

TMS320C55x DSP Programmer's Guide

Preliminary Draft

This document contains preliminary data current as of the publication date and is subject to change without notice.

SPRU376
April 2000



IMPORTANT NOTICE

Texas Instruments and its subsidiaries (TI) reserve the right to make changes to their products or to discontinue any product or service without notice, and advise customers to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgement, including those pertaining to warranty, patent infringement, and limitation of liability.

TI warrants performance of its semiconductor products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

CERTAIN APPLICATIONS USING SEMICONDUCTOR PRODUCTS MAY INVOLVE POTENTIAL RISKS OF DEATH, PERSONAL INJURY, OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE ("CRITICAL APPLICATIONS"). TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS. INCLUSION OF TI PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE FULLY AT THE CUSTOMER'S RISK.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards must be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance or customer product design. TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used. TI's publication of information regarding any third party's products or services does not constitute TI's approval, warranty or endorsement thereof.

Read This First

About This Manual

This manual describes ways to optimize C and assembly code for the TMS320C55x™ DSPs and recommends ways to write TMS320C55x code for specific applications.

Notational Conventions

This document uses the following conventions.

- The device number TMS320C55x is often abbreviated as C55x.
- Program listings, program examples, and interactive displays are shown in a special typeface similar to a typewriter's. Examples use a **bold version** of the special typeface for emphasis; interactive displays use a **bold version** of the special typeface to distinguish commands that you enter from items that the system displays (such as prompts, command output, error messages, etc.).

Here is a sample program listing:

```
0011 0005 0001      .field    1, 2
0012 0005 0003      .field    3, 4
0013 0005 0006      .field    6, 3
0014 0006           .even
```

Here is an example of a system prompt and a command that you might enter:

```
C:  csr -a /user/ti/simuboard/utilities
```

- In syntax descriptions, the instruction, command, or directive is in a **bold typeface** font and parameters are in an *italic typeface*. Portions of a syntax that are in **bold** should be entered as shown; portions of a syntax that are in *italics* describe the type of information that should be entered. Here is an example of a directive syntax:

.asect *"section name", address*

.asect is the directive. This directive has two parameters, indicated by *section name* and *address*. When you use **.asect**, the first parameter must be an actual section name, enclosed in double quotes; the second parameter must be an address.

- ❑ Some directives can have a varying number of parameters. For example, the `.byte` directive can have up to 100 parameters. The syntax for this directive is:

.byte *value*₁ [, ... , *value*_{*n*}]

This syntax shows that `.byte` must have at least one value parameter, but you have the option of supplying additional value parameters, separated by commas.

- ❑ In most cases, hexadecimal numbers are shown with the suffix `h`. For example, the following number is a hexadecimal 40 (decimal 64):

40h

Similarly, binary numbers usually are shown with the suffix `b`. For example, the following number is the decimal number 4 shown in binary form:

0100b

- ❑ Bits are sometimes referenced with the following notation:

Notation	Description	Example
Register(<i>n</i> – <i>m</i>)	Bits <i>n</i> through <i>m</i> of Register	AC0(15