupts	5-22
5.6.3	Manual Interrupt Processing	5-24
5.6.4	Traps	5-25
A	Mapping Between Instruction and Functional Unit	A-1
	<i>Lists the instructions that execute on each functional unit.</i>	

Figures

1-1	TMS320C62x DSP Block Diagram	1-6
2-1	TMS320C62x CPU Data Paths	2-3
2-2	Storage Scheme for 40-Bit Data in a Register Pair	2-4
2-3	Addressing Mode Register (AMR)	2-10
2-4	Control Status Register (CSR)	2-13
2-5	PWRD Field of Control Status Register (CSR)	2-13
2-6	Interrupt Clear Register (ICR)	2-16
2-7	Interrupt Enable Register (IER)	2-17
2-8	Interrupt Flag Register (IFR)	2-18
2-9	Interrupt Return Pointer Register (IRP)	2-19
2-10	Interrupt Set Register (ISR)	2-20
2-11	Interrupt Service Table Pointer Register (ISTP)	2-21
2-12	NMI Return Pointer Register (NRP)	2-22
2-13	E1 Phase Program Counter (PCE1)	2-22
3-1	Basic Format of a Fetch Packet	3-7
3-2	Examples of the Detectability of Write Conflicts by the Assembler	3-15
4-1	Pipeline Stages	4-2
4-2	Fetch Phases of the Pipeline	4-3
4-3	Decode Phases of the Pipeline	4-4
4-4	Execute Phases of the Pipeline	4-5
4-5	Pipeline Phases	4-6
4-6	Pipeline Operation: One Execute Packet per Fetch Packet	4-6
4-7	Pipeline Phases Block Diagram	4-8
4-8	Single-Cycle Instruction Phases	4-12
4-9	Single-Cycle Instruction Execution Block Diagram	4-12
4-10	Two-Cycle Instruction Phases	4-13
4-11	Single 16 × 16 Multiply Instruction Execution Block Diagram	4-13
4-12	Store Instruction Phases	4-14
4-13	Store Instruction Execution Block Diagram	4-14
4-14	Load Instruction Phases	4-15
4-15	Load Instruction Execution Block Diagram	4-15
4-16	Branch Instruction Phases	4-17
4-17	Branch Instruction Execution Block Diagram	4-17
4-18	Pipeline Operation: Fetch Packets With Different Numbers of Execute Packets	4-19
4-19	Multicycle NOP in an Execute Packet	4-20
4-20	Branching and Multicycle NOPs	4-21

4-21	Pipeline Phases Used During Memory Accesses	4-22
4-22	Program and Data Memory Stalls	4-23
4-23	4-Bank Interleaved Memory	4-24
4-24	4-Bank Interleaved Memory With Two Memory Spaces	4-25
5-1	Interrupt Service Table	5-5
5-2	Interrupt Service Fetch Packet	5-6
5-3	Interrupt Service Table With Branch to Additional Interrupt Service Code Located Outside the IST	5-7
5-4	Nonreset Interrupt Detection and Processing: Pipeline Operation	5-16
5-5	$\overline{\text{RESET}}$ Interrupt Detection and Processing: Pipeline Operation	5-18

Tables

1-1	Typical Applications for the TMS320 DSPs	1-3
2-1	40-Bit/64-Bit Register Pairs	2-4
2-2	Functional Units and Operations Performed	2-5
2-3	Control Registers	2-7
2-4	Register Addresses for Accessing the Control Registers	2-8
2-5	Addressing Mode Register (AMR) Field Descriptions	2-10
2-6	Block Size Calculations	2-12
2-7	Control Status Register (CSR) Field Descriptions	2-14
2-8	Interrupt Clear Register (ICR) Field Descriptions	2-16
2-9	Interrupt Enable Register (IER) Field Descriptions	2-17
2-10	Interrupt Flag Register (IFR) Field Descriptions	2-18
2-11	Interrupt Set Register (ISR) Field Descriptions	2-20
2-12	Interrupt Service Table Pointer Register (ISTP) Field Descriptions	2-21
3-1	Instruction Operation and Execution Notations	3-2
3-2	Instruction Syntax and Opcode Notations	3-5
3-3	Delay Slots	3-6
3-4	Registers That Can Be Tested by Conditional Operations	3-10
3-5	Indirect Address Generation for Load/Store	3-19
3-6	Address Generator Options for Load/Store	3-19
3-7	Relationships Between Operands, Operand Size, Signed/Unsigned, Functional Units, and Opfields for Example Instruction (ADD)	3-22
3-8	Program Counter Values for Example Branch Using a Displacement	3-44
3-9	Program Counter Values for Example Branch Using a Register	3-46
3-10	Program Counter Values for B IRP Instruction	3-48
3-11	Program Counter Values for B NRP Instruction	3-50
3-12	Data Types Supported by LDB(U) Instruction	3-73
3-13	Data Types Supported by LDB(U) Instruction (15-Bit Offset)	3-76
3-14	Data Types Supported by LDH(U) Instruction	3-78
3-15	Data Types Supported by LDH(U) Instruction (15-Bit Offset)	3-81
3-16	Register Addresses for Accessing the Control Registers	3-112
4-1	Operations Occurring During Pipeline Phases	4-7
4-2	Execution Stage Length Description for Each Instruction Type	4-11
4-3	Program Memory Accesses Versus Data Load Accesses	4-22
4-4	Loads in Pipeline From Example 4-2	4-25
5-1	Interrupt Priorities	5-3
5-2	Interrupt Control Registers	5-9
A-1	Instruction to Functional Unit Mapping	A-1

Examples

3-1	Fully Serial p-Bit Pattern in a Fetch Packet	3-8
3-2	Fully Parallel p-Bit Pattern in a Fetch Packet	3-8
3-3	Partially Serial p-Bit Pattern in a Fetch Packet	3-9
3-4	LDW Instruction in Circular Mode	3-17
3-5	ADDAH Instruction in Circular Mode	3-18
4-1	Execute Packet in Figure 4-7	4-10
4-2	Load From Memory Banks	4-24
5-1	Relocation of Interrupt Service Table	5-8
5-2	Interrupts Versus Writes to GIE	5-10
5-3	Code Sequence to Disable Maskable Interrupts Globally	5-11
5-4	Code Sequence to Enable Maskable Interrupts Globally	5-11
5-5	Code Sequence to Enable an Individual Interrupt (INT9)	5-12
5-6	Code Sequence to Disable an Individual Interrupt (INT9)	5-12
5-7	Code to Set an Individual Interrupt (INT6) and Read the Flag Register	5-13
5-8	Code to Clear an Individual Interrupt (INT6) and Read the Flag Register	5-13
5-9	Code to Return From NMI	5-14
5-10	Code to Return from a Maskable Interrupt	5-14
5-11	Code Without Single Assignment: Multiple Assignment of A1	5-21
5-12	Code Using Single Assignment	5-22
5-13	Assembly Interrupt Service Routine That Allows Nested Interrupts	5-23
5-14	C Interrupt Service Routine That Allows Nested Interrupts	5-24
5-15	Manual Interrupt Processing	5-24
5-16	Code Sequence to Invoke a Trap	5-25
5-17	Code Sequence for Trap Return	5-25

Introduction

The TMS320C6000™ digital signal processor (DSP) platform is part of the TMS320™ DSP family. The TMS320C62x™ DSP generation and the TMS320C64x™ DSP generation comprise fixed-point devices in the C6000™ DSP platform, and the TMS320C67x™ DSP generation comprises floating-point devices in the C6000 DSP platform. The C62x™ and C64x™ DSPs are code-compatible. The C62x and C67x™ DSPs are code-compatible. All three DSPs use the VelociTI™ architecture, a high-performance, advanced very long instruction word (VLIW) architecture, making these DSPs excellent choices for multichannel and multifunction applications.

Topic	Page
1.1 TMS320 DSP Family Overview	1-2
1.2 TMS320C6000 DSP Family Overview	1-2
1.3 TMS320C62x DSP Features and Options	1-4
1.4 TMS320C62x DSP Architecture	1-6

1.1 TMS320 DSP Family Overview

The TMS320™ DSP family consists of fixed-point, floating-point, and multiprocessor digital signal processors (DSPs). TMS320™ DSPs have an architecture designed specifically for real-time signal processing.

Table 1-1 lists some typical applications for the TMS320™ family of DSPs. The TMS320™ DSPs offer adaptable approaches to traditional signal-processing problems. They also support complex applications that often require multiple operations to be performed simultaneously.

1.2 TMS320C6000 DSP Family Overview

With a performance of up to 8000 million instructions per second (MIPS) and an efficient C compiler, the TMS320C6000 DSPs give system architects unlimited possibilities to differentiate their products. High performance, ease of use, and affordable pricing make the C6000 generation the ideal solution for multichannel, multifunction applications, such as:

- Pooled modems
- Wireless local loop base stations
- Remote access servers (RAS)
- Digital subscriber loop (DSL) systems
- Cable modems
- Multichannel telephony systems

The C6000 generation is also an ideal solution for exciting new applications; for example:

- Personalized home security with face and hand/fingerprint recognition
- Advanced cruise control with global positioning systems (GPS) navigation and accident avoidance
- Remote medical diagnostics
- Beam-forming base stations
- Virtual reality 3-D graphics
- Speech recognition
- Audio
- Radar
- Atmospheric modeling
- Finite element analysis
- Imaging (examples: fingerprint recognition, ultrasound, and MRI)

Table 1-1. Typical Applications for the TMS320 DSPs

Automotive	Consumer	Control
Adaptive ride control	Digital radios/TVs	Disk drive control
Antiskid brakes	Educational toys	Engine control
Cellular telephones	Music synthesizers	Laser printer control
Digital radios	Pagers	Motor control
Engine control	Power tools	Robotics control
Global positioning	Radar detectors	Servo control
Navigation	Solid-state answering machines	
Vibration analysis		
Voice commands		
General-Purpose	Graphics/Imaging	Industrial
Adaptive filtering	3-D transformations	Numeric control
Convolution	Animation/digital maps	Power-line monitoring
Correlation	Homomorphic processing	Robotics
Digital filtering	Image compression/transmission	Security access
Fast Fourier transforms	Image enhancement	
Hilbert transforms	Pattern recognition	
Waveform generation	Robot vision	
Windowing	Workstations	
Instrumentation	Medical	Military
Digital filtering	Diagnostic equipment	Image processing
Function generation	Fetal monitoring	Missile guidance
Pattern matching	Hearing aids	Navigation
Phase-locked loops	Patient monitoring	Radar processing
Seismic processing	Prosthetics	Radio frequency modems
Spectrum analysis	Ultrasound equipment	Secure communications
Transient analysis		Sonar processing
Telecommunications		Voice/Speech
1200- to 56 600-bps modems	Faxing	Speaker verification
Adaptive equalizers	Future terminals	Speech enhancement
ADPCM transcoders	Line repeaters	Speech recognition
Base stations	Personal communications systems (PCS)	Speech synthesis
Cellular telephones	Personal digital assistants (PDA)	Speech vocoding
Channel multiplexing	Speaker phones	Text-to-speech
Data encryption	Spread spectrum communications	Voice mail
Digital PBXs	Digital subscriber loop (xDSL)	
Digital speech interpolation (DSI)	Video conferencing	
DTMF encoding/decoding	X.25 packet switching	
Echo cancellation		

1.3 TMS320C62x DSP Features and Options

The C6000 devices execute up to eight 32-bit instructions per cycle. The C62x CPU consists of 32 general-purpose 32-bit registers and eight functional units. These eight functional units contain:

- Two multipliers
- Six ALUs

The C6000 generation has a complete set of optimized development tools, including an efficient C compiler, an assembly optimizer for simplified assembly-language programming and scheduling, and a Windows™ based debugger interface for visibility into source code execution characteristics. A hardware emulation board, compatible with the TI XDS510™ and XDS560™ emulator interface, is also available. This tool complies with IEEE Standard 1149.1-1990, IEEE Standard Test Access Port and Boundary-Scan Architecture.

Features of the C6000 devices include:

- Advanced VLIW CPU with eight functional units, including two multipliers and six arithmetic units
 - Executes up to eight instructions per cycle for up to ten times the performance of typical DSPs
 - Allows designers to develop highly effective RISC-like code for fast development time
- Instruction packing
 - Gives code size equivalence for eight instructions executed serially or in parallel
 - Reduces code size, program fetches, and power consumption
- Conditional execution of all instructions
 - Reduces costly branching
 - Increases parallelism for higher sustained performance
- Efficient code execution on independent functional units
 - Industry's most efficient C compiler on DSP benchmark suite
 - Industry's first assembly optimizer for fast development and improved parallelization
- 8/16/32-bit data support, providing efficient memory support for a variety of applications
- 40-bit arithmetic options add extra precision for vocoders and other computationally intensive applications

- Saturation and normalization provide support for key arithmetic operations
- Field manipulation and instruction extract, set, clear, and bit counting support common operation found in control and data manipulation applications.

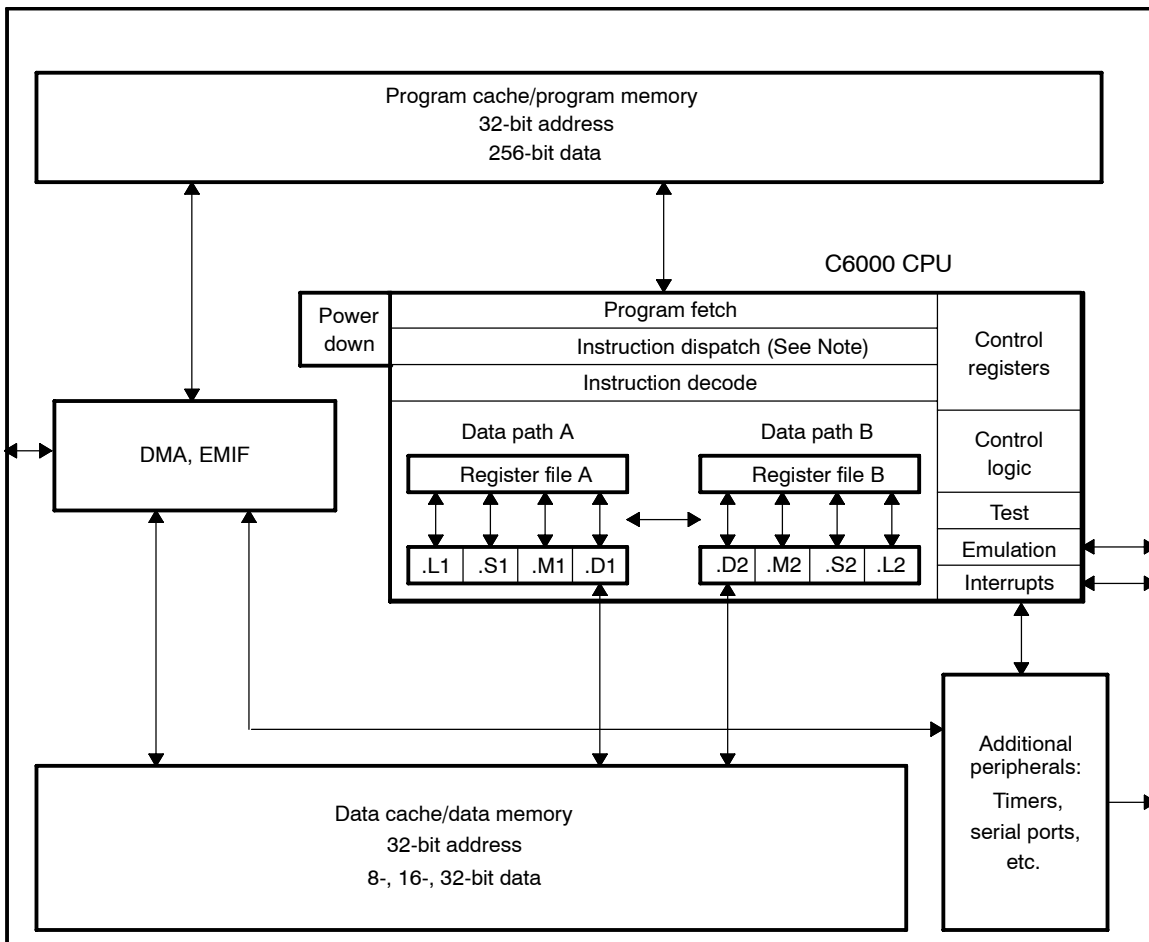
The VelociTI architecture of the C6000 platform of devices make them the first off-the-shelf DSPs to use advanced VLIW to achieve high performance through increased instruction-level parallelism. A traditional VLIW architecture consists of multiple execution units running in parallel, performing multiple instructions during a single clock cycle. Parallelism is the key to extremely high performance, taking these DSPs well beyond the performance capabilities of traditional superscalar designs. VelociTI is a highly deterministic architecture, having few restrictions on how or when instructions are fetched, executed, or stored. It is this architectural flexibility that is key to the breakthrough efficiency levels of the TMS320C6000 Optimizing C compiler. VelociTI's advanced features include:

- Instruction packing: reduced code size
- All instructions can operate conditionally: flexibility of code
- Variable-width instructions: flexibility of data types
- Fully pipelined branches: zero-overhead branching.

1.4 TMS320C62x DSP Architecture

Figure 1-1 is the block diagram for the C62x DSP. The C6000 devices come with program memory, which, on some devices, can be used as a program cache. The devices also have varying sizes of data memory. Peripherals such as a direct memory access (DMA) controller, power-down logic, and external memory interface (EMIF) usually come with the CPU, while peripherals such as serial ports and host ports are on only certain devices. Check your data manual for your device to determine the specific peripheral configurations.

Figure 1-1. TMS320C62x DSP Block Diagram



1.4.1 Central Processing Unit (CPU)

The C62x CPU, in Figure 1-1, contains:

- Program fetch unit
- Instruction dispatch unit
- Instruction decode unit
- Two data paths, each with four functional units
- 32 32-bit registers
- Control registers
- Control logic
- Test, emulation, and interrupt logic

The program fetch, instruction dispatch, and instruction decode units can deliver up to eight 32-bit instructions to the functional units every CPU clock cycle. The processing of instructions occurs in each of the two data paths (A and B), each of which contains four functional units (.L, .S, .M, and .D) and 16 32-bit general-purpose registers. The data paths are described in more detail in Chapter 2. A control register file provides the means to configure and control various processor operations. To understand how instructions are fetched, dispatched, decoded, and executed in the data path, see Chapter 4.

1.4.2 Internal Memory

The C62x DSP has a 32-bit, byte-addressable address space. Internal (on-chip) memory is organized in separate data and program spaces. When off-chip memory is used, these spaces are unified on most devices to a single memory space via the external memory interface (EMIF).

The C62x DSP has two 32-bit internal ports to access internal data memory. The C62x DSP has a single internal port to access internal program memory, with an instruction-fetch width of 256 bits.

1.4.3 Memory and Peripheral Options

A variety of memory and peripheral options are available for the C6000 platform:

- Large on-chip RAM, up to 7M bits
- Program cache
- 2-level caches
- 32-bit external memory interface supports SDRAM, SBSRAM, SRAM, and other asynchronous memories for a broad range of external memory requirements and maximum system performance.

- ❑ The direct memory access (DMA) controller transfers data between address ranges in the memory map without intervention by the CPU. The DMA controller has four programmable channels and a fifth auxiliary channel.
- ❑ The enhanced direct memory access (EDMA) controller (C6211 DSP only) performs the same functions as the DMA controller. The EDMA has 16 programmable channels, as well as a RAM space to hold multiple configurations for future transfers.
- ❑ The host port interface (HPI) is a parallel port through which a host processor can directly access the CPU memory space. The host device functions as a master to the interface, which increases ease of access. The host and CPU can exchange information via internal or external memory. The host also has direct access to memory-mapped peripherals. Connectivity to the CPU memory space is provided through the DMA/EDMA controller.
- ❑ The expansion bus is a replacement for the HPI, as well as an expansion of the EMIF. The expansion provides two distinct areas of functionality (host port and I/O port) that can co-exist in a system. The host port of the expansion bus can operate in either asynchronous slave mode, similar to the HPI, or in synchronous master/slave mode. This allows the device to interface to a variety of host bus protocols. Synchronous FIFOs and asynchronous peripheral I/O devices may interface to the expansion bus.
- ❑ The peripheral component interconnect (PCI) port supports connection of the C62x DSP to a PCI host via the integrated PCI master/slave bus interface.
- ❑ The multichannel buffered serial port (McBSP) is based on the standard serial port interface found on the TMS320C2000™ and TMS320C5000™ devices. In addition, the port can buffer serial samples in memory automatically with the aid of the DMA/EDMA controller. It also has multichannel capability compatible with the T1, E1, SCSA, and MVIP networking standards.
- ❑ Timers in the C6000 devices are two 32-bit general-purpose timers used for these functions:
 - Time events
 - Count events
 - Generate pulses
 - Interrupt the CPU
 - Send synchronization events to the DMA/EDMA controller.

- Power-down logic allows reduced clocking to reduce power consumption. Most of the operating power of CMOS logic dissipates during circuit switching from one logic state to another. By preventing some or all of the chip's logic from switching, you can realize significant power savings without losing any data or operational context.

For an overview of the peripherals available on the C6000 DSP, refer to the *TMS320C6000 DSP Peripherals Overview Reference Guide* (SPRU190) or to your device-specific data manual.

CPU Data Paths and Control

This chapter focuses on the CPU, providing information about the data paths and control registers. The two register files and the data cross paths are described.

Topic	Page
2.1 Introduction	2-2
2.2 General-Purpose Register Files	2-2
2.3 Functional Units	2-5
2.4 Register File Cross Paths	2-6
2.5 Memory, Load, and Store Paths	2-6
2.6 Data Address Paths	2-7
2.7 Control Register File	2-7

2.1 Introduction

The components of the data path for the TMS320C62x CPU are shown in Figure 2-1. These components consist of:

- Two general-purpose register files (A and B)
- Eight functional units (.L1, .L2, .S1, .S2, .M1, .M2, .D1, and .D2)
- Two load-from-memory data paths (LD1 and LD2)
- Two store-to-memory data paths (ST1 and ST2)
- Two data address paths (DA1 and DA2)
- Two register file data cross paths (1X and 2X)

2.2 General-Purpose Register Files

There are two general-purpose register files (A and B) in the C62x CPU data paths. Each of these files contains 16 32-bit registers (A0–A15 for file A and B0–B15 for file B), as shown in Table 2-1. The general-purpose registers can be used for data, data address pointers, or condition registers.

The C62x DSP general-purpose register files support data ranging in size from packed 16-bit through 40-bit fixed-point data. Values larger than 32 bits, such as 40-bit long quantities, are stored in register pairs. The 32 LSBs of data are placed in an even-numbered register and the remaining 8 MSBs in the next upper register (that is always an odd-numbered register). Packed data types store two 16-bit values in a single 32-bit register.

There are 16 valid register pairs for 40-bit data in the C62x DSP cores. In assembly language syntax, a colon between the register names denotes the register pairs, and the odd-numbered register is specified first.

Figure 2-2 shows the register storage scheme for 40-bit long data. Operations requiring a long input ignore the 24 MSBs of the odd-numbered register. Operations producing a long result zero-fill the 24 MSBs of the odd-numbered register. The even-numbered register is encoded in the opcode.

Figure 2-1. TMS320C62x CPU Data Paths

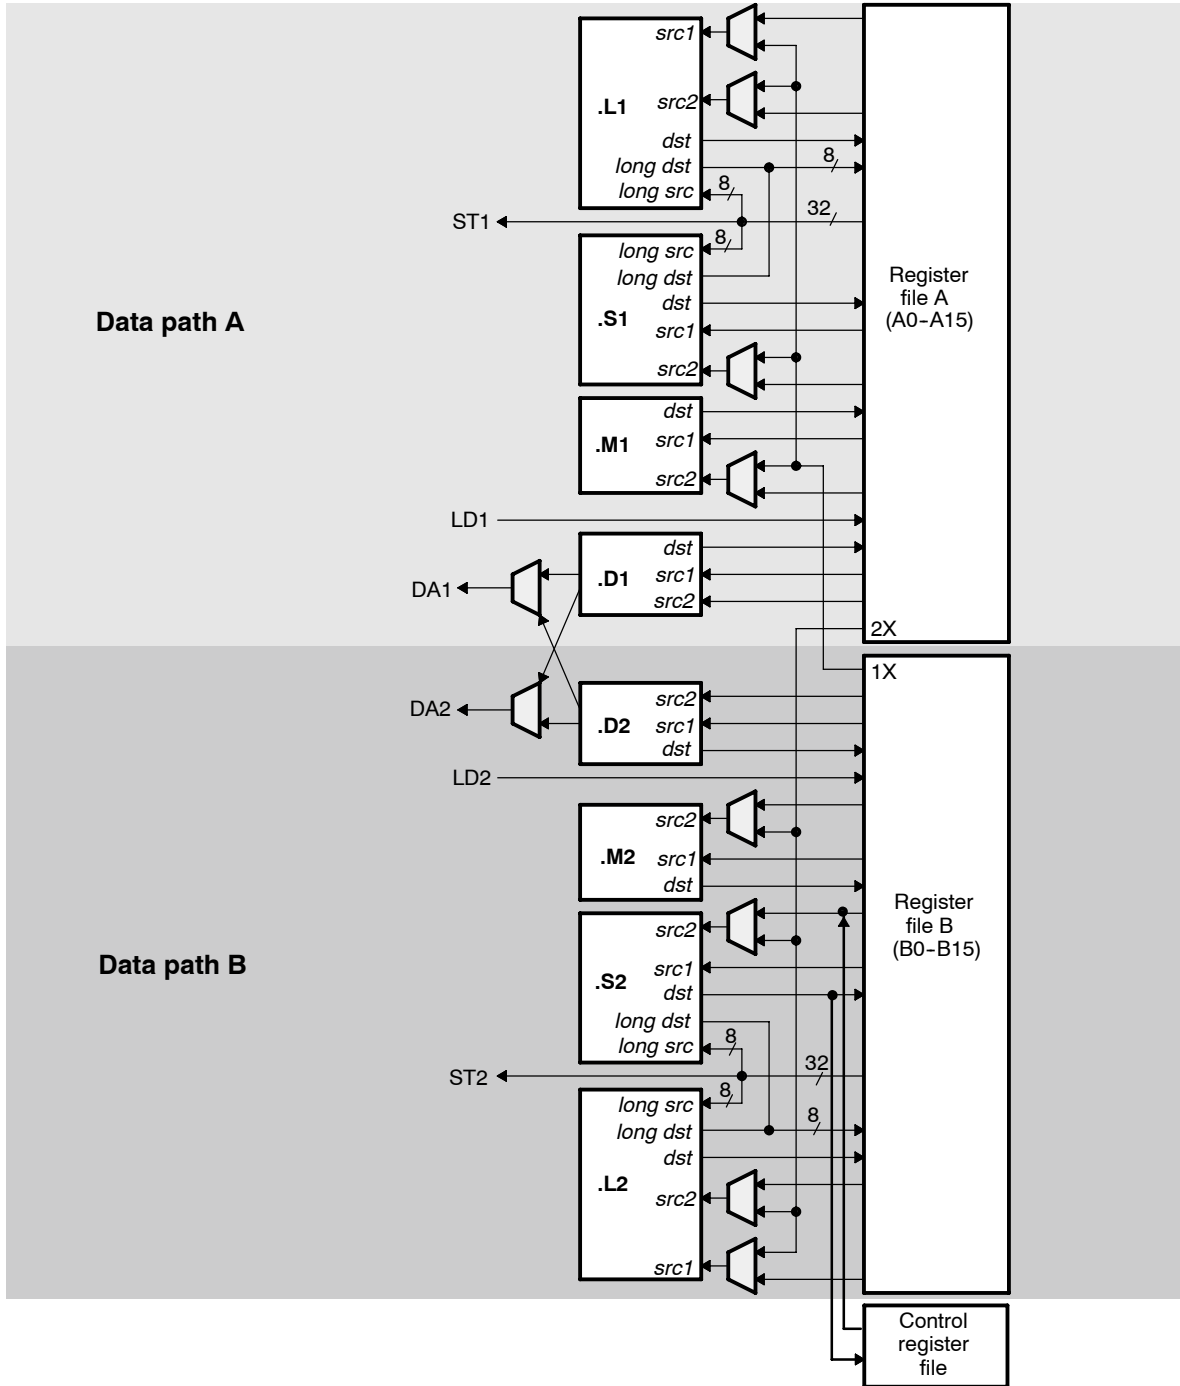
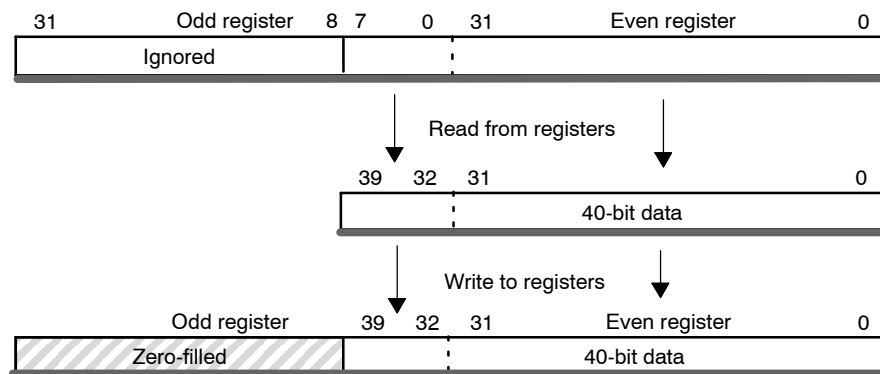


Table 2-1. 40-Bit/64-Bit Register Pairs

Register Files	
A	B
A1:A0	B1:B0
A3:A2	B3:B2
A5:A4	B5:B4
A7:A6	B7:B6
A9:A8	B9:B8
A11:A10	B11:B10
A13:A12	B13:B12
A15:A14	B15:B14

Figure 2-2. Storage Scheme for 40-Bit Data in a Register Pair



2.3 Functional Units

The eight functional units in the C6000 data paths can be divided into two groups of four; each functional unit in one data path is almost identical to the corresponding unit in the other data path. The functional units are described in Table 2-2.

Most data lines in the CPU support 32-bit operands, and some support long (40-bit) operands. Each functional unit has its own 32-bit write port into a general-purpose register file (refer to Figure 2-1). All units ending in 1 (for example, .L1) write to register file A, and all units ending in 2 write to register file B. Each functional unit has two 32-bit read ports for source operands *src1* and *src2*. Four units (.L1, .L2, .S1, and .S2) have an extra 8-bit-wide port for 40-bit long writes, as well as an 8-bit input for 40-bit long reads. Each unit has its own 32-bit write port, so all eight units can be used in parallel every cycle.

See Appendix A for a list of the instructions that execute on each functional unit.

Table 2-2. Functional Units and Operations Performed

Functional Unit	Fixed-Point Operations
.L unit (.L1, .L2)	32/40-bit arithmetic and compare operations 32-bit logical operations Leftmost 1 or 0 counting for 32 bits Normalization count for 32 and 40 bits
.S unit (.S1, .S2)	32-bit arithmetic operations 32/40-bit shifts and 32-bit bit-field operations 32-bit logical operations Branches Constant generation Register transfers to/from control register file (.S2 only)
.M unit (.M1, .M2)	16 × 16-bit multiply operations
.D unit (.D1, .D2)	32-bit add, subtract, linear and circular address calculation Loads and stores with 5-bit constant offset Loads and stores with 15-bit constant offset (.D2 only)

2.4 Register File Cross Paths

Each functional unit reads directly from and writes directly to the register file within its own data path. That is, the .L1, .S1, .D1, and .M1 units write to register file A and the .L2, .S2, .D2, and .M2 units write to register file B. The register files are connected to the opposite-side register file's functional units via the 1X and 2X cross paths. These cross paths allow functional units from one data path to access a 32-bit operand from the opposite side register file. The 1X cross path allows the functional units of data path A to read their source from register file B, and the 2X cross path allows the functional units of data path B to read their source from register file A.

On the C62x DSP, six of the eight functional units have access to the register file on the opposite side, via a cross path. The *src2* inputs of .M1, .M2, .S1, and .S2 units are selectable between the cross path and the same side register file. In the case of the .L1 and .L2, both *src1* and *src2* inputs are selectable between the cross path and the same-side register file.

Only two cross paths, 1X and 2X, exist in the C6000 architecture. Thus, the limit is one source read from each data path's opposite register file per cycle, or a total of two cross path source reads per cycle. In the C62x DSP, only one functional unit per data path, per execute packet, can get an operand from the opposite register file.

2.5 Memory, Load, and Store Paths

The C62x DSP has two 32-bit paths for loading data from memory to the register file: LD1 for register file A, and LD2 for register file B. There are also two 32-bit paths, ST1 and ST2, for storing register values to memory from each register file.

On the C6000 architecture, some of the ports for long operands are shared between functional units. This places a constraint on which long operations can be scheduled on a data path in the same execute packet. See section 3.6.4.

2.6 Data Address Paths

The data address paths (DA1 and DA2) are each connected to the .D units in both data paths. This allows data addresses generated by any one path to access data to or from any register.

The DA1 and DA2 resources and their associated data paths are specified as T1 and T2, respectively. T1 consists of the DA1 address path and the LD1 and ST1 data paths. Similarly, T2 consists of the DA2 address path and the LD2 and ST2 data paths.

The T1 and T2 designations appear in the functional unit fields for load and store instructions. For example, the following load instruction uses the .D1 unit to generate the address but is using the LD2 path resource from DA2 to place the data in the B register file. The use of the DA2 resource is indicated with the T2 designation.

```
LDW  .D1T2  *A0[3],B1
```

2.7 Control Register File

Table 2-3 lists the control registers contained in the control register file.

Table 2-3. Control Registers

Acronym	Register Name	Section
AMR	Addressing mode register	2.7.3
CSR	Control status register	2.7.4
ICR	Interrupt clear register	2.7.5
IER	Interrupt enable register	2.7.6
IFR	Interrupt flag register	2.7.7
IRP	Interrupt return pointer register	2.7.8
ISR	Interrupt set register	2.7.9
ISTP	Interrupt service table pointer register	2.7.10
NRP	Nonmaskable interrupt return pointer register	2.7.11
PCE1	Program counter, E1 phase	2.7.12

2.7.1 Register Addresses for Accessing the Control Registers

Table 2-4 lists the register addresses for accessing the control register file. One unit (.S2) can read from and write to the control register file. Each control register is accessed by the **MVC** instruction. See the **MVC** instruction description, page 3-110, for information on how to use this instruction.

Additionally, some of the control register bits are specially accessed in other ways. For example, arrival of a maskable interrupt on an external interrupt pin, $INTm$, triggers the setting of flag bit $IFRm$. Subsequently, when that interrupt is processed, this triggers the clearing of $IFRm$ and the clearing of the global interrupt enable bit, GIE. Finally, when that interrupt processing is complete, the **B IRP** instruction in the interrupt service routine restores the pre-interrupt value of the GIE. Similarly, saturating instructions like **SADD** set the SAT (saturation) bit in the control status register (CSR).

Table 2-4. Register Addresses for Accessing the Control Registers

Acronym	Register Name	Address	Read/ Write
AMR	Addressing mode register	00000	R, W
CSR	Control status register	00001	R, W
ICR	Interrupt clear register	00011	W
IER	Interrupt enable register	00100	R, W
IFR	Interrupt flag register	00010	R
IRP	Interrupt return pointer	00110	R, W
ISR	Interrupt set register	00010	W
ISTP	Interrupt service table pointer	00101	R, W
NRP	Nonmaskable interrupt return pointer	00111	R, W
PCE1	Program counter, E1 phase	10000	R

Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction

2.7.2 Pipeline/Timing of Control Register Accesses

All **MVC** instructions are single-cycle instructions that complete their access of the explicitly named registers in the E1 pipeline phase. This is true whether **MVC** is moving a general register to a control register, or conversely. In all cases, the source register content is read, moved through the .S2 unit, and written to the destination register in the E1 pipeline phase.

Pipeline Stage	E1
Read	<i>src2</i>
Written	<i>dst</i>
Unit in use	.S2

Even though **MVC** modifies the particular target control register in a single cycle, it can take extra clocks to complete modification of the non-explicitly named register. For example, the **MVC** cannot modify bits in the IFR directly. Instead, **MVC** can only write 1's into the ISR or the ICR to specify setting or clearing, respectively, of the IFR bits. **MVC** completes this ISR/ICR write in a single (E1) cycle but the modification of the IFR bits occurs one clock later. For more information on the manipulation of ISR, ICR, and IFR, see section 2.7.9, section 2.7.5, and section 2.7.7.


Saturating instructions, such as **SADD**, set the saturation flag bit (SAT) in CSR indirectly. As a result, several of these instructions update the SAT bit one full clock cycle after their primary results are written to the register file. For example, the **SMPY** instruction writes its result at the end of pipeline stage E2; its primary result is available after one delay slot. In contrast, the SAT bit in CSR is updated one cycle later than the result is written; this update occurs after two delay slots. (For the specific behavior of an instruction, refer to the description of that individual instruction).

The **B IRP** and **B NRP** instructions directly update the GIE and NMIE bits, respectively. Because these branches directly modify CSR and IER, respectively, there are no delay slots between when the branch is issued and when the control register updates take effect.

2.7.3 Addressing Mode Register (AMR)

For each of the eight registers (A4–A7, B4–B7) that can perform linear or circular addressing, the addressing mode register (AMR) specifies the addressing mode. A 2-bit field for each register selects the address modification mode: linear (the default) or circular mode. With circular addressing, the field also specifies which BK (block size) field to use for a circular buffer. In addition, the buffer must be aligned on a byte boundary equal to the block size. The mode select fields and block size fields are shown in Figure 2-3 and described in Table 2-5.

Figure 2-3. Addressing Mode Register (AMR)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -n = value after reset

Table 2-5. Addressing Mode Register (AMR) Field Descriptions

Bit	Field	Value	Description
31-26	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
25-21	BK1	0-1Fh	Block size field 1. A 5-bit value used in calculating block sizes for circular addressing. Table 2-6 shows block size calculations for all 32 possibilities. <i>Block size (in bytes) = 2^(N+1), where N is the 5-bit value in BK1</i>
20-16	BK0	0-1Fh	Block size field 0. A 5-bit value used in calculating block sizes for circular addressing. Table 2-6 shows block size calculations for all 32 possibilities. <i>Block size (in bytes) = 2^(N+1), where N is the 5-bit value in BK0</i>
15-14	B7 MODE	0-3h	Address mode selection for register file B7. 0 Linear modification (default at reset) 1h Circular addressing using the BK0 field 2h Circular addressing using the BK1 field 3h Reserved

Table 2-5. Addressing Mode Register (AMR) Field Descriptions (Continued)

Bit	Field	Value	Description
13-12	B6 MODE	0-3h	Address mode selection for register file B6.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved
11-10	B5 MODE	0-3h	Address mode selection for register file B5.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved
9-8	B4 MODE	0-3h	Address mode selection for register file B4.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved
7-6	A7 MODE	0-3h	Address mode selection for register file A7.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved
5-4	A6 MODE	0-3h	Address mode selection for register file A6.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved

Table 2-5. Addressing Mode Register (AMR) Field Descriptions (Continued)

Bit	Field	Value	Description
3-2	A5 MODE	0-3h	Address mode selection for register file a5.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved
1-0	A4 MODE	0-3h	Address mode selection for register file A4.
		0	Linear modification (default at reset)
		1h	Circular addressing using the BK0 field
		2h	Circular addressing using the BK1 field
		3h	Reserved

Table 2-6. Block Size Calculations

BK _n Value	Block Size	BK _n Value	Block Size
00000	2	10000	131 072
00001	4	10001	262 144
00010	8	10010	524 288
00011	16	10011	1 048 576
00100	32	10100	2 097 152
00101	64	10101	4 194 304
00110	128	10110	8 388 608
00111	256	10111	16 777 216
01000	512	11000	33 554 432
01001	1 024	11001	67 108 864
01010	2 048	11010	134 217 728
01011	4 096	11011	268 435 456
01100	8 192	11100	536 870 912
01101	16 384	11101	1 073 741 824
01110	32 768	11110	2 147 483 648
01111	65 536	11111	4 294 967 296

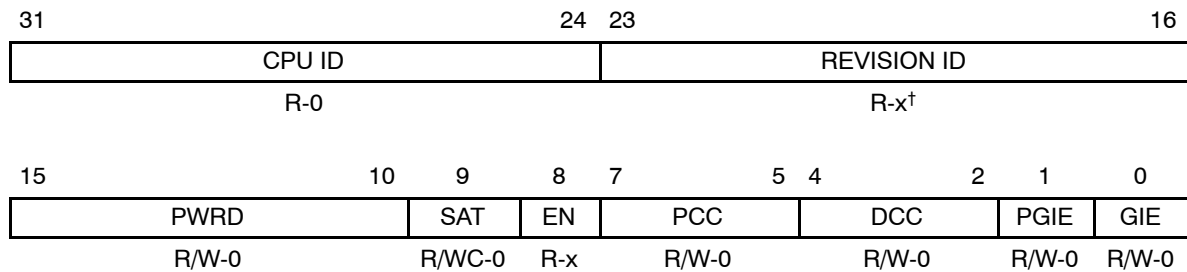
Note: When *n* is 11111, the behavior is identical to linear addressing.

2.7.4 Control Status Register (CSR)

The control status register (CSR) contains control and status bits. The CSR is shown in Figure 2-4 and described in Table 2-7. For the PWRD, EN, PCC, and DCC fields, see the device-specific datasheet to see if it supports the options that these fields control.

The power-down modes and their wake-up methods are programmed by the PWRD field (bits 15-10) of CSR. The PWRD field of CSR is shown in Figure 2-5. When writing to CSR, all bits of the PWRD field should be configured at the same time. A logic 0 should be used when writing to the reserved bit (bit 15) of the PWRD field.


Figure 2-4. Control Status Register (CSR)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; WC = Bit is cleared on write; -n = value after reset; -x = value is indeterminate after reset

[†] See the device-specific data manual for the default value of this field.

Figure 2-5. PWRD Field of Control Status Register (CSR)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -n = value after reset

Table 2-7. Control Status Register (CSR) Field Descriptions

Bit	Field	Value	Description
31-24	CPU ID	0-FFh	Identifies the CPU of the device. Not writable by the MVC instruction.
		0	C62x DSP
		1h-FFh	Reserved
23-16	REVISION ID	0-FFh	Identifies silicon revision of the CPU. For the most current silicon revision information, see the device-specific datasheet. Not writable by the MVC instruction.
15-10	PWRD	0-3Fh	Power-down mode field. See Figure 2-5. Writable by the MVC instruction.
		0	No power-down.
		1h-8h	Reserved
		9h	Power-down mode PD1; wake by an enabled interrupt.
		Ah-10h	Reserved
		11h	Power-down mode PD1; wake by an enabled or nonenabled interrupt.
		12h-19h	Reserved
		1Ah	Power-down mode PD2; wake by a device reset.
		1Bh	Reserved
		1Ch	Power-down mode PD3; wake by a device reset.
9	SAT	1D-3Fh	Reserved
			Saturate bit. Can be cleared only by the MVC instruction and can be set only by a functional unit. The set by a functional unit has priority over a clear (by the MVC instruction), if they occur on the same cycle. The SAT bit is set one full cycle (one delay slot) after a saturate occurs. The SAT bit will not be modified by a conditional instruction whose condition is false.
		0	No functional units generated saturated results.
		1	One or more functional units performed an arithmetic operation which resulted in saturation.
8	EN		Endian mode. Not writable by the MVC instruction.
		0	Big endian
		1	Little endian

Table 2-7. Control Status Register (CSR) Field Descriptions (Continued)

Bit	Field	Value	Description
7-5	PCC	0-7h	Program cache control mode. Writable by the MVC instruction. See the <i>TMS320C621x/C671x DSP Two-Level Internal Memory Reference Guide</i> (SPRU609).
		0	Direct-mapped cache enabled
		1h	Reserved
		2h	Direct-mapped cache enabled
		3h-7h	Reserved
4-2	DCC	0-7h	Data cache control mode. Writable by the MVC instruction. See the <i>TMS320C621x/C671x DSP Two-Level Internal Memory Reference Guide</i> (SPRU609).
		0	2-way cache enabled
		1h	Reserved
		2h	2-way cache enabled
		3h-7h	Reserved
1	PGIE		Previous GIE (global interrupt enable). This bit contains a copy of the GIE bit at the point when interrupt is taken. Writeable by the MVC instruction.
		0	Interrupts will be disabled after return from interrupt.
		1	Interrupts will be enabled after return from interrupt.
0	GIE		Global interrupt enable. Physically the same bit as GIE bit in the task state register (TSR). Writable by the MVC instruction.
		0	Disables all interrupts, except the reset interrupt and NMI (nonmaskable interrupt).
		1	Enables all interrupts.

2.7.5 Interrupt Clear Register (ICR)

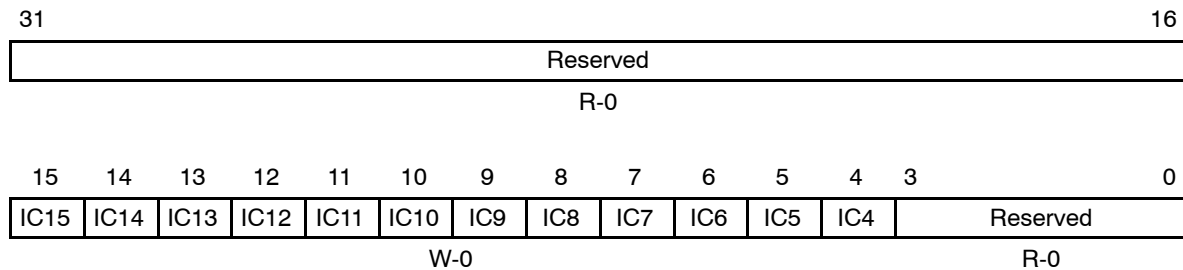
The interrupt clear register (ICR) allows you to manually clear the maskable interrupts (INT15–INT4) in the interrupt flag register (IFR). Writing a 1 to any of the bits in ICR causes the corresponding interrupt flag (IF n) to be cleared in IFR. Writing a 0 to any bit in ICR has no effect. Incoming interrupts have priority and override any write to ICR. You cannot set any bit in ICR to affect NMI or reset. The ISR is shown in Figure 2-6 and described in Table 2-8.

Note:

Any write to ICR (by the **MVC** instruction) effectively has one delay slot because the results cannot be read (by the **MVC** instruction) in IFR until two cycles after the write to ICR.

Any write to ICR is ignored by a simultaneous write to the same bit in the interrupt set register (ISR).

Figure 2-6. Interrupt Clear Register (ICR)



Legend: R = Read only; W = Writeable by the **MVC** instruction; - n = value after reset

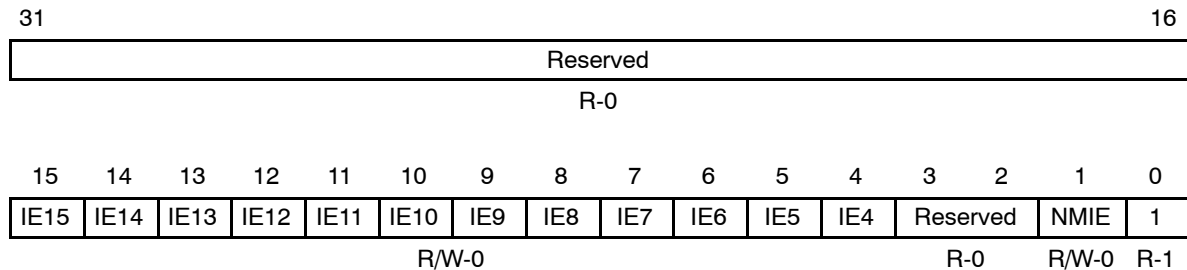
Table 2-8. Interrupt Clear Register (ICR) Field Descriptions

Bit	Field	Value	Description
31-16	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
15-4	IC n	0	Corresponding interrupt flag (IF n) in IFR is not cleared.
		1	Corresponding interrupt flag (IF n) in IFR is cleared.
3-0	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.

2.7.6 Interrupt Enable Register (IER)

The interrupt enable register (IER) enables and disables individual interrupts. The IER is shown in Figure 2-7 and described in Table 2-9.

Figure 2-7. Interrupt Enable Register (IER)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -n = value after reset

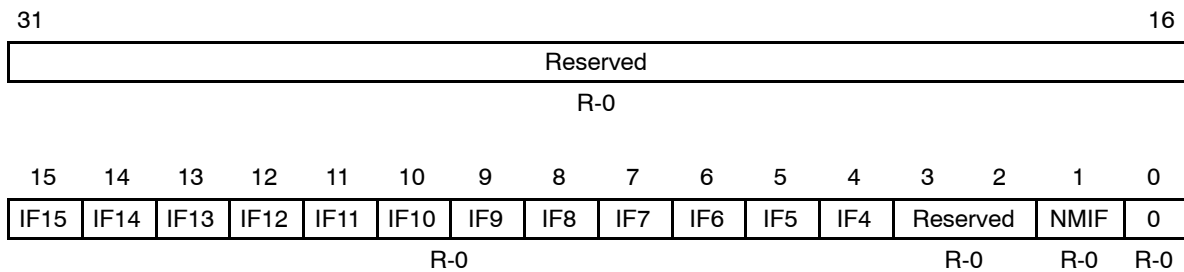
Table 2-9. Interrupt Enable Register (IER) Field Descriptions

Bit	Field	Value	Description
31-16	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
15-4	IE _n	0	Interrupt enable. An interrupt triggers interrupt processing only if the corresponding bit is set to 1.
		0	Interrupt is disabled.
		1	Interrupt is enabled.
3-2	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
1	NMIE	0	Nonmaskable interrupt enable. An interrupt triggers interrupt processing only if the bit is set to 1.
		1	The NMIE bit is cleared at reset. After reset, you must set the NMIE bit to enable the NMI and to allow INT15-INT4 to be enabled by the GIE bit in CSR and the corresponding IER bit. You cannot manually clear the NMIE bit; a write of 0 has no effect. The NMIE bit is also cleared by the occurrence of an NMI.
		0	All nonreset interrupts are disabled.
		1	All nonreset interrupts are enabled. The NMIE bit is set only by completing a B NRP instruction or by a write of 1 to the NMIE bit.
0	1	1	Reset interrupt enable. You cannot disable the reset interrupt.

2.7.7 Interrupt Flag Register (IFR)

The interrupt flag register (IFR) contains the status of INT4–INT15 and NMI interrupt. Each corresponding bit in the IFR is set to 1 when that interrupt occurs; otherwise, the bits are cleared to 0. If you want to check the status of interrupts, use the **MVC** instruction to read the IFR. (See the **MVC** instruction description, page 3-110, for information on how to use this instruction.) The IFR is shown in Figure 2-8 and described in Table 2-10.

Figure 2-8. Interrupt Flag Register (IFR)



Legend: R = Readable by the **MVC** instruction; -n = value after reset

Table 2-10. Interrupt Flag Register (IFR) Field Descriptions


Bit	Field	Value	Description
31-16	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
15-4	IFn	0	Interrupt flag. Indicates the status of the corresponding maskable interrupt. An interrupt flag may be manually set by setting the corresponding bit (ISn) in the interrupt set register (ISR) or manually cleared by setting the corresponding bit (ICn) in the interrupt clear register (ICR).
		0	Interrupt has not occurred.
		1	Interrupt has occurred.
3-2	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
1	NMIF	0	Nonmaskable interrupt flag.
		0	Interrupt has not occurred.
		1	Interrupt has occurred.
0	0	0	Reset interrupt flag.

2.7.8 Interrupt Return Pointer Register (IRP)

The interrupt return pointer register (IRP) contains the return pointer that directs the CPU to the proper location to continue program execution after processing a maskable interrupt. A branch using the address in IRP (**B IRP**) in your interrupt service routine returns to the program flow when interrupt servicing is complete. The IRP is shown in Figure 2-9.

The IRP contains the 32-bit address of the first execute packet in the program flow that was not executed because of a maskable interrupt. Although you can write a value to IRP, any subsequent interrupt processing may overwrite that value.

Figure 2-9. Interrupt Return Pointer Register (IRP)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -x = value is indeterminate after reset

2.7.9 Interrupt Set Register (ISR)

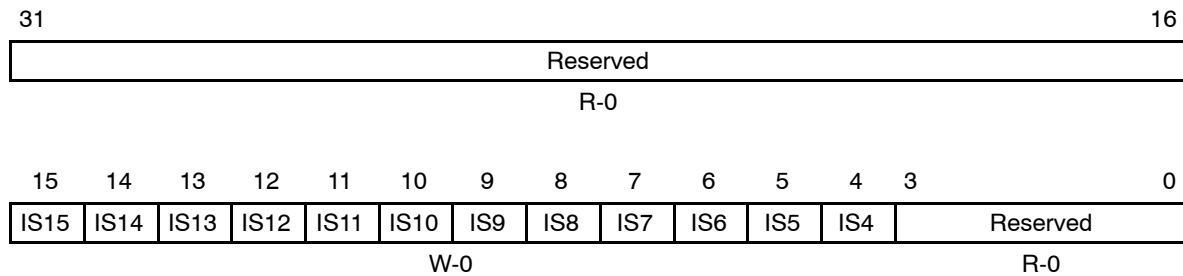
The interrupt set register (ISR) allows you to manually set the maskable interrupts (INT15–INT4) in the interrupt flag register (IFR). Writing a 1 to any of the bits in ISR causes the corresponding interrupt flag (IF n) to be set in IFR. Writing a 0 to any bit in ISR has no effect. You cannot set any bit in ISR to affect NMI or reset. The ISR is shown in Figure 2-10 and described in Table 2-11.

Note:

Any write to ISR (by the **MVC** instruction) effectively has one delay slot because the results cannot be read (by the **MVC** instruction) in IFR until two cycles after the write to ISR.

Any write to the interrupt clear register (ICR) is ignored by a simultaneous write to the same bit in ISR.

Figure 2-10. Interrupt Set Register (ISR)



Legend: R = Read only; W = Writeable by the **MVC** instruction; - n = value after reset

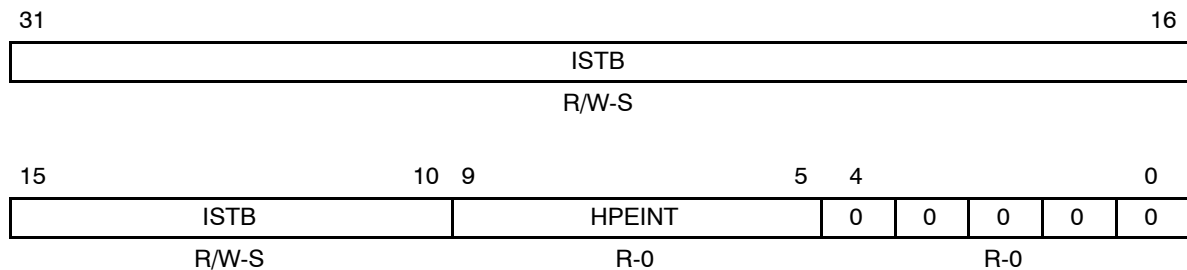
Table 2-11. Interrupt Set Register (ISR) Field Descriptions

Bit	Field	Value	Description
31-16	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.
15-4	IS n	0	Corresponding interrupt flag (IF n) in IFR is not set.
		1	Corresponding interrupt flag (IF n) in IFR is set.
3-0	Reserved	0	Reserved. The reserved bit location is always read as 0. A value written to this field has no effect.

2.7.10 Interrupt Service Table Pointer Register (ISTP)

The interrupt service table pointer register (ISTP) is used to locate the interrupt service routine (ISR). The ISTB field identifies the base portion of the address of the interrupt service table (IST) and the HPEINT field identifies the specific interrupt and locates the specific fetch packet within the IST. The ISTP is shown in Figure 2-11 and described in Table 2-12. See section 5.1.2.2 on page 5-8 for a discussion of the use of the ISTP.

Figure 2-11. Interrupt Service Table Pointer Register (ISTP)



Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -n = value after reset; S = See the device-specific data manual for the default value of this field after reset

Table 2-12. Interrupt Service Table Pointer Register (ISTP) Field Descriptions


Bit	Field	Value	Description
31-10	ISTB	0-3F FFFFh	Interrupt service table base portion of the IST address. This field is cleared to a device-specific default value on reset; therefore, upon startup the IST must reside at this specific address. See the device-specific data manual for more information. After reset, you can relocate the IST by writing a new value to ISTB. If relocated, the first ISFP (corresponding to RESET) is never executed via interrupt processing, because reset clears the ISTB to its default value. See Example 5-1 on page 5-8.
9-5	HPEINT	0-1Fh	Highest priority enabled interrupt that is currently pending. This field indicates the number (related bit position in the IFR) of the highest priority interrupt (as defined in Table 5-1 on page 5-3) that is enabled by its bit in the IER. Thus, the ISTP can be used for manual branches to the highest priority enabled interrupt. If no interrupt is pending and enabled, HPEINT contains the value 0. The corresponding interrupt need not be enabled by NMIE (unless it is NMI) or by GIE.
4-0	-	-	Cleared to 0 (fetch packets must be aligned on 8-word (32-byte) boundaries).

2.7.11 Nonmaskable Interrupt (NMI) Return Pointer Register (NRP)

The NMI return pointer register (NRP) contains the return pointer that directs the CPU to the proper location to continue program execution after processing of a nonmaskable interrupt (NMI) or exception. A branch using the address in NRP (**B NRP**) in your interrupt service routine or exception service routine returns to the program flow when NMI or exception servicing is complete. The NRP is shown in Figure 2-12.

The NRP contains the 32-bit address of the first execute packet in the program flow that was not executed because of a nonmaskable interrupt. Although you can write a value to NRP, any subsequent interrupt processing may overwrite that value.

Figure 2-12. NMI Return Pointer Register (NRP)




Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction; -x = value is indeterminate after reset

2.7.12 E1 Phase Program Counter (PCE1)

The E1 phase program counter (PCE1), shown in Figure 2-13, contains the 32-bit address of the fetch packet in the E1 pipeline phase.

Figure 2-13. E1 Phase Program Counter (PCE1)



Legend: R = Readable by the **MVC** instruction; -x = value is indeterminate after reset

Instruction Set

This chapter describes the assembly language instructions of the TMS320C62x DSP. Also described are parallel operations, conditional operations, resource constraints, and addressing modes.

Topic	Page
3.1 Instruction Operation and Execution Notations	3-2
3.2 Instruction Syntax and Opcode Notations	3-5
3.3 Delay Slots	3-6
3.4 Parallel Operations	3-7
3.5 Conditional Operations	3-10
3.6 Resource Constraints	3-11
3.7 Addressing Modes	3-16
3.8 Instruction Compatibility	3-20
3.9 Instruction Descriptions	3-20

3.1 Instruction Operation and Execution Notations

Table 3-1 explains the symbols used in the instruction descriptions.

Table 3-1. Instruction Operation and Execution Notations

Symbol	Meaning
abs(x)	Absolute value of x
and	Bitwise AND
-a	Perform 2s-complement subtraction using the addressing mode defined by the AMR
+a	Perform 2s-complement addition using the addressing mode defined by the AMR
b _{y..z}	Selection of bits y through z of bit string b
cond	Check for either <i>creg</i> equal to 0 or <i>creg</i> not equal to 0
<i>creg</i>	3-bit field specifying a conditional register, see section 3.5
<i>cstn</i>	n-bit constant field (for example, <i>cst5</i>)
int	32-bit integer value
lmb0(x)	Leftmost 0 bit search of x
lmb1(x)	Leftmost 1 bit search of x
long	40-bit integer value
lsbn or LSBn	n least-significant bits (for example, <i>lsb16</i>)
msbn or MSBn	n most-significant bits (for example, <i>msb16</i>)
nop	No operation
norm(x)	Leftmost nonredundant sign bit of x
not	Bitwise logical complement
op	Opfields
or	Bitwise OR
R	Any general-purpose register
scstn	n-bit signed constant field
sint	Signed 32-bit integer value
slong	Signed 40-bit integer value

Table 3-1. Instruction Operation and Execution Notations (Continued)

Symbol	Meaning
slsb16	Signed 16-bit integer value in lower half of 32-bit register
smsb16	Signed 16-bit integer value in upper half of 32-bit register
-s	Perform 2s-complement subtraction and saturate the result to the result size, if an overflow occurs
+s	Perform 2s-complement addition and saturate the result to the result size, if an overflow occurs
ucstrn	n-bit unsigned constant field (for example, ucst5)
uint	Unsigned 32-bit integer value
ulong	Unsigned 40-bit integer value
ulsb16	Unsigned 16-bit integer value in lower half of 32-bit register
umsb16	Unsigned 16-bit integer value in upper half of 32-bit register
x clear b,e	Clear a field in x, specified by b (beginning bit) and e (ending bit)
x ext l,r	Extract and sign-extend a field in x, specified by l (shift left value) and r (shift right value)
x extu l,r	Extract an unsigned field in x, specified by l (shift left value) and r (shift right value)
x set b,e	Set field in x to all 1s, specified by b (beginning bit) and e (ending bit)
xint	32-bit integer value that can optionally use cross path
xor	Bitwise exclusive-OR
xsint	Signed 32-bit integer value that can optionally use cross path
xslsb16	Signed 16 LSB of register that can optionally use cross path
xsmsb16	Signed 16 MSB of register that can optionally use cross path
xuint	Unsigned 32-bit integer value that can optionally use cross path
xulsb16	Unsigned 16 LSB of register that can optionally use cross path
xumsb16	Unsigned 16 MSB of register that can optionally use cross path
→	Assignment
+	Addition
++	Increment by 1
x	Multiplication

Table 3-1. Instruction Operation and Execution Notations (Continued)

Symbol	Meaning
-	Subtraction
==	Equal to
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
<<	Shift left
>>	Shift right
>>s	Shift right with sign extension
>>z	Shift right with a zero fill
~	Logical inverse
&	Logical AND

3.2 Instruction Syntax and Opcode Notations

Table 3–2 explains the syntaxes and opcode fields used in the instruction descriptions.

Table 3–2. *Instruction Syntax and Opcode Notations*

Symbol	Meaning
<i>baseR</i>	base address register
<i>creg</i>	3-bit field specifying a conditional register, see section 3.5
<i>cst</i>	constant
<i>csta</i>	constant a
<i>cstb</i>	constant b
<i>cstn</i>	n-bit constant field
<i>dst</i>	destination
<i>mode</i>	addressing mode, see section 3.7
<i>offsetR</i>	register offset
<i>op</i>	opfield; field within opcode that specifies a unique instruction
<i>p</i>	parallel execution; 0 = next instruction is not executed in parallel, 1 = next instruction is executed in parallel
<i>r</i>	Load/store instruction
<i>s</i>	side A or B for destination; 0 = side A, 1 = side B.
<i>scstn</i>	n-bit signed constant field
<i>src</i>	source
<i>src1</i>	source 1
<i>src2</i>	source 2
<i>ucstn</i>	n-bit unsigned constant field
<i>unit</i>	unit decode
<i>x</i>	cross path for <i>src2</i> ; 0 = do not use cross path, 1 = use cross path
<i>y</i>	.D1 or .D2 unit; 0 = .D1 unit, 1 = .D2 unit
<i>z</i>	test for equality with zero or nonzero

3.3 Delay Slots

The execution of fixed-point instructions can be defined in terms of delay slots. The number of delay slots is equivalent to the number of additional cycles required after the source operands are read for the result to be available for reading. For a single-cycle type instruction (such as **ADD**), source operands read in cycle i produce a result that can be read in cycle $i + 1$. For a multiply (**MPY**) instruction, source operands read in cycle i produce a result that can be read in cycle $i + 2$. Table 3–3 shows the number of delay slots associated with each type of instruction.

Delay slots are equivalent to an execution or result latency. All of the instructions in the C62x DSP have a functional unit latency of 1. This means that a new instruction can be started on the functional unit each cycle. Single-cycle throughput is another term for single-cycle functional unit latency.

Table 3–3. Delay Slots

Instruction Type	Delay Slots	Read Cycles [†]	Write Cycles [†]	Branch Taken [†]
NOP (no operation)	0			
Store	0	i	i	
Single cycle	0	i	i	
Multiply (16×16)	1	i	$i + 1$	
Load	4	i	$i, i + 4$ [§]	
Branch	5	i [‡]		$i + 5$

[†] Cycle i is in the E1 pipeline phase.


[‡] The branch to label, branch to IRP, and branch to NRP instructions do not read any general-purpose registers.

[§] The write on cycle $i + 4$ uses a separate write port from other .D unit instructions.

3.4 Parallel Operations

Instructions are always fetched eight words at a time. This constitutes a *fetch packet*. The basic format of a fetch packet is shown in Figure 3-1. Fetch packets are aligned on 256-bit (8-word) boundaries.

Figure 3-1. Basic Format of a Fetch Packet



The execution of the individual instructions is partially controlled by a bit in each instruction, the *p*-bit. The *p*-bit (bit 0) determines whether the instruction executes in parallel with another instruction. The *p*-bits are scanned from left to right (lower to higher address). If the *p*-bit of instruction *i* is 1, then instruction *i* + 1 is to be executed in parallel with (in the the same cycle as) instruction *i*. If the *p*-bit of instruction *i* is 0, then instruction *i* + 1 is executed in the cycle after instruction *i*. All instructions executing in parallel constitute an *execute packet*. An execute packet can contain up to eight instructions. Each instruction in an execute packet must use a different functional unit.


An execute packet cannot cross an 8-word boundary. Therefore, the last *p*-bit in a fetch packet is always cleared to 0, and each fetch packet starts a new execute packet. There are three types of *p*-bit patterns for fetch packets. These three *p*-bit patterns result in the following execution sequences for the eight instructions:

- Fully serial
- Fully parallel
- Partially serial

Example 3-1 through Example 3-3 show the conversion of a *p*-bit sequence into a cycle-by-cycle execution stream of instructions.

Example 3-1. Fully Serial p-Bit Pattern in a Fetch Packet

This p-bit pattern:




results in this execution sequence:

Cycle/Execute Packet	Instructions
1	A
2	B
3	C
4	D
5	E
6	F
7	G
8	H

The eight instructions are executed sequentially.

Example 3-2. Fully Parallel p-Bit Pattern in a Fetch Packet

This p-bit pattern:




results in this execution sequence:

Cycle/Execute Packet	Instructions							
1	A	B	C	D	E	F	G	H

All eight instructions are executed in parallel.

Example 3-3. Partially Serial p-Bit Pattern in a Fetch Packet

This p -bit pattern:



results in this execution sequence:

Cycle/Execute Packet	Instructions
1	A
2	B
3	C D E
4	F G H

Note: Instructions C, D, and E do not use any of the same functional units, cross paths, or other data path resources. This is also true for instructions F, G, and H.

3.4.1 Example Parallel Code

The vertical bars || signify that an instruction is to execute in parallel with the previous instruction. The code for the fetch packet in Example 3-3 would be represented as this:

```

instruction A
instruction B
instruction C
|| instruction D
|| instruction E

instruction F
|| instruction G
|| instruction H

```

3.4.2 Branching Into the Middle of an Execute Packet

If a branch into the middle of an execute packet occurs, all instructions at lower addresses are ignored. In Example 3-3, if a branch to the address containing instruction D occurs, then only D and E execute. Even though instruction C is in the same execute packet, it is ignored. Instructions A and B are also ignored because they are in earlier execute packets. If your result depends on executing A, B, or C, the branch to the middle of the execute packet will produce an erroneous result.

3.5 Conditional Operations

All instructions can be conditional. The condition is controlled by a 3-bit opcode field (*creg*) that specifies the condition register tested, and a 1-bit field (*z*) that specifies a test for zero or nonzero. The four MSBs of every opcode are *creg* and *z*. The specified condition register is tested at the beginning of the E1 pipeline stage for all instructions. For more information on the pipeline, see Chapter 4. If *z* = 1, the test is for equality with zero; if *z* = 0, the test is for nonzero. The case of *creg* = 0 and *z* = 0 is treated as always true to allow instructions to be executed unconditionally. The *creg* field is encoded in the instruction opcode as shown in Table 3-4.

Table 3-4. Registers That Can Be Tested by Conditional Operations

Specified Conditional Register	<i>creg</i>			<i>z</i>	
	Bit	31	30	29	28
Unconditional		0	0	0	0
Reserved†		0	0	0	1
B0		0	0	1	<i>z</i>
B1		0	1	0	<i>z</i>
B2		0	1	1	<i>z</i>
A1		1	0	0	<i>z</i>
A2		1	0	1	<i>z</i>
Reserved		1	1	<i>x</i> ‡	<i>x</i> ‡

† This value is reserved for software breakpoints that are used for emulation purposes.

‡ *x* can be any value.

Conditional instructions are represented in code by using square brackets, [], surrounding the condition register name. The following execute packet contains two **ADD** instructions in parallel. The first **ADD** is conditional on B0 being nonzero. The second **ADD** is conditional on B0 being zero. The character ! indicates the inverse of the condition.

```
[B0]  ADD  .L1  A1,A2,A3
|| [!B0] ADD  .L2  B1,B2,B3
```

The above instructions are mutually exclusive, only one will execute. If they are scheduled in parallel, mutually exclusive instructions are constrained as described in section 3.6. If mutually exclusive instructions share any resources as described in section 3.6, they cannot be scheduled in parallel (put in the same execute packet), even though only one will execute.

3.6 Resource Constraints

No two instructions within the same execute packet can use the same resources. Also, no two instructions can write to the same register during the same cycle. The following sections describe how an instruction can use each of the resources.

3.6.1 Constraints on Instructions Using the Same Functional Unit

Two instructions using the same functional unit cannot be issued in the same execute packet.

The following execute packet is invalid:

```

      ADD  .S1   A0, A1, A2  ; .S1 is used for
|| SHR   .S1   A3, 15, A4  ; both instructions

```

The following execute packet is valid:

```

      ADD  .L1   A0, A1, A2  ; Two different functional
|| SHR   .S1   A3, 15, A4  ; units are used

```

3.6.2 Constraints on Cross Paths (1X and 2X)

Only one unit (.S, .L, or .M unit) per data path, per execute packet, can read a source operand from its opposite register file via the cross paths (1X and 2X), provided that each unit is reading the same operand.

For example, the .S1 unit can read both its operands from the A register file; or it can read an operand from the B register file using the 1X cross path and the other from the A register file. The use of a cross path is denoted by an X following the functional unit name in the instruction syntax (as in S1X).

The following execute packet is invalid because the 1X cross path is being used for two different B register operands:

```

      MV .S1X   B0,A0  ;Invalid: Instructions are using the 1X
|| MV .L1X   B1,A1  ;cross path with different B registers

```

The following execute packet is valid because all uses of the 1X cross path are for the same B register operand, and all uses of the 2X cross path are for the same A register operand:

```

      ADD.L1X  A0,B1,A1 ;Instructions use the 1X
|| MPY.M2X   B4,A4,B2 ;and 2X cross paths

```

The operand comes from a register file opposite of the destination, if the x bit in the instruction field is set.

3.6.3 Constraints on Loads and Stores

Load and store instructions can use an address pointer from one register file while loading to or storing from the other register file. Two load and store instructions using a destination/source from the same register file cannot be issued in the same execute packet. The address register must be on the same side as the .D unit used.

The following execute packet is invalid:

```
LDW.D1    *A0,A1 ; \ .D2 unit must use the address
|| LDW.D2    *A2,B2 ; / register from the B register file
```

The following execute packet is valid:

```
LDW.D1    *A0,A1 ; \ Address registers from correct
|| LDW.D2    *B0,B2 ; / register files
```

Two loads and/or stores loading to and/or storing from the same register file cannot be issued in the same execute packet.

The following execute packet is invalid:

```
LDW.D1    *A4,A5 ; \ Loading to and storing from the
|| STW.D2    A6,*B4 ; / same register file
```

The following execute packets are valid:

```
LDW.D1    *A4,B5 ; \ Loading to, and storing from
|| STW.D2    A6,*B4 ; / different register files
```

```
LDW.D1    *A0,B2 ; \ Loading to
|| LDW.D2    *B0,A1 ; / different register files
```


3.6.4 Constraints on Long (40-Bit) Data

Because the .S and .L units share a read register port for long source operands and a write register port for long results, only one long result may be issued per register file in an execute packet. All instructions with a long result on the .S and .L units have zero delay slots. See section 2.2 for the order for long pairs.

The following execute packet is invalid:

```

    ADD.L1    A5:A4,A1,A3:A2    ; \ Two long writes
|| SHL.S1    A8,A9,A7:A6      ; / on A register file

```

The following execute packet is valid:

```

    ADD.L1    A5:A4,A1,A3:A2    ; \ One long write for
|| SHL.S2    B8,B9,B7:B6      ; / each register file

```

Because the .L and .S units share their long read port with the store port, operations that read a long value cannot be issued on the .L and/or .S units in the same execute packet as a store.

The following execute packet is invalid:

```

    ADD.L1    A5:A4,A1,A3:A2    ; \ Long read operation and a
|| STW.D1    A8,*A9           ; / store

```

The following execute packets are valid:

```

    ADD.L1    A4, A1, A3:A2     ; \ No long read with
|| STW.D1    A8,*A9           ; / the store

```

```

    ADD.L1    A1,A5:A4,A3:A2    ; \ One long write for
|| SHL.S2    B8,B9,B7:B6      ; / each register file

```

```

    ADD.L1    A4, A1, A3:A2     ; \ No long read with the
|| STW.D1T1  A8,*A9           ; / store on T1 path of .D1

```

3.6.5 Constraints on Register Reads

More than four reads of the same register cannot occur on the same cycle. Conditional registers are not included in this count.

The following execute packets are invalid:

```
    MPY  .M1  A1, A1, A4 ; five reads of register A1
|| ADD  .L1  A1, A1, A5
|| SUB  .D1  A1, A2, A3
```

```
    MPY  .M1  A1, A1, A4 ; five reads of register A1
|| ADD  .L1  A1, A1, A5
|| SUB  .D2x A1, B2, B3
```

The following execute packet is valid:

```
|| [A1]    MPY  .M1  A1, A1, A4 ; only four reads of A1
||         ADD  .L1  A0, A1, A5
||         SUB  .D1  A1, A2, A3
```

3.6.6 Constraints on Register Writes

Two instructions cannot write to the same register on the same cycle. Two instructions with the same destination can be scheduled in parallel as long as they do not write to the destination register on the same cycle. For example, an **MPY** issued on cycle i followed by an **ADD** on cycle $i + 1$ cannot write to the same register because both instructions write a result on cycle $i + 1$. Therefore, the following code sequence is invalid unless a branch occurs after the **MPY**, causing the **ADD** not to be issued.

```
MPY .M1  A0, A1, A2
ADD .L1  A4, A5, A2
```

However, this code sequence is valid:

```
    MPY .M1  A0, A1, A2
||  ADD .L1  A4, A5, A2
```

Figure 3–2 shows different multiple-write conflicts. For example, **ADD** and **SUB** in execute packet L1 write to the same register. This conflict is easily detectable.

MPY in packet L2 and **ADD** in packet L3 might both write to B2 simultaneously; however, if a branch instruction causes the execute packet after L2 to be something other than L3, a conflict would not occur. Thus, the potential conflict in L2 and L3 might not be detected by the assembler. The instructions in L4 do not constitute a write conflict because they are mutually exclusive. In contrast, because the instructions in L5 may or may not be mutually exclusive, the assembler cannot determine a conflict. If the pipeline does receive commands to perform multiple writes to the same register, the result is undefined.

Figure 3–2. Examples of the Detectability of Write Conflicts by the Assembler

```
L1:      ADD.L2  B5,B6,B7 ; \ detectable, conflict
||      SUB.S2  B8,B9,B7 ; /
L2:      MPY.M2  B0,B1,B2 ; \ not detectable
L3:      ADD.L2  B3,B4,B2 ; /
L4:[!B0] ADD.L2  B5,B6,B7 ; \ detectable, no conflict
|| [B0]  SUB.S2  B8,B9,B7 ; /
L5:[!B1] ADD.L2  B5,B6,B7 ; \ not detectable
|| [B0]  SUB.S2  B8,B9,B7 ; /
```

3.7 Addressing Modes

The addressing modes on the C62x DSP are linear, circular using BK0, and circular using BK1. The addressing mode is specified by the addressing mode register (AMR), described in section 2.7.3.

All registers can perform linear addressing. Only eight registers can perform circular addressing: A4–A7 are used by the .D1 unit and B4–B7 are used by the .D2 unit. No other units can perform circular addressing. **LDB(U)/LDH(U)/LDW**, **STB/STH/STW**, **ADDAB/ADDAH/ADDAW**, and **SUBAB/SUBAH/SUBAW** instructions all use AMR to determine what type of address calculations are performed for these registers.

3.7.1 Linear Addressing Mode

3.7.1.1 LD and ST Instructions

For load and store instructions, linear mode simply shifts the *offsetR/cst* operand to the left by 2, 1, or 0 for word, halfword, or byte access, respectively; and then performs an add or a subtract to *baseR* (depending on the operation specified).

For the preincrement, predecrement, positive offset, and negative offset address generation options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is the address to be accessed from memory.

3.7.1.2 ADDA and SUBA Instructions

For integer addition and subtraction instructions, linear mode simply shifts the *src1/cst* operand to the left by 2, 1, or 0 for word, halfword, or byte data sizes, respectively, and then performs the add or subtract specified.

3.7.2 Circular Addressing Mode

The BK0 and BK1 fields in AMR specify the block sizes for circular addressing, see section 2.7.3.

3.7.2.1 LD and ST Instructions

As with linear address arithmetic, *offsetR/cst* is shifted left by 2, 1, or 0 according to the data size, and is then added to or subtracted from *baseR* to produce the final address. Circular addressing modifies this slightly by only allowing bits N through 0 of the result to be updated, leaving bits 31 through N + 1 unchanged after address arithmetic. The resulting address is bounded to $2^{(N+1)}$ range, regardless of the size of the *offsetR/cst*.

The circular buffer size in AMR is not scaled; for example, a block-size of 8 is 8 bytes, not 8 times the data size (byte, halfword, word). So, to perform circular addressing on an array of 8 words, a size of 32 should be specified, or $N = 4$. Example 3-4 shows an **LDW** performed with register A4 in circular mode and $BK0 = 4$, so the buffer size is 32 bytes, 16 halfwords, or 8 words. The value in AMR for this example is 0004 0001h.

Example 3-4. LDW Instruction in Circular Mode

```
LDW    .D1    *++A4[9],A1
```

	Before LDW	1 cycle after LDW	5 cycles after LDW
A4	0000 0100h	0000 0104h	0000 0104h
A1	XXXX XXXXh	XXXX XXXXh	1234 5678h
mem 104h	1234 5678h	1234 5678h	1234 5678h

Note: 9h words is 24h bytes. 24h bytes is 4 bytes beyond the 32-byte (20h) boundary 100h-11Fh; thus, it is wrapped around to (124h - 20h = 104h).


3.7.2.2 ADDA and SUBA Instructions

As with linear address arithmetic, *offsetR/cst* is shifted left by 2, 1, or 0 according to the data size, and is then added to or subtracted from *baseR* to produce the final address. Circular addressing modifies this slightly by only allowing bits N through 0 of the result to be updated, leaving bits 31 through N + 1 unchanged after address arithmetic. The resulting address is bounded to $2^{(N+1)}$ range, regardless of the size of the *offsetR/cst*.

The circular buffer size in AMR is not scaled; for example, a block size of 8 is 8 bytes, not 8 times the data size (byte, halfword, word). So, to perform circular addressing on an array of 8 words, a size of 32 should be specified, or $N = 4$. Example 3-5 shows an **ADDAH** performed with register A4 in circular mode and $BK0 = 4$, so the buffer size is 32 bytes, 16 halfwords, or 8 words. The value in AMR for this example is 0004 0001h.

Example 3-5. ADDAH Instruction in Circular Mode

```
ADDAH    .D1    A4, A1, A4
```



Note: 13h halfwords is 26h bytes. 26h bytes is 6 bytes beyond the 32-byte (20h) boundary 100h-11Fh; thus, it is wrapped around to (126h - 20h = 106h).

3.7.3 Syntax for Load/Store Address Generation

The C62x DSP has a load/store architecture, which means that the only way to access data in memory is with a load or store instruction. Table 3-5 shows the syntax of an indirect address to a memory location. Sometimes a large offset is required for a load/store. In this case, you can use the B14 or B15 register as the base register, and use a 15-bit constant (*ucst15*) as the offset.

Table 3-6 describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*).

Table 3-5. Indirect Address Generation for Load/Store

Addressing Type	No Modification of Address Register	Preincrement or Predecrement of Address Register	Postincrement or Postdecrement of Address Register
Register indirect	*R	*++R *- -R	*R++ *R- -
Register relative	*+R[ucst5] *-R[ucst5]	*++R[ucst5] *- -R[ucst5]	*R++[ucst5] *R- -[ucst5]
Register relative with 15-bit constant offset	*+B14/B15[ucst15]	not supported	not supported
Base + index	*+R[offsetR] *-R[offsetR]	*++R[offsetR] *- -R[offsetR]	*R++[offsetR] *R- -[offsetR]

Table 3-6. Address Generator Options for Load/Store

Mode Field				Syntax	Modification Performed
0	0	0	0	*-R[ucst5]	Negative offset
0	0	0	1	*+R[ucst5]	Positive offset
0	1	0	0	*-R[offsetR]	Negative offset
0	1	0	1	*+R[offsetR]	Positive offset
1	0	0	0	*- -R[ucst5]	Predecrement
1	0	0	1	*++R[ucst5]	Preincrement
1	0	1	0	*R- -[ucst5]	Postdecrement
1	0	1	1	*R++[ucst5]	Postincrement
1	1	0	0	*--R[offsetR]	Predecrement
1	1	0	1	*++R[offsetR]	Preincrement
1	1	1	0	*R- -[offsetR]	Postdecrement
1	1	1	1	*R++[offsetR]	Postincrement

3.8 Instruction Compatibility

The C62x, C64x, and C67x DSPs share an instruction set. All of the instructions valid for the C62x DSP are also valid for the C64x and C67x DSPs.

3.9 Instruction Descriptions

This section gives detailed information on the instruction set. Each instruction may present the following information:

- Assembler syntax
- Functional units
- Operands
- Opcode
- Description
- Execution
- Pipeline
- Instruction type
- Delay slots
- Examples

The **ADD** instruction is used as an example to familiarize you with the way each instruction is described. The example describes the kind of information you will find in each part of the individual instruction description and where to obtain more information.

Example*The way each instruction is described.***Syntax**

EXAMPLE (.unit) *src*, *dst*
 .unit = .L1, .L2, .S1, .S2, .D1, .D2

src and *dst* indicate source and destination, respectively. The (.unit) dictates which functional unit the instruction is mapped to (.L1, .L2, .S1, .S2, .M1, .M2, .D1, or .D2).

A table is provided for each instruction that gives the opcode map fields, units the instruction is mapped to, types of operands, and the opcode.

The opcode shows the various fields that make up each instruction. These fields are described in Table 3-2 on page 3-5.

There are instructions that can be executed on more than one functional unit. Table 3-7 shows how this is documented for the **ADD** instruction. This instruction has three opcode map fields: *src1*, *src2*, and *dst*. In the fifth group, the operands have the types *cst5*, *long*, and *long* for *src1*, *src2*, and *dst*, respectively. The ordering of these fields implies *cst5* + *long* → *long*, where + represents the operation being performed by the **ADD**. This operation can be done on .L1 or .L2 (both are specified in the unit column). The *s* in front of each operand signifies that *src1* (*scst5*), *src2* (*slong*), and *dst* (*slong*) are all signed values.

In the ninth group, *src1*, *src2*, and *dst* are *int*, *cst5*, and *int*, respectively. The *u* in front of the *cst5* operand signifies that *src1* (*ucst5*) is an unsigned value. Any operand that begins with *x* can be read from a register file that is different from the destination register file. The operand comes from the register file opposite the destination, if the *x* bit in the instruction is set (shown in the opcode map).

Example *The way each instruction is described*

Table 3-7. Relationships Between Operands, Operand Size, Signed/Unsigned, Functional Units, and Opfields for Example Instruction (ADD)

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	sint	.L1, .L2	000 0011
<i>src2</i>	xsint		
<i>dst</i>	sint		
<i>src1</i>	sint	.L1, .L2	010 0011
<i>src2</i>	xsint		
<i>dst</i>	slong		
<i>src1</i>	xsint	.L1, .L2	010 0001
<i>src2</i>	slong		
<i>dst</i>	slong		
<i>src1</i>	scst5	.L1, .L2	000 0010
<i>src2</i>	xsint		
<i>dst</i>	sint		
<i>src1</i>	scst5	.L1, .L2	010 0000
<i>src2</i>	slong		
<i>dst</i>	slong		
<i>src1</i>	sint	.S1, .S2	00 0111
<i>src2</i>	xsint		
<i>dst</i>	sint		
<i>src1</i>	scst5	.S1, .S2	00 0110
<i>src2</i>	xsint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	01 0000
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	01 0010
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description Instruction execution and its effect on the rest of the processor or memory contents are described. Any constraints on the operands imposed by the processor or the assembler are discussed. The description parallels and supplements the information given by the execution block.

Execution for .L1, .L2 and .S1, .S2 Opcodes

if (cond) $src1 + src2 \rightarrow dst$
 else nop

Execution for .D1, .D2 Opcodes

if (cond) $src2 + src1 \rightarrow dst$
 else nop

The execution describes the processing that takes place when the instruction is executed. The symbols are defined in Table 3-1 (page 3-2).

Pipeline This section contains a table that shows the sources read from, the destinations written to, and the functional unit used during each execution cycle of the instruction.

Instruction Type This section gives the type of instruction. See section 4.2 (page 4-11) for information about the pipeline execution of this type of instruction.


Delay Slots This section gives the number of delay slots the instruction takes to execute. See section 3.3 (page 3-6) for an explanation of delay slots.

Example Examples of instruction execution. If applicable, register and memory values are given before and after instruction execution.

ABS Absolute Value With Saturation

Syntax **ABS** (.unit) *src2*, *dst*

.unit = .L1 or .L2

Opcode


Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i> <i>dst</i>	xsint sint	.L1, .L2	001 1010
<i>src2</i> <i>dst</i>	slong slong	.L1, L2	011 1000

Description The absolute value of *src2* is placed in *dst*.

Execution if (cond) $\text{abs}(\text{src2}) \rightarrow \text{dst}$
 else nop

 The absolute value of *src2* when *src2* is an sint is determined as follows:

- 1) If $\text{src2} \geq 0$, then $\text{src2} \rightarrow \text{dst}$
- 2) If $\text{src2} < 0$ and $\text{src2} \neq -2^{31}$, then $-\text{src2} \rightarrow \text{dst}$
- 3) If $\text{src2} = -2^{31}$, then $2^{31} - 1 \rightarrow \text{dst}$

 The absolute value of *src2* when *src2* is an slong is determined as follows:

- 1) If $\text{src2} \geq 0$, then $\text{src2} \rightarrow \text{dst}$
- 2) If $\text{src2} < 0$ and $\text{src2} \neq -2^{39}$, then $-\text{src2} \rightarrow \text{dst}$
- 3) If $\text{src2} = -2^{39}$, then $2^{39} - 1 \rightarrow \text{dst}$

Pipeline

Pipeline Stage	E1
Read	<i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

Example 1 ABS .L1 A1,A5

	Before instruction		1 cycle after instruction
A1	8000 4E3Dh	-2147463619	A1 8000 4E3Dh -2147463619
A5	xxxxx xxxxxh		A5 7FFF B1C3h 2147463619

Example 2 ABS .L1 A1,A5


	Before instruction		1 cycle after instruction
A1	3FF6 0010h	1073086480	A1 3FF6 0010h 1073086480
A5	xxxxx xxxxxh		A5 3FF6 0010h 1073086480

ADD *Add Two Signed Integers Without Saturation*

ADD *Add Two Signed Integers Without Saturation*

Syntax **ADD** (.unit) *src1*, *src2*, *dst*
 or
ADD (.D1 or .D2) *src2*, *src1*, *dst*
.unit = .L1, .L2, .S1, .S2


Opcode .L unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint sint	.L1, .L2	000 0011
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint slong	.L1, .L2	010 0011
<i>src1</i> <i>src2</i> <i>dst</i>	xsint slong slong	.L1, .L2	010 0001
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xsint sint	.L1, .L2	000 0010
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 slong slong	.L1, .L2	010 0000

Opcode

.S unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	sint	.S1, .S2	00 0111
<i>src2</i>	xsint		
<i>dst</i>	sint		
<i>src1</i>	scst5	.S1, .S2	00 0110
<i>src2</i>	xsint		
<i>dst</i>	sint		

Description for .L1, .L2 and .S1, .S2 Opcodes

src2 is added to *src1*. The result is placed in *dst*.


Execution for .L1, .L2 and .S1, .S2 Opcodes

if (cond)
 $src1 + src2 \rightarrow dst$
 else nop

ADD Add Two Signed Integers Without Saturation

Opcode

.D unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	01 0000
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	01 0010
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description for .D1, .D2 Opcodes

src1 is added to *src2*. The result is placed in *dst*.

Execution for .D1, .D2 Opcodes

if (cond)
 $src2 + src1 \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.L, .S, or .D

Instruction Type Single-cycle

Delay Slots 0

See Also ADDK, ADDU, ADD2, SADD, SUB

Example 1 `ADD .L2X A1,B1,B2`

	Before instruction		1 cycle after instruction		
A1	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table> 12890	0000 325Ah		A1 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table>	0000 325Ah
0000 325Ah					
0000 325Ah					
B1	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>FFFF FF12h</td></tr></table> -238	FFFF FF12h		B1 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>FFFF FF12h</td></tr></table>	FFFF FF12h
FFFF FF12h					
FFFF FF12h					
B2	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		B2 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 316Ch</td></tr></table> 12652	0000 316Ch
xxxx xxxxh					
0000 316Ch					

Example 2 `ADD .L1 A1,A3:A2,A5:A4`

	Before instruction		1 cycle after instruction				
A1	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table> 12890	0000 325Ah		A1 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table>	0000 325Ah		
0000 325Ah							
0000 325Ah							
A3:A2	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 00FFh</td><td>FFFF FF12h</td></tr></table> -228 [§]	0000 00FFh	FFFF FF12h		A3:A2 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 00FFh</td><td>FFFF FF12h</td></tr></table>	0000 00FFh	FFFF FF12h
0000 00FFh	FFFF FF12h						
0000 00FFh	FFFF FF12h						
A5:A4	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 0000h</td><td>0000 0000h</td></tr></table> 0 [§]	0000 0000h	0000 0000h		A5:A4 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 0000h</td><td>0000 316Ch</td></tr></table> 12652 [§]	0000 0000h	0000 316Ch
0000 0000h	0000 0000h						
0000 0000h	0000 316Ch						

[§] Signed 40-bit (long) integer

Example 3 `ADD .L1 -13,A1,A6`

	Before instruction		1 cycle after instruction		
A1	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table> 12890	0000 325Ah		A1 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table>	0000 325Ah
0000 325Ah					
0000 325Ah					
A6	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		A6 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 324Dh</td></tr></table> 12877	0000 324Dh
xxxx xxxxh					
0000 324Dh					

Example 4 `ADD .D1 A1,26,A6`

	Before instruction		1 cycle after instruction		
A1	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table> 12890	0000 325Ah		A1 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 325Ah</td></tr></table>	0000 325Ah
0000 325Ah					
0000 325Ah					
A6	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		A6 <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>0000 3274h</td></tr></table> 12916	0000 3274h
xxxx xxxxh					
0000 3274h					


ADDAB *Add Using Byte Addressing Mode*

ADDAB *Add Using Byte Addressing Mode*

Syntax **ADDAB** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 0000
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 0010
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description *src1* is added to *src2* using the byte addressing mode specified for *src2*. The addition defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). The result is placed in *dst*.

Execution if (cond) *src2* +a *src1* → *dst*
 else nop

Pipeline

Pipeline stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, ADDAH, ADDAW**

Example 1

ADDAB .D1 A4,A2,A4

	Before instruction	1 cycle after instruction
A2	0000 000Bh	0000 000Bh
A4	0000 0100h	0000 0103h
AMR	0002 0001h	0002 0001h

BK0 = 2 → size = 8
A4 in circular addressing mode using BK0

ADDAH *Add Using Halfword Addressing Mode*

ADDAH *Add Using Halfword Addressing Mode*

Syntax **ADDAH** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode

31	29	28	27	23	22	18	17	13	12	7	6	5	4	3	2	1	0						
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>			<i>op</i>			1	0	0	0	0	0	<i>s</i>	<i>p</i>
3			1	5			5			5			6			1	1			1	1		

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 0100
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 0110
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description *src1* is added to *src2* using the halfword addressing mode specified for *src2*. The addition defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). If circular addressing is enabled, *src1* is left shifted by 1. The result is placed in *dst*.

Execution if (cond) *src2* +a *src1* → *dst*
 else nop

Pipeline stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, ADDAB, ADDAW**

Example 1

ADDAH .D1 A4,A2,A4

Before instruction

A2	0000 000Bh
A4	0000 0100h
AMR	0002 0001h

1 cycle after instruction

A2	0000 000Bh
A4	0000 0106h
AMR	0002 0001h

BK0 = 2 → size = 8
A4 in circular addressing mode using BK0


ADDAW *Add Using Word Addressing Mode*

ADDAW *Add Using Word Addressing Mode*

Syntax **ADDAW** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 1000
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 1010
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description *src1* is added to *src2* using the word addressing mode specified for *src2*. The addition defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). If circular addressing is enabled, *src1* is left shifted by 2. The result is placed in *dst*.

Execution if (cond) *src2* +a *src1* → *dst*
 else nop

Pipeline

Pipeline stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, ADDAB, ADDAH**

Example 1

ADDAW .D1 A4,2,A4

	Before instruction	1 cycle after instruction		
A4	<table border="1"><tr><td>0002 0000h</td></tr></table>	0002 0000h	<table border="1"><tr><td>0002 0000h</td></tr></table>	0002 0000h
0002 0000h				
0002 0000h				
AMR	<table border="1"><tr><td>0002 0001h</td></tr></table>	0002 0001h	<table border="1"><tr><td>0002 0001h</td></tr></table>	0002 0001h
0002 0001h				
0002 0001h				

BK0 = 2 → size = 8
A4 in circular addressing mode using BK0


ADDK *Add Signed 16-Bit Constant to Register*

ADDK *Add Signed 16-Bit Constant to Register*

Syntax **ADDK** (.unit) *cst*, *dst*

.unit = .S1 or .S2

Opcode



Opcode map field used...	For operand type...	Unit
<i>cst16</i>	scst16	.S1, .S2
<i>dst</i>	uint	

Description A 16-bit signed constant, *cst16*, is added to the *dst* register specified. The result is placed in *dst*.

Execution if (cond) $cst + dst \rightarrow dst$
else nop

Pipeline

Pipeline Stage	E1
Read	<i>cst16</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type Single-cycle


Delay Slots 0

Example ADDK .S1 15401, A1

	Before instruction		1 cycle after instruction		
A1	0021 37E1h	2176993	A1	0021 740Ah	2192394

ADDU*Add Two Unsigned Integers Without Saturation*

Syntax **ADDU** (.unit) *src1*, *src2*, *dst*
 .unit = .L1 or .L2

Opcode

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	010 1011
<i>src2</i>	xuint		
<i>dst</i>	ulong		
<i>src1</i>	xuint	.L1, .L2	010 1001
<i>src2</i>	ulong		
<i>dst</i>	ulong		

Description *src2* is added to *src1*. The result is placed in *dst*.

Execution if (cond)
 $src1 + src2 \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L


Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, SADD, SUBU**

ADDU *Add Two Unsigned Integers Without Saturation*


Example 1 `ADDU .L1 A1,A2,A5:A4`



[†] Unsigned 32-bit integer

[‡] Unsigned 40-bit (long) integer

Example 2 `ADDU .L1 A1,A3:A2,A5:A4`



[†] Unsigned 32-bit integer

[‡] Unsigned 40-bit (long) integer

ADD2


Add Two 16-Bit Integers on Upper and Lower Register Halves

Syntax

ADD2 (.unit) *src1*, *src2*, *dst*

.unit = .S1 or .S2

Opcode

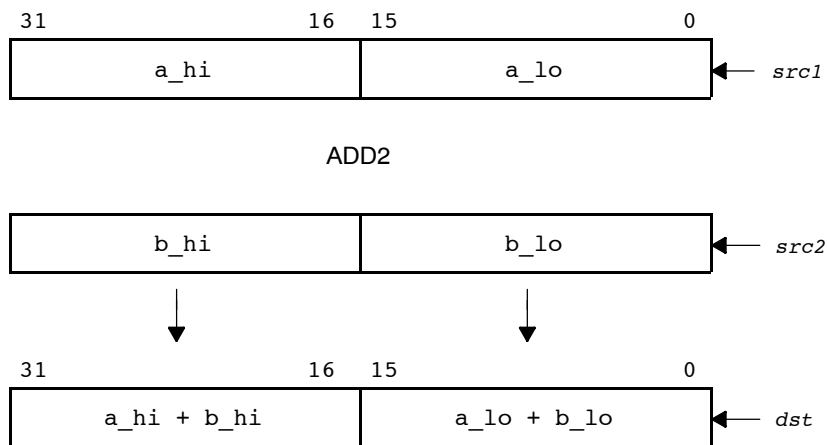


Opcode map field used...	For operand type...	Unit
<i>src1</i>	sint	.S1, .S2
<i>src2</i>	xsint	
<i>dst</i>	sint	

Description

The upper and lower halves of the *src1* operand are added to the upper and lower halves of the *src2* operand. The values in *src1* and *src2* are treated as signed, packed 16-bit data and the results are written in signed, packed 16-bit format into *dst*.

For each pair of signed packed 16-bit values found in the *src1* and *src2*, the sum between the 16-bit value from *src1* and the 16-bit value from *src2* is calculated to produce a 16-bit result. The result is placed in the corresponding positions in the *dst*. The carry from the lower half add does not affect the upper half add.



ADD2 *Add Two 16-Bit Integers on Upper and Lower Register Halves*

Execution

```

if (cond) {
    msb16(src1) + msb16(src2) → msb16(dst);
    lsb16(src1) + lsb16(src2) → lsb16(dst);
}
else nop
    
```

Pipeline


Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type Single-cycle

Delay Slots 0


See Also **ADD, ADDU, SUB2**

Example `ADD2 .S1X A1,B1,A2`




AND*Bitwise AND***Syntax** **AND** (.unit) *src1*, *src2*, *dst*

.unit = .L1, .L2, .S1, .S2

Opcode .L unit

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	111 1011
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	111 1010
<i>src2</i>	xuint		
<i>dst</i>	uint		

Opcode .S unit

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.S1, .S2	01 1111
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.S1, .S2	01 1110
<i>src2</i>	xuint		
<i>dst</i>	uint		

Description Performs a bitwise **AND** operation between *src1* and *src2*. The result is placed in *dst*. The *scst5* operands are sign extended to 32 bits.**Execution** if (cond) *src1* AND *src2* → *dst*
else nop


Pipeline	Pipeline Stage	E1
	Read	<i>src1, src2</i>
	Written	<i>dst</i>
	Unit in use	.L or .S

Instruction Type Single-cycle


Delay Slots 0

See Also **OR, XOR**

Example 1 `AND .L1X A1,B1,A2`




Example 2 `AND .L1 15,A1,A3`



B Branch Using a Displacement**Syntax****B** (.unit) label

.unit = .S1 or .S2

Opcode

Opcode map field used...	For operand type...	Unit
<i>cst21</i>	<i>scst21</i>	.S1, .S2

Description

A 21-bit signed constant, *cst21*, is shifted left by 2 bits and is added to the address of the first instruction of the fetch packet that contains the branch instruction. The result is placed in the program fetch counter (PFC). The assembler/linker automatically computes the correct value for *cst21* by the following formula:

$$cst21 = (\text{label} - \text{PCE1}) \gg 2$$

If two branches are in the same execute packet and both are taken, behavior is undefined.

Two conditional branches can be in the same execute packet if one branch uses a displacement and the other uses a register, IRP, or NRP. As long as only one branch has a true condition, the code executes in a well-defined way.

Execution

if (cond) $cst21 \ll 2 + \text{PCE1} \rightarrow \text{PFC}$
 else nop

Notes:

- 1) PCE1 (program counter) represents the address of the first instruction in the fetch packet in the E1 stage of the pipeline. PFC is the program fetch counter.
- 2) The execute packets in the delay slots of a branch cannot be interrupted. This is true regardless of whether the branch is taken.
- 3) See section 3.4.2 on page 3-9 for information on branching into the middle of an execute packet.

B Branch Using a Displacement

Pipeline Stage	Target Instruction						
	E1	PS	PW	PR	DP	DC	E1
Read							
Written							
Branch Taken							✓
Unit in use	.S						

Instruction Type Branch

Delay Slots 5

Example Table 3-8 gives the program counter values and actions for the following code example.

```

0000 0000          B      .S1 LOOP
0000 0004          ADD    .L1 A1, A2, A3
0000 0008          ||    ADD    .L2 B1, B2, B3
0000 000C  LOOP:  MPY    .M1X A3, B3, A4
0000 0010          ||    SUB    .D1 A5, A6, A6
0000 0014          MPY    .M1 A3, A6, A5
0000 0018          MPY    .M1 A6, A7, A8
0000 001C          SHR    .S1 A4, 15, A4
0000 0020          ADD    .D1 A4, A6, A4


```

Table 3-8. Program Counter Values for Example Branch Using a Displacement

Cycle	Program Counter Value	Action
Cycle 0	0000 0000h	Branch command executes (target code fetched)
Cycle 1	0000 0004h	
Cycle 2	0000 000Ch	
Cycle 3	0000 0014h	
Cycle 4	0000 0018h	
Cycle 5	0000 001Ch	
Cycle 6	0000 000Ch	Branch target code executes
Cycle 7	0000 0014h	

B *Branch Using a Register*

Syntax **B** (.unit) *src2*
 .unit = .S2

Opcode

Opcode map field used...	For operand type...	Unit
<i>src2</i>	xuint	.S2

Description

src2 is placed in the program fetch counter (PFC).

If two branches are in the same execute packet and are both taken, behavior is undefined.

Two conditional branches can be in the same execute packet if one branch uses a displacement and the other uses a register, IRP, or NRP. As long as only one branch has a true condition, the code executes in a well-defined way.

Execution

if (cond) *src2* → PFC
 else nop

Notes:

- 1) This instruction executes on .S2 only. PFC is program fetch counter.
- 2) The execute packets in the delay slots of a branch cannot be interrupted. This is true regardless of whether the branch is taken.
- 3) See section 3.4.2 on page 3-9 for information on branching into the middle of an execute packet.

Pipeline

Pipeline Stage	Target Instruction						
	E1	PS	PW	PR	DP	DC	E1
Read	<i>src2</i>						
Written							
Branch Taken	↗						
Unit in use	.S2						

B Branch Using a Register

Instruction Type Branch

Delay Slots 5

Example Table 3–9 gives the program counter values and actions for the following code example. In this example, the B10 register holds the value 1000 000Ch.


```
          B10 1000 000Ch
1000 0000      B      .S2 B10
1000 0004      ADD   .L1 A1, A2, A3
1000 0008      ||  ADD   .L2 B1, B2, B3
1000 000C      MPY   .M1X A3, B3, A4
1000 0010      ||  SUB   .D1 A5, A6, A6
1000 0014      MPY   .M1 A3, A6, A5
1000 0018      MPY   .M1 A6, A7, A8
1000 001C      SHR   .S1 A4, 15, A4
1000 0020      ADD   .D1 A4, A6, A4
```

Table 3–9. Program Counter Values for Example Branch Using a Register

Cycle	Program Counter Value	Action
Cycle 0	1000 0000h	Branch command executes (target code fetched)
Cycle 1	1000 0004h	
Cycle 2	1000 000Ch	
Cycle 3	1000 0014h	
Cycle 4	1000 0018h	
Cycle 5	1000 001Ch	
Cycle 6	1000 000Ch	Branch target code executes
Cycle 7	1000 0014h	

B IRP*Branch Using an Interrupt Return Pointer***Syntax****B** (.unit) IRP

.unit = .S2

Opcode

Opcode map field used...	For operand type...	Unit
<i>src2</i>	xsint	.S2

Description

IRP is placed in the program fetch counter (PFC). This instruction also moves the PGIE bit value to the GIE bit. The PGIE bit is unchanged.

If two branches are in the same execute packet and are both taken, behavior is undefined.

Two conditional branches can be in the same execute packet if one branch uses a displacement and the other uses a register, IRP, or NRP. As long as only one branch has a true condition, the code executes in a well-defined way.

Execution

if (cond) IRP → PFC
else nop

Notes:

- 1) This instruction executes on .S2 only. PFC is the program fetch counter.
- 2) Refer to Chapter 5, *Interrupts*, for more information on IRP, PGIE, and GIE.
- 3) The execute packets in the delay slots of a branch cannot be interrupted. This is true regardless of whether the branch is taken.
- 4) See section 3.4.2 on page 3-9 for information on branching into the middle of an execute packet.

B IRP *Branch Using an Interrupt Return Pointer*

Pipeline

Pipeline Stage	Target Instruction						
	E1	PS	PW	PR	DP	DC	E1
Read	IRP						
Written							
Branch Taken							✓
Unit in use	.S2						

Instruction Type

Branch

Delay Slots

5

Example

Table 3-10 gives the program counter values and actions for the following code example. Given that an interrupt occurred at

```

PC = [0000 1000]   IRP = [0000 1000]

0000 0020   B       .S2 IRP
0000 0024   ADD    .S1 A0, A2, A1
0000 0028   MPY    .M1 A1, A0, A1
0000 002C   NOP
0000 0030   SHR    .S1 A1, 15, A1
0000 0034   ADD    .L1 A1, A2, A1
0000 0038   ADD    .L2 B1, B2, B3
    
```

Table 3-10. Program Counter Values for B IRP Instruction

Cycle	Program Counter Value	Action
Cycle 0	0000 0020	Branch command executes (target code fetched)
Cycle 1	0000 0024	
Cycle 2	0000 0028	
Cycle 3	0000 002C	
Cycle 4	0000 0030	
Cycle 5	0000 0034	
Cycle 6	0000 1000	Branch target code executes

B NRP*Branch Using NMI Return Pointer***Syntax****B** (.unit) **NRP**

.unit = .S2

Opcode

31	29	28	27	23	22	18	17	13	12	11	6	5	4	3	2	1	0											
<i>creg</i>	<i>z</i>	<i>dst</i>			0	0	1	1	1	0	0	0	0	0	x	0	0	0	0	1	1	1	0	0	0	0	<i>s</i>	<i>p</i>
3	1	5										1											1	1				

Opcode map field used...	For operand type...	Unit
<i>src2</i>	xsint	.S2

Description

NRP is placed in the program fetch counter (PFC). This instruction also sets the NMIE bit. The PGIE bit is unchanged.

If two branches are in the same execute packet and are both taken, behavior is undefined.

Two conditional branches can be in the same execute packet if one branch uses a displacement and the other uses a register, IRP, or NRP. As long as only one branch has a true condition, the code executes in a well-defined way.

Execution

if (cond) NRP → PFC
else nop

Notes:

- 1) This instruction executes on .S2 only. PFC is program fetch counter.
- 2) Refer to Chapter 5, *Interrupts*, for more information on NRP and NMIE.
- 3) The execute packets in the delay slots of a branch cannot be interrupted. This is true regardless of whether the branch is taken.
- 4) See section 3.4.2 on page 3-9 for information on branching into the middle of an execute packet.

B NRP *Branch Using NMI Return Pointer*

Pipeline	Target Instruction						
	E1	PS	PW	PR	DP	DC	E1
Read	NRP						
Written							
Branch Taken							✓
Unit in use	.S2						

Instruction Type Branch

Delay Slots 5

Example Table 3-11 gives the program counter values and actions for the following code example. Given that an interrupt occurred at

```

PC = 0000 1000   NRP = 0000 1000

0000 0020   B      .S2 NRP
0000 0024   ADD   .S1 A0, A2, A1
0000 0028   MPY   .M1 A1, A0, A1
0000 002C   NOP
0000 0030   SHR   .S1 A1, 15, A1
0000 0034   ADD   .L1 A1, A2, A1
0000 0038   ADD   .L2 B1, B2, B3
    
```

Table 3-11. Program Counter Values for B NRP Instruction

Cycle	Program Counter Value	Action
Cycle 0	0000 0020	Branch command executes (target code fetched)
Cycle 1	0000 0024	
Cycle 2	0000 0028	
Cycle 3	0000 002C	
Cycle 4	0000 0030	
Cycle 5	0000 0034	
Cycle 6	0000 1000	Branch target code executes

CLR*Clear a Bit Field***Syntax****CLR** (.unit) *src2*, *csta*, *cstb*, *dst*


or

CLR (.unit) *src2*, *src1*, *dst*

.unit = .S1 or .S2

Opcode


Constant form



Opcode map field used...	For operand type...	Unit
<i>src2</i>	uint	.S1, .S2
<i>csta</i>	ucst5	
<i>cstb</i>	ucst5	
<i>dst</i>	uint	

Opcode


Register form



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xuint	.S1, .S2
<i>src1</i>	uint	
<i>dst</i>	uint	

Description

The field in *src2*, specified by *csta* and *cstb*, is cleared to zero. *csta* and *cstb* may be specified as constants or as the ten LSBs of the *src1* registers, with *cstb* being bits 0–4 and *csta* bits 5–9. *csta* signifies the bit location of the LSB in the field and *cstb* signifies the bit location of the MSB in the field. In other words, *csta* and *cstb* represent the beginning and ending bits, respectively, of the field to be cleared. The LSB location of *src2* is 0 and the MSB location of *src2* is 31. In the example below, *csta* is 15 and *cstb* is 23. Only the ten LSBs are valid for the register version of the instruction. If any of the 22 MSBs are non-zero, the result is invalid.



Execution

If the constant form is used:

if (cond) *src2* clear *csta*, *cstb* → *dst*
 else nop

If the register form is used:

if (cond) *src2* clear *src1*_{9..5}, *src1*_{4..0} → *dst*
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type Single-cycle

Delay Slots 0

See Also SET

Example 1

CLR .S1 A1,4,19,A2

	Before instruction	1 cycle after instruction
A1	07A4 3F2Ah	07A4 3F2Ah
A2	xxxx xxxxh	07A0 000Ah

Example 2

CLR .S2 B1,B3,B2

	Before instruction	1 cycle after instruction
B1	03B6 E7D5h	03B6 E7D5h
B2	xxxx xxxxh	03B0 0001h
B3	0000 0052h	0000 0052h

CMPEQ Compare for Equality, Signed Integers

CMPEQ Compare for Equality, Signed Integers

Syntax **CMPEQ** (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode

31	29	28	27	23	22	18	17	13	12	11	5	4	3	2	1	0				
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>		<i>x</i>	<i>op</i>			1	1	0	<i>s</i>	<i>p</i>
3			1	5			5			5		1	7			1	1	1		1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint uint	.L1, .L2	101 0011
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xsint uint	.L1, .L2	101 0010
<i>src1</i> <i>src2</i> <i>dst</i>	xsint slong uint	.L1, .L2	101 0001
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 slong uint	.L1, .L2	101 0000

Description Compares *src1* to *src2*. If *src1* equals *src2*, then 1 is written to *dst*; otherwise, 0 is written to *dst*.

Execution

```
if (cond) {
    if (src1 == src2) 1 → dst
    else 0 → dst
}
else nop
```

Pipeline


Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle


Delay Slots 0

See Also **CMPGT, CMPLT**


Example 1 `CMPEQ .L1X A1,B1,A2`



Example 2 `CMPEQ .L1 Ch,A1,A2`



Example 3 `CMPEQ .L2X A1,B3:B2,B1`



CMPGT *Compare for Greater Than, Signed Integers*

CMPGT


Compare for Greater Than, Signed Integers

Syntax

CMPGT (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	sint	.L1, .L2	100 0111
<i>src2</i>	xsint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	100 0110
<i>src2</i>	xsint		
<i>dst</i>	uint		
<i>src1</i>	xsint	.L1, .L2	100 0101
<i>src2</i>	slong		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	100 0100
<i>src2</i>	slong		
<i>dst</i>	uint		

Description Performs a signed comparison of *src1* to *src2*. If *src1* is greater than *src2*, then a 1 is written to *dst*; otherwise, a 0 is written to *dst*.

Note:

The **CMPGT** instruction allows using a 5-bit constant as *src1*. If *src2* is a 5-bit constant, as in

```
CMPGT .L1 A4, 5, A0
```

Then to implement this operation, the assembler converts this instruction to

```
CMPLT .L1 5, A4, A0
```

These two instructions are equivalent, with the second instruction using the conventional operand types for *src1* and *src2*.

Similarly, the **CMPGT** instruction allows a cross path operand to be used as *src2*. If *src1* is a cross path operand as in

```
CMPGT .L1x B4, A5, A0
```

Then to implement this operation the assembler converts this instruction to

```
CMPLT .L1x A5, B4, A0
```

In both of these operations the listing file (.lst) will have the first implementation, and the second implementation will appear in the debugger.

Execution

```
if (cond) {
    if (src1 > src2) 1 → dst
    else 0 → dst
}
else nop
```

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

See Also **CMPEQ, CMPGTU, CMPLT**

Example 1

CMPGT .L1X A1,B1,A2

Before instruction		1 cycle after instruction			
A1	0000 01B6h	438	A1	0000 01B6h	
A2	xxxx xxxxh		A2	0000 0000h	false
B1	0000 08BDh	2237	B1	0000 08BDh	

Example 2

CMPGT .L1X A1,B1,A2

Before instruction		1 cycle after instruction			
A1	FFFF FE91h	-367	A1	FFFF FE91h	
A2	xxxx xxxxh		A2	0000 0001h	true
B1	FFFF FDC4h	-572	B1	FFFF FDC4h	

Example 3

CMPGT .L1 8,A1,A2

Before instruction		1 cycle after instruction			
A1	0000 0023h	35	A1	0000 0023h	
A2	xxxx xxxxh		A2	0000 0000h	false

Example 4

CMPGT .L1X A1,B1,A2

Before instruction		1 cycle after instruction			
A1	0000 00EBh	235	A1	0000 00EBh	
A2	xxxx xxxxh		A2	0000 0000h	false
B1	0000 00EBh	235	B1	0000 00EBh	

CMPGTU*Compare for Greater Than, Unsigned Integers***Syntax****CMPGTU** (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode

31	29	28	27	23	22	18	17	13	12	11	5	4	3	2	1	0			
<i>creg</i>	<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>			<i>x</i>	<i>op</i>			1	1	0	<i>s</i>	<i>p</i>
3	1	5			5			5			1	7			1	1	1	1	1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	100 1111
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	ucst4	.L1, .L2	100 1110
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	xuint	.L1, .L2	100 1101
<i>src2</i>	ulong		
<i>dst</i>	uint		
<i>src1</i>	ucst4	.L1, .L2	100 1100
<i>src2</i>	ulong		
<i>dst</i>	uint		

Description

Performs an unsigned comparison of *src1* to *src2*. If *src1* is greater than *src2*, then a 1 is written to *dst*; otherwise, a 0 is written to *dst*. Only the four LSBs are valid in the 5-bit *dst* field when the *ucst4* operand is used. If the MSB of the *dst* field is nonzero, the result is invalid.

Execution

```
if (cond) {
    if (src1 > src2) 1 → dst
    else 0 → dst
}
else nop
```

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L


CMPGTU *Compare for Greater Than, Unsigned Integers*

Instruction Type Single-cycle

Delay Slots 0


See Also **CMPGT, CMPLTU**

Example 1 `CMPGTU .L1 A1,A2,A3`




[†] Unsigned 32-bit integer

Example 2 `CMPGTU .L1 0Ah,A1,A2`



[†] Unsigned 32-bit integer

Example 3 `CMPGTU .L1 0Eh,A3:A2,A4`



[‡] Unsigned 40-bit (long) integer

CMPLT


Compare for Less Than, Signed Integers

Syntax

CMPLT (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	sint	.L1, .L2	101 0111
<i>src2</i>	xsint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	101 0110
<i>src2</i>	xsint		
<i>dst</i>	uint		
<i>src1</i>	xsint	.L1, .L2	101 0101
<i>src2</i>	slong		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	101 0100
<i>src2</i>	slong		
<i>dst</i>	uint		

CMPLT *Compare for Less Than, Signed Integers*

Description Performs a signed comparison of *src1* to *src2*. If *src1* is less than *src2*, then 1 is written to *dst*; otherwise, 0 is written to *dst*.

Note:

The **CMPLT** instruction allows using a 5-bit constant as *src1*. If *src2* is a 5-bit constant, as in

```
CMPLT .L1 A4, 5, A0
```

Then to implement this operation, the assembler converts this instruction to

```
CMPGT .L1 5, A4, A0
```

These two instructions are equivalent, with the second instruction using the conventional operand types for *src1* and *src2*.

Similarly, the **CMPLT** instruction allows a cross path operand to be used as *src2*. If *src1* is a cross path operand as in

```
CMPLT .L1x B4, A5, A0
```

Then to implement this operation, the assembler converts this instruction to

```
CMPGT .L1x A5, B4, A0
```

In both of these operations the listing file (.lst) will have the first implementation, and the second implementation will appear in the debugger.

Execution

```
if (cond) {
    if (src1 < src2) 1 → dst
    else 0 → dst
}
else nop
```

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

See Also **CMPEQ, CMPGT, CMPLTU**

Example 1

CMPLT .L1 A1,A2,A3

Before instruction		1 cycle after instruction			
A1	0000 07E2h	2018	A1	0000 07E2h	
A2	0000 0F6Bh	3947	A2	0000 0F6Bh	
A3	xxxx xxxxh		A3	0000 0001h	true

Example 2

CMPLT .L1 A1,A2,A3

Before instruction		1 cycle after instruction			
A1	FFFF FED6h	-298	A1	FFFF FED6h	
A2	0000 000Ch	12	A2	0000 000Ch	
A3	xxxx xxxxh		A3	0000 0001h	true

Example 3

CMPLT .L1 9,A1,A2

Before instruction		1 cycle after instruction			
A1	0000 0005h	5	A1	0000 0005h	
A2	xxxx xxxxh		A2	0000 0000h	false

CMPLTU *Compare for Less Than, Unsigned Integers*

CMPLTU

Compare for Less Than, Unsigned Integers

Syntax

CMPLTU (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode

31	29	28	27	23	22	18	17	13	12	11	5	4	3	2	1	0			
<i>creg</i>	<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>			<i>x</i>	<i>op</i>			1	1	0	<i>s</i>	<i>p</i>
3	1	5			5			5			1	7			1	1	1		1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	uint xuint uint	.L1, .L2	101 1111
<i>src1</i> <i>src2</i> <i>dst</i>	ucst4 xuint uint	.L1, .L2	101 1110
<i>src1</i> <i>src2</i> <i>dst</i>	xuint ulong uint	.L1, .L2	101 1101
<i>src1</i> <i>src2</i> <i>dst</i>	ucst4 ulong uint	.L1, .L2	101 1100

Description

Performs an unsigned comparison of *src1* to *src2*. If *src1* is less than *src2*, then 1 is written to *dst*; otherwise, 0 is written to *dst*.

Execution

```
if (cond) {
    if (src1 < src2) 1 → dst
    else 0 → dst
}
else nop
```

Pipeline

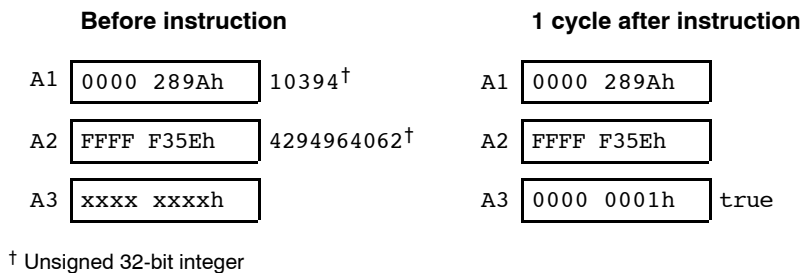
Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle


Delay Slots 0

See Also **CMPGTU, CMPLT**


Example 1 `CMPLTU .L1 A1,A2,A3`



Example 2 `CMPLTU .L1 14,A1,A2`



Example 3 `CMPLTU .L1 A1,A5:A4,A2`



[†] Unsigned 32-bit integer

[‡] Unsigned 40-bit (long) integer

EXT *Extract and Sign-Extend a Bit Field*

EXT


Extract and Sign-Extend a Bit Field

Syntax

EXT (.unit) *src2*, *csta*, *cstb*, *dst*
 or
EXT (.unit) *src2*, *src1*, *dst*
 .unit = .S1 or .S2

Opcode


Constant form



Opcode map field used...	For operand type...	Unit
<i>src2</i>	sint	.S1, .S2
<i>csta</i>	ucst5	
<i>cstb</i>	ucst5	
<i>dst</i>	sint	

Opcode

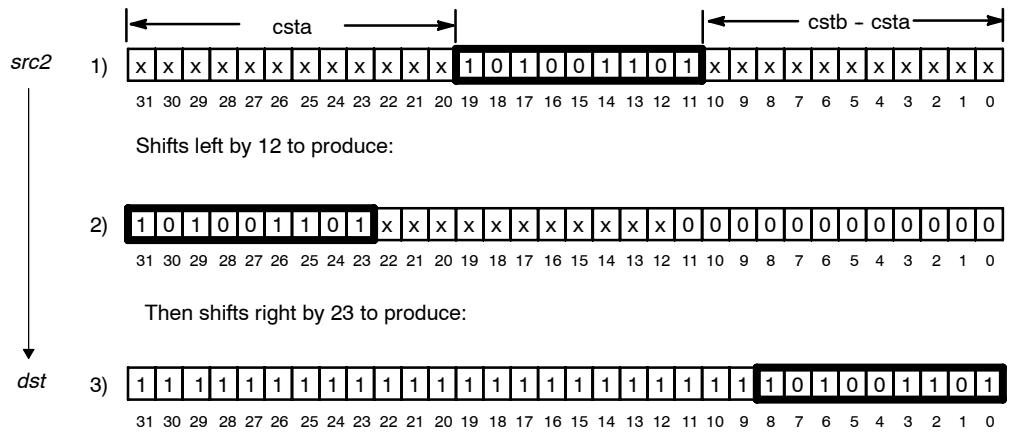
Register form



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xsint	.S1, .S2
<i>src1</i>	uint	
<i>dst</i>	sint	

Description

The field in *src2*, specified by *csta* and *cstb*, is extracted and sign-extended to 32 bits. The extract is performed by a shift left followed by a signed shift right. *csta* and *cstb* are the shift left amount and shift right amount, respectively. This can be thought of in terms of the LSB and MSB of the field to be extracted. Then $csta = 31 - \text{MSB of the field}$ and $cstb = csta + \text{LSB of the field}$. The shift left and shift right amounts may also be specified as the ten LSBs of the *src1* register with *cstb* being bits 0–4 and *csta* bits 5–9. In the example below, *csta* is 12 and *cstb* is $11 + 12 = 23$. Only the ten LSBs are valid for the register version of the instruction. If any of the 22 MSBs are non-zero, the result is invalid.



Execution

If the constant form is used:

if (cond) $src2 \text{ ext } csta, cstb \rightarrow dst$
 else nop

If the register form is used:

if (cond) $src2 \text{ ext } src1_{9..5}, src1_{4..0} \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.S


EXT *Extract and Sign-Extend a Bit Field*

Instruction Type Single-cycle


Delay Slots 0

See Also **EXTU**

Example 1 `EXT .S1 A1,10,19,A2`



Example 2 `EXT .S1 A1,A2,A3`



EXTU


Extract and Zero-Extend a Bit Field

Syntax

EXTU (.unit) *src2*, *csta*, *cstb*, *dst*
 or
EXTU (.unit) *src2*, *src1*, *dst*
 .unit = .S1 or .S2

Opcode


Constant width and offset form:



Opcode map field used...	For operand type...	Unit
<i>src2</i>	uint	.S1, .S2
<i>csta</i>	ucst5	
<i>cstb</i>	ucst5	
<i>dst</i>	uint	

Opcode

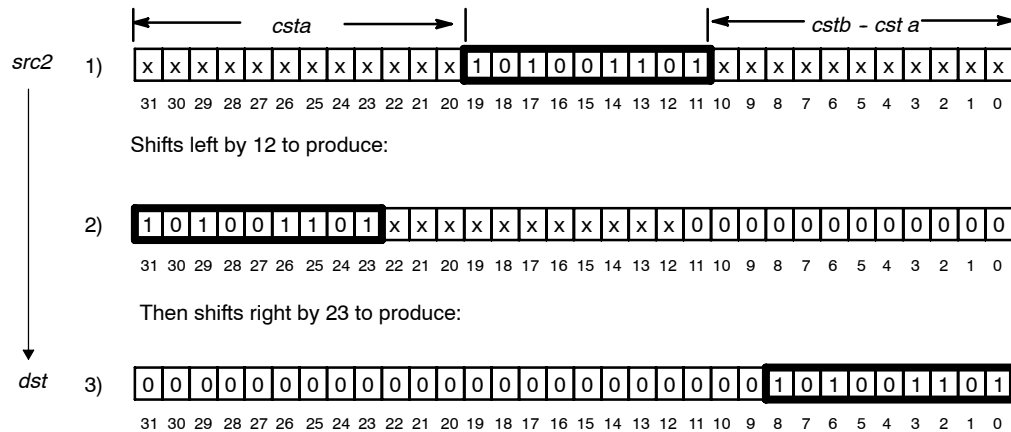
Register width and offset form:



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xuint	.S1, .S2
<i>src1</i>	uint	
<i>dst</i>	uint	

Description

The field in *src2*, specified by *csta* and *cstb*, is extracted and zero extended to 32 bits. The extract is performed by a shift left followed by an unsigned shift right. *csta* and *cstb* are the amounts to shift left and shift right, respectively. This can be thought of in terms of the LSB and MSB of the field to be extracted. Then $csta = 31 - \text{MSB of the field}$ and $cstb = csta + \text{LSB of the field}$. The shift left and shift right amounts may also be specified as the ten LSBs of the *src1* register with *cstb* being bits 0-4 and *csta* bits 5-9. In the example below, *csta* is 12 and *cstb* is $11 + 12 = 23$. Only the ten LSBs are valid for the register version of the instruction. If any of the 22 MSBs are non-zero, the result is invalid.



Execution

If the constant form is used:

```
if (cond)  src2 extu csta, cstb → dst
else nop
```

If the register width and offset form is used:

```
if (cond)  src2 extu src19..5, src14..0 → dst
else nop
```

Pipeline


Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type Single-cycle


Delay Slots 0

See Also **EXT**

Example 1 `EXTU .S1 A1,10,19,A2`



Example 2 `EXTU .S1 A1,A2,A3`




IDLE *Multicycle NOP With No Termination Until Interrupt*

IDLE *Multicycle NOP With No Termination Until Interrupt*

Syntax **IDLE**
.unit = none

Opcode



Description Performs an infinite multicycle **NOP** that terminates upon servicing an interrupt, or a branch occurs due to an **IDLE** instruction being in the delay slots of a branch.

Instruction Type NOP

Delay Slots 0

LDB(U)

Load Byte From Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Syntax

Register Offset

Unsigned Constant Offset

LDB (.unit) *+baseR[offsetR], dst

LDB (.unit) *+baseR[ucst5], dst

or

or

LDBU (.unit) *+baseR[offsetR], dst

LDBU (.unit) *+baseR[ucst5], dst

.unit = .D1 or .D2

Opcode

31	29	28	27	23	22	18	17	13	12	9	8	7	6	4	3	2	1	0			
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>baseR</i>		<i>offsetR/ucst5</i>			<i>mode</i>		0	<i>y</i>	<i>op</i>		0	1	<i>s</i>	<i>p</i>
3			1	5			5		5			4		1	3	1		1	1		1

Description

Loads a byte from memory to a general-purpose register (*dst*). Table 3-12 summarizes the data types supported by loads. Table 3-6 (page 3-19) describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*). If an offset is not given, the assembler assigns an offset of zero.

offsetR and *baseR* must be in the same register file and on the same side as the .D unit used. The *y* bit in the opcode determines the .D unit and register file used: *y* = 0 selects the .D1 unit and *baseR* and *offsetR* from the A register file, and *y* = 1 selects the .D2 unit and *baseR* and *offsetR* from the B register file.

offsetR/ucst5 is scaled by a left-shift of 0 bits. After scaling, *offsetR/ucst5* is added to or subtracted from *baseR*. For the preincrement, predecrement, positive offset, and negative offset address generator options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is the address to be accessed in memory.

Table 3-12. Data Types Supported by LDB(U) Instruction

Mnemonic	<i>op</i> Field	Load Data Type	Size	Left Shift of Offset
LDB	0 1 0	Load byte	8	0 bits
LDBU	0 0 1	Load byte unsigned	8	0 bits

LDB(U) *Load Byte From Memory With a 5-Bit Unsigned Constant Offset or Register Offset*

The addressing arithmetic that performs the additions and subtractions defaults to linear mode. However, for A4–A7 and for B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10).

For **LDB(U)**, the values are loaded into the 8 LSBs of *dst*. For **LDB**, the upper 24 bits of *dst* values are sign-extended; for **LDBU**, the upper 24 bits of *dst* are zero-filled. The *s* bit determines which file *dst* will be loaded into: *s* = 0 indicates *dst* will be loaded in the A register file and *s* = 1 indicates *dst* will be loaded in the B register file. The *r* bit should be cleared to 0.

Increments and decrements default to 1 and offsets default to 0 when no bracketed register or constant is specified. Loads that do no modification to the *baseR* can use the syntax *R. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 0. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Execution if (cond) mem → *dst*
 else nop

Pipeline	Pipeline Stage					
	E1	E2	E3	E4	E5	
Read	<i>baseR</i> <i>offsetR</i>					
Written	<i>baseR</i>				<i>dst</i>	
Unit in use	.D					

Instruction Type Load

Delay Slots 4 for loaded value
 0 for address modification from pre/post increment/decrement
 For more information on delay slots for a load, see Chapter 4.

See Also **LDH, LDW**

Example

LDB .D1 *-A5[4],A7

Before LDB

A5 0000 0204h
 A7 1951 1970h
 AMR 0000 0000h
 mem 200h E1h

1 cycle after LDB

A5 0000 0204h
 A7 1951 1970h
 AMR 0000 0000h
 mem 200h E1h

5 cycles after LDB

A5 0000 0204h
 A7 FFFF FFE1h
 AMR 0000 0000h
 mem 200h E1h

LDB(U) *Load Byte From Memory With a 15-Bit Unsigned Constant Offset*


LDB(U)

Load Byte From Memory With a 15-Bit Unsigned Constant Offset

Syntax

LDB (.unit) *+B14/B15[*ucst15*], *dst*
 or
LDBU (.unit) *+B14/B15[*ucst15*], *dst*
 .unit = .D2

Opcode



Description

Loads a byte from memory to a general-purpose register (*dst*). Table 3–13 summarizes the data types supported by loads. The memory address is formed from a base address register B14 (*y* = 0) or B15 (*y* = 1) and an offset, which is a 15-bit unsigned constant (*ucst15*). The assembler selects this format only when the constant is larger than five bits in magnitude. This instruction operates only on the .D2 unit.

The offset, *ucst15*, is scaled by a left shift of 0 bits. After scaling, *ucst15* is added to *baseR*. Subtraction is not supported. The result of the calculation is the address sent to memory. The addressing arithmetic is always performed in linear mode.

For **LDB(U)**, the values are loaded into the 8 LSBs of *dst*. For **LDB**, the upper 24 bits of *dst* values are sign-extended; for **LDBU**, the upper 24 bits of *dst* are zero-filled. The *s* bit determines which file *dst* will be loaded into: *s* = 0 indicates *dst* will be loaded in the A register file and *s* = 1 indicates *dst* will be loaded in the B register file.

Square brackets, [], indicate that the *ucst15* offset is left-shifted by 0. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Table 3–13. *Data Types Supported by LDB(U) Instruction (15-Bit Offset)*

Mnemonic	<i>op</i> Field	Load Data Type	Size	Left Shift of Offset
LDB	0 1 0	Load byte	8	0 bits
LDBU	0 0 1	Load byte unsigned	8	0 bits

Execution if (cond) mem → *dst*
 else nop

Note:

This instruction executes only on the B side (.D2).

Pipeline

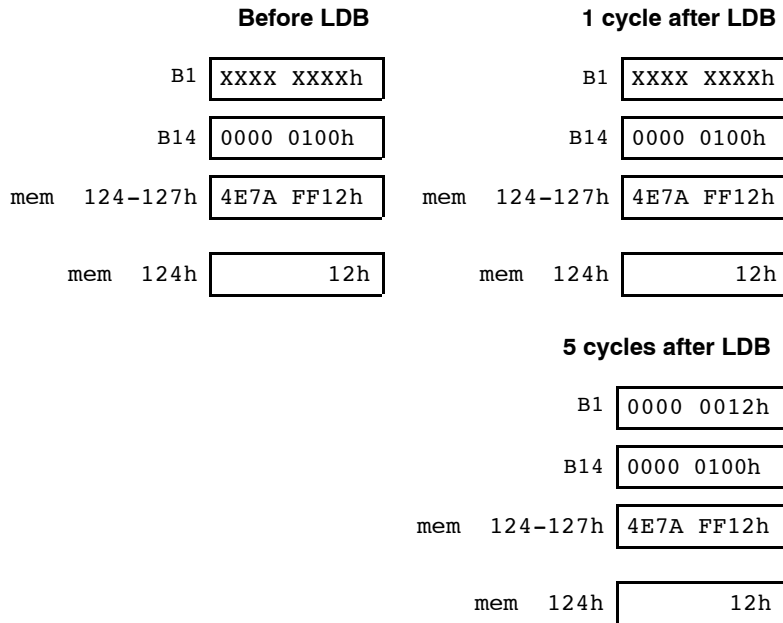
Pipeline Stage	E1	E2	E3	E4	E5
Read	B14 / B15				
Written	<i>dst</i>				
Unit in use	.D2				

Instruction Type Load

Delay Slots 4

See Also **LDH, LDW**

Example LDB .D2 **/B14[36],B1*



LDH(U) Load Halfword From Memory With a 5-Bit Unsigned Constant Offset or Register Offset


LDH(U) Load Halfword From Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Syntax

Register Offset	Unsigned Constant Offset
LDH (.unit) <i>*+baseR[offsetR], dst</i>	LDH (.unit) <i>*+baseR[ucst5], dst</i>
or	or
LDHU (.unit) <i>*+baseR[offsetR], dst</i>	LDHU (.unit) <i>*+baseR[ucst5], dst</i>

.unit = .D1 or .D2

Opcode



Description

Loads a halfword from memory to a general-purpose register (*dst*). Table 3–14 summarizes the data types supported by halfword loads. Table 3–6 (page 3-19) describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*). If an offset is not given, the assembler assigns an offset of zero.

offsetR and *baseR* must be in the same register file and on the same side as the .D unit used. The *y* bit in the opcode determines the .D unit and register file used: *y* = 0 selects the .D1 unit and *baseR* and *offsetR* from the A register file, and *y* = 1 selects the .D2 unit and *baseR* and *offsetR* from the B register file.

offsetR/ucst5 is scaled by a left-shift of 1 bit. After scaling, *offsetR/ucst5* is added to or subtracted from *baseR*. For the preincrement, predecrement, positive offset, and negative offset address generator options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is the address to be accessed in memory.

Table 3–14. Data Types Supported by LDH(U) Instruction

Mnemonic	<i>op</i> Field	Load Data Type	Size	Left Shift of Offset
LDH	1 0 0	Load halfword	16	1 bit
LDHU	0 0 0	Load halfword unsigned	16	1 bit

The addressing arithmetic that performs the additions and subtractions defaults to linear mode. However, for A4–A7 and for B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10).

For **LDH(U)**, the values are loaded into the 16 LSBs of *dst*. For **LDH**, the upper 16 bits of *dst* are sign-extended; for **LDHU**, the upper 16 bits of *dst* are zero-filled. The *s* bit determines which file *dst* will be loaded into: *s* = 0 indicates *dst* will be loaded in the A register file and *s* = 1 indicates *dst* will be loaded in the B register file. The *r* bit should be cleared to 0.

Increments and decrements default to 1 and offsets default to 0 when no bracketed register or constant is specified. Loads that do no modification to the *baseR* can use the syntax *R. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 1. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Halfword addresses must be aligned on halfword (LSB is 0) boundaries.

Execution if (cond) mem → *dst*
else nop

Pipeline Stage	E1	E2	E3	E4	E5
Read	<i>baseR</i> <i>offsetR</i>				
Written	<i>baseR</i>				<i>dst</i>
Unit in use	.D				

Instruction Type Load

Delay Slots 4 for loaded value
0 for address modification from pre/post increment/decrement
For more information on delay slots for a load, see Chapter 4.

See Also **LDB, LDW**

LDH(U) *Load Halfword From Memory With a 5-Bit Unsigned Constant Offset or Register Offset*

Example

LDH .D1 *++A4[A1],A8

Before LDH

A1	0000 0002h
A4	0000 0020h
A8	1103 51FFh
AMR	0000 0000h

1 cycle after LDH

A1	0000 0002h
A4	0000 0024h
A8	1103 51FFh
AMR	0000 0000h

5 cycles after LDH

A1	0000 0002h
A4	0000 0024h
A8	FFFF A21Fh
AMR	0000 0000h

mem 24h A21Fh

mem 24h A21Fh

mem 24h A21Fh


LDH(U)

Load Halfword From Memory With a 15-Bit Unsigned Constant Offset

Syntax

LDH (.unit) *+B14/B15[*ucst15*], *dst*
 or
LDHU (.unit) *+B14/B15[*ucst15*], *dst*
 .unit = .D2

Opcode



Description

Loads a halfword from memory to a general-purpose register (*dst*). Table 3-15 summarizes the data types supported by loads. The memory address is formed from a base address register B14 (*y* = 0) or B15 (*y* = 1) and an offset, which is a 15-bit unsigned constant (*ucst15*). The assembler selects this format only when the constant is larger than five bits in magnitude. This instruction operates only on the .D2 unit.

The offset, *ucst15*, is scaled by a left shift of 1 bit. After scaling, *ucst15* is added to *baseR*. Subtraction is not supported. The result of the calculation is the address sent to memory. The addressing arithmetic is always performed in linear mode.

For **LDH(U)**, the values are loaded into the 16 LSBs of *dst*. For **LDH**, the upper 16 bits of *dst* are sign-extended; for **LDHU**, the upper 16 bits of *dst* are zero-filled. The *s* bit determines which file *dst* will be loaded into: *s* = 0 indicates *dst* will be loaded in the A register file and *s* = 1 indicates *dst* will be loaded in the B register file.

Square brackets, [], indicate that the *ucst15* offset is left-shifted by 1. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Halfword addresses must be aligned on halfword (LSB is 0) boundaries.

Table 3-15. Data Types Supported by LDH(U) Instruction (15-Bit Offset)

Mnemonic	<i>op</i> Field	Load Data Type	Size	Left Shift of Offset
LDH	1 0 0	Load halfword	16	1 bit
LDHU	0 0 0	Load halfword unsigned	16	1 bit

LDH(U) *Load Halfword From Memory With a 15-Bit Unsigned Constant Offset*

Execution if (cond) mem → *dst*
 else nop

Note:

This instruction executes only on the B side (.D2).

Pipeline

Pipeline Stage	E1	E2	E3	E4	E5
Read	B14 / B15				
Written					<i>dst</i>
Unit in use	.D2				

Instruction Type Load

Delay Slots 4

See Also **LDB, LDW**

LDW

Load Word From Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Syntax


Register Offset

Unsigned Constant Offset

LDW (.unit) **+baseR[offsetR], dst* **LDW** (.unit) **+baseR[ucst5], dst*

.unit = .D1 or .D2

Opcode



Description

Loads a word from memory to a general-purpose register (*dst*). Table 3-6 (page 3-19) describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*). If an offset is not given, the assembler assigns an offset of zero.

offsetR and *baseR* must be in the same register file and on the same side as the .D unit used. The *y* bit in the opcode determines the .D unit and register file used: *y* = 0 selects the .D1 unit and *baseR* and *offsetR* from the A register file, and *y* = 1 selects the .D2 unit and *baseR* and *offsetR* from the B register file.

offsetR/ucst5 is scaled by a left-shift of 2 bits. After scaling, *offsetR/ucst5* is added to or subtracted from *baseR*. For the preincrement, predecrement, positive offset, and negative offset address generator options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is the address to be accessed in memory.

The addressing arithmetic that performs the additions and subtractions defaults to linear mode. However, for A4-A7 and for B4-B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10).

For **LDW**, the entire 32 bits fills *dst*. *dst* can be in either register file, regardless of the .D unit or *baseR* or *offsetR* used. The *s* bit determines which file *dst* will be loaded into: *s* = 0 indicates *dst* will be loaded in the A register file and *s* = 1 indicates *dst* will be loaded in the B register file. The *r* bit should be cleared to 0.

LDW *Load Word From Memory With a 5-Bit Unsigned Constant Offset or Register Offset*

Increments and decrements default to 1 and offsets default to 0 when no bracketed register or constant is specified. Loads that do no modification to the *baseR* can use the syntax **R*. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 2. Parentheses, (), can be used to set a nonscaled, constant offset. For example, **LDW** (.unit) **+baseR* (12) *dst* represents an offset of 12 bytes; whereas, **LDW** (.unit) **+baseR* [12] *dst* represents an offset of 12 words, or 48 bytes. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Word addresses must be aligned on word (two LSBs are 0) boundaries.

Execution if (cond) mem → *dst*
 else nop

Pipeline Stage	E1	E2	E3	E4	E5
Read	<i>baseR</i> <i>offsetR</i>				
Written	<i>baseR</i>				<i>dst</i>
Unit in use	.D				

Instruction Type Load

Delay Slots 4 for loaded value
 0 for address modification from pre/post increment/decrement
 For more information on delay slots for a load, see Chapter 4.

See Also **LDB, LDH**

Example 1 LDW .D1 *A10,B1

Before LDW

B1 0000 0000h
 A10 0000 0100h
 mem 100h 21F3 1996h

1 cycle after LDW

B1 0000 0000h
 A10 0000 0100h
 mem 100h 21F3 1996h

5 cycles after LDW

B1 21F3 1996h
 A10 0000 0100h
 mem 100h 21F3 1996h

Example 2 LDW .D1 *A4++[1],A6

Before LDW

A4 0000 0100h
 A6 1234 4321h
 AMR 0000 0000h
 mem 100h 0798 F25Ah
 mem 104h 1970 19F3h

1 cycle after LDW

A4 0000 0104h
 A6 1234 4321h
 AMR 0000 0000h
 mem 100h 0798 F25Ah
 mem 104h 1970 19F3h

5 cycles after LDW

A4 0000 0104h
 A6 0798 F25Ah
 AMR 0000 0000h
 mem 100h 0798 F25Ah
 mem 104h 1970 19F3h

Example 3 LDW .D1 *++A4[1],A6

Before LDW

A4 0000 0100h
 A6 1234 5678h
 AMR 0000 0000h
 mem 104h 0217 6991h

1 cycle after LDW

A4 0000 0104h
 A6 1234 5678h
 AMR 0000 0000h
 mem 104h 0217 6991h

5 cycles after LDW

A4 0000 0104h
 A6 0217 6991h
 AMR 0000 0000h
 mem 104h 0217 6991h


LDW Load Word From Memory With a 15-Bit Unsigned Constant Offset

LDW Load Word From Memory With a 15-Bit Unsigned Constant Offset

Syntax LDW (.unit) *+B14/B15[*ucst15*], *dst*

.unit = .D2

Opcode



Description

Load a word from memory to a general-purpose register (*dst*). The memory address is formed from a base address register B14 ($y = 0$) or B15 ($y = 1$) and an offset, which is a 15-bit unsigned constant (*ucst15*). The assembler selects this format only when the constant is larger than five bits in magnitude. This instruction operates only on the .D2 unit.

The offset, *ucst15*, is scaled by a left shift of 2 bits. After scaling, *ucst15* is added to *baseR*. Subtraction is not supported. The result of the calculation is the address sent to memory. The addressing arithmetic is always performed in linear mode.

For **LDW**, the entire 32 bits fills *dst*. *dst* can be in either register file. The *s* bit determines which file *dst* will be loaded into: $s = 0$ indicates *dst* will be loaded in the A register file and $s = 1$ indicates *dst* will be loaded in the B register file.

Square brackets, [], indicate that the *ucst15* offset is left-shifted by 2. Parentheses, (), can be used to set a nonscaled, constant offset. For example, **LDW** (.unit) *+B14/B15(60), *dst* represents an offset of 60 bytes; whereas, **LDW** (.unit) *+B14/B15[60], *dst* represents an offset of 60 words, or 240 bytes. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Word addresses must be aligned on word (two LSBs are 0) boundaries.

Execution

if (cond) mem → *dst*
else nop

Note:

This instruction executes only on the B side (.D2).

Pipeline	Pipeline					
	Stage	E1	E2	E3	E4	E5
Read	B14 / B15					
Written	<i>dst</i>					
Unit in use	.D2					

Instruction Type Load

Delay Slots 4


See Also **LDB, LDH**

LMBD Leftmost Bit Detection

LMBD Leftmost Bit Detection

Syntax **LMBD** (.unit) *src1*, *src2*, *dst*
.unit = .L1 or .L2

Opcode




Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	110 1011
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	cst5	.L1, .L2	110 1010
<i>src2</i>	xuint		
<i>dst</i>	uint		


Description The LSB of the *src1* operand determines whether to search for a leftmost 1 or 0 in *src2*. The number of bits to the left of the first 1 or 0 when searching for a 1 or 0, respectively, is placed in *dst*.

The following diagram illustrates the operation of **LMBD** for several cases.


When searching for 0 in *src2*, **LMBD** returns 0:



When searching for 1 in *src2*, **LMBD** returns 4:



When searching for 0 in *src2*, **LMBD** returns 32:



Execution

```

if (cond) {
    if (src10 == 0) lmb0(src2) → dst
    if (src10 == 1) lmb1(src2) → dst
}
else nop
    
```


Pipeline

Pipeline Stage	E1
Read	src1, src2
Written	dst
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

Example LMBD .L1 A1,A2,A3



MPY *Multiply Signed 16 LSB x Signed 16 LSB*

MPY *Multiply Signed 16 LSB x Signed 16 LSB*

Syntax **MPY** (.unit) *src1, src2, dst*

.unit = .M1 or .M2

Opcode

31	29	28	27	23	22	18	17	13	12	11	7	6	5	4	3	2	1	0	
<i>creg</i>			<i>z</i>	<i>dst</i>		<i>src2</i>		<i>src1/cst</i>		<i>x</i>	<i>op</i>		0	0	0	0	0	<i>s</i>	<i>p</i>
3			1	5		5		5		1	5							1	1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	slsb16	.M1, .M2	11001
<i>src2</i>	xslsb16		
<i>dst</i>	sint		
<i>src1</i>	scst5	.M1, .M2	11000
<i>src2</i>	xslsb16		
<i>dst</i>	sint		

Description The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are signed by default.

Execution if (cond) $\text{lsb16}(\text{src1}) \times \text{lsb16}(\text{src2}) \rightarrow \text{dst}$
else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1, src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPYU, MPYSU, MPYUS, SMPY**

Example 1

MPY .M1 A1,A2,A3

Before instruction

A1 0000 0123h 291[†]
 A2 01E0 FA81h -1407[†]
 A3 xxxx xxxxh

[†] Signed 16-LSB integer

2 cycles after instruction

A1 0000 0123h
 A2 01E0 FA81h
 A3 FFF9 C0A3 -409437

Example 2

MPY .M1 13,A1,A2

Before instruction

A1 3497 FFF3h -13[†]
 A2 xxxx xxxxh

[†] Signed 16-LSB integer

2 cycles after instruction

A1 3497 FFF3h
 A2 FFFF FF57h -163

MPYH *Multiply Signed 16 MSB x Signed 16 MSB*

MPYH


Multiply Signed 16 MSB x Signed 16 MSB

Syntax

MPYH (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	smsb16	.M1, .M2
<i>src2</i>	xsmsb16	
<i>dst</i>	sint	

Description

The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are signed by default.

Execution

if (cond) $\text{msb16}(\text{src1}) \times \text{msb16}(\text{src2}) \rightarrow \text{dst}$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type

Multiply (16 × 16)

Delay Slots

1

See Also

MPYHU, MPYHSU, MPYHUS, SMPYH

Example

MPYH .M1 A1, A2, A3

	Before instruction		2 cycles after instruction
A1	0023 0000h	35 [†]	0023 0000h
A2	FFA7 1234h	-89 [†]	FFA7 1234h
A3	xxxx xxxxh		FFFF F3D5h

[†] Signed 16-MSB integer


MPYHL

Multiply Signed 16 MSB x Signed 16 LSB

Syntax **MPYHL** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	smsb16	.M1, .M2
<i>src2</i>	xslsb16	
<i>dst</i>	sint	

Description The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are signed by default.

Execution if (cond) msb16(*src1*) × lsb16(*src2*) → *dst*
else nop

Pipeline


Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also MPYHLU, MPYHSLU, MPYHULS, SMPYHL

Example MPYHL .M1 A1,A2,A3



[†] Signed 16-MSB integer

[‡] Signed 16-LSB integer

MPYHLU *Multiply Unsigned 16 MSB x Unsigned 16 LSB*

MPYHLU

Multiply Unsigned 16 MSB x Unsigned 16 LSB

Syntax **MPYHLU** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode

31	29	28	27	23	22	18	17	13	12	11	7	6	5	4	3	2	1	0								
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1</i>			x	0	1	1	1	1	0	0	0	0	0	0	s	p
3			1	5			5			5			1											1	1	

Opcode map field used...	For operand type...	Unit
<i>src1</i>	umsb16	.M1, .M2
<i>src2</i>	xulsb16	
<i>dst</i>	uint	

Description The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are unsigned by default.

Execution if (cond) $\text{msb16}(\text{src1}) \times \text{lsb16}(\text{src2}) \rightarrow \text{dst}$
else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPYHL**, **MPYHSLU**, **MPYHULS**


MPYHSLU

Multiply Signed 16 MSB x Unsigned 16 LSB

Syntax **MPYHSLU** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	smsb16	.M1, .M2
<i>src2</i>	xulsb16	
<i>dst</i>	sint	

Description The signed operand *src1* is multiplied by the unsigned operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $msb16(src1) \times lsb16(src2) \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written	<i>dst</i>	
Unit in use	.M	

Instruction Type Multiply (16 x 16)

Delay Slots 1

See Also **MPYHL, MPYHLU, MPYHULS**

MPYHSU *Multiply Signed 16 MSB x Unsigned 16 MSB*


MPYHSU

Multiply Signed 16 MSB x Unsigned 16 MSB

Syntax **MPYHSU** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	smsb16	.M1, .M2
<i>src2</i>	xmsb16	
<i>dst</i>	sint	

Description The signed operand *src1* is multiplied by the unsigned operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $msb16(src1) \times msb16(src2) \rightarrow dst$
 else nop

Pipeline


Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPYH, MPYHU, MPYHUS**

Example MPYHSU .M1 A1, A2, A3



[†] Signed 16-MSB integer
[‡] Unsigned 16-MSB integer

MPYHU


Multiply Unsigned 16 MSB x Unsigned 16 MSB

Syntax

MPYHU (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	umsb16	.M1, .M2
<i>src2</i>	xumsb16	
<i>dst</i>	uint	

Description

The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are unsigned by default.

Execution

if (cond) msb16(*src1*) × msb16(*src2*) → *dst*
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type

Multiply (16 × 16)

Delay Slots


1

See Also

MPYH, MPYHSU, MPYHUS

Example

MPYHU .M1 A1, A2, A3



[‡] Unsigned 16-MSB integer

[§] Unsigned 32-bit integer

MPYHULS *Multiply Unsigned 16 MSB x Signed 16 LSB*


MPYHULS

Multiply Unsigned 16 MSB x Signed 16 LSB

Syntax **MPYHULS** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	umsb16	.M1, .M2
<i>src2</i>	xslsb16	
<i>dst</i>	sint	

Description The unsigned operand *src1* is multiplied by the signed operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $msb16(src1) \times lsb16(src2) \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 x 16)

Delay Slots 1

See Also **MPYHL, MPYHLU, MPYHSLU**


MPYHUS

Multiply Unsigned 16 MSB x Signed 16 MSB

Syntax **MPYHUS** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	umsb16	.M1, .M2
<i>src2</i>	xmsb16	
<i>dst</i>	sint	

Description The unsigned operand *src1* is multiplied by the signed operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $msb16(src1) \times msb16(src2) \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written	<i>dst</i>	
Unit in use	.M	

Instruction Type Multiply (16 x 16)

Delay Slots 1

See Also **MPYH, MPYHU, MPYHSU**

MPYLH *Multiply Signed 16 LSB x Signed 16 MSB*

MPYLH

Multiply Signed 16 LSB x Signed 16 MSB

Syntax

MPYLH (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode

31	29	28	27	23	22	18	17	13	12	11	7	6	5	4	3	2	1	0								
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1</i>			<i>x</i>	1	0	0	0	1	0	0	0	0	0	0	<i>s</i>	<i>p</i>
3			1	5			5			5			1	1 1												

Opcode map field used...	For operand type...	Unit
<i>src1</i>	slsb16	.M1, .M2
<i>src2</i>	xmsb16	
<i>dst</i>	sint	

Description

The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are signed by default.

Execution

if (cond) $\text{lsb16}(\text{src1}) \times \text{msb16}(\text{src2}) \rightarrow \text{dst}$
else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type

Multiply (16 × 16)

Delay Slots

1

See Also

MPYLHU, MPYLSHU, MPYLUHS, SMPYLH

Example

MPYLH .M1 A1,A2,A3

	Before instruction		2 cycles after instruction
A1	0900 000Eh 14 [†]	A1	0900 000Eh
A2	0029 00A7h 41 [‡]	A2	0029 00A7h
A3	xxxx xxxxh	A3	0000 023Eh 574


[†] Signed 16-LSB integer

[‡] Signed 16-MSB integer

MPYLHU *Multiply Unsigned 16 LSB × Unsigned 16 MSB*

Syntax **MPYLHU** (.unit) *src1*, *src2*, *dst*
 .unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	ulsb16	.M1, .M2
<i>src2</i>	xumsb16	
<i>dst</i>	uint	

Description The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are unsigned by default.

Execution if (cond) $\text{lsb16}(\text{src1}) \times \text{msb16}(\text{src2}) \rightarrow \text{dst}$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPYLH**, **MPYLSHU**, **MPYLUHS**


MPYLSHU *Multiply Signed 16 LSB x Unsigned 16 MSB*

MPYLSHU

Multiply Signed 16 LSB x Unsigned 16 MSB

Syntax **MPYLSHU** (.unit) *src1*, *src2*, *dst*
 .unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	s1sb16	.M1, .M2
<i>src2</i>	xum1sb16	
<i>dst</i>	sint	

Description The signed operand *src1* is multiplied by the unsigned operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $lsb16(src1) \times msb16(src2) \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written	<i>dst</i>	
Unit in use	.M	

Instruction Type Multiply (16 x 16)

Delay Slots 1


See Also **MPYLH, MPYLHU, MPYLUHS**

MPYLUHS

Multiply Unsigned 16 LSB x Signed 16 MSB

Syntax **MPYLUHS** (.unit) *src1*, *src2*, *dst*
 .unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	ulsb16	.M1, .M2
<i>src2</i>	xmsb16	
<i>dst</i>	sint	

Description The unsigned operand *src1* is multiplied by the signed operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $lsb16(src1) \times msb16(src2) \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written	<i>dst</i>	
Unit in use	.M	

Instruction Type Multiply (16 x 16)

Delay Slots 1

See Also **MPYLH, MPYLHU, MPYLSHU**


MPYSU *Multiply Signed 16 LSB x Unsigned 16 LSB*

MPYSU *Multiply Signed 16 LSB x Unsigned 16 LSB*

Syntax **MPYSU** (.unit) *src1, src2, dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	slsb16 xulsb16 sint	.M1, .M2	11011
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xulsb16 sint	.M1, .M2	11110

Description The signed operand *src1* is multiplied by the unsigned operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution if (cond) $lsb16(src1) \times lsb16(src2) \rightarrow dst$
else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1, src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPY, MPYU, MPYUS**

Example

MPYSU .M1 13,A1,A2

Before instruction

A1 3497 FFF3h 65523[‡]

A2 xxxx xxxxh

2 cycles after instruction

A1 3497 FFF3h

A2 000C FF57h 851779

[‡] Unsigned 16-LSB integer


MPYU *Multiply Unsigned 16 LSB x Unsigned 16 LSB*

MPYU

Multiply Unsigned 16 LSB x Unsigned 16 LSB

Syntax **MPYU** (.unit) *src1*, *src2*, *dst*
 .unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	ulsb16	.M1, .M2
<i>src2</i>	xulsb16	
<i>dst</i>	uint	

Description The *src1* operand is multiplied by the *src2* operand. The result is placed in *dst*. The source operands are unsigned by default.

Execution if (cond) $lsb16(src1) \times lsb16(src2) \rightarrow dst$
 else nop

Pipeline


Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Multiply (16 × 16)

Delay Slots 1

See Also **MPY**, **MPYSU**, **MPYUS**

Example `MPYU .M1 A1,A2,A3`



[‡] Unsigned 16-LSB integer

MPYUS


Multiply Unsigned 16 LSB x Signed 16 LSB

Syntax

MPYUS (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	ulsb16	.M1, .M2
<i>src2</i>	xslsb16	
<i>dst</i>	sint	

Description

The unsigned operand *src1* is multiplied by the signed operand *src2*. The result is placed in *dst*. The **S** is needed in the mnemonic to specify a signed operand when both signed and unsigned operands are used.

Execution

if (cond) lsb16(*src1*) × lsb16(*src2*) → *dst*
 else nop

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type

Multiply (16 × 16)

Delay Slots


1

See Also

MPY, MPYU, MPYSU

Example

MPYUS .M1 A1,A2,A3




[†] Signed 16-LSB integer
[‡] Unsigned 16-LSB integer

MV *Move From Register to Register*

MV *Move From Register to Register*


Syntax **MV** (.unit) *src2*, *dst*
 .unit = .L1, .L2, .S1, .S2, .D1, .D2

Opcode .L unit




Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	xsint	.L1, .L2	000 0010
<i>dst</i>	sint		
<i>src2</i>	slong	.L1, .L2	010 0000
<i>dst</i>	slong		

Opcode .S unit



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xsint	.S1, .S2
<i>dst</i>	sint	

Opcode .D unit



Opcode map field used...	For operand type...	Unit
<i>src2</i>	sint	.D1, .D2
<i>dst</i>	sint	

Description	The MV pseudo-operation moves a value from one register to another. The assembler uses the ADD (.unit) 0, <i>src2</i> , <i>dst</i> operation to perform this task.
Execution	if (cond) 0 + <i>src2</i> → <i>dst</i> else nop
Instruction Type	Single-cycle
Delay Slots	0

MVC Move Between Control File and Register File


MVC

Move Between Control File and Register File

Syntax MVC (.unit) *src2*, *dst*

.unit = .S2

Opcode



Operands when moving from the control file to the register file:

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	uint	.S2	00 1111
<i>dst</i>	uint		

Description

The *src2* register is moved from the control register file to the register file. Valid values for *src2* are any register listed in the control register file.

Register addresses for accessing the control registers are in Table 3-16 (page 3-112).

Operands when moving from the register file to the control file:

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	xuint	.S2	00 1110
<i>dst</i>	uint		

Description

The *src2* register is moved from the register file to the control register file. Valid values for *src2* are any register listed in the control register file.

Register addresses for accessing the control registers are in Table 3-16 (page 3-112).

Execution if (cond) *src2* → *dst*
 else nop

Note:

The **MVC** instruction executes only on the B side (.S2).

Refer to the individual control register descriptions for specific behaviors and restrictions in accesses via the **MVC** instruction.

Pipeline

Pipeline Stage	E1
Read	<i>src2</i>
Written	<i>dst</i>
Unit in use	.S2

Instruction Type Single-cycle

Any write to the ISR or ICR (by the **MVC** instruction) effectively has one delay slot because the results cannot be read (by the **MVC** instruction) in the IFR until two cycles after the write to the ISR or ICR.

Delay Slots 0

Example MVC .S2 B1,AMR

	Before instruction	1 cycle after instruction
B1	F009 0001h	F009 0001h
AMR	0000 0000h	0009 0001h

Note:

The six MSBs of the AMR are reserved and therefore are not written to.


Table 3-16. Register Addresses for Accessing the Control Registers

Acronym	Register Name	Address	Read/ Write
AMR	Addressing mode register	00000	R, W
CSR	Control status register	00001	R, W
ICR	Interrupt clear register	00011	W
IER	Interrupt enable register	00100	R, W
IFR	Interrupt flag register	00010	R
IRP	Interrupt return pointer	00110	R, W
ISR	Interrupt set register	00010	W
ISTP	Interrupt service table pointer	00101	R, W
NRP	Nonmaskable interrupt return pointer	00111	R, W
PCE1	Program counter, E1 phase	10000	R

Legend: R = Readable by the **MVC** instruction; W = Writeable by the **MVC** instruction

MVK*Move Signed Constant Into Register and Sign Extend***Syntax****MVK** (.unit) *cst*, *dst*

.unit = .S1 or .S2

Opcode

Opcode map field used...	For operand type...	Unit
<i>cst16</i>	scst16	.S1, .S2
<i>dst</i>	sint	

Description

The 16-bit signed constant, *cst*, is sign extended and placed in *dst*.

In most cases, the C6000 assembler and linker issue a warning or an error when a constant is outside the range supported by the instruction. In the case of **MVK** .S, a warning is issued whenever the constant is outside the signed 16-bit range, -32768 to 32767 (or FFFF 8000h to 0000 7FFFh).

For example:

```
MVK .S1 0x00008000X, A0
```

will generate a warning; whereas:

```
MVK .S1 0xFFFF8000, A0
```

will not generate a warning.

Execution

if (cond) scst → *dst*
else nop

Pipeline

Pipeline Stage	E1
Read	
Written	<i>dst</i>
Unit in use	.S

MVK *Move Signed Constant Into Register and Sign Extend*

Instruction Type Single cycle

Delay Slots 0

See Also **MVKH, MVKL, MVKLH**

Example 1 MVK .S2 -5, B8

Before instruction

B8

xxxx xxxxh

1 cycle after instruction

B8

FFFF FFFBh

Example 2 MVK .S2 14, B8

Before instruction

B8

xxxx xxxxh

1 cycle after instruction

B8

0000 000Eh

MVKH/MVKLH*Move 16-Bit Constant Into Upper Bits of Register***Syntax**

MVKH (.unit) *cst*, *dst*
 or
MVKLH (.unit) *cst*, *dst*
 .unit = .S1 or .S2

Opcode

31	29	28	27	23	22	7	6	5	4	3	2	1	0				
<i>creg</i>			<i>z</i>	<i>dst</i>		<i>cst16</i>					1	1	0	1	0	<i>s</i>	<i>p</i>
3			1	5		16					1 1						

Opcode map field used...	For operand type...	Unit
<i>cst16</i>	uscst16	.S1, .S2
<i>dst</i>	sint	

Description

The 16-bit constant, *cst16*, is loaded into the upper 16 bits of *dst*. The 16 LSBs of *dst* are unchanged. For the **MVKH** instruction, the assembler encodes the 16 MSBs of a 32-bit constant into the *cst16* field of the opcode. For the **MVKLH** instruction, the assembler encodes the 16 LSBs of a constant into the *cst16* field of the opcode.

Execution

For the **MVKLH** instruction:

if (cond)((*cst*_{15..0}) << 16) or (*dst*_{15..0}) → *dst*
 else nop

For the **MVKH** instruction:

if (cond)((*cst*_{31..16}) << 16) or (*dst*_{15..0}) → *dst*
 else nop

Pipeline

Pipeline Stage	E1
Read	
Written	<i>dst</i>
Unit in use	.S

MVKH/MVKLH *Move 16-Bit Constant Into Upper Bits of Register*

Instruction Type Single-cycle

Delay Slots 0

Note:

Use the **MVK** instruction (page 3-113) to load 16-bit constants. The assembler generates a warning for any constant over 16 bits. To load 32-bit constants, such as 1234 5678h, use the following pair of instructions:

```
MVKL 0x12345678
MVKH 0x12345678
```

If you are loading the address of a label, use:

```
MVKL label
MVKH label
```

See Also **MVK, MVKL**

Example 1 `MVKH .S1 0A329123h,A1`

Before instruction

A1 0000 7634h

1 cycle after instruction

A1 0A32 7634h

Example 2 `MVKLH .S1 7A8h,A1`

Before instruction

A1 FFFF F25Ah

1 cycle after instruction


A1 07A8 F25Ah

MVKL

Move Signed Constant Into Register and Sign Extend

Syntax **MVKL** (.unit) *cst*, *dst*
 .unit = .S1 or .S2

Opcode



Opcode map field used...	For operand type...	Unit
<i>cst16</i>	scst16	.S1, .S2
<i>dst</i>	sint	

Description

The **MVKL** pseudo-operation sign extends the 16-bit constant, *cst16*, and places it in *dst*.

The **MVKL** instruction is equivalent to the **MVK** instruction (page 3-113), except that the **MVKL** instruction disables the constant range checking normally performed by the assembler/linker. This allows the **MVKL** instruction to be paired with the **MVKH** instruction (page 3-115) to generate 32-bit constants.

To load 32-bit constants, such as 1234 ABCDh, use the following pair of instructions:

```
MVKL .S1 0x0ABCD, A4
MVKLH .S1 0x1234, A4
```

This could also be used:

```
MVKL .S1 0x1234ABCD, A4
MVKH .S1 0x1234ABCD, A4
```

Use this to load the address of a label:

```
MVKL .S2 label, B5
MVKH .S2 label, B5
```

Execution

if (cond) *scst* → *dst*
 else nop

MVKL *Move Signed Constant Into Register and Sign Extend—Used with MVKH*


Pipeline	Pipeline Stage	E1
	Read	
	Written	<i>dst</i>
	Unit in use	<i>.S</i>

Instruction Type Single cycle


Delay Slots 0


See Also **MVK, MVKH, MVKLH**

Example 1 `MVKL .S1 5678h, A8`




Example 2 `MVKL .S1 0C678h, A8`



NEG *Negate***Syntax** **NEG** (.unit) *src2*, *dst**.unit* = .L1, .L2, .S1, .S2**Opcode** .S unit

Opcode map field used...	For operand type...	Unit
<i>src2</i>	xsint	.S1, .S2
<i>dst</i>	sint	

Opcode .L unit

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	xsint	.L1, .L2	000 0110
<i>dst</i>	sint		
<i>src2</i>	slong	.L1, .L2	010 0100
<i>dst</i>	slong		


Description The **NEG** pseudo-operation negates *src2* and places the result in *dst*. The assembler uses **SUB** (.unit) 0, *src2*, *dst* to perform this operation.**Execution** if (cond) 0 -s *src2* → *dst*
else nop**Instruction Type** Single-cycle**Delay Slots** 0

NOP *No Operation*

NOP *No Operation*

Syntax **NOP** [*count*]
 .unit = none

Opcode



Opcode map field used...	For operand type...	Unit
<i>src</i>	ucst4	none

Description *src* is encoded as *count* - 1. For *src* + 1 cycles, no operation is performed. The maximum value for *count* is 9. **NOP** with no operand is treated like **NOP 1** with *src* encoded as 0000.

A multicycle **NOP** will not finish if a branch is completed first. For example, if a branch is initiated on cycle *n* and a **NOP 5** instruction is initiated on cycle *n* + 3, the branch is complete on cycle *n* + 6 and the **NOP** is executed only from cycle *n* + 3 to cycle *n* + 5. A single-cycle **NOP** in parallel with other instructions does not affect operation.


Execution No operation for *count* cycles

Instruction Type **NOP**

Delay Slots 0


Example 1

```
NOP
MVK .S1 125h,A1
```



Example 2

```
MVK .S1 1,A1
MVKLH .S1 0,A1
NOP 5
ADD .L1 A1,A2,A1
```




Execution if (cond) norm(*src*) → *dst*
 else nop

Pipeline Stage	E1
Read	<i>src2</i>
Written	<i>dst</i>
Unit in use	.L


Instruction Type Single-cycle

Delay Slots 0

Example 1 NORM .L1 A1,A2



Example 2 NORM .L1 A1,A2



NOT Bitwise NOT

NOT

Bitwise NOT


Syntax

NOT (.unit) *src2*, *dst*

.unit = .L1, .L2, .S1, .S2

Opcode


.L unit



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xuint	.L1, .L2
<i>dst</i>	uint	

Opcode

.S unit



Opcode map field used...	For operand type...	Unit
<i>src2</i>	xuint	.S1, .S2
<i>dst</i>	uint	

Description

The **NOT** pseudo-operation performs a bitwise **NOT** on the *src2* operand and places the result in *dst*. The assembler uses **XOR** (.unit) -1, *src2*, *dst* to perform this operation.

Execution

if (cond) -1 XOR *src2* → *dst*
else nop

Instruction Type

Single-cycle

Delay Slots


0

OR*Bitwise OR***Syntax****OR** (.unit) *src1*, *src2*, *dst*

.unit = .L1, .L2, .S1, .S2

Opcode


.L unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	111 1111
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	111 1110
<i>src2</i>	xuint		
<i>dst</i>	uint		

Opcode

.S unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.S1, .S2	01 1011
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.S1, .S2	01 1010
<i>src2</i>	xuint		
<i>dst</i>	uint		

Description

Performs a bitwise **OR** operation between *src1* and *src2*. The result is placed in *dst*. The *scst5* operands are sign extended to 32 bits.

Execution

if (cond) *src1* OR *src2* → *dst*
 else nop

OR *Bitwise OR*

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.L or .S

Instruction Type

Single-cycle

Delay Slots

0

See Also

AND, XOR

Example 1

OR .S1 A3, A4, A5

Before instruction

A3	08A3 A49Fh
A4	00FF 375Ah
A5	xxxx xxxxh

1 cycle after instruction

A3	08A3 A49Fh
A4	00FF 375Ah
A5	08FF B7DFh

Example 2

OR .L2 -12, B2, B8

Before instruction


B2	0000 3A41h
B8	xxxx xxxxh

1 cycle after instruction

B2	0000 3A41h
B8	FFFF FFF5h

SADD*Add Two Signed Integers With Saturation***Syntax** **SADD** (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint sint	.L1, .L2	001 0011
<i>src1</i> <i>src2</i> <i>dst</i>	xsint slong slong	.L1, .L2	011 0001
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xsint sint	.L1, .L2	001 0010
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 slong slong	.L1, .L2	011 0000

Description

src1 is added to *src2* and saturated, if an overflow occurs according to the following rules:

- 1) If the *dst* is an int and $src1 + src2 > 2^{31} - 1$, then the result is $2^{31} - 1$.
- 2) If the *dst* is an int and $src1 + src2 < -2^{31}$, then the result is -2^{31} .
- 3) If the *dst* is a long and $src1 + src2 > 2^{39} - 1$, then the result is $2^{39} - 1$.
- 4) If the *dst* is a long and $src1 + src2 < -2^{39}$, then the result is -2^{39} .

The result is placed in *dst*. If a saturate occurs, the SAT bit in the control status register (CSR) is set one cycle after *dst* is written.

Execution

if (cond) $src1 +s src2 \rightarrow dst$
else nop

SADD *Add Two Signed Integers With Saturation*

Pipeline	Pipeline Stage	E1
	Read	<i>src1, src2</i>
	Written	<i>dst</i>
	Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, SSUB**

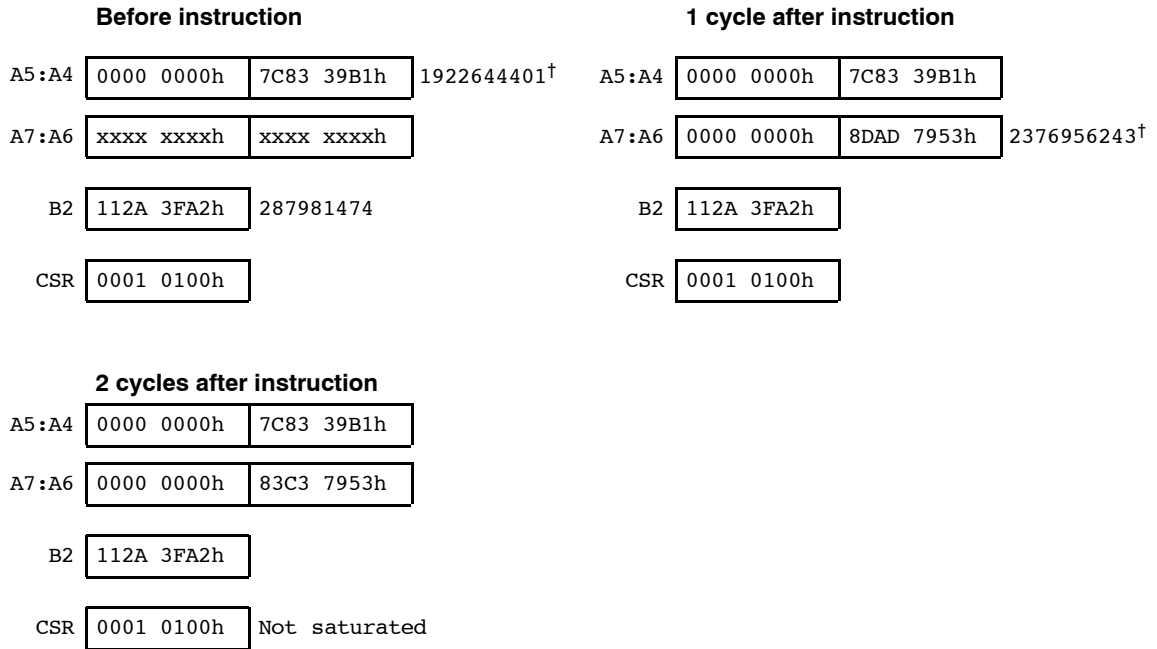
Example 1 SADD .L1 A1,A2,A3

	Before instruction		1 cycle after instruction		2 cycles after instruction
A1	5A2E 51A3h 1512984995	A1	5A2E 51A3h	A1	5A2E 51A3h
A2	012A 3FA2h 19546018	A2	012A 3FA2h	A2	012A 3FA2h
A3	xxxx xxxxh	A3	5B58 9145h 1532531013	A3	5B58 9145h
CSR	0001 0100h	CSR	0001 0100h	CSR	0001 0100h Not saturated

Example 2 SADD .L1 A1,A2,A3

	Before instruction		1 cycle after instruction		2 cycles after instruction
A1	4367 71F2h 1130852850	A1	4367 71F2h	A1	4367 71F2h
A2	5A2E 51A3h 1512984995	A2	5A2E 51A3h	A2	5A2E 51A3h
A3	xxxx xxxxh	A3	7FFF FFFFh 2147483647	A3	7FFF FFFFh
CSR	0001 0100h	CSR	0001 0100h	CSR	0001 0300h Saturated

Example 3 SADD .L1X B2,A5:A4,A7:A6



[†] Signed 40-bit (long) integer

SAT Saturate a 40-Bit Integer to a 32-Bit Integer

SAT


Saturate a 40-Bit Integer to a 32-Bit Integer

Syntax

SAT (.unit) *src2*, *dst*

.unit = .L1 or .L2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src2</i>	slong	.L1, .L2
<i>dst</i>	sint	

Description

A 40-bit *src2* value is converted to a 32-bit value. If the value in *src2* is greater than what can be represented in 32-bits, *src2* is saturated. The result is placed in *dst*. If a saturate occurs, the SAT bit in the control status register (CSR) is set one cycle after *dst* is written.

Execution

```

if (cond) {
    if (src2 > (231 - 1) )
        (231 - 1) → dst
    else if (src2 < -231)
        -231 → dst
    else src231..0 → dst
}
else nop
    
```

Pipeline

Pipeline Stage	E1
Read	<i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type

Single-cycle

Delay Slots

0

Example 1 SAT .L2 B1:B0,B5

	Before instruction	1 cycle after instruction	2 cycles after instruction
B1:B0	0000 001Fh 3413 539Ah	0000 001Fh 3413 539Ah	0000 001Fh 3413 539Ah
B5	xxxx xxxxh	7FFF FFFFh	7FFF FFFFh
CSR	0001 0100h	0001 0100h	0001 0300h Saturated

Example 2 SAT .L2 B1:B0,B5

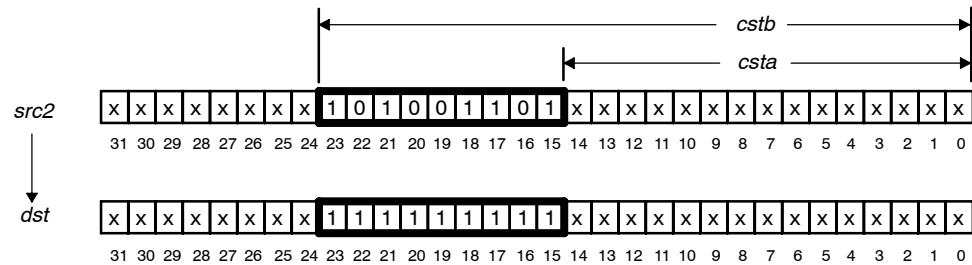
	Before instruction	1 cycle after instruction	2 cycles after instruction
B1:B0	0000 0000h A190 7321h	0000 0000h A190 7321h	0000 0000h A190 7321h
B5	xxxx xxxxh	7FFF FFFFh	7FFF FFFFh
CSR	0001 0100h	0001 0100h	0001 0300h Saturated

Example 3 SAT .L2 B1:B0,B5

	Before instruction	1 cycle after instruction	2 cycles after instruction
B1:B0	0000 00FFh A190 7321h	0000 00FFh A190 7321h	0000 00FFh A190 7321h
B5	xxxx xxxxh	A190 7321h	A190 7321h
CSR	0001 0100h	0001 0100h	0001 0100h Not saturated

Description

The field in *src2*, specified by *csta* and *cstb*, is set to all 1s. The *csta* and *cstb* operands may be specified as constants or in the ten LSBs of the *src1* register, with *cstb* being bits 0–4 and *csta* bits 5–9. *csta* signifies the bit location of the LSB of the field and *cstb* signifies the bit location of the MSB of the field. In other words, *csta* and *cstb* represent the beginning and ending bits, respectively, of the field to be set to all 1s. The LSB location of *src2* is 0 and the MSB location of *src2* is 31. In the example below, *csta* is 15 and *cstb* is 23. Only the ten LSBs are valid for the register version of the instruction. If any of the 22 MSBs are non-zero, the result is invalid.

**Execution**

If the constant form is used:

```
if (cond)  src2 SET csta, cstb → dst
else nop
```

If the register form is used:

```
if (cond)  src2 SET src19..5, src14..0 → dst
else nop
```

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type

Single-cycle

Delay Slots

0

See Also

CLR

Example 1

SET .S1 A0,7,21,A1

Before instruction

A0 4B13 4A1Eh
A1 xxxx xxxxxh

1 cycle after instruction

A0 4B13 4A1Eh
A1 4B3F FF9Eh

Example 2

SET .S2 B0,B1,B2

Before instruction

B0 9ED3 1A31h
B1 0000 C197h
B2 xxxx xxxxxh

1 cycle after instruction

B0 9ED3 1A31h
B1 0000 C197h
B2 9EFF FA31h

SHL Arithmetic Shift Left**Syntax** **SHL** (.unit) *src2*, *src1*, *dst*

.unit = .S1 or .S2

Opcode

31	29	28	27	23	22	18	17	13	12	11	6	5	4	3	2	1	0				
<i>creg</i>	<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>			<i>x</i>	<i>op</i>				1	0	0	0	<i>s</i>	<i>p</i>
3	1	5			5			5			1	6				1	1			1	1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i> <i>src1</i> <i>dst</i>	xsint uint sint	.S1, .S2	11 0011
<i>src2</i> <i>src1</i> <i>dst</i>	slong uint slong	.S1, .S2	11 0001
<i>src2</i> <i>src1</i> <i>dst</i>	xuint uint ulong	.S1, .S2	01 0011
<i>src2</i> <i>src1</i> <i>dst</i>	xsint ucst5 sint	.S1, .S2	11 0010
<i>src2</i> <i>src1</i> <i>dst</i>	slong ucst5 slong	.S1, .S2	11 0000
<i>src2</i> <i>src1</i> <i>dst</i>	xuint ucst5 ulong	.S1, .S2	01 0010

Description

The *src2* operand is shifted to the left by the *src1* operand. The result is placed in *dst*. When a register is used, the six LSBs specify the shift amount and valid values are 0–40. When an immediate is used, valid shift amounts are 0–31.

If $39 < src1 < 64$, *src2* is shifted to the left by 40. Only the six LSBs of *src1* are used by the shifter, so any bits set above bit 5 do not affect execution.

Execution

if (cond) $src2 \ll src1 \rightarrow dst$
else nop

SHL *Arithmetic Shift Left*

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type Single-cycle

Delay Slots 0

See Also SHR, SSHL

Example 1 SHL .S1 A0,4,A1

	Before instruction	1 cycle after instruction
A0	29E3 D31Ch	29E3 D31Ch
A1	xxxx xxxxh	9E3D 31C0h

Example 2 SHL .S2 B0,B1,B2

	Before instruction	1 cycle after instruction
B0	4197 51A5h	4197 51A5h
B1	0000 0009h	0000 0009h
B2	xxxx xxxxh	2EA3 4A00h

Example 3 SHL .S2 B1:B0,B2,B3:B2

	Before instruction	1 cycle after instruction
B1:B0	0000 0009h 4197 51A5h	0000 0009h 4197 51A5h
B2	0000 0022h	0000 0000h
B3:B2	xxxx xxxxh xxxx xxxxh	0000 0094h 0000 0000h

SHR*Arithmetic Shift Right***Syntax****SHR** (.unit) *src2*, *src1*, *dst*

.unit = .S1 or .S2

Opcode

31	29	28	27	23	22	18	17	13	12	11	6	5	4	3	2	1	0
<i>creg</i>	<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>			<i>x</i>	<i>op</i>				<i>s</i>	<i>p</i>
3	1	5			5			5			1	6				1	1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i> <i>src1</i> <i>dst</i>	xsint uint sint	.S1, .S2	11 0111
<i>src2</i> <i>src1</i> <i>dst</i>	slong uint slong	.S1, .S2	11 0101
<i>src2</i> <i>src1</i> <i>dst</i>	xsint ucst5 sint	.S1, .S2	11 0110
<i>src2</i> <i>src1</i> <i>dst</i>	slong ucst5 slong	.S1, .S2	11 0100

Description

The *src2* operand is shifted to the right by the *src1* operand. The sign-extended result is placed in *dst*. When a register is used, the six LSBs specify the shift amount and valid values are 0–40. When an immediate value is used, valid shift amounts are 0–31.

If $39 < src1 < 64$, *src2* is shifted to the right by 40. Only the six LSBs of *src1* are used by the shifter, so any bits set above bit 5 do not affect execution.

Execution

if (cond) *src2* >>s *src1* → *dst*
else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.S

SHR *Arithmetic Shift Right*

Instruction Type Single-cycle

Delay Slots 0

See Also SHL, SHRU

Example 1 SHR .S1 A0,8,A1

	Before instruction	1 cycle after instruction
A0	F123 63D1h	F123 63D1h
A1	xxxx xxxxh	FFF1 2363h

Example 2 SHR .S2 B0,B1,B2

	Before instruction	1 cycle after instruction
B0	1492 5A41h	1492 5A41h
B1	0000 0012h	0000 0012h
B2	xxxx xxxxh	0000 0524h

Example 3 SHR .S2 B1:B0,B2,B3:B2

	Before instruction	1 cycle after instruction
B1:B0	0000 0012h 1492 5A41h	0000 0012h 1492 5A41h
B2	0000 0019h	0000 090Ah
B3:B2	xxxx xxxxh xxxx xxxxh	0000 0000h 0000 090Ah

SHRU
Logical Shift Right
Syntax
SHRU (.unit) *src2*, *src1*, *dst*

.unit = .S1 or .S2

Opcode

31	29	28	27	23	22	18	17	13	12	11	6	5	4	3	2	1	0				
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>		<i>x</i>	<i>op</i>			1	0	0	0	<i>s</i>	<i>p</i>
3			1	5			5			5		1	6			1	1			1	1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	xuint	.S1, .S2	10 0111
<i>src1</i>	uint		
<i>dst</i>	uint		
<i>src2</i>	ulong	.S1, .S2	10 0101
<i>src1</i>	uint		
<i>dst</i>	ulong		
<i>src2</i>	xuint	.S1, .S2	10 0110
<i>src1</i>	ucst5		
<i>dst</i>	uint		
<i>src2</i>	ulong	.S1, .S2	10 0100
<i>src1</i>	ucst5		
<i>dst</i>	ulong		

Description

The *src2* operand is shifted to the right by the *src1* operand. The zero-extended result is placed in *dst*. When a register is used, the six LSBs specify the shift amount and valid values are 0–40. When an immediate value is used, valid shift amounts are 0–31.

If $39 < src1 < 64$, *src2* is shifted to the right by 40. Only the six LSBs of *src1* are used by the shifter, so any bits set above bit 5 do not affect execution.

Execution

if (cond) $src2 \gg z src1 \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.S

SHRU *Logical Shift Right*

Instruction Type Single-cycle

Delay Slots 0

See Also **SHL, SHR**

Example SHRU .S1 A0,8,A1

	Before instruction	1 cycle after instruction		
A0	<table border="1"><tr><td>F123 63D1h</td></tr></table>	F123 63D1h	<table border="1"><tr><td>F123 63D1h</td></tr></table>	F123 63D1h
F123 63D1h				
F123 63D1h				
A1	<table border="1"><tr><td>xxxx xxxxxh</td></tr></table>	xxxx xxxxxh	<table border="1"><tr><td>00F1 2363h</td></tr></table>	00F1 2363h
xxxx xxxxxh				
00F1 2363h				


SMPY

Multiply Signed 16 LSB × Signed 16 LSB With Left Shift and Saturation

Syntax **SMPY** (.unit) *src1*, *src2*, *dst*

.unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	s1sb16	.M1, .M2
<i>src2</i>	xs1sb16	
<i>dst</i>	sint	

Description The least significant 16 bits of *src1* operand is multiplied by the least significant 16 bits of the *src2* operand. The result is left shifted by 1 and placed in *dst*. If the left-shifted result is 8000 0000h, then the result is saturated to 7FFF FFFFh. If a saturate occurs, the SAT bit in CSR is set one cycle after *dst* is written. The source operands are signed by default.

Execution

```

if (cond) {
    if (((lsb16(src1) × lsb16(src2)) << 1) != 8000 0000h)
        ((lsb16(src1) × lsb16(src2)) << 1) → dst
    else
        7FFF FFFFh → dst
}
else nop
    
```

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Single-cycle (16 × 16)

Delay Slots 1

See Also **MPY, SMPYH, SMPYHL, SMPYLH**

SMPY *Multiply Signed 16 LSB x Signed 16 LSB With Left Shift and Saturation*

Example

SMPY .M1 A1,A2,A3

	Before instruction		2 cycle after instruction		
A1	<table border="1"><tr><td>0000 0123h</td></tr></table> 291 [‡]	0000 0123h		A1 <table border="1"><tr><td>0000 0123h</td></tr></table>	0000 0123h
0000 0123h					
0000 0123h					
A2	<table border="1"><tr><td>01E0 FA81h</td></tr></table> -1407 [‡]	01E0 FA81h		A2 <table border="1"><tr><td>01E0 FA81h</td></tr></table>	01E0 FA81h
01E0 FA81h					
01E0 FA81h					
A3	<table border="1"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		A3 <table border="1"><tr><td>FFF3 8146h</td></tr></table> -818874	FFF3 8146h
xxxx xxxxh					
FFF3 8146h					
CSR	<table border="1"><tr><td>0001 0100h</td></tr></table>	0001 0100h		CSR <table border="1"><tr><td>0001 0100h</td></tr></table> Not saturated	0001 0100h
0001 0100h					
0001 0100h					


[‡] Signed 16-LSB integer

SMPYHL *Multiply Signed 16 MSB x Signed 16 LSB With Left Shift and Saturation*

SMPYHL *Multiply Signed 16 MSB x Signed 16 LSB With Left Shift and Saturation*

Syntax **SMPYHL** (.unit) *src1*, *src2*, *dst*
 .unit = .M1 or .M2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	smsb16	.M1, .M2
<i>src2</i>	xslsb16	
<i>dst</i>	sint	

Description The most significant 16 bits of the *src1* operand is multiplied by the least significant bits of the *src2* operand. The result is left shifted by 1 and placed in *dst*. If the left-shifted result is 8000 0000h, then the result is saturated to 7FFF FFFFh. If a saturation occurs, the SAT bit in CSR is set one cycle after *dst* is written.

Execution

```

if (cond) {
    if (((msb16(src1) × lsb16(src2)) << 1) != 8000 0000h)
        ((msb16(src1) × lsb16(src2)) << 1) → dst
    else
        7FFF FFFFh → dst
}
else nop
    
```

Pipeline

Pipeline Stage	E1	E2
Read	<i>src1</i> , <i>src2</i>	
Written		<i>dst</i>
Unit in use	.M	

Instruction Type Single-cycle (16 × 16)

Delay Slots 1

See Also **MPYHL, SMPY, SMPYH, SMPYLH**

Example

SMPYHL .M1 A1,A2,A3

Before instruction		2 cycles after instruction			
A1	008A 0000h	138 [†]	A1	008A 0000h	
A2	0000 00A7h	167 [‡]	A2	0000 00A7h	
A3	xxxx xxxxh		A3	0000 B40Ch	46092
CSR	0001 0100h		CSR	0001 0100h	Not saturated

[†] Signed 16-MSB integer
[‡] Signed 16-LSB integer

Example

SMPYLH .M1 A1,A2,A3

	Before instruction		2 cycles after instruction
A1	0000 8000h -32768‡	A1	0000 8000h
A2	8000 0000h -32768†	A2	8000 0000h
A3	xxxx xxxxh	A3	7FFF FFFFh 2147483647
CSR	0001 0100h	CSR	0001 0300h saturated

† Signed 16-MSB integer

‡ Signed 16-LSB integer

Instruction Type Single-cycle

Delay Slots 0

See Also **SHL, SHR**

Example 1 `SSHL .S1 A0,2,A1`

	Before instruction	1 cycle after instruction	2 cycles after instruction
A0	02E3 031Ch	02E3 031Ch	02E3 031Ch
A1	xxxx xxxxh	0B8C 0C70h	0B8C 0C70h
CSR	0001 0100h	0001 0100h	0001 0100h Not saturated

Example 2 `SSHL .S1 A0,A1,A2`

	Before instruction	1 cycle after instruction	2 cycles after instruction
A0	4719 1925h	4719 1925h	4719 1925h
A1	0000 0006h	0000 0006h	0000 0006h
A2	xxxx xxxxh	7FFF FFFFh	7FFF FFFFh
CSR	0001 0100h	0001 0100h	0001 0300h Saturated

SSUB Subtract Two Signed Integers With Saturation

SSUB

Subtract Two Signed Integers With Saturation

Syntax

SSUB (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode

31	29	28	27	23	22	18	17	13	12	11	5	4	3	2	1	0				
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>			<i>src1/cst</i>		<i>x</i>	<i>op</i>			1	1	0	<i>s</i>	<i>p</i>
3			1	5			5			5		1	7			1	1			1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint sint	.L1, .L2	000 1111
<i>src1</i> <i>src2</i> <i>dst</i>	xsint sint sint	.L1, .L2	001 1111
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xsint sint	.L1, .L2	000 1110
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 slong slong	.L1, .L2	010 1100

Description

src2 is subtracted from *src1* and is saturated to the result size according to the following rules:

- 1) If the result is an int and $src1 - src2 > 2^{31} - 1$, then the result is $2^{31} - 1$.
- 2) If the result is an int and $src1 - src2 < -2^{31}$, then the result is -2^{31} .
- 3) If the result is a long and $src1 - src2 > 2^{39} - 1$, then the result is $2^{39} - 1$.
- 4) If the result is a long and $src1 - src2 < -2^{39}$, then the result is -2^{39} .

The result is placed in *dst*. If a saturate occurs, the SAT bit in CSR is set one cycle after *dst* is written.

Execution

if (cond) $src1 -s src2 \rightarrow dst$
else nop

Pipeline	Pipeline Stage	E1
	Read	<i>src1, src2</i>
	Written	<i>dst</i>
	Unit in use	<i>.L</i>

Instruction Type Single-cycle

Delay Slots 0

See Also **SUB**

Example 1 `SSUB .L2 B1, B2, B3`

	Before instruction	1 cycle after instruction	2 cycles after instruction
B1	5A2E 51A3h 1512984995	B1 5A2E 51A3h	B1 5A2E 51A3h
B2	802A 3FA2h -2144714846	B2 802A 3FA2h	B2 802A 3FA2h
B3	xxxx xxxxh	B3 7FFF FFFFh 2147483647	B3 7FFF FFFFh
CSR	0001 0100h	CSR 0001 0100h	CSR 0001 0300h Saturated

Example 2 `SSUB .L1 A0, A1, A2`

	Before instruction	1 cycle after instruction	2 cycles after instruction
A0	4367 71F2h 1130852850	A0 4367 71F2h	A0 4367 71F2h
A1	5A2E 51A3h 1512984995	A1 5A2E 51A3h	A1 5A2E 51A3h
A2	xxxx xxxxh	A2 E939 204Fh -382132145	A2 E939 204Fh
CSR	0001 0100h	CSR 0001 0100h	CSR 0001 0100h Not saturated

STB

Store Byte to Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Syntax

Register Offset


Unsigned Constant Offset

STB (.unit) *src*, *+*baseR*[*offsetR*]

STB (.unit) *src*, *+*baseR*[*ucst5*]

.unit = .D1 or .D2

Opcode



Description

Stores a byte to memory from a general-purpose register (*src*). Table 3-6 (page 3-19) describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*).

offsetR and *baseR* must be in the same register file and on the same side as the .D unit used. The *y* bit in the opcode determines the .D unit and register file used: *y* = 0 selects the .D1 unit and *baseR* and *offsetR* from the A register file, and *y* = 1 selects the .D2 unit and *baseR* and *offsetR* from the B register file.

offsetR/ucst5 is scaled by a left-shift of 0 bits. After scaling, *offsetR/ucst5* is added to or subtracted from *baseR*. For the preincrement, predecrement, positive offset, and negative offset address generator options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is sent to memory.

The addressing arithmetic that performs the additions and subtractions defaults to linear mode. However, for A4-A7 and for B4-B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10).

For **STB**, the 8 LSBs of the *src* register are stored. *src* can be in either register file, regardless of the .D unit or *baseR* or *offsetR* used. The *s* bit determines which file *src* is read from: *s* = 0 indicates *src* will be in the A register file and *s* = 1 indicates *src* will be in the B register file. The *r* bit should be cleared to 0.

Increments and decrements default to 1 and offsets default to zero when no bracketed register or constant is specified. Stores that do no modification to the *baseR* can use the syntax **R*. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 0. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Execution if (cond) *src* → mem
 else nop


Pipeline Stage	E1
Read	<i>baseR, offsetR, src</i>
Written	<i>baseR</i>
Unit in use	.D2

Instruction Type Store

Delay Slots 0
 For more information on delay slots for a store, see Chapter 4.

See Also **STH, STW**

Example STB .D1 A1, *A10



STB


Store Byte to Memory With a 15-Bit Unsigned Constant Offset

Syntax

STB (.unit) *src*, *+B14/B15[*ucst15*]

.unit = .D2

Opcode



Description

Stores a byte to memory from a general-purpose register (*src*). The memory address is formed from a base address register B14 (*y* = 0) or B15 (*y* = 1) and an offset, which is a 15-bit unsigned constant (*ucst15*). The assembler selects this format only when the constant is larger than five bits in magnitude. This instruction executes only on the .D2 unit.

The offset, *ucst15*, is scaled by a left-shift of 0 bits. After scaling, *ucst15* is added to *baseR*. The result of the calculation is the address that is sent to memory. The addressing arithmetic is always performed in linear mode.

For **STB**, the 8 LSBs of the *src* register are stored. *src* can be in either register file. The *s* bit determines which file *src* is read from: *s* = 0 indicates *src* is in the A register file and *s* = 1 indicates *src* is in the B register file.

Square brackets, [], indicate that the *ucst15* offset is left-shifted by 0. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Execution

if (cond) *src* → mem
else nop

Note:
This instruction executes only on the B side (.D2).

Pipeline


Pipeline Stage	E1
Read	B14/B15, <i>src</i>
Written	
Unit in use	.D2

Instruction Type Store

Delay Slots 0

See Also **STH, STW**

Example `STB .D2 B1, *+B14[40]`



Increments and decrements default to 1 and offsets default to zero when no bracketed register or constant is specified. Stores that do no modification to the *baseR* can use the syntax *R. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 1. Parentheses, (), can be used to set a nonscaled, constant offset. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Halfword addresses must be aligned on halfword (LSB is 0) boundaries.

Execution if (cond) *src* → mem
 else nop


Pipeline Stage	E1
Read	<i>baseR, offsetR, src</i>
Written	<i>baseR</i>
Unit in use	.D2

Instruction Type Store

Delay Slots 0
 For more information on delay slots for a store, see Chapter 4.

See Also **STB, STW**

Example 1 `STH .D1 A1, *+A10(4)`



STH Store Halfword to Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Example 2 `STH .D1 A1, *A10--[A11]`


	Before instruction	1 cycle after instruction	3 cycles after instruction
A1	9A32 2634h	9A32 2634h	9A32 2634h
A10	0000 0100h	0000 00F8h	0000 00F8h
A11	0000 0004h	0000 0004h	0000 0004h
mem F8h	0000h	0000h	0000h
mem 100h	0000	0000h	2634h

STH *Store Halfword to Memory With a 15-Bit Unsigned Constant Offset*

Instruction Type	Store
Delay Slots	0
See Also	STB, STW

STW*Store Word to Memory With a 5-Bit Unsigned Constant Offset or Register Offset***Syntax****Register Offset****Unsigned Constant Offset****STW** (.unit) *src*, *+*baseR*[*offsetR*]**STW** (.unit) *src*, *+*baseR*[*ucst5*]

.unit = .D1 or .D2

Opcode**Description**

Stores a word to memory from a general-purpose register (*src*). Table 3–6 (page 3-19) describes the addressing generator options. The memory address is formed from a base address register (*baseR*) and an optional offset that is either a register (*offsetR*) or a 5-bit unsigned constant (*ucst5*).

offsetR and *baseR* must be in the same register file and on the same side as the .D unit used. The *y* bit in the opcode determines the .D unit and register file used: *y* = 0 selects the .D1 unit and *baseR* and *offsetR* from the A register file, and *y* = 1 selects the .D2 unit and *baseR* and *offsetR* from the B register file.

offsetR/ucst5 is scaled by a left-shift of 2 bits. After scaling, *offsetR/ucst5* is added to or subtracted from *baseR*. For the preincrement, predecrement, positive offset, and negative offset address generator options, the result of the calculation is the address to be accessed in memory. For postincrement or postdecrement addressing, the value of *baseR* before the addition or subtraction is sent to memory.

The addressing arithmetic that performs the additions and subtractions defaults to linear mode. However, for A4–A7 and for B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10).

For **STW**, the entire 32-bits of the *src* register are stored. *src* can be in either register file, regardless of the .D unit or *baseR* or *offsetR* used. The *s* bit determines which file *src* is read from: *s* = 0 indicates *src* will be in the A register file and *s* = 1 indicates *src* will be in the B register file. The *r* bit should be cleared to 0.

STW Store Word to Memory With a 5-Bit Unsigned Constant Offset or Register Offset

Increments and decrements default to 1 and offsets default to zero when no bracketed register or constant is specified. Stores that do no modification to the *baseR* can use the syntax **R*. Square brackets, [], indicate that the *ucst5* offset is left-shifted by 2. Parentheses, (), can be used to set a nonscaled, constant offset. For example, **STW** (.unit) *src*, *++baseR*(12) represents an offset of 12 bytes; whereas, **STW** (.unit) *src*, *++baseR*[12] represents an offset of 12 words, or 48 bytes. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Word addresses must be aligned on word (two LSBs are 0) boundaries.

Execution if (cond) *src* → mem
else nop

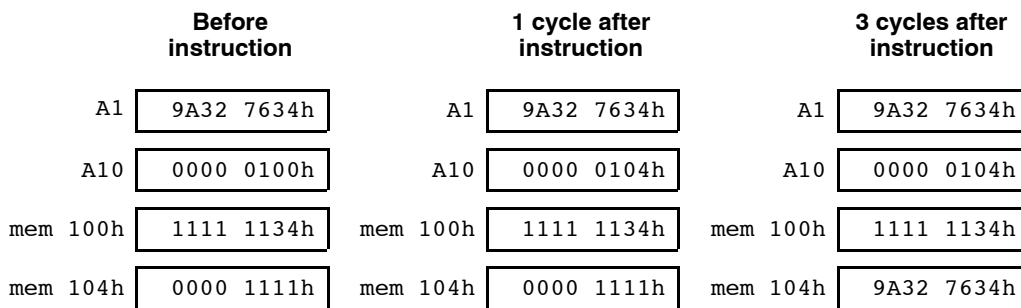
Pipeline Stage	E1
Read	<i>baseR</i> , <i>offsetR</i> , <i>src</i>
Written	<i>baseR</i>
Unit in use	.D2

Instruction Type Store

Delay Slots 0
For more information on delay slots for a store, see Chapter 4.


See Also STB, STH

Example STW .D1 A1, ++A10[1]



STW**Store Word to Memory With a 15-Bit Unsigned Constant Offset****Syntax****STW** (.unit) *src*, *+B14/B15[*ucst15*]

.unit = .D2

Opcode**Description**

Stores a word to memory from a general-purpose register (*src*). The memory address is formed from a base address register B14 ($y = 0$) or B15 ($y = 1$) and an offset, which is a 15-bit unsigned constant (*ucst15*). The assembler selects this format only when the constant is larger than five bits in magnitude. This instruction executes only on the .D2 unit.

The offset, *ucst15*, is scaled by a left-shift of 2 bits. After scaling, *ucst15* is added to *baseR*. The result of the calculation is the address that is sent to memory. The addressing arithmetic is always performed in linear mode.

For **STW**, the entire 32-bits of the *src* register are stored. *src* can be in either register file. The *s* bit determines which *src* is read from: $s = 0$ indicates *src* is in the A register file and $s = 1$ indicates *src* is in the B register file.

Square brackets, [], indicate that the *ucst15* offset is left-shifted by 2. Parentheses, (), can be used to set a nonscaled, constant offset. For example, **STW** (.unit) *src*, *+B14/B15(60) represents an offset of 12 bytes; whereas, **STW** (.unit) *src*, *+B14/B15[60] represents an offset of 60 words, or 240 bytes. You must type either brackets or parentheses around the specified offset, if you use the optional offset parameter.

Word addresses must be aligned on word (two LSBs are 0) boundaries.

Execution

```
if (cond)   src → mem
else nop
```

Note:

This instruction executes only on the B side (.D2).

STW *Store Word to Memory With a 15-Bit Unsigned Constant Offset*

Pipeline	Pipeline Stage	E1
	Read	B14/B15, <i>src</i>
	Written	
	Unit in use	.D2

Instruction Type Store

Delay Slots 0

See Also STB, STH

SUB


Subtract Two Signed Integers Without Saturation

Syntax

SUB (.unit) *src1*, *src2*, *dst*
 or
SUB (.D1 or .D2) *src2*, *src1*, *dst*
 .unit = .L1, .L2, .S1, .S2

Opcode


.L unit



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint sint	.L1, .L2	000 0111
<i>src1</i> <i>src2</i> <i>dst</i>	xsint sint sint	.L1, .L2	001 0111
<i>src1</i> <i>src2</i> <i>dst</i>	sint xsint slong	.L1, .L2	010 0111
<i>src1</i> <i>src2</i> <i>dst</i>	xsint sint slong	.L1, .L2	011 0111
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 xsint sint	.L1, .L2	000 0110
<i>src1</i> <i>src2</i> <i>dst</i>	scst5 slong slong	.L1, .L2	010 0100

SUB Subtract Two Signed Integers Without Saturation

Opcode .S unit



Opcode map field used...	For operand type...	Unit	Opfield
src1	sint	.S1, .S2	01 0111
src2	xsint		
dst	sint		
src1	scst5	.S1, .S2	01 0110
src2	xsint		
dst	sint		


Description for .L1, .L2 and .S1, .S2 Opcodes

src2 is subtracted from *src1*. The result is placed in *dst*.

Execution for .L1, .L2 and .S1, .S2 Opcodes

if (cond)
 $src1 - src2 \rightarrow dst$
 else nop

Opcode .D unit



Opcode map field used...	For operand type...	Unit	Opfield
src2	sint	.D1, .D2	01 0001
src1	sint		
dst	sint		
src2	sint	.D1, .D2	01 0011
src1	ucst5		
dst	sint		

Description for .D1, .D2 Opcodes

src1 is subtracted from *src2*. The result is placed in *dst*.

Execution for .D1, .D2 Opcodes

if (cond)
 $src2 - src1 \rightarrow dst$
 else nop

Note:

Subtraction with a signed constant on the .L and .S units allows either the first or the second operand to be the signed 5-bit constant.

SUB (.unit) *src1, scst5, dst* is encoded as **ADD** (.unit) *-scst5, src2, dst* where the *src1* register is now *src2* and *scst5* is now *-scst5*.

However, the .D unit provides only the second operand as a constant since it is an unsigned 5-bit constant. *ucst5* allows a greater offset for addressing with the .D unit.

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.L, .S, or .D

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, SSUB, SUBC, SUBU, SUB2**

Example SUB .L1 A1, A2, A3

	Before instruction		1 cycle after instruction
A1	0000 325Ah 12810	A1	0000 325Ah
A2	FFFF FF12h -238	A2	FFFF FF12h
A3	xxxx xxxxh	A3	0000 3348h 13128

SUBAB *Subtract Using Byte Addressing Mode*


SUBAB

Subtract Using Byte Addressing Mode

Syntax **SUBAB** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 0001
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 0011
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description *src1* is subtracted from *src2* using the byte addressing mode specified for *src2*. The subtraction defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). The result is placed in *dst*.

Execution if (cond) *src2* -a *src1* → *dst*
else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type Single-cycle

Delay Slots 0

See Also **SUB, SUBAH, SUBAW**

Example

SUBAB .D1 A5,A0,A5

	Before instruction		1 cycle after instruction
A0	0000 0004h	A0	0000 0004h
A5	0000 4000h	A5	0000 400Ch
AMR	0003 0004h	AMR	0003 0004h

BK0 = 3 → size = 16

A5 in circular addressing mode using BK0

SUBAH *Subtract Using Halfword Addressing Mode*


SUBAH

Subtract Using Halfword Addressing Mode

Syntax **SUBAH** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 0101
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 0111
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description *src1* is subtracted from *src2* using the halfword addressing mode specified for *src2*. The subtraction defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). If circular addressing is enabled, *src1* is left shifted by 1. The result is placed in *dst*.

Execution if (cond) *src2* -a *src1* → *dst*
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type Single-cycle

Delay Slots 0

See Also **SUB, SUBAB, SUBAW**

SUBAW*Subtract Using Word Addressing Mode***Syntax****SUBAW** (.unit) *src2*, *src1*, *dst*

.unit = .D1 or .D2

Opcode

31	29	28	27	23	22	18	17	13	12	7	6	5	4	3	2	1	0		
<i>creg</i>			<i>z</i>	<i>dst</i>			<i>src2</i>		<i>src1/cst</i>		<i>op</i>		1	0	0	0	0	<i>s</i>	<i>p</i>
3			1	5			5		5		6						1		1

Opcode map field used...	For operand type...	Unit	Opfield
<i>src2</i>	sint	.D1, .D2	11 1001
<i>src1</i>	sint		
<i>dst</i>	sint		
<i>src2</i>	sint	.D1, .D2	11 1011
<i>src1</i>	ucst5		
<i>dst</i>	sint		

Description

src1 is subtracted from *src2* using the word addressing mode specified for *src2*. The subtraction defaults to linear mode. However, if *src2* is one of A4–A7 or B4–B7, the mode can be changed to circular mode by writing the appropriate value to the AMR (see section 2.7.3, page 2-10). If circular addressing is enabled, *src1* is left shifted by 2. The result is placed in *dst*.

Execution

if (cond) $src2 -a src1 \rightarrow dst$
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.D

Instruction Type

Single-cycle

Delay Slots

0

See Also**SUB, SUBAB, SUBAH**

SUBAW *Subtract Using Word Addressing Mode*

Example

SUBAW .D1 A5,2,A3

	Before instruction	1 cycle after instruction
A3	xxxx xxxxh	0000 0108h
A5	0000 0100h	0000 0100h
AMR	0003 0004h	0003 0004h

BK0 = 3 → size = 16

A5 in circular addressing mode using BK0


SUBC

Subtract Conditionally and Shift—Used for Division

Syntax **SUBC** (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode



Opcode map field used...	For operand type...	Unit
<i>src1</i>	uint	.L1, .L2
<i>src2</i>	xuint	
<i>dst</i>	uint	

Description Subtract *src2* from *src1*. If result is greater than or equal to 0, left shift result by 1, add 1 to it, and place it in *dst*. If result is less than 0, left shift *src1* by 1, and place it in *dst*. This step is commonly used in division.

Execution if (cond) {
 if (*src1* - *src2* ≥ 0)
 ((*src1* - *src2*) << 1) + 1 → *dst*
 else *src1* << 1 → *dst*
 }
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

See Also **ADD, SSUB, SUB, SUBU, SUB2**

SUBC *Subtract Conditionally and Shift-Used for Division*

Example 1

SUBC .L1 A0,A1,A0

	Before instruction	1 cycle after instruction
A0	0000 125Ah 4698	0000 024B4h 9396
A1	0000 1F12h 7954	0000 1F12h

Example 2

SUBC .L1 A0,A1,A0

	Before instruction	1 cycle after instruction
A0	0002 1A31h 137777	0000 47E5h 18405
A1	0001 F63Fh 128575	0001 F63Fh


SUBU

Subtract Two Unsigned Integers Without Saturation

Syntax **SUBU** (.unit) *src1*, *src2*, *dst*

.unit = .L1 or .L2

Opcode



Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	010 1111
<i>src2</i>	xuint		
<i>dst</i>	ulong		
<i>src1</i>	xuint	.L1, .L2	011 1111
<i>src2</i>	uint		
<i>dst</i>	ulong		

Description *src2* is subtracted from *src1*. The result is placed in *dst*.

Execution if (cond)
 src1 - *src2* → *dst*
 else nop

Pipeline

Pipeline Stage	E1
Read	<i>src1</i> , <i>src2</i>
Written	<i>dst</i>
Unit in use	.L

Instruction Type Single-cycle

Delay Slots 0

See Also **ADDU, SSUB, SUB, SUBC, SUB2**

SUBU Subtract Two Unsigned Integers Without Saturation

Example

SUBU .L1 A1,A2,A5:A4

	Before instruction		1 cycle after instruction				
A1	<table border="1"><tr><td>0000 325Ah</td></tr></table> 12810 [†]	0000 325Ah		A1 <table border="1"><tr><td>0000 325Ah</td></tr></table>	0000 325Ah		
0000 325Ah							
0000 325Ah							
A2	<table border="1"><tr><td>FFFF FF12h</td></tr></table> 4294967058 [†]	FFFF FF12h		A2 <table border="1"><tr><td>FFFF FF12h</td></tr></table>	FFFF FF12h		
FFFF FF12h							
FFFF FF12h							
A5:A4	<table border="1"><tr><td>xxxx xxxxh</td><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh	xxxx xxxxh		A5:A4 <table border="1"><tr><td>0000 00FFh</td><td>0000 3348h</td></tr></table> -4294954168 [‡]	0000 00FFh	0000 3348h
xxxx xxxxh	xxxx xxxxh						
0000 00FFh	0000 3348h						

[†] Unsigned 32-bit integer

[‡] Signed 40-bit (long) integer

SUB2 Subtract Two 16-Bit Integers on Upper and Lower Register Halves

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.S

Instruction Type

Single-cycle

Delay Slots

0

See Also

ADD2, SSUB, SUB, SUBC, SUBU

Example 1

SUB2 .S1 A3, A4, A5

	Before instruction		1 cycle after instruction		
A3	<table border="1"><tr><td>1105 6E30h</td></tr></table> 4357 28208	1105 6E30h		A3 <table border="1"><tr><td>1105 6E30h</td></tr></table> 4357 28208	1105 6E30h
1105 6E30h					
1105 6E30h					
A4	<table border="1"><tr><td>1105 6980h</td></tr></table> 4357 27008	1105 6980h		A4 <table border="1"><tr><td>1105 6980h</td></tr></table> 4357 27008	1105 6980h
1105 6980h					
1105 6980h					
A5	<table border="1"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		A5 <table border="1"><tr><td>0000 04B0h</td></tr></table> 0 1200	0000 04B0h
xxxx xxxxh					
0000 04B0h					

Example 2

SUB2 .S2X B1,A0,B2


	Before instruction		1 cycle after instruction		
A0	<table border="1"><tr><td>0021 3271h</td></tr></table> †33 12913‡	0021 3271h		A0 <table border="1"><tr><td>0021 3271h</td></tr></table>	0021 3271h
0021 3271h					
0021 3271h					
B1	<table border="1"><tr><td>003A 1B48h</td></tr></table> †58 6984‡	003A 1B48h		B1 <table border="1"><tr><td>003A 1B48h</td></tr></table>	003A 1B48h
003A 1B48h					
003A 1B48h					
B2	<table border="1"><tr><td>xxxx xxxxh</td></tr></table>	xxxx xxxxh		B2 <table border="1"><tr><td>0019 E8D7h</td></tr></table> 25† -5929‡	0019 E8D7h
xxxx xxxxh					
0019 E8D7h					

† Signed 16-MSB integer


‡ Signed 16-LSB integer

XOR*Bitwise Exclusive OR***Syntax** **XOR** (.unit) *src1*, *src2*, *dst*

.unit = .L1, .L2, .S1, .S2

Opcode .L unit

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.L1, .L2	110 1111
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.L1, .L2	110 1110
<i>src2</i>	xuint		
<i>dst</i>	uint		

Opcode .S unit

Opcode map field used...	For operand type...	Unit	Opfield
<i>src1</i>	uint	.S1, .S2	00 1011
<i>src2</i>	xuint		
<i>dst</i>	uint		
<i>src1</i>	scst5	.S1, .S2	00 1010
<i>src2</i>	xuint		
<i>dst</i>	uint		

Description Performs a bitwise exclusive-OR (**XOR**) operation between *src1* and *src2*. The result is placed in *dst*. The *scst5* operands are sign extended to 32 bits.**Execution** if (cond) *src1* XOR *src2* → *dst*
else nop

XOR *Bitwise Exclusive OR*

Pipeline

Pipeline Stage	E1
Read	<i>src1, src2</i>
Written	<i>dst</i>
Unit in use	.L or .S

Instruction Type Single-cycle

Delay Slots 0

See Also **AND, OR**

Example 1 XOR .S1 A3, A4, A5

Before instruction

A3	0721 325Ah
A4	0019 0F12h
A5	xxxx xxxxh

1 cycle after instruction

A3	0721 325Ah
A4	0019 0F12h
A5	0738 3D48h

Example 2 XOR .L2 B1, 0dh, B8

Before instruction

B1	0000 1023h
B8	xxxx xxxxh

1 cycle after instruction

B1	0000 1023h
B8	0000 102Eh

ZERO*Zero a Register***Syntax****ZERO** (.unit) *dst*

.unit = .L1, .L2, .D1, .D2, .S1, .S2

Opcode

Opcode map field used...	For operand type...	Unit	Opfield
<i>dst</i>	sint	.L1, .L2	001 0111
<i>dst</i>	sint	.D1, .D2	01 0001
<i>dst</i>	sint	.S1, .S2	01 0111
<i>dst</i>	slong	.L1, .L2	011 0111

Description

The **ZERO** pseudo-operation fills the *dst* register with 0s by subtracting the *dst* from itself and placing the result in the *dst*.

In the case where *dst* is sint, the assembler uses the **MVK** (.unit) 0, *dst* instruction.

In the case where *dst* is slong, the assembler uses the **SUB** (.unit) *src1*, *src2*, *dst* instruction.

Execution

if (cond) $dst - dst \rightarrow dst$
 else nop

Instruction Type

Single-cycle

Delay Slots

0

See Also**MVK, SUB****Example**

ZERO .D1 A1

Before instructionA1 B174 6CA1h**1 cycle after instruction**A1 0000 0000h

The C62x DSP pipeline provides flexibility to simplify programming and improve performance. These two factors provide this flexibility:

- 1) Control of the pipeline is simplified by eliminating pipeline interlocks.
- 2) Increased pipelining eliminates traditional architectural bottlenecks in program fetch, data access, and multiply operations. This provides single-cycle throughput.

This chapter starts with a description of the pipeline flow. Highlights are:

- The pipeline can dispatch eight parallel instructions every cycle.
- Parallel instructions proceed simultaneously through each pipeline phase.
- Serial instructions proceed through the pipeline with a fixed relative phase difference between instructions.
- Load and store addresses appear on the CPU boundary during the same pipeline phase, eliminating read-after-write memory conflicts.

All instructions require the same number of pipeline phases for fetch and decode, but require a varying number of execute phases. This chapter contains a description of the number of execution phases for each type of instruction.

Finally, this chapter contains performance considerations for the pipeline. These considerations include the occurrence of fetch packets that contain multiple execute packets, execute packets that contain multicycle **NOPs**, and memory considerations for the pipeline. For more information about fully optimizing a program and taking full advantage of the pipeline, see the *TMS320C6000 Programmer's Guide* (SPRU198).

Topic	Page
4.1 Pipeline Operation Overview	4-2
4.2 Pipeline Execution of Instruction Types	4-11
4.3 Performance Considerations	4-18


4.1 Pipeline Operation Overview

The pipeline phases are divided into three stages:

- Fetch
- Decode
- Execute

All instructions in the C62x DSP instruction set flow through the fetch, decode, and execute stages of the pipeline. The fetch stage of the pipeline has four phases for all instructions, and the decode stage has two phases for all instructions. The execute stage of the pipeline requires a varying number of phases, depending on the type of instruction. The stages of the C62x DSP pipeline are shown in Figure 4-1.

Figure 4-1. Pipeline Stages



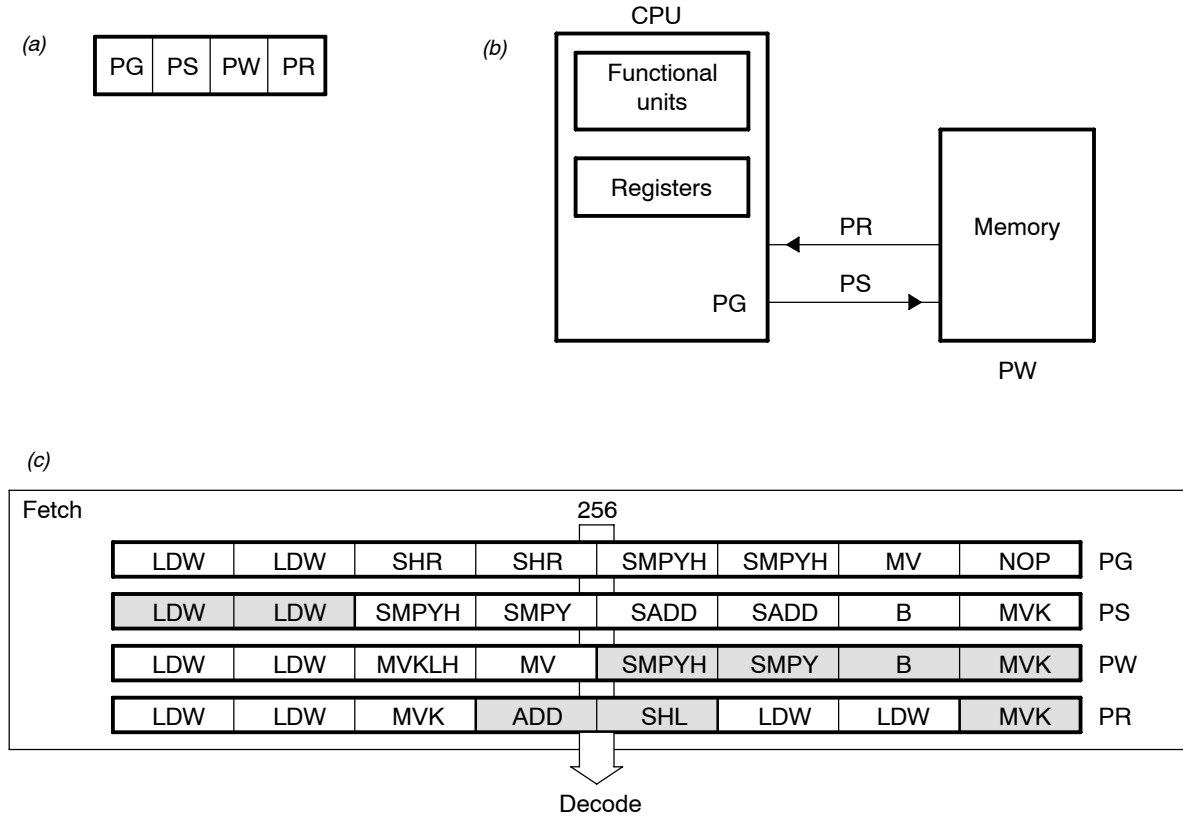
4.1.1 Fetch

The fetch phases of the pipeline are:

- PG:** Program address generate
- PS:** Program address send
- PW:** Program access ready wait
- PR:** Program fetch packet receive

The C62x DSP uses a fetch packet (FP) of eight words. All eight of the words proceed through fetch processing together, through the PG, PS, PW, and PR phases. Figure 4-2(a) shows the fetch phases in sequential order from left to right. Figure 4-2(b) is a functional diagram of the flow of instructions through the fetch phases. During the PG phase, the program address is generated in the CPU. In the PS phase, the program address is sent to memory. In the PW phase, a memory read occurs. Finally, in the PR phase, the fetch packet is received at the CPU. Figure 4-2(c) shows fetch packets flowing through the phases of the fetch stage of the pipeline. In Figure 4-2(c), the first fetch packet (in PR) is made up of four execute packets, and the second and third fetch packets (in PW and PS) contain two execute packets each. The last fetch packet (in PG) contains a single execute packet of eight instructions.

Figure 4-2. Fetch Phases of the Pipeline



4.1.2 Decode

The decode phases of the pipeline are:


- DP:** Instruction dispatch
- DC:** Instruction decode

In the DP phase of the pipeline, the fetch packets are split into execute packets. Execute packets consist of one instruction or from two to eight parallel instructions. During the DP phase, the instructions in an execute packet are assigned to the appropriate functional units. In the DC phase, the source registers, destination registers, and associated paths are decoded for the execution of the instructions in the functional units.

Figure 4-3(a) shows the decode phases in sequential order from left to right. Figure 4-3(b) shows a fetch packet that contains two execute packets as they are processed through the decode stage of the pipeline. The last six instructions of the fetch packet (FP) are parallel and form an execute packet (EP). This EP is in the dispatch phase (DP) of the decode stage. The arrows indicate each instruction's assigned functional unit for execution during the same cycle. The **NOP** instruction in the eighth slot of the FP is not dispatched to a functional unit because there is no execution associated with it.

The first two slots of the fetch packet (shaded below) represent an execute packet of two parallel instructions that were dispatched on the previous cycle. This execute packet contains two **MPY** instructions that are now in decode (DC) one cycle before execution. There are no instructions decoded for the .L, .S, and .D functional units for the situation illustrated.

Figure 4-3. Decode Phases of the Pipeline




† NOP is not dispatched to a functional unit.

4.1.3 Execute

The execute portion of the pipeline is subdivided into five phases (E1–E5). Different types of instructions require different numbers of these phases to complete their execution. These phases of the pipeline play an important role in your understanding the device state at CPU cycle boundaries. The execution of different types of instructions in the pipeline is described in section 4.2, *Pipeline Execution of Instruction Types*. Figure 4–4(a) shows the execute phases of the pipeline in sequential order from left to right. Figure 4–4(b) shows the portion of the functional block diagram in which execution occurs.

Figure 4–4. Execute Phases of the Pipeline



4.1.4 Pipeline Operation Summary

Figure 4-5 shows all the phases in each stage of the C62x DSP pipeline in sequential order, from left to right.

Figure 4-5. Pipeline Phases




Figure 4-6 shows an example of the pipeline flow of consecutive fetch packets that contain eight parallel instructions. In this case, where the pipeline is full, all instructions in a fetch packet are in parallel and split into one execute packet per fetch packet. The fetch packets flow in lockstep fashion through each phase of the pipeline.

For example, examine cycle 7 in Figure 4-6. When the instructions from FPN reach E1, the instructions in the execute packet from FPN + 1 are being decoded. FP n + 2 is in dispatch while FPs n + 3, n + 4, n + 5, and n + 6 are each in one of four phases of program fetch. See section 4.3, page 4-18, for additional detail on code flowing through the pipeline. Table 4-1 summarizes the pipeline phases and what happens in each phase.

Figure 4-6. Pipeline Operation: One Execute Packet per Fetch Packet

Fetch packet	Clock cycle												
	1	2	3	4	5	6	7	8	9	10	11	12	13
n	PG	PS	PW	PR	DP	DC	E1	E2	E3	E4	E5		
n+1		PG	PS	PW	PR	DP	DC	E1	E2	E3	E4	E5	
n+2			PG	PS	PW	PR	DP	DC	E1	E2	E3	E4	E5
n+3				PG	PS	PW	PR	DP	DC	E1	E2	E3	E4
n+4					PG	PS	PW	PR	DP	DC	E1	E2	E3
n+5						PG	PS	PW	PR	DP	DC	E1	E2
n+6							PG	PS	PW	PR	DP	DC	E1
n+7								PG	PS	PW	PR	DP	DC
n+8									PG	PS	PW	PR	DP
n+9										PG	PS	PW	PR
n+10											PG	PS	PW

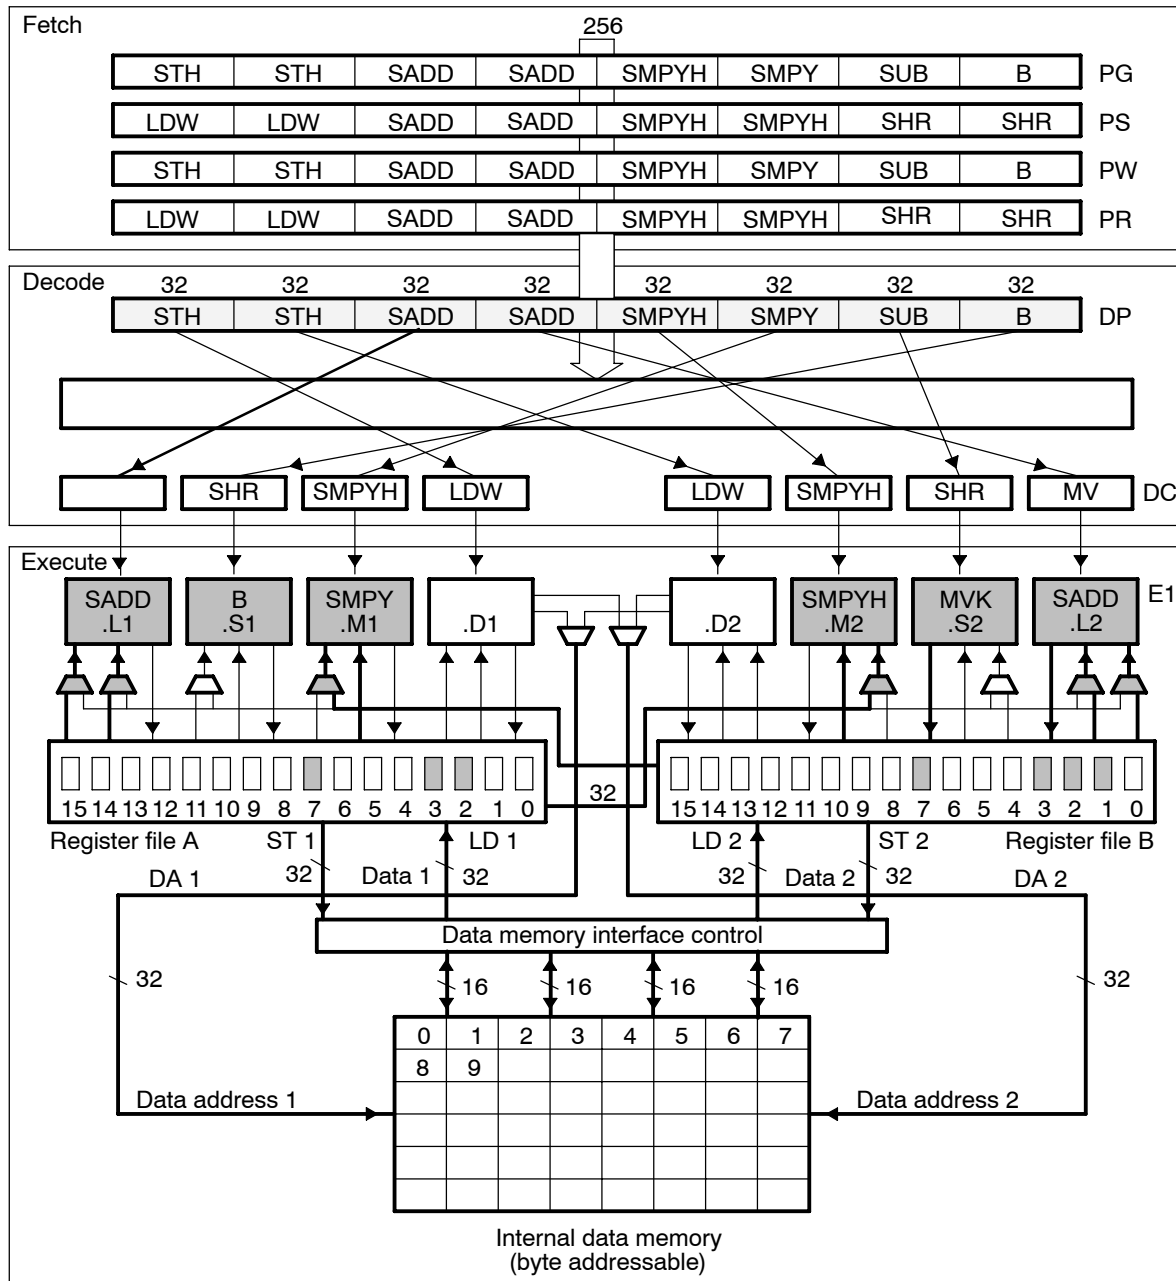
Table 4-1. Operations Occurring During Pipeline Phases

Stage	Phase	Symbol	During This Phase	Instruction Type Completed
Program fetch	Program address generate	PG	The address of the fetch packet is determined.	
	Program address send	PS	The address of the fetch packet is sent to memory.	
	Program wait	PW	A program memory access is performed.	
	Program data receive	PR	The fetch packet is at the CPU boundary.	
Program decode	Dispatch	DP	The next execute packet in the fetch packet is determined and sent to the appropriate functional units to be decoded.	
	Decode	DC	Instructions are decoded in functional units.	
Execute	Execute 1	E1	For all instruction types, the conditions for the instructions are evaluated and operands are read. For load and store instructions, address generation is performed and address modifications are written to a register file. [†] For branch instructions, branch fetch packet in PG phase is affected. [†] For single-cycle instructions, results are written to a register file. [†]	Single cycle
	Execute 2	E2	For load instructions, the address is sent to memory. For store instructions, the address and data are sent to memory. [†] Single-cycle instructions that saturate results set the SAT bit in the control status register (CSR) if saturation occurs. [†] For single 16 × 16 multiply instructions, results are written to a register file. [†]	Multiply
	Execute 3	E3	Data memory accesses are performed. Any multiply instruction that saturates results sets the SAT bit in the control status register (CSR) if saturation occurs. [†]	Store
	Execute 4	E4	For load instructions, data is brought to the CPU. [†]	Load
	Execute 5	E5	For load instructions, data is written into a register. [†]	Load

[†] This assumes that the conditions for the instructions are evaluated as true. If the condition is evaluated as false, the instruction does not write any results or have any pipeline operation after E1.

Figure 4-7 shows a functional block diagram of the pipeline stages.

Figure 4-7. Pipeline Phases Block Diagram



The pipeline operation is based on CPU cycles. A CPU cycle is the period during which a particular execute packet is in a particular pipeline phase. CPU cycle boundaries always occur at clock cycle boundaries.

As code flows through the pipeline phases, it is processed by different parts of the C62x DSP. Figure 4-7 shows a full pipeline with a fetch packet in every phase of fetch. One execute packet of eight instructions is being dispatched at the same time that a 7-instruction execute packet is in decode. The arrows between DP and DC correspond to the functional units identified in the code in Example 4-1.

In the DC phase portion of Figure 4-7, one box is empty because a **NOP** was the eighth instruction in the fetch packet in DC and no functional unit is needed for a **NOP**. Finally, Figure 4-7 shows six functional units processing code during the same cycle of the pipeline.

Registers used by the instructions in E1 are shaded in Figure 4-7. The multiplexers used for the input operands to the functional units are also shaded in the figure. The bold crosspaths are used by the **MPY** instructions.

Most C62x DSP instructions are single-cycle instructions, which means they have only one execution phase (E1). A small number of instructions require more than one execute phase. The types of instructions, each of which require different numbers of execute phases, are described in section 4.2.

Example 4-1. Execute Packet in Figure 4-7

	SADD	.L1	A2,A7,A2	; E1 Phase
	SADD	.L2	B2,B7,B2	
	SMPYH	.M2X	B3,A3,B2	
	SMPY	.M1X	B3,A3,A2	
	B	.S1	LOOP1	
	MVK	.S2	117,B1	
	LDW	.D2	*B4++,B3	; DC Phase
	LDW	.D1	*A4++,A3	
	MV	.L2X	A1,B0	
	SMPYH	.M1	A2,A2,A0	
	SMPYH	.M2	B2,B2,B10	
	SHR	.S1	A2,16,A5	
	SHR	.S2	B2,16,B5	
LOOP1:				
	STH	.D1	A5,*A8++[2]	; DP, PW, and PG Phases
	STH	.D2	B5,*B8++[2]	
	SADD	.L1	A2,A7,A2	
	SADD	.L2	B2,B7,B2	
	SMPYH	.M2X	B3,A3,B2	
	SMPY	.M1X	B3,A3,A2	
	[B1] B	.S1	LOOP1	
	[B1] SUB	.S2	B1,1,B1	
	LDW	.D2	*B4++,B3	: PR and PS Phases
	LDW	.D1	*A4++,A3	
	SADD	.L1	A0,A1,A1	
	SADD	.L2	B10,B0,B0	
	SMPYH	.M1	A2,A2,A0	
	SMPYH	.M2	B2,B2,B10	
	SHR	.S1	A2,16,A5	
	SHR	.S2	B2,16,B5	

4.2 Pipeline Execution of Instruction Types

The pipeline operation of the C62x DSP instructions can be categorized into six instruction types. Five of these are shown in Table 4-2 (**NOP** is not included in the table), which is a mapping of operations occurring in each execution phase for the different instruction types. The delay slots associated with each instruction type are also listed.

The execution of instructions is defined in terms of delay slots. A delay slot is a CPU cycle that occurs after the first execution phase (E1) of an instruction. Results from instructions with delay slots are not available until the end of the last delay slot. For example, a multiply instruction has one delay slot, which means that one CPU cycle elapses before the results of the multiply are available for use by a subsequent instruction. However, results are available from other instructions finishing execution during the same CPU cycle in which the multiply is in a delay slot.

Table 4-2. Execution Stage Length Description for Each Instruction Type

Execution phases	Instruction Type				
	Single Cycle	16 × 16 Single Multiply	Store	Load	Branch
E1	Compute result and write to register	Read operands and start computations	Compute address	Compute address	Target code in PG [‡]
E2		Compute result and write to register	Send address and data to memory	Send address to memory	
E3			Access memory	Access memory	
E4				Send data back to CPU	
E5				Write data into register	
Delay slots	0	1	0 [†]	4 [†]	5 [‡]

[†] See sections 4.2.3 and 4.2.4 for more information on execution and delay slots for stores and loads.

[‡] See section 4.2.5 for more information on branches.

- Notes:**
- 1) This table assumes that the condition for each instruction is evaluated as true. If the condition is evaluated as false, the instruction does not write any results or have any pipeline operation after E1.
 - 2) **NOP** is not shown and has no operation in any of the execution phases.

4.2.1 Single-Cycle Instructions

Single-cycle instructions complete execution during the E1 phase of the pipeline. Figure 4-8 shows the fetch, decode, and execute phases of the pipeline that the single-cycle instructions use.

Figure 4-9 shows the single-cycle execution diagram. The operands are read, the operation is performed, and the results are written to a register, all during E1. Single-cycle instructions have no delay slots.

Figure 4-8. Single-Cycle Instruction Phases





Figure 4-9. Single-Cycle Instruction Execution Block Diagram



4.2.2 Two-Cycle Instructions

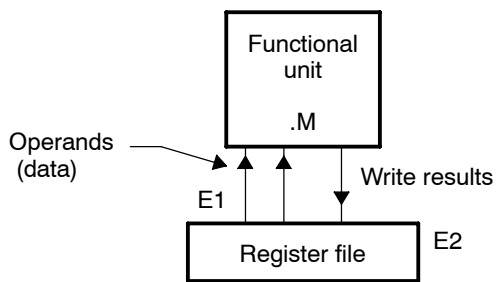
Two-cycle or multiply instructions use both the E1 and E2 phases of the pipeline to complete their operations. Figure 4-10 shows the fetch, decode, and execute phases of the pipeline that the two-cycle instructions use.

Figure 4-11 shows the operations occurring in the pipeline for a multiply instruction. In the E1 phase, the operands are read and the multiply begins. In the E2 phase, the multiply finishes, and the result is written to the destination register. Multiply instructions have one delay slot.

Figure 4-10. Two-Cycle Instruction Phases



Figure 4-11. Single 16 × 16 Multiply Instruction Execution Block Diagram



4.2.3 Store Instructions

Store instructions require phases E1 through E3 of the pipeline to complete their operations. Figure 4-12 shows the fetch, decode, and execute phases of the pipeline that the store instructions use.

Figure 4-13 shows the operations occurring in the pipeline phases for a store instruction. In the E1 phase, the address of the data to be stored is computed. In the E2 phase, the data and destination addresses are sent to data memory. In the E3 phase, a memory write is performed. The address modification is performed in the E1 stage of the pipeline. Even though stores finish their execution in the E3 phase of the pipeline, they have no delay slots. There is additional explanation of why stores have zero delay slots in section 4.2.4.

Figure 4-12. Store Instruction Phases





Figure 4-13. Store Instruction Execution Block Diagram



When you perform a load and a store to the same memory location, these rules apply ($i = \text{cycle}$):

- When a load is executed before a store, the old value is loaded and the new value is stored.

i	LDW
$i + 1$	STW
- When a store is executed before a load, the new value is stored and the new value is loaded.

i	STW
$i + 1$	LDW
- When the instructions are executed in parallel, the old value is loaded first and then the new value is stored, but both occur in the same phase.

i	STW
i	LDW

4.2.4 Load Instructions

Data loads require all five, E1–E5, of the pipeline execute phases to complete their operations. Figure 4–14 shows the fetch, decode, and execute phases of the pipeline that the load instructions use.

Figure 4–15 shows the operations occurring in the pipeline phases for a load. In the E1 phase, the data address pointer is modified in its register. In the E2 phase, the data address is sent to data memory. In the E3 phase, a memory read at that address is performed.

Figure 4–14. Load Instruction Phases





Figure 4–15. Load Instruction Execution Block Diagram



In the E4 stage of a load, the data is received at the CPU core boundary. Finally, in the E5 phase, the data is loaded into a register. Because data is not written to the register until E5, load instructions have four delay slots. Because pointer results are written to the register in E1, there are no delay slots associated with the address modification.

In the following code, pointer results are written to the A4 register in the first execute phase of the pipeline and data is written to the A3 register in the fifth execute phase.

```
LDW  .D1  *A4++,A3
```

Because a store takes three execute phases to write a value to memory and a load takes three execute phases to read from memory, a load following a store accesses the value placed in memory by that store in the cycle after the store is completed. This is why the store is considered to have zero delay slots.

4.2.5 Branch Instructions

Although branch instructions take one execute phase, there are five delay slots between the execution of the branch and execution of the target code. Figure 4-16 shows the pipeline phases used by the branch instruction and branch target code. The delay slots are shaded.

Figure 4-17 shows a branch instruction execution block diagram. If a branch is in the E1 phase of the pipeline (in the .S2 unit in the figure), its branch target is in the fetch packet that is in PG during that same cycle (shaded in the figure). Because the branch target has to wait until it reaches the E1 phase to begin execution, the branch takes five delay slots before the branch target code executes.

Figure 4-16. Branch Instruction Phases





Figure 4-17. Branch Instruction Execution Block Diagram



4.3 Performance Considerations

The C62x DSP pipeline is most effective when it is kept as full as the algorithms in the program allow it to be. It is useful to consider some situations that can affect pipeline performance.

A fetch packet (FP) is a grouping of eight instructions. Each FP can be split into from one to eight execute packets (EPs). Each EP contains instructions that execute in parallel. Each instruction executes in an independent functional unit. The effect on the pipeline of combinations of EPs that include varying numbers of parallel instructions, or just a single instruction that executes serially with other code, is considered here.

In general, the number of execute packets in a single FP defines the flow of instructions through the pipeline. Another defining factor is the instruction types in the EP. Each type of instruction has a fixed number of execute cycles that determines when this instruction's operations are complete. Section 4.3.2 covers the effect of including a multicycle **NOP** in an individual EP.

Finally, the effect of the memory system on the operation of the pipeline is considered. The access of program and data memory is discussed, along with memory stalls.

4.3.1 Pipeline Operation With Multiple Execute Packets in a Fetch Packet

Referring to Figure 4-6 on page 4-6, pipeline operation is shown with eight instructions in every fetch packet. Figure 4-18, however, shows the pipeline operation with a fetch packet that contains multiple execute packets. Code for Figure 4-18 might have this layout:

```

    instruction A ; EP k           FP n
|| instruction B ;

    instruction C ; EP k + 1     FP n
|| instruction D
|| instruction E

    instruction F ; EP k + 2     FP n
|| instruction G
|| instruction H

    instruction I ; EP k + 3     FP n + 1
|| instruction J
|| instruction K
|| instruction L
|| instruction M
|| instruction N
|| instruction O
|| instruction P

```

... continuing with EPs $k + 4$ through $k + 8$, which have eight instructions in parallel, like $k + 3$.

Figure 4-18. Pipeline Operation: Fetch Packets With Different Numbers of Execute Packets

		Clock cycle												
Fetch packet (FP)	Execute packet (EP)	1	2	3	4	5	6	7	8	9	10	11	12	13
n	k	PG	PS	PW	PR	DP	DC	E1	E2	E3	E4	E5		
n	k+1					DP	DC	E1	E2	E3	E4	E5		
n	k+2						DP	DC	E1	E2	E3	E4	E5	
n+1	k+3	PG	PS	PW	PR			DP	DC	E1	E2	E3	E4	
n+2	k+4		PG	PS	PW	Pipeline		PR	DP	DC	E1	E2	E3	
n+3	k+5			PG	PS	stall		PW	PR	DP	DC	E1	E2	
n+4	k+6				PG			PS	PW	PR	DP	DC	E1	
n+5	k+7							PG	PS	PW	PR	DP	DC	
n+6	k+8								PG	PS	PW	PR	DP	

In Figure 4-18, fetch packet n, which contains three execute packets, is shown followed by six fetch packets (n + 1 through n + 6), each with one execute packet (containing eight parallel instructions). The first fetch packet (n) goes through the program fetch phases during cycles 1-4. During these cycles, a program fetch phase is started for each of the fetch packets that follow.

In cycle 5, the program dispatch (DP) phase, the CPU scans the *p*-bits and detects that there are three execute packets (k through k + 2) in fetch packet n. This forces the pipeline to stall, which allows the DP phase to start for execute packets k + 1 and k + 2 in cycles 6 and 7. Once execute packet k + 2 is ready to move on to the DC phase (cycle 8), the pipeline stall is released.

The fetch packets n + 1 through n + 4 were all stalled so the CPU could have time to perform the DP phase for each of the three execute packets (k through k + 2) in fetch packet n. Fetch packet n + 5 was also stalled in cycles 6 and 7: it was not allowed to enter the PG phase until after the pipeline stall was released in cycle 8. The pipeline continues operation as shown with fetch packets n + 5 and n + 6 until another fetch packet containing multiple execution packets enters the DP phase, or an interrupt occurs.

4.3.2 Multicycle NOPs

The **NOP** instruction has an optional operand, *count*, that allows you to issue a single instruction for multicycle **NOPs**. A **NOP 2**, for example, fills in extra delay slots for the instructions in its execute packet and for all previous execute packets. If a **NOP 2** is in parallel with an **MPY** instruction, the **MPY** result is available for use by instructions in the next execute packet.

Figure 4-19 shows how a multicycle **NOP** drives the execution of other instructions in the same execute packet. Figure 4-19(a) shows a **NOP** in an execute packet (in parallel) with other code. The results of the **LD**, **ADD**, and **MPY** is available during the proper cycle for each instruction. Hence, **NOP** has no effect on the execute packet.

Figure 4-19(b) shows the replacement of the single-cycle **NOP** with a multicycle **NOP (NOP 5)** in the same execute packet. The **NOP 5** causes no operation to perform other than the operations from the instructions inside its execute packet. The results of the **LD**, **ADD**, and **MPY** cannot be used by any other instructions until the **NOP 5** period has completed.

Figure 4-19. Multicycle NOP in an Execute Packet

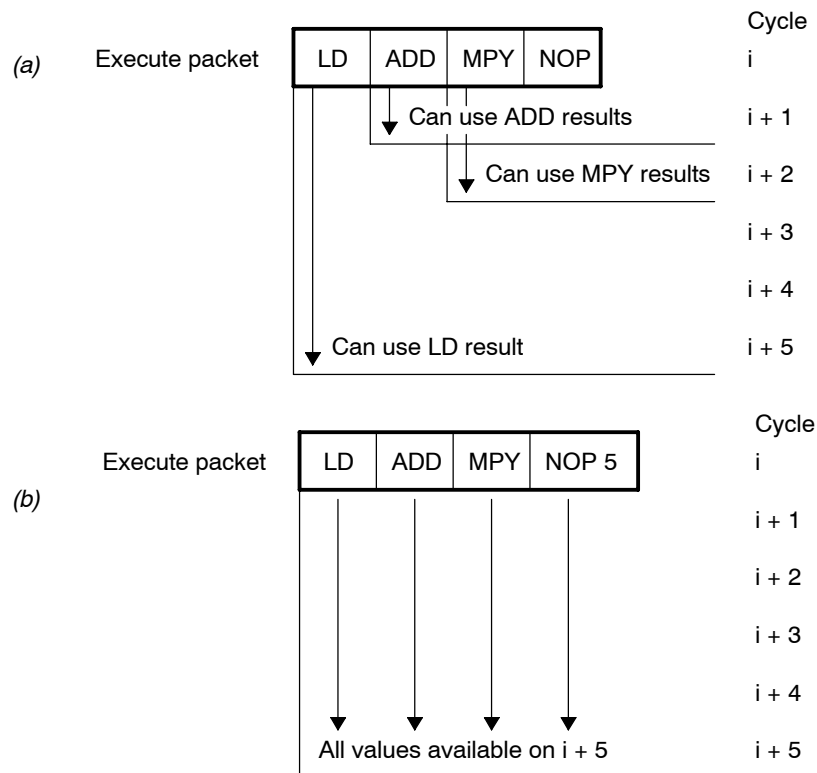
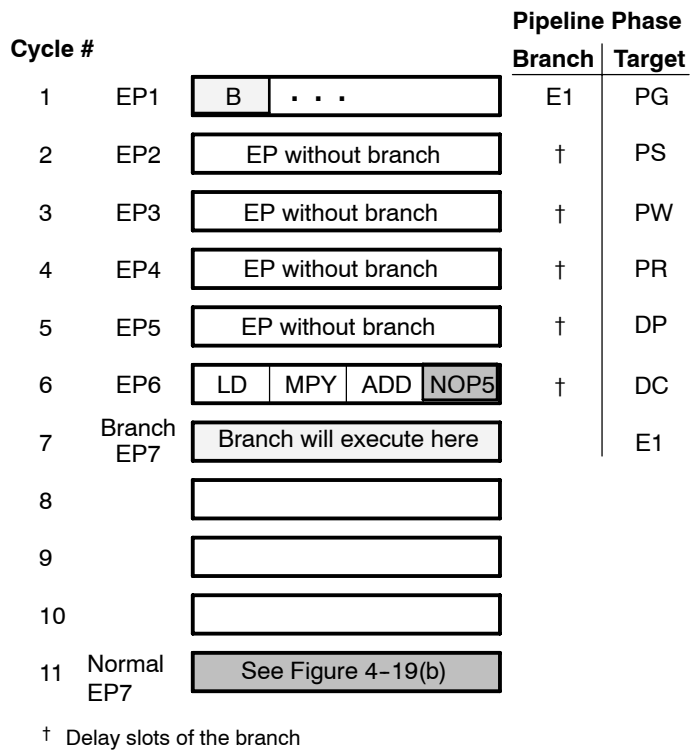


Figure 4-20 shows how a multicycle **NOP** can be affected by a branch. If the delay slots of a branch finish while a multicycle **NOP** is still dispatching **NOPs** into the pipeline, the branch overrides the multicycle **NOP** and the branch target begins execution five delay slots after the branch was issued.

Figure 4-20. Branching and Multicycle NOPs



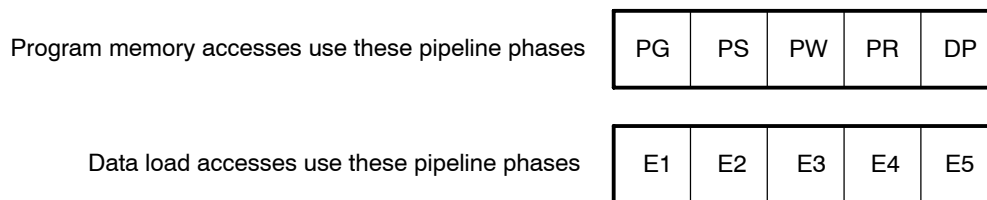
In one case, execute packet 1 (EP1) does not have a branch. The **NOP 5** in EP6 forces the CPU to wait until cycle 11 to execute EP7.

In the other case, EP1 does have a branch. The delay slots of the branch coincide with cycles 2 through 6. Once the target code reaches E1 in cycle 7, it executes.

4.3.3 Memory Considerations

The C62x DSP has a memory configuration with program memory in one physical space and data memory in another physical space. Data loads and program fetches have the same operation in the pipeline, they just use different phases to complete their operations. With both data loads and program fetches, memory accesses are broken into multiple phases. This enables the C62x DSP to access memory at a high speed. These phases are shown in Figure 4-21.

Figure 4-21. Pipeline Phases Used During Memory Accesses



To understand the memory accesses, compare data loads and instruction fetches/dispatches. The comparison is valid because data loads and program fetches operate on internal memories of the same speed on the C62x DSP and perform the same types of operations (listed in Table 4-3) to accommodate those memories. Table 4-3 shows the operation of program fetches pipeline versus the operation of a data load.

Table 4-3. Program Memory Accesses Versus Data Load Accesses

Operation	Program Memory Access Phase	Data Load Access Phase
Compute address	PG	E1
Send address to memory	PS	E2
Memory read/write	PW	E3
Program memory: receive fetch packet at CPU boundary Data load: receive data at CPU boundary	PR	E4
Program memory: send instruction to functional units Data load: send data to register	DP	E5

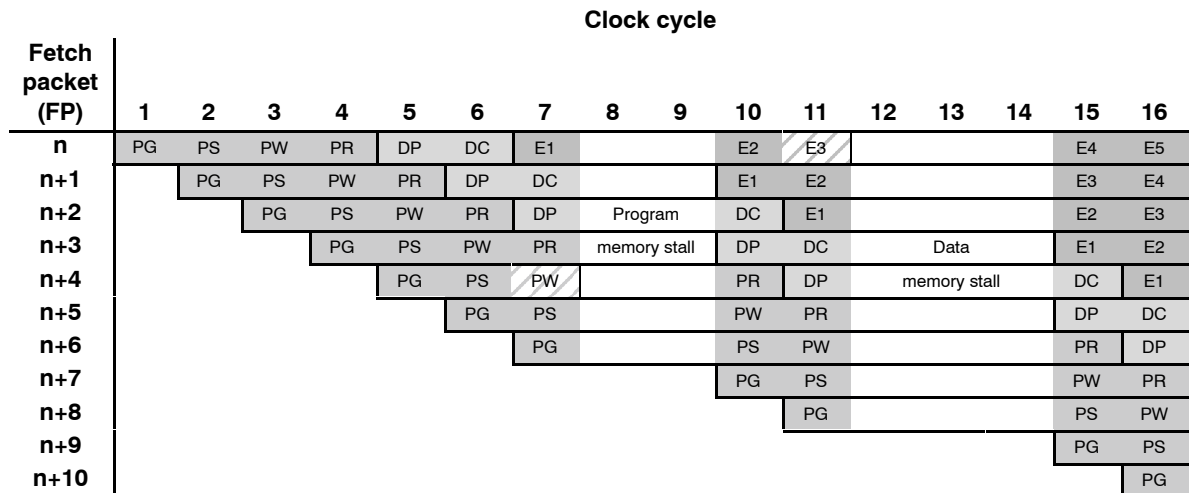
Depending on the type of memory and the time required to complete an access, the pipeline may stall to ensure proper coordination of data and instructions. This is discussed in section 4.3.3.1.

In the instance where multiple accesses are made to a single ported memory, the pipeline stalls to allow the extra access to occur. This is called a memory bank hit and is discussed in section 4.3.3.2.

4.3.3.1 Memory Stalls

A memory stall occurs when memory is not ready to respond to an access from the CPU. This access occurs during the PW phase for a program memory access and during the E3 phase for a data memory access. The memory stall causes all of the pipeline phases to lengthen beyond a single clock cycle, causing execution to take additional clock cycles to finish. The results of the program execution are identical whether a stall occurs or not. Figure 4-22 illustrates this point.


Figure 4-22. Program and Data Memory Stalls



4.3.3.2 Memory Bank Hits

Most C62x devices use an interleaved memory bank scheme, as shown in Figure 4-23; however, the C6211 DSP uses a two-level cache memory scheme. Each number in Figure 4-23 represents a byte address. A load byte (**LDB**) instruction from address 0 loads byte 0 in bank 0. A load halfword (**LDH**) from address 0 loads the halfword value in bytes 0 and 1, which are also in bank 0. A load word (**LDW**) from address 0 loads bytes 0 through 3 in banks 0 and 1.

Figure 4-23. 4-Bank Interleaved Memory



Because each of these banks is single-ported memory, only one access to each bank is allowed per cycle. Two accesses to a single bank in a given cycle result in a memory stall that halts all pipeline operation for one cycle, while the second value is read from memory. Two memory operations per cycle are allowed without any stall, as long as they do not access the same bank.

Consider the code in Example 4-2. Because both loads are trying to access the same bank at the same time, one load must wait. The first **LDW** accesses bank 0 on cycle $i + 2$ (in the E3 phase) and the second **LDW** accesses bank 0 on cycle $i + 3$ (in the E3 phase). See Table 4-4 for identification of cycles and phases. The E4 phase for both LDW instructions is in cycle $i + 4$. To eliminate this extra phase, the loads must access data from different banks (B4 address would need to be in bank 1). For more information on programming topics, see the *TMS320C6000 Programmer's Guide* (SPRU198).

Example 4-2. Load From Memory Banks

	LDW	.D1	*A4++,A5	;	load 1, A4 address is in bank 0
	LDW	.D2	*B4++,B5	;	load 2, B4 address is in bank 0

Table 4-4. Loads in Pipeline From Example 4-2

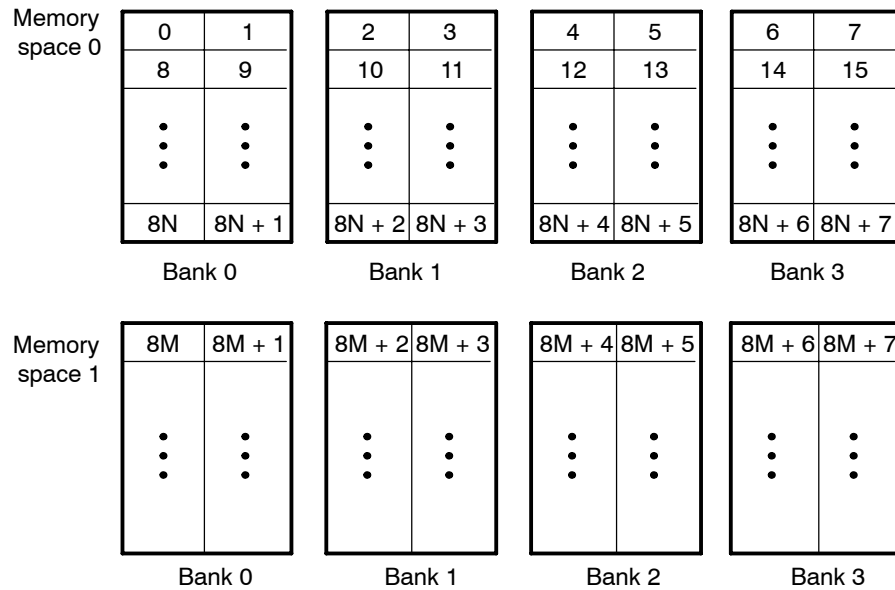
	i	$i + 1$	$i + 2$	$i + 3$	$i + 4$	$i + 5$
LDW .D1 Bank 0	E1	E2	E3	†	E4	E5
LDW .D2 Bank 0	E1	E2	†	E3	E4	E5

† Stall due to memory bank hit

For devices that have more than one memory space (see Figure 4-24), an access to bank 0 in one space does not interfere with an access to bank 0 in another memory space, and no pipeline stall occurs.

The internal memory of the C62x family of DSPs varies from device to device. See your device-specific datasheet to determine the memory spaces in your device.

Figure 4-24. 4-Bank Interleaved Memory With Two Memory Spaces



Interrupts

This chapter describes CPU interrupts, including reset and the nonmaskable interrupt (NMI). It details the related CPU control registers and their functions in controlling interrupts. It also describes interrupt processing, the method the CPU uses to detect automatically the presence of interrupts and divert program execution flow to your interrupt service code. Finally, the chapter describes the programming implications of interrupts.

Topic	Page
5.1 Overview	5-2
5.2 Globally Enabling and Disabling Interrupts	5-10
5.3 Individual Interrupt Control	5-12
5.4 Interrupt Detection and Processing	5-15
5.5 Performance Considerations	5-20
5.6 Programming Considerations	5-21

5.1 Overview

Typically, DSPs work in an environment that contains multiple external asynchronous events. These events require tasks to be performed by the DSP when they occur. An interrupt is an event that stops the current process in the CPU so that the CPU can attend to the task needing completion because of the event. These interrupt sources can be on chip or off chip, such as timers, analog-to-digital converters, or other peripherals.

Servicing an interrupt involves saving the context of the current process, completing the interrupt task, restoring the registers and the process context, and resuming the original process. There are eight registers that control servicing interrupts.

An appropriate transition on an interrupt pin sets the pending status of the interrupt within the interrupt flag register (IFR). If the interrupt is properly enabled, the CPU begins processing the interrupt and redirecting program flow to the interrupt service routine.

5.1.1 Types of Interrupts and Signals Used

There are three types of interrupts on the C6000 CPU.

- Reset
- Maskable
- Nonmaskable

These three types are differentiated by their priorities, as shown in Table 5-1. The reset interrupt has the highest priority and corresponds to the $\overline{\text{RESET}}$ signal. The nonmaskable interrupt (NMI) has the second highest priority and corresponds to the NMI signal. The lowest priority interrupts are interrupts 4-15 corresponding to the INT4-INT15 signals. $\overline{\text{RESET}}$, NMI, and some of the INT4-INT15 signals are mapped to pins on C6000 devices. Some of the INT4-INT15 interrupt signals are used by internal peripherals and some may be unavailable or can be used under software control. Check your device-specific data manual to see your interrupt specifications.

Table 5-1. Interrupt Priorities

Priority	Interrupt Name	Interrupt Type
Highest	Reset	Reset
	NMI	Nonmaskable
	INT4	Maskable
	INT5	Maskable
	INT6	Maskable
	INT7	Maskable
	INT8	Maskable
	INT9	Maskable
	INT10	Maskable
	INT11	Maskable
	INT12	Maskable
	INT13	Maskable
	INT14	Maskable
	Lowest	INT15

5.1.1.1 Reset ($\overline{\text{RESET}}$)

Reset is the highest priority interrupt and is used to halt the CPU and return it to a known state. The reset interrupt is unique in a number of ways:

- $\overline{\text{RESET}}$ is an active-low signal. All other interrupts are active-high signals.
- $\overline{\text{RESET}}$ must be held low for 10 clock cycles before it goes high again to reinitialize the CPU properly.
- The instruction execution in progress is aborted and all registers are returned to their default states.
- The reset interrupt service fetch packet must be located at a specific address which is specific to the specific device. See the device data manual for more information.
- $\overline{\text{RESET}}$ is not affected by branches.

5.1.1.2 Nonmaskable Interrupt (NMI)

NMI is the second-highest priority interrupt and is generally used to alert the CPU of a serious hardware problem such as imminent power failure.

For NMI processing to occur, the nonmaskable interrupt enable (NMIE) bit in the interrupt enable register must be set to 1. If NMIE is set to 1, the only condition that can prevent NMI processing is if the NMI occurs during the delay slots of a branch (whether the branch is taken or not).

NMIE is cleared to 0 at reset to prevent interruption of the reset. It is cleared at the occurrence of an NMI to prevent another NMI from being processed. You cannot manually clear NMIE, but you can set NMIE to allow nested NMIs. While NMI is cleared, all maskable interrupts (INT4–INT15) are disabled.

5.1.1.3 Maskable Interrupts (INT4-INT15)

The CPUs of the C6000 DSPs have 12 interrupts that are maskable. These have lower priority than the NMI and reset interrupts. These interrupts can be associated with external devices, on-chip peripherals, software control, or not be available.

Assuming that a maskable interrupt does not occur during the delay slots of a branch (this includes conditional branches that do not complete execution due to a false condition), the following conditions must be met to process a maskable interrupt:

- The global interrupt enable bit (GIE) bit in the control status register (CSR) is set to 1.
- The NMIE bit in the interrupt enable register (IER) is set to 1.
- The corresponding interrupt enable (IE) bit in the IER is set to 1.
- The corresponding interrupt occurs, which sets the corresponding bit in the interrupt flags register (IFR) to 1 and there are no higher priority interrupt flag (IF) bits set in the IFR.

5.1.1.4 Interrupt Acknowledgment (IACK) and Interrupt Number (INUM_n)

The IACK and INUM_n signals alert hardware external to the C6000 that an interrupt has occurred and is being processed. The IACK signal indicates that the CPU has begun processing an interrupt. The INUM_n signal (INUM3–INUM0) indicates the number of the interrupt (bit position in the IFR) that is being processed. For example:

INUM3 = 0 (MSB)
INUM2 = 1
INUM1 = 1
INUM0 = 1 (LSB)

Together, these signals provide the 4-bit value 0111, indicating INT7 is being processed.

5.1.2 Interrupt Service Table (IST)

When the CPU begins processing an interrupt, it references the interrupt service table (IST). The IST is a table of fetch packets that contain code for servicing the interrupts. The IST consists of 16 consecutive fetch packets. Each interrupt service fetch packet (ISFP) contains eight instructions. A simple interrupt service routine may fit in an individual fetch packet.

The addresses and contents of the IST are shown in Figure 5-1. Because each fetch packet contains eight 32-bit instruction words (or 32 bytes), each address in the table is incremented by 32 bytes (20h) from the one adjacent to it.

Figure 5-1. Interrupt Service Table

000h	RESET ISFP
020h	NMI ISFP
040h	Reserved
060h	Reserved
080h	INT4 ISFP
0A0h	INT5 ISFP
0C0h	INT6 ISFP
0E0h	INT7 ISFP
100h	INT8 ISFP
120h	INT9 ISFP
140h	INT10 ISFP
160h	INT11 ISFP
180h	INT12 ISFP
1A0h	INT13 ISFP
1C0h	INT14 ISFP
1E0h	INT15 ISFP

Program memory

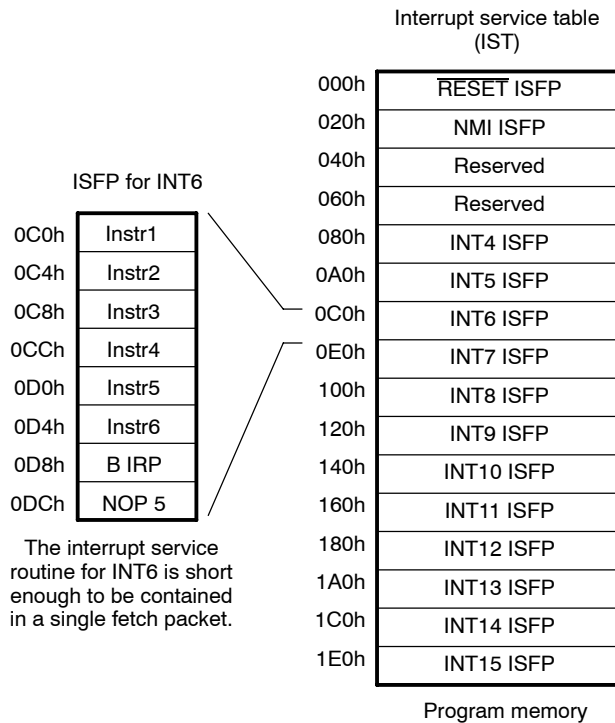
5.1.2.1 Interrupt Service Fetch Packet (ISFP)

An ISFP is a fetch packet used to service an interrupt. Figure 5–2 shows an ISFP that contains an interrupt service routine small enough to fit in a single fetch packet (FP). To branch back to the main program, the FP contains a branch to the interrupt return pointer instruction (**B IRP**). This is followed by a **NOP 5** instruction to allow the branch target to reach the execution stage of the pipeline.

Note:

If the **NOP 5** was not in the routine, the CPU would execute the next five execute packets (some of which are likely to be associated with the next ISFP) because of the delay slots associated with the **B IRP** instruction. See section 4.2.5 for more information.

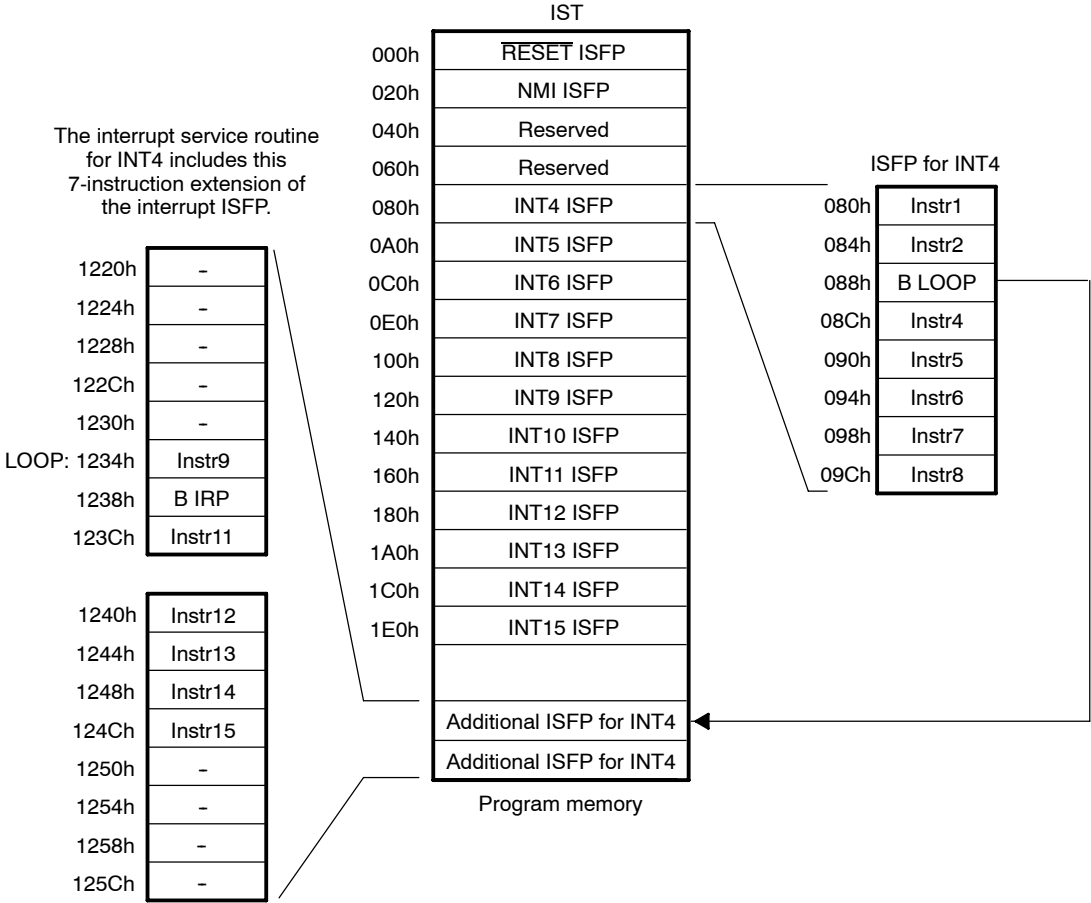
Figure 5-2. Interrupt Service Fetch Packet



If the interrupt service routine for an interrupt is too large to fit in a single fetch packet, a branch to the location of additional interrupt service routine code is required. Figure 5-3 shows that the interrupt service routine for INT4 was too large for a single fetch packet, and a branch to memory location 1234h is required to complete the interrupt service routine.

Note:
The instruction **B LOOP** branches into the middle of a fetch packet and processes code starting at address 1234h. The CPU ignores code from address 1220h-1230h, even if it is in parallel to code at address 1234h.

Figure 5-3. Interrupt Service Table With Branch to Additional Interrupt Service Code Located Outside the IST



5.1.2.2 Interrupt Service Table Pointer (ISTP)

The reset fetch packet must be located at the default location (see device data manual for more information), but the rest of the IST can be at any program memory location that is on a 256-word boundary (that is, any 1K byte boundary). The location of the IST is determined by the interrupt service table base (ISTB) field of the interrupt service table pointer register (ISTP). The ISTP is shown in Figure 2-11 (page 2-21) and described in Table 2-12 (page 2-21). Example 5-1 shows the relationship of the ISTB to the table location.

Example 5-1. Relocation of Interrupt Service Table

(a) Relocating the IST to 800h

1) Copy IST, located between 0h and 200h, to the memory location between 800h and A00h.

2) Write 800h to ISTP: MVK 800h, A2
 MVC A2, ISTP

ISTP = 800h = 1000 0000 0000b

(b) How the ISTP directs the CPU to the appropriate ISFP in the relocated IST

Assume: IFR = BBC0h = 1011 1011 1100 0000b
 IER = 1230h = 0001 0010 0011 0001b

2 enabled interrupts pending: INT9 and INT12

The 1s in IFR indicate pending interrupts; the 1s in IER indicate the interrupts that are enabled. INT9 has a higher priority than INT12, so HPEINT is encoded with the value for INT9, 01001b.

HPEINT corresponds to bits 9-5 of the ISTP:
ISTP = 1001 0010 0000b = 920h = address of INT9

IST	
0	RESET ISFP
800h	RESET ISFP
820h	NMI ISFP
840h	Reserved
860h	Reserved
880h	INT4 ISFP
8A0h	INT5 ISFP
8C0h	INT6 ISFP
8E0h	INT7 ISFP
900h	INT8 ISFP
920h	INT9 ISFP
940h	INT10 ISFP
96h0	INT11 ISFP
980h	INT12 ISFP
9A0h	INT13 ISFP
9C0h	INT14 ISFP
9E0h	INT15 ISFP

Program memory

5.1.3 Summary of Interrupt Control Registers

Table 5–2 lists the interrupt control registers on the C62x CPU.

Table 5-2. Interrupt Control Registers

Acronym	Register Name	Description	Page
CSR	Control status register	Allows you to globally set or disable interrupts	2-13
ICR	Interrupt clear register	Allows you to clear flags in the IFR manually	2-16
IER	Interrupt enable register	Allows you to enable interrupts	2-17
IFR	Interrupt flag register	Shows the status of interrupts	2-18
IRP	Interrupt return pointer register	Contains the return address used on return from a maskable interrupt. This return is accomplished via the B IRP instruction.	2-19
ISR	Interrupt set register	Allows you to set flags in the IFR manually	2-20
ISTP	Interrupt service table pointer register	Pointer to the beginning of the interrupt service table	2-21
NRP	Nonmaskable interrupt return pointer register	Contains the return address used on return from a nonmaskable interrupt. This return is accomplished via the B NRP instruction.	2-22

5.2 Globally Enabling and Disabling Interrupts

The control status register (CSR) contains two fields that control interrupts: GIE and PGIE, as shown in Figure 2-4 (page 2-13) and described in Table 2-7 (page 2-14). The global interrupt enable (GIE) bit allows you to enable or disable all maskable interrupts:

- GIE = 1 enables the maskable interrupts so that they are processed.
- GIE = 0 disables the maskable interrupts so that they are not processed.

The CPU detects interrupts in parallel with instruction execution. As a result, the CPU may begin interrupt processing in the same cycle that an **MVC** instruction writes 0 to GIE to disable interrupts. The PGIE bit (bit 1 of CSR) records the value of GIE after the CPU begins interrupt processing, recording whether the program was in the process of disabling interrupts.

During maskable interrupt processing, the CPU finishes executing the current execute packet. The CPU then copies the current value of GIE to PGIE, overwriting the previous value of PGIE. The CPU then clears GIE to prevent another maskable interrupt from occurring before the handler saves the machine's state. (Section 5.6.2 discusses nesting interrupts.)

When the interrupt handler returns to the interrupted code with the **B IRP** instruction, the CPU copies PGIE back to GIE. When the interrupted code resumes, GIE reflects the last value written by the interrupted code.

Because interrupt detection occurs in parallel with CPU execution, the CPU can take an interrupt in the cycle immediately following an **MVC** instruction that clears GIE. The behavior of PGIE and the **B IRP** instruction ensures, however, that interrupts do not occur after subsequent execute packets. Consider the code in Example 5-2.

Example 5-2. Interrupts Versus Writes to GIE

```

;Assume GIE = 1
MVC   CSR,B0           ;(1) Get CSR
AND   -2,B0,B0        ;(2) Get ready to clear GIE
MVC   B0,CSR          ;(3) Clear GIE
ADD   A0,A1,A2        ;(4)
ADD   A3,A4,A5        ;(5)
```

In Example 5-2, the CPU may service an interrupt between instructions 1 and 2, between instructions 2 and 3, or between instructions 3 and 4. The CPU will not service an interrupt between instructions 4 and 5.

If the CPU services an interrupt between instructions 1 and 2 or between instructions 2 and 3, the PGIE bit will hold the value 1 when arriving at the interrupt service routine. If the CPU services an interrupt between instructions 3 and 4, the PGIE bit will hold the value 0. Thus, when the interrupt service routine resumes the interrupted code, it will resume with GIE set as the interrupted code intended.

Programs must directly manipulate the GIE bit in CSR to disable and enable interrupts. Example 5-3 and Example 5-4 show code examples for disabling and enabling maskable interrupts globally, respectively.

Example 5-3. Code Sequence to Disable Maskable Interrupts Globally

```
MVC   CSR,B0      ; get CSR
AND   -2,B0,B0   ; get ready to clear GIE
MVC   B0,CSR      ; clear GIE
```

Example 5-4. Code Sequence to Enable Maskable Interrupts Globally

```
MVC   CSR,B0      ; get CSR
OR    1,B0,B0     ; get ready to set GIE
MVC   B0,CSR      ; set GIE
```

5.3 Individual Interrupt Control

Servicing interrupts effectively requires individual control of all three types of interrupts: reset, nonmaskable, and maskable. Enabling and disabling individual interrupts is done with the interrupt enable register (IER). The status of pending interrupts is stored in the interrupt flag register (IFR). Manual interrupt processing can be accomplished through the use of the interrupt set register (ISR) and interrupt clear register (ICR). The interrupt return pointers restore context after servicing nonmaskable and maskable interrupts.

5.3.1 Enabling and Disabling Interrupts

You can enable and disable individual interrupts by setting and clearing bits in the IER that correspond to the individual interrupts. An interrupt can trigger interrupt processing only if the corresponding bit in the IER is set. Bit 0, corresponding to reset, is not writeable and is always read as 1, so the reset interrupt is always enabled. You cannot disable the reset interrupt. Bits IE4–IE15 can be written as 1 or 0, enabling or disabling the associated interrupt, respectively. The IER is shown in Figure 2–7 (page 2-17) and described in Table 2–9.

When NMIE = 0, all nonreset interrupts are disabled, preventing interruption of an NMI. The NMIE bit is cleared at reset to prevent any interruption of process or initialization until you enable NMI. After reset, you must set the NMIE bit to enable the NMI and to allow INT15–INT4 to be enabled by the GIE bit in CSR and the corresponding IER bit. You cannot manually clear the NMIE bit; the NMIE bit is unaffected by a write of 0. The NMIE bit is also cleared by the occurrence of an NMI. If cleared, the NMIE bit is set only by completing a **B NRP** instruction or by a write of 1 to the NMIE bit. Example 5–5 and Example 5–6 show code for enabling and disabling individual interrupts, respectively.

Example 5-5. Code Sequence to Enable an Individual Interrupt (INT9)

MVK	200h,B1	; set bit 9
MVC	IER,B0	; get IER
OR	B1,B0,B0	; get ready to set IE9
MVC	B0,IER	; set bit 9 in IER

Example 5-6. Code Sequence to Disable an Individual Interrupt (INT9)

MVK	FDFh,B1	; clear bit 9
MVC	IER,B0	
AND	B1,B0,B0	; get ready to clear IE9
MVC	B0,IER	; clear bit 9 in IER

5.3.2 Status of Interrupts

The interrupt flag register (IFR) contains the status of INT4–INT15 and NMI. Each interrupt's corresponding bit in IFR is set to 1 when that interrupt occurs; otherwise, the bits have a value of 0. If you want to check the status of interrupts, use the **MVC** instruction to read IFR. The IFR is shown in Figure 2–8 (page 2-18) and described in Table 2–10.

5.3.3 Setting and Clearing Interrupts

The interrupt set register (ISR) and the interrupt clear register (ICR) allow you to set or clear maskable interrupts manually in IFR. Writing a 1 to IS4–IS15 in ISR causes the corresponding interrupt flag to be set in IFR. Similarly, writing a 1 to a bit in ICR causes the corresponding interrupt flag to be cleared. Writing a 0 to any bit of either ISR or ICR has no effect. Incoming interrupts have priority and override any write to ICR. You cannot set or clear any bit in ISR or ICR to affect NMI or reset. The ISR is shown in Figure 2–10 (page 2-20) and described in Table 2–11. The ICR is shown in Figure 2–6 (page 2-16) and described in Table 2–8.

Note:

Any write to the ISR or ICR (by the **MVC** instruction) effectively has one delay slot because the results cannot be read (by the **MVC** instruction) in IFR until two cycles after the write to ISR or ICR.

Any write to ICR is ignored by a simultaneous write to the same bit in ISR.

Example 5–7 and Example 5–8 show code examples to set and clear individual interrupts, respectively.

Example 5–7. Code to Set an Individual Interrupt (INT6) and Read the Flag Register

```
MVK    40h, B3
MVC    B3, ISR
NOP
MVC    IFR, B4
```

Example 5–8. Code to Clear an Individual Interrupt (INT6) and Read the Flag Register

```
MVK    40h, B3
MVC    B3, ICR
NOP
MVC    IFR, B4
```

5.3.4 Returning From Interrupt Servicing

After $\overline{\text{RESET}}$ goes high, the control registers are brought to a known value and program execution begins at address 0h. After nonmaskable and maskable interrupt servicing, use a branch to the corresponding return pointer register to continue the previous program execution.

5.3.4.1 CPU State After $\overline{\text{RESET}}$

After $\overline{\text{RESET}}$, the control registers and bits contain the following values:

- AMR, ISR, ICR, and IFR = 0h
- ISTP = Default value varies by device (See data manual for correct value)
- IER = 1h
- IRP and NRP = undefined
- CSR bits 15-0 = 100h in little-endian mode
= 000h in big-endian mode

The program execution begins at the address specified by the ISTB field in ISTP.

5.3.4.2 Returning From Nonmaskable Interrupts

The NMI return pointer register (NRP), shown in Figure 2-12 (page 2-22), contains the return pointer that directs the CPU to the proper location to continue program execution after NMI processing. A branch using the address in NRP (**B NRP**) in your interrupt service routine returns to the program flow when NMI servicing is complete. Example 5-9 shows how to return from an NMI.

Example 5-9. Code to Return From NMI

B	NRP	; return, sets NMIE
NOP	5	; delay slots

5.3.4.3 Returning From Maskable Interrupts

The interrupt return pointer register (IRP), shown in Figure 2-9 (page 2-19), contains the return pointer that directs the CPU to the proper location to continue program execution after processing a maskable interrupt. A branch using the address in IRP (**B IRP**) in your interrupt service routine returns to the program flow when interrupt servicing is complete. Example 5-10 shows how to return from a maskable interrupt.

Example 5-10. Code to Return from a Maskable Interrupt

B	IRP	; return, moves PGIE to GIE
NOP	5	; delay slots

5.4 Interrupt Detection and Processing

When an interrupt occurs, it sets a flag in the interrupt flag register (IFR). Depending on certain conditions, the interrupt may or may not be processed. This section discusses the mechanics of setting the flag bit, the conditions for processing an interrupt, and the order of operation for detecting and processing an interrupt. The similarities and differences between reset and nonreset interrupts are also discussed.

5.4.1 Setting the Nonreset Interrupt Flag

Figure 5-4 shows the processing of a nonreset interrupt (INT_m). The flag (IF_m) for INT_m in the IFR is set following the low-to-high transition of the INT_m signal on the CPU boundary. This transition is detected on a clock-cycle by clock-cycle basis and is not affected by memory stalls that might extend a CPU cycle. Once there is a low-to-high transition on an external interrupt pin (cycle 1), it takes two clock cycles for the signal to reach the CPU boundary (cycle 3). When the interrupt signal enters the CPU, it is has been detected (cycle 4). Two clock cycles after detection, the interrupt's corresponding flag bit in the IFR is set (cycle 6).

In Figure 5-4, IF_m is set during CPU cycle 6. You could attempt to clear IF_m by using an **MVC** instruction to write a 1 to bit *m* of the ICR in execute packet *n* + 3 (during CPU cycle 4). However, in this case, the automated write by the interrupt detection logic takes precedence and IF_m remains set.

Figure 5-4 assumes INT_m is the highest-priority pending interrupt and is enabled by the GIE and NMIE bits, as necessary. If it is not the highest-priority pending interrupt, IF_m remains set until either you clear it by writing a 1 to bit *m* of ICR or the processing of INT_m occurs.

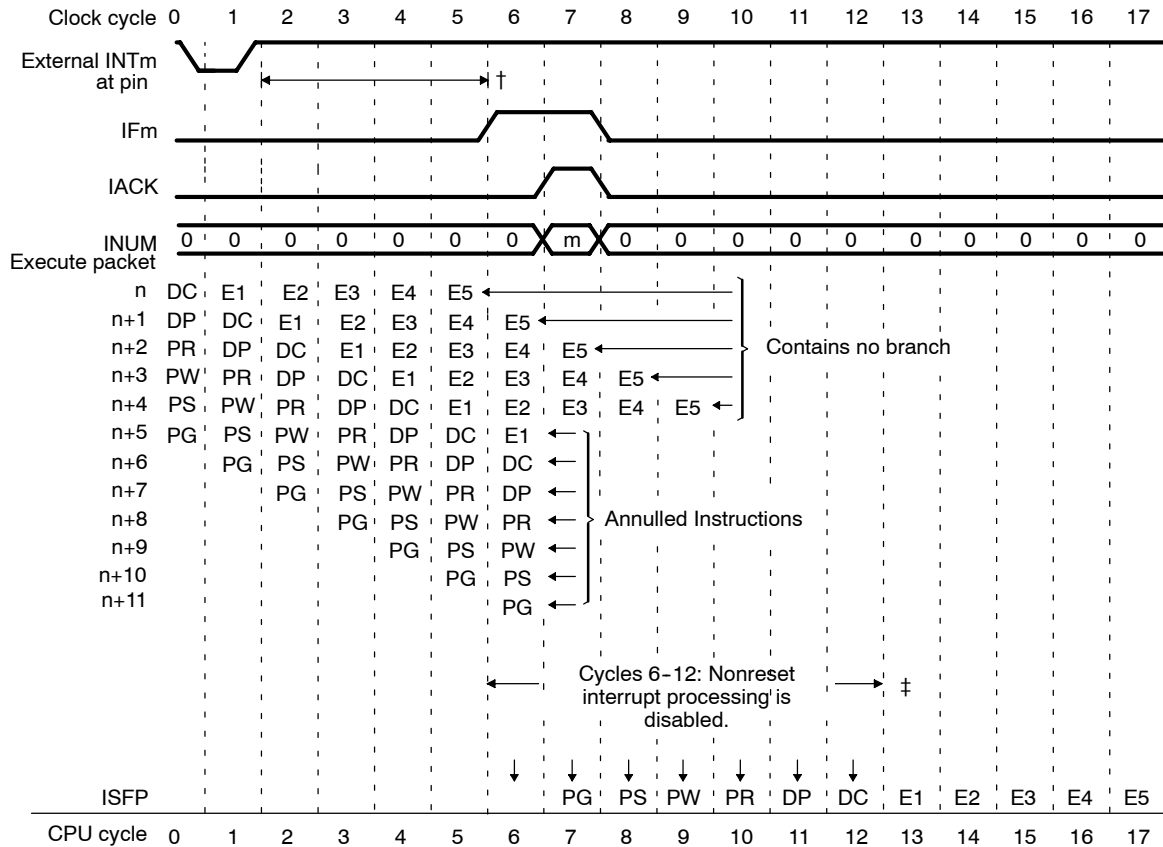
5.4.2 Conditions for Processing a Nonreset Interrupt

In clock cycle 4 of Figure 5-4, a nonreset interrupt in need of processing is detected. For this interrupt to be processed, the following conditions must be valid on the same clock cycle and are evaluated every clock cycle:

- IF_m is set during CPU cycle 6. (This determination is made in CPU cycle 4 by the interrupt logic.)
- There is not a higher priority IF_m bit set in the IFR.
- The corresponding bit in the IER is set (IE_m = 1).
- GIE = 1
- NMIE = 1
- The five previous execute packets (*n* through *n* + 4) do not contain a branch (even if the branch is not taken) and are not in the delay slots of a branch.

Any pending interrupt will be taken as soon as pending branches are completed.

Figure 5-4. Nonreset Interrupt Detection and Processing: Pipeline Operation



† IFm is set on the next CPU cycle boundary after a 4-clock cycle delay after the rising edge of INTm.

‡ After this point, interrupts are still disabled. All nonreset interrupts are disabled when NMIE = 0. All maskable interrupts are disabled when GIE = 0.

5.4.3 Actions Taken During Nonreset Interrupt Processing


During CPU cycles 6 through 12 of Figure 5–4, the following interrupt processing actions occur:

- Processing of subsequent nonreset interrupts is disabled.
- For all interrupts except NMI, the PGIE bit is set to the value of the GIE bit and then the GIE bit is cleared.
- For NMI, the NMIE bit is cleared.
- The next execute packets (from $n + 5$ on) are annulled. If an execute packet is annulled during a particular pipeline stage, it does not modify any CPU state. Annulling also forces an instruction to be annulled in future pipeline stages.
- The address of the first annulled execute packet ($n + 5$) is loaded in NRP (in the case of NMI) or IRP (for all other interrupts).
- A branch to the address held in ISTP (the pointer to the ISFP for INTm) is forced into the E1 phase of the pipeline during cycle 7.
- During cycle 7, IACK is asserted and the proper INUM n signals are asserted to indicate which interrupt is being processed. The timings for these signals in Figure 5–4 represent only the signals' characteristics inside the CPU. The external signals may be delayed and be longer in duration to handle external devices. Check the device-specific datasheet for your timing values.
- IFm is cleared during cycle 8.

5.4.4 Setting the $\overline{\text{RESET}}$ Interrupt Flag

$\overline{\text{RESET}}$ must be held low for a minimum of 10 clock cycles. Four clock cycles after $\overline{\text{RESET}}$ goes high, processing of the reset vector begins. The flag for $\overline{\text{RESET}}$ (IF0) in the IFR is set by the low-to-high transition of the $\overline{\text{RESET}}$ signal on the CPU boundary. In Figure 5-5, IF0 is set during CPU cycle 15. This transition is detected on a clock-cycle by clock-cycle basis and is not affected by memory stalls that might extend a CPU cycle.

Figure 5-5. $\overline{\text{RESET}}$ Interrupt Detection and Processing: Pipeline Operation



† IF0 is set on the next CPU cycle boundary after a 4-clock cycle delay after the rising edge of $\overline{\text{RESET}}$.

‡ After this point, interrupts are still disabled. All nonreset interrupts are disabled when NMIE = 0. All maskable interrupts are disabled when GIE = 0.

5.4.5 Actions Taken During $\overline{\text{RESET}}$ Interrupt Processing

A low signal on the $\overline{\text{RESET}}$ pin is the only requirement to process a reset. Once $\overline{\text{RESET}}$ makes a high-to-low transition, the pipeline is flushed and CPU registers are returned to their reset values. The GIE bit, the NMIE bit, and the ISTB bits in ISTP are cleared. For the CPU state after reset, see section 5.3.4.1.

During CPU cycles 15 through 21 of Figure 5-5, the following reset processing actions occur:

- Processing of subsequent nonreset interrupts is disabled because the GIE and NMIE bits are cleared.
- A branch to the address held in ISTP (the pointer to the ISFP for INT0) is forced into the E1 phase of the pipeline during cycle 16.
- During cycle 16, IACK is asserted and the proper INUM n signals are asserted to indicate a reset is being processed.
- IF0 is cleared during cycle 17.

Note:

Code that starts running after reset must explicitly enable the GIE bit, the NMIE bit, and IER to allow interrupts to be processed.

5.5 Performance Considerations

The interaction of the C62x CPU and sources of interrupts present performance issues for you to consider when you are developing your code.

5.5.1 General Performance

- ❑ **Overhead.** Overhead for all CPU interrupts is 7 cycles. You can see this in Figure 5–4, where no new instructions are entering the E1 pipeline phase during CPU cycles 6 through 12.
- ❑ **Latency.** Interrupt latency is 11 cycles (21 cycles for `RESET`). In Figure 5–4, although the interrupt is active in cycle 2, execution of interrupt service code does not begin until cycle 13.
- ❑ **Frequency.** The logic clears the nonreset interrupt (IFm) on cycle 8, with any incoming interrupt having highest priority. Thus, an interrupt is can be recognized every second cycle. Also, because a low-to-high transition is necessary, an interrupt can occur only every second cycle. However, the frequency of interrupt processing depends on the time required for interrupt service and whether you reenables interrupts during processing, thereby allowing nested interrupts. Effectively, only two occurrences of a specific interrupt can be recognized in two cycles.

5.5.2 Pipeline Interaction

Because the serial or parallel encoding of fetch packets does not affect the DC and subsequent phases of the pipeline, no conflicts between code parallelism and interrupts exist. There are three operations or conditions that can affect or are affected by interrupts:

- ❑ **Branches.** Nonreset interrupts are delayed, if any execute packets n through $n + 4$ in Figure 5–4 contain a branch or are in the delay slots of a branch.
- ❑ **Memory stalls.** Memory stalls delay interrupt processing, because they inherently extend CPU cycles.
- ❑ **Multicycle NOPs.** Multicycle NOPs (including the `IDLE` instruction) operate like other instructions when interrupted, except when an interrupt causes annulment of any but the first cycle of a multicycle `NOP`. In that case, the address of the next execute packet in the pipeline is saved in NRP or IRP. This prevents returning to an `IDLE` instruction or a multicycle `NOP` that was interrupted.

5.6 Programming Considerations

The interaction of the C62x CPUs and sources of interrupts present programming issues for you to consider when you are developing your code.

5.6.1 Single Assignment Programming

Using the same register to store different variables (called here: multiple assignment) can result in unpredictable operation when the code can be interrupted.

To avoid unpredictable operation, you must employ the single assignment method in code that can be interrupted. When an interrupt occurs, all instructions entering E1 prior to the beginning of interrupt processing are allowed to complete execution (through E5). All other instructions are annulled and refetched upon return from interrupt. The instructions encountered after the return from the interrupt do not experience any delay slots from the instructions prior to processing the interrupt. Thus, instructions with delay slots prior to the interrupt can appear, to the instructions after the interrupt, to have fewer delay slots than they actually have.

Example 5-11 shows a code fragment which stores two variables into A1 using multiple assignment. Example 5-12 shows equivalent code using the single assignment programming method which stores the two variables into two different registers.

For example, suppose that register A1 contains 0 and register A0 points to a memory location containing a value of 10 before reaching the code in Example 5-11. The **ADD** instruction, which is in a delay slot of the **LDW**, sums A2 with the value in A1 (0) and the result in A3 is just a copy of A2. If an interrupt occurred between the **LDW** and **ADD**, the **LDW** would complete the update of A1 (10), the interrupt would be processed, and the **ADD** would sum A1 (10) with A2 and place the result in A3 (equal to A2 + 10). Obviously, this situation produces incorrect results.

In Example 5-12, the single assignment method is used. The register A1 is assigned only to the **ADD** input and not to the result of the **LDW**. Regardless of the value of A6 with or without an interrupt, A1 does not change before it is summed with A2. Result A3 is equal to A2.

Example 5-11. Code Without Single Assignment: Multiple Assignment of A1

LDW	.D1	*A0, A1	
ADD	.L1	A1, A2, A3	
NOP	3		
MPY	.M1	A1, A4, A5	; uses new A1

Example 5-12. Code Using Single Assignment

LDW	.D1	*A0,A6	
ADD	.L1	A1,A2,A3	
NOP	3		
MPY	.M1	A6,A4,A5	; uses A6

5.6.2 Nested Interrupts

Generally, when the CPU enters an interrupt service routine, interrupts are disabled. However, when the interrupt service routine is for one of the maskable interrupts (INT4–INT15), an NMI can interrupt processing of the maskable interrupt. In other words, an NMI can interrupt a maskable interrupt, but neither an NMI nor a maskable interrupt can interrupt an NMI.

There may be times when you want to allow an interrupt service routine to be interrupted by another (particularly higher priority) interrupt. Even though the processor by default does not allow interrupt service routines to be interrupted unless the source is an NMI, it is possible to nest interrupts under software control. To allow nested interrupts, the interrupt service routine must perform the following initial steps in addition to its normal work of saving any registers (including control registers) that it modifies:

- 1) The contents of IRP (or NRP) must be saved
- 2) The contents of the PGIE bit must be saved
- 3) The GIE bit must be set to 1

Prior to returning from the interrupt service routine, the code must restore the registers saved above as follows:

- 1) The GIE bit must be first cleared to 0
- 2) The PGIE bit saved value must be restored
- 3) The IRP (or NRP) saved value must be restored

Although steps 2 and 3 above may be performed in either order, it is important that the GIE bit is cleared first. This means that the GIE and PGIE bits must be restored with separate writes to CSR. If these bits are not restored separately, then it is possible that the PGIE bit is overwritten by nested interrupt processing just as interrupts are being disabled.

Example 5-13 shows a simple assembly interrupt handler that allows nested interrupts. This example saves its context to the compiler's stack, pointed to by B15. This assumes that the C runtime conventions are being followed. The example code is not optimized, to aid in readability.

Example 5-13. Assembly Interrupt Service Routine That Allows Nested Interrupts

```

_isr:
    STW    B0, *B15--[4]      ; Save B0, allocate 4 words of stack
    STW    B1, *B15[1]       ; Save B1 on stack

    MVC    IRP, B0
    STW    B0, *B15[2]       ; Save IRP on stack

    MVC    CSR, B0
    STW    B0, *B15[3]       ; Save CSR (and thus PGIE) on stack

    OR     B0, 1, B1
    MVC    B1, CSR           ; Enable interrupts

    ; Interrupt service code goes here.
    ; Interrupts may occur while this code executes.

    MVC    CSR, B0           ;\
    AND    B0, -2, B1        ; |-- Disable interrupts.
    MVC    B1, CSR           ;/   (Set GIE to 0)

    LDW    *B15[3], B0       ; get saved value of CSR into B0
    NOP    4                 ; wait for LDW *B15[3] to finish
    MVC    B0, CSR           ; Restore PGIE

    LDW    *B15[2], B0       ; get saved value of IRP into B1
    NOP    4                 ;
    MVC    B0, IRP          ; Restore IRP

    B      IRP               ; Return from interrupt
|| LDW    *B15[1], B1       ; Restore B1

    LDW    *++B15[4], B0     ; Restore B0, release stack.
    NOP    4                 ; wait for B IRP and LDW to complete.

```

Example 5-14 shows a C-based interrupt handler that allows nested interrupts. The steps are similar, although the compiler takes care of allocating the stack and saving CPU registers. For more information on using C to access control registers and write interrupt handlers, see the *TMS320C6000 Optimizing C Compiler Users Guide*, SPRU187.

Example 5-14 uses the interrupt keyword along with explicit context save and restore code. An alternative is to use the DSP/BIOS interrupt dispatcher that also provides an easy way to nest interrupt service routines.

Example 5-14. C Interrupt Service Routine That Allows Nested Interrupts

```

/* c6x.h contains declarations of the C6x control registers          */
#include <c6x.h>

interrupt void isr(void)
{
    unsigned old_csr;
    unsigned old_irp;

    old_irp = IRP          ;/* Save IRP          */
    old_csr = CSR          ;/* Save CSR (and thus PGIE) */

    CSR = old_csr | 1     ;/* Enable interrupts */

    /* Interrupt service code goes here. */
    /* Interrupts may occur while this code executes */

    CSR = CSR & -2       ;/* Disable interrupts */
    CSR = old_csr        ;/* Restore CSR (and thus PGIE) */
    IRP = old_irp        ;/* Restore IRP          */
}

```

5.6.3 Manual Interrupt Processing

You can poll the IFR and IER to detect interrupts manually and then branch to the value held in the ISTP as shown below in Example 5-15.

The code sequence begins by copying the address of the highest priority interrupt from the ISTP to the register B2. The next instruction extracts the number of the interrupt, which is used later to clear the interrupt. The branch to the interrupt service routine comes next with a parallel instruction to set up the ICR word.

The last five instructions fill the delay slots of the branch. First, the 32-bit return address is stored in the B2 register and then copied to the interrupt return pointer (IRP). Finally, the number of the highest priority interrupt, stored in B1, is used to shift the ICR word in B1 to clear the interrupt.

Example 5-15. Manual Interrupt Processing

```

MVC      ISTP,B2          ; get related ISFP address
EXTU    B2,23,27,B1     ; extract HPEINT
[B1]    B                ; branch to interrupt
|| [B1]  MVK 1,A0        ; setup ICR word
[B1]    MVK RET_ADR,B2   ; create return address
[B1]    MVKH RET_ADR,B2  ;
[B1]    MVC B2,IRP       ; save return address
[B1]    SHL A0,B1,B1     ; create ICR word
[B1]    MVC B1,ICR       ; clear interrupt flag
RET_ADR: (Post interrupt service routine Code)

```

5.6.4 Traps

A trap behaves like an interrupt, but is created and controlled with software. The trap condition can be stored in any one of the conditional registers: A1, A2, B0, B1, or B2. If the trap condition is valid, a branch to the trap handler routine processes the trap and the return.

Example 5-16 and Example 5-17 show a trap call and the return code sequence, respectively. In the first code sequence, the address of the trap handler code is loaded into register B0 and the branch is called. In the delay slots of the branch, the context is saved in the B0 register, the GIE bit is cleared to disable maskable interrupts, and the return pointer is stored in the B1 register. If the trap handler were within the 21-bit offset for a branch using a displacement, the **MVKH** instructions could be eliminated, thus shortening the code sequence.

The trap is processed with the code located at the address pointed to by the label TRAP_HANDLER. If the B0 or B1 registers are needed in the trap handler, their contents must be stored to memory and restored before returning. The code shown in Example 5-17 should be included at the end of the trap handler code to restore the context prior to the trap and return to the TRAP_RETURN address.

Example 5-16. Code Sequence to Invoke a Trap

[A1]	MVK	TRAP_HANDLER,B0	; load 32-bit trap address
[A1]	MVKH	TRAP_HANDLER,B0	
[A1]	B	B0	; branch to trap handler
[A1]	MVC	CSR,B0	; read CSR
[A1]	AND	-2,B0,B1	; disable interrupts: GIE=0
[A1]	MVC	B1,CSR	; write to CSR
[A1]	MVK	TRAP_RETURN,B1	; load 32-bit return address
[A1]	MVKH	TRAP_RETURN,B1	
TRAP_RETURN:		(post-trap code)	

Note: A1 contains the trap condition.

Example 5-17. Code Sequence for Trap Return

B	B1	; return
MVC	B0,CSR	; restore CSR
NOP	4	; delay slots

Mapping Between Instruction and Functional Unit

Table A-1 lists the instructions that execute on each functional unit.

Table A-1. Instruction to Functional Unit Mapping

Instruction	Page	Functional Unit			
		.L Unit	.M Unit	.S Unit	.D Unit
ABS	3-24	✓			
ADD	3-26	✓		✓	✓
ADDAB	3-30				✓
ADDAH	3-32				✓
ADDAW	3-34				✓
ADDK	3-36			✓	
ADDU	3-37	✓			
ADD2	3-39			✓	
AND	3-41	✓		✓	
B displacement	3-43			✓	
B register	3-45			✓†	
B IRP	3-47			✓†	
B NRP	3-49			✓†	
CLR	3-51			✓	
CMPEQ	3-54	✓			
CMPGT	3-56	✓			

† S2 only
‡ D2 only

Table A-1. Instruction to Functional Unit Mapping (Continued)

Instruction	Page	Functional Unit			
		.L Unit	.M Unit	.S Unit	.D Unit
CMPGTU	3-59	✓			
CMPLT	3-61	✓			
CMPLTU	3-64	✓			
EXT	3-66			✓	
EXTU	3-69			✓	
IDLE	3-72		No unit		
LDB memory	3-73				✓
LDB memory (15-bit offset)	3-76				✓†
LDBU memory	3-73				✓
LDBU memory (15-bit offset)	3-76				✓†
LDH memory	3-78				✓
LDH memory (15-bit offset)	3-81				✓†
LDHU memory	3-78				✓
LDHU memory (15-bit offset)	3-81				✓†
LDW memory	3-83				✓
LDW memory (15-bit offset)	3-86				✓†
LMBD	3-88	✓			
MPY	3-90		✓		
MPYH	3-92		✓		
MPYHL	3-93		✓		
MPYHLU	3-94		✓		
MPYHSLU	3-95		✓		
MPYHSU	3-96		✓		
MPYHU	3-97		✓		

† S2 only

‡ D2 only

Table A-1. Instruction to Functional Unit Mapping (Continued)

Instruction	Page	Functional Unit			
		.L Unit	.M Unit	.S Unit	.D Unit
MPYHULS	3-98		✓		
MPYHUS	3-99		✓		
MPYLH	3-100		✓		
MPYLHU	3-101		✓		
MPYLSHU	3-102		✓		
MPYLUHS	3-103		✓		
MPYSU	3-104		✓		
MPYU	3-106		✓		
MPYUS	3-107		✓		
MV	3-108	✓		✓	✓
MVC	3-110			✓†	
MVK	3-113			✓	
MVKH	3-115			✓	
MVKL	3-117			✓	
MVKLH	3-115			✓	
NEG	3-119	✓		✓	
NOP	3-120			No unit	
NORM	3-122	✓			
NOT	3-124	✓		✓	
OR	3-125	✓		✓	
SADD	3-127	✓			
SAT	3-130	✓			
SET	3-132			✓	

† S2 only
‡ D2 only

Table A-1. Instruction to Functional Unit Mapping (Continued)

Instruction	Page	Functional Unit			
		.L Unit	.M Unit	.S Unit	.D Unit
SHL	3-135			✓	
SHR	3-137			✓	
SHRU	3-139			✓	
SMPY	3-141		✓		
SMPYH	3-143		✓		
SMPYHL	3-144		✓		
SMPYLH	3-146		✓		
SSHL	3-148			✓	
SSUB	3-150	✓			
STB memory	3-152				✓
STB memory (15-bit offset)	3-154				✓†
STH memory	3-156				✓
STH memory (15-bit offset)	3-159				✓†
STW memory	3-161				✓
STW memory (15-bit offset)	3-163				✓†
SUB	3-165	✓		✓	✓
SUBAB	3-168				✓
SUBAH	3-170				✓
SUBAW	3-171				✓
SUBC	3-173	✓			
SUBU	3-175	✓			
SUB2	3-177			✓	
XOR	3-179	✓		✓	✓
ZERO	3-181	✓		✓	✓

† S2 only

‡ D2 only

1X and 2X paths 2-6

A

A4 MODE bits 2-10

A5 MODE bits 2-10

A6 MODE bits 2-10

A7 MODE bits 2-10

ABS instruction 3-24

absolute value, with saturation (ABS) 3-24

actions taken during nonreset interrupt processing 5-17

actions taken during $\overline{\text{RESET}}$ interrupt processing 5-19

add

signed 16-bit constant to register (ADDK) 3-36

two 16-bit integers on upper and lower register halves (ADD2) 3-39

using byte addressing mode (ADDAB) 3-30

using halfword addressing mode (ADDAH) 3-32

using word addressing mode (ADDAW) 3-34

with saturation, two signed integers (SADD) 3-127

without saturation

two signed integers (ADD) 3-26

two unsigned integers (ADDU) 3-37

ADD instruction 3-26

add instructions

using circular addressing 3-18

using linear addressing 3-16

ADD2 instruction 3-39

ADDAB instruction 3-30

ADDAH instruction 3-32

ADDAW instruction 3-34

ADDK instruction 3-36

address generation for load/store 3-18

address paths 2-7

addressing mode 3-16

circular mode 3-17

linear mode 3-16

addressing mode register (AMR) 2-10

ADDU instruction 3-37

AMR 2-10

AND instruction 3-41

applications, TMS320 DSP family 1-3

architecture, TMS320C62x DSP 1-6

arithmetic shift left (SHL) 3-135

arithmetic shift right (SHR) 3-137

B

B instruction

using a displacement 3-43

using a register 3-45

B IRP instruction 3-47

B NRP instruction 3-49

B4 MODE bits 2-10

B5 MODE bits 2-10

B6 MODE bits 2-10

B7 MODE bits 2-10

bit field

clear (CLR) 3-51

extract and sign-extend a bit field (EXT) 3-66

extract and zero-extend a bit field (EXTU) 3-69

set (SET) 3-132

bitwise AND (AND) 3-41

bitwise exclusive OR (XOR) 3-179

bitwise NOT (NOT) 3-124

bitwise OR (OR) 3-125

BK0 bits 2-10

BK1 bits 2-10

- block diagram
 - branch instructions 4-17
 - decode pipeline phases 4-4
 - execute pipeline phases 4-5
 - fetch pipeline phases 4-3
 - load instructions 4-15
 - multiply instructions 4-13
 - pipeline phases 4-8
 - single-cycle instructions 4-12
 - store instructions 4-14
 - TMS320C62x CPU data path 2-3
 - TMS320C62x DSP 1-6
 - block size calculations 2-12
 - branch
 - using a displacement (B) 3-43
 - using a register (B) 3-45
 - using an interrupt return pointer (B IRP) 3-47
 - using NMI return pointer (B NRP) 3-49
 - branch instruction
 - block diagram 4-17
 - pipeline operation 4-17
 - branching
 - into the middle of an execute packet 3-9
 - performance considerations 5-20
 - to additional interrupt service routine 5-7
- C**
- circular addressing, block size calculations 2-12
 - circular addressing mode
 - add instructions 3-18
 - block size specification 3-17
 - load instructions 3-17
 - store instructions 3-17
 - subtract instructions 3-18
 - clear a bit field (CLR) 3-51
 - clear an individual interrupt 5-13
 - clearing interrupts 5-13
 - CLR instruction 3-51
 - CMPEQ instruction 3-54
 - CMPGT instruction 3-56
 - CMPGTU instruction 3-59
 - CMPLT instruction 3-61
 - CMPLTU instruction 3-64
 - compare
 - for equality, signed integers (CMPEQ) 3-54
 - for greater than
 - signed integers (CMPGT) 3-56
 - unsigned integers (CMPGTU) 3-59
 - for less than
 - signed integers (CMPLT) 3-61
 - unsigned integers (CMPLTU) 3-64
 - compare for equality, signed integers (CMPEQ) 3-54
 - compare for greater than
 - signed integers (CMPGT) 3-56
 - unsigned integers (CMPGTU) 3-59
 - compare for less than
 - signed integers (CMPLT) 3-61
 - unsigned integers (CMPLTU) 3-64
 - conditional operations 3-10
 - conditional subtract and shift (SUBC) 3-173
 - conditions for processing a nonreset interrupt 5-15
 - constraints
 - on cross paths 3-11
 - on instructions using the same functional unit 3-11
 - on loads and stores 3-12
 - on long data 3-13
 - on register reads 3-14
 - on register writes 3-15
 - control, individual interrupts 5-12
 - control register, interrupts 5-9
 - control status register (CSR) 2-13
 - CPU
 - control register file 2-7
 - data paths 2-3
 - functional units 2-5
 - general-purpose register files 2-2
 - introduction 1-7
 - load and store paths 2-6
 - CPU data paths
 - relationship to register files 2-6
 - TMS320C62x DSP 2-3
 - CPU ID bits 2-13
 - cross paths 2-6
 - CSR 2-13

D

DA1 and DA2 2-7
 data address paths 2-7
 DC pipeline phase 4-3
 DCC bits 2-13
 decoding instructions 4-3
 delay slots 3-6
 disabling an individual interrupt 5-12
 disabling maskable interrupts globally 5-11
 DP pipeline phase 4-3

E

E1 phase program counter (PCE1) 2-22
 E1-E5 pipeline phases 4-5
 EN bit 2-13
 enabling an individual interrupt 5-12
 enabling maskable interrupts globally 5-11
 execute packet, pipeline operation 4-18
 execution notations 3-2
 EXT instruction 3-66
 extract and sign-extend a bit field (EXT) 3-66
 extract and zero-extend a bit field (EXTU) 3-69
 EXTU instruction 3-69

F

features, TMS320C62x DSP 1-4
 fetch packet 3-7, 5-6
 fetch packets
 fully parallel 3-8
 fully serial 3-8
 partially serial 3-9
 fetch pipeline phase 4-2
 functional unit to instruction mapping A-1
 functional units 2-5

G

general-purpose register files
 cross paths 2-6
 data address paths 2-7
 description 2-2
 memory, load, and store paths 2-6
 GIE bit 2-13

H

HPEINT bits 2-21

I

ICn bit 2-16
 ICR 2-16
 IDLE instruction 3-72
 IEn bit 2-17
 IER 2-17
 IFn bit 2-18
 IFR 2-18
 individual interrupt control 5-12
 instruction compatibility 3-20
 instruction descriptions 3-20
 instruction operation, notations 3-2
 instruction to functional unit mapping A-1
 instruction types
 branch instructions 4-17
 load instructions 4-15
 multiply instructions 4-13
 single-cycle 4-12
 store instructions 4-13
 two-cycle 4-13
 interleaved memory bank scheme 4-24
 interrupt clear register (ICR) 2-16
 interrupt detection and processing 5-15
 interrupt enable register (IER) 2-17
 interrupt flag register (IFR) 2-18
 interrupt return pointer register (IRP) 2-19
 interrupt service fetch packet (ISFP) 5-6
 interrupt service table (IST) 5-5
 interrupt service table pointer (ISTP), overview 5-8
 interrupt service table pointer register (ISTP) 2-21
 interrupt set register (ISR) 2-20

interrupts

- clearing 5-13
- control 5-12
- control registers 5-9
- detection 5-15
- detection and processing
 - actions taken during nonreset interrupt processing* 5-17
 - actions taken during RESET interrupt processing* 5-19
 - conditions for processing a nonreset interrupt* 5-15
 - setting the nonreset interrupt flag* 5-15
 - setting the RESET interrupt flag* 5-18
- disabling 5-12
- enabling 5-12
- global control 5-10
- globally disabling 5-10
- globally enabling 5-10
- manual interrupt processing 5-24
- overview 5-2
- performance considerations 5-20
 - frequency* 5-20
 - latency* 5-20
 - overhead* 5-20
 - pipeline interaction* 5-20
- pipeline interaction
 - branches* 5-20
 - code parallelism* 5-20
 - memory stalls* 5-20
 - multicycle NOPs* 5-20
- priorities 5-3
- processing 5-15
- programming considerations 5-21
 - nested interrupts* 5-22
 - single assignment* 5-21
 - traps* 5-25
- returning from interrupt servicing 5-14
- setting 5-13
- signals used 5-2
- status 5-13
- types of 5-2

invoking a trap 5-25

IRP 2-19

IRP bits 2-19

ISFP 5-6

ISn bit 2-20

ISR 2-20

IST 5-5

ISTB bits 2-21

ISTP 2-21

L

latency 3-6

LDB instruction

- 5-bit unsigned constant offset or register offset 3-73
- 15-bit unsigned constant offset 3-76

LDBU instruction

- 5-bit unsigned constant offset or register offset 3-73
- 15-bit constant offset 3-76

LDH instruction

- 5-bit unsigned constant offset or register offset 3-78
- 15-bit unsigned constant offset 3-81

LDHU instruction

- 5-bit unsigned constant offset or register offset 3-78
- 15-bit unsigned constant offset 3-81

LDW instruction

- 5-bit unsigned constant offset or register offset 3-83
- 15-bit unsigned constant offset 3-86

leftmost bit detection (LMBD) 3-88

linear addressing mode 3-16

- add instructions 3-16
- load instructions 3-16
- store instructions 3-16
- subtract instructions 3-16

LMBD instruction 3-88

load

- byte
 - from memory with a 5-bit unsigned constant offset or register offset (LDB and LDBU)* 3-73
 - from memory with a 15-bit unsigned constant offset (LDB and LDBU)* 3-76
- halfword
 - from memory with a 5-bit unsigned constant offset or register offset (LDH and LDHU)* 3-78
 - from memory with a 15-bit unsigned constant offset (LDH and LDHU)* 3-81

- load (continued)
 - word
 - from memory with a 5-bit unsigned constant offset or register offset (LDW)* 3-83
 - from memory with a 15-bit unsigned constant offset (LDW)* 3-86
 - load and store paths CPU 2-6
 - load instructions
 - block diagram 4-15
 - conflicts 3-12
 - pipeline operation 4-15
 - syntax for indirect addressing 3-18
 - using circular addressing 3-17
 - using linear addressing 3-16
 - load or store to the same memory location, rules 4-14
 - load paths 2-6
 - logical shift right (SHRU) 3-139
- M**
- memory
 - introduction 1-7
 - paths 2-6
 - memory bank hits 4-24
 - memory considerations 4-22
 - memory bank hits 4-24
 - memory stalls 4-23
 - memory paths 2-6
 - memory stalls 4-23
 - move
 - 16-bit constant into upper bits of register (MVKH and MVKLH) 3-115
 - between control file and register file (MVC) 3-110
 - from register to register (MV) 3-108
 - signed constant into register and sign extend (MVK) 3-113
 - signed constant into register and sign extend (MVKL) 3-117
 - MPY instruction 3-90
 - MPYH instruction 3-92
 - MPYHL instruction 3-93
 - MPYHLU instruction 3-94
 - MPYHSLU instruction 3-95
 - MPYHSU instruction 3-96
 - MPYHU instruction 3-97
 - MPYHULS instruction 3-98
 - MPYHUS instruction 3-99
 - MPYLH instruction 3-100
 - MPYLHU instruction 3-101
 - MPYLSHU instruction 3-102
 - MPYLUHS instruction 3-103
 - MPYSU instruction 3-104
 - MPYU instruction 3-106
 - MPYUS instruction 3-107
 - multicycle NOP with no termination until interrupt (IDLE) 3-72
 - multicycle NOPs 4-20
 - multiply
 - signed by signed
 - signed 16 LSB by signed 16 LSB (MPY)* 3-90
 - signed 16 LSB by signed 16 LSB with left shift and saturation (SMPY)* 3-141
 - signed 16 LSB by signed 16 MSB (MPYLH)* 3-100
 - signed 16 LSB by signed 16 MSB with left shift and saturation (SMPYLH)* 3-146
 - signed 16 MSB by signed 16 LSB (MPYHL)* 3-93
 - signed 16 MSB by signed 16 LSB with left shift and saturation (SMPYHL)* 3-144
 - signed 16 MSB by signed 16 MSB (MPYH)* 3-92
 - signed 16 MSB by signed 16 MSB with left shift and saturation (SMPYH)* 3-143
 - signed by unsigned
 - signed 16 LSB by unsigned 16 LSB (MPYSU)* 3-104
 - signed 16 LSB by unsigned 16 MSB (MPYLSHU)* 3-102
 - signed 16 MSB by unsigned 16 LSB (MPYHSLU)* 3-95
 - signed 16 MSB by unsigned 16 MSB (MPYHSU)* 3-96
 - unsigned by signed
 - unsigned 16 LSB by signed 16 LSB (MPYUS)* 3-107
 - unsigned 16 LSB by signed 16 MSB (MPYLUHS)* 3-103
 - unsigned 16 MSB by signed 16 LSB (MPYHULS)* 3-98
 - unsigned 16 MSB by signed 16 MSB (MPYHUS)* 3-99

- multiply (continued)
 - unsigned by unsigned
 - unsigned 16 LSB by unsigned 16 LSB (MPYU)* 3-106
 - unsigned 16 LSB by unsigned 16 MSB (MPYLHU)* 3-101
 - unsigned 16 MSB by unsigned 16 LSB (MPYHLU)* 3-94
 - unsigned 16 MSB by unsigned 16 MSB (MPYHU)* 3-97
- multiply instructions
 - block diagram 4-13
 - pipeline operation 4-13
- MV instruction 3-108
- MVC instruction 3-110
- MVK instruction 3-113
- MVKH instruction 3-115
- MVKL instruction 3-117
- MVKLH instruction 3-115

N

- NEG instruction 3-119
- negate (NEG) 3-119
- nested interrupts 5-22
- NMI return pointer register (NRP) 2-22
- NMIE bit 2-17
- NMIF bit 2-18
- no operation (NOP) 3-120
- NOP instruction 3-120
- NORM instruction 3-122
- normalize integer (NORM) 3-122
- NOT instruction 3-124
- notational conventions iii
- NRP 2-22
- NRP bits 2-22

O

- opcode, fields and meanings 3-5
- operands, examples 3-21
- options, TMS320C62x DSP 1-4
- OR instruction 3-125

- overview
 - interrupts 5-2
 - pipeline 4-2
 - TMS320 DSP family 1-2
 - TMS320C6000 DSP family 1-2

P

- parallel code 3-9
- parallel fetch packets 3-8
- parallel operations 3-7
 - branch into the middle of an execute packet 3-9
 - parallel code 3-9
- partially serial fetch packets 3-9
- PCC bits 2-13
- PCE1 2-22
- PCE1 bits 2-22
- performance considerations
 - interrupts 5-20
 - pipeline 4-18
- PG pipeline phase 4-2
- PGIE bit 2-13
- pipeline
 - decode stage 4-3
 - execute stage 4-5
 - execution 4-11
 - fetch stage 4-2
 - overview 4-2
 - performance considerations 4-18
 - phases 4-2
 - stages 4-2
 - summary 4-6
- pipeline execution 4-11
- pipeline operation
 - branch instructions 4-17
 - load instructions 4-15
 - multiple execute packets in a fetch packet 4-18
 - multiply instructions 4-13
 - one execute packet per fetch packet 4-6
 - single-cycle instructions 4-12
 - store instructions 4-13
 - two-cycle instructions 4-13
- pipeline phases
 - block diagram 4-8
 - used during memory accesses 4-22
- PR pipeline phase 4-2
- programming considerations, interrupts 5-21
- PS pipeline phase 4-2

PW pipeline phase 4-2
 PWRD bits 2-13

R

register files
 cross paths 2-6
 data address paths 2-7
 general-purpose 2-2
 memory, load, and store paths 2-6
 relationship to data paths 2-6

registers
 addresses for accessing 2-8
 addressing mode register (AMR) 2-10
 control register file 2-7
 control status register (CSR) 2-13
 E1 phase program counter (PCE1) 2-22
 interrupt clear register (ICR) 2-16
 interrupt enable register (IER) 2-17
 interrupt flag register (IFR) 2-18
 interrupt return pointer register (IRP) 2-19
 interrupt service table pointer register (ISTP) 2-21
 interrupt set register (ISR) 2-20
 NMI return pointer register (NRP) 2-22
 read constraints 3-14
 write constraints 3-15

related documentation from Texas Instruments iii

resource constraints 3-11
 cross paths 3-11
 on loads and stores 3-12
 on long data 3-13
 on register reads 3-14
 on register writes 3-15
 using the same functional unit 3-11

return from NMI 5-14
 returning from a trap 5-25
 returning from interrupt servicing 5-14
 returning from maskable interrupts 5-14
 returning from nonmaskable interrupts 5-14
 REVISION ID bits 2-13

S

SADD instruction 3-127
 SAT bit 2-13
 SAT instruction 3-130

saturate a 40-bit integer to a 32-bit integer (SAT) 3-130

serial fetch packets 3-8

set a bit field (SET) 3-132

set an individual interrupt 5-13

SET instruction 3-132

setting interrupts 5-13

setting the nonreset interrupt flag 5-15

setting the $\overline{\text{RESET}}$ interrupt flag 5-18

shift
 arithmetic shift left (SHL) 3-135
 arithmetic shift right (SHR) 3-137
 logical shift right (SHRU) 3-139
 shift left with saturation (SSHL) 3-148

shift left with saturation (SSHL) 3-148

SHL instruction 3-135

SHR instruction 3-137

SHRU instruction 3-139

single-cycle instructions
 block diagram 4-12
 pipeline operation 4-12

SMPY instruction 3-141

SMPYH instruction 3-143

SMPYHL instruction 3-144

SMPYLH instruction 3-146

SSHL instruction 3-148

SSUB instruction 3-150

STB instruction
 5-bit unsigned constant offset or register offset 3-152
 15-bit unsigned constant offset 3-154

STH instruction
 5-bit unsigned constant offset or register offset 3-156
 15-bit unsigned constant offset 3-159

store
 byte
 to memory with a 5-bit unsigned constant offset or register offset (STB) 3-152
 to memory with a 15-bit unsigned constant offset (STB) 3-154
 halfword
 to memory with a 5-bit unsigned constant offset or register offset (STH) 3-156
 to memory with a 15-bit unsigned constant offset (STH) 3-159

store (continued)
word
 to memory with a 5-bit unsigned constant offset or register offset (STW) 3-161
 to memory with a 15-bit unsigned constant offset (STW) 3-163

store instructions
 block diagram 4-14
 conflicts 3-12
 pipeline operation 4-13
 syntax for indirect addressing 3-18
 using circular addressing 3-17
 using linear addressing 3-16

store or load to the same memory location, rules 4-14

store paths 2-6

STW instruction
 5-bit unsigned constant offset or register offset 3-161
 15-bit unsigned constant offset 3-163

SUB instruction 3-165

SUB2 instruction 3-177

SUBAB instruction 3-168

SUBAH instruction 3-170

SUBAW instruction 3-171

SUBC instruction 3-173

subtract
 conditionally and shift (SUBC) 3-173
 two 16-bit integers on upper and lower register halves (SUB2) 3-177
 using byte addressing mode (SUBAB) 3-168
 using halfword addressing mode (SUBAH) 3-170
 using word addressing mode (SUBAW) 3-171
 with saturation, two signed integers (SSUB) 3-150
 without saturation
 two signed integers (SUB) 3-165
 two unsigned integers (SUBU) 3-175

subtract instructions
 using circular addressing 3-18
 using linear addressing 3-16

SUBU instruction 3-175

syntax, fields and meanings 3-5

T

TMS320 DSP family
 applications 1-3
 overview 1-2

TMS320C6000 DSP family, overview 1-2

TMS320C62x DSP
 architecture 1-6
 block diagram 1-6
 features 1-4
 options 1-4

trademarks iv

two 16-bit integers
 add on upper and lower register halves (ADD2) 3-39
 subtract on upper and lower register halves (SUB2) 3-177

two-cycle instructions, pipeline operation 4-13

V

VelociTI architecture 1-1

VLIW (very long instruction word) architecture 1-1

X

XOR instruction 3-179

Z

zero a register (ZERO) 3-181

ZERO instruction 3-181

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products		Applications	
Amplifiers	amplifier.ti.com	Audio	www.ti.com/audio
Data Converters	dataconverter.ti.com	Automotive	www.ti.com/automotive
DLP® Products	www.dlp.com	Communications and Telecom	www.ti.com/communications
DSP	dsp.ti.com	Computers and Peripherals	www.ti.com/computers
Clocks and Timers	www.ti.com/clocks	Consumer Electronics	www.ti.com/consumer-apps
Interface	interface.ti.com	Energy	www.ti.com/energy
Logic	logic.ti.com	Industrial	www.ti.com/industrial
Power Mgmt	power.ti.com	Medical	www.ti.com/medical
Microcontrollers	microcontroller.ti.com	Security	www.ti.com/security
RFID	www.ti-rfid.com	Space, Avionics & Defense	www.ti.com/space-avionics-defense
RF/IF and ZigBee® Solutions	www.ti.com/lprf	Video and Imaging	www.ti.com/video
		Wireless	www.ti.com/wireless-apps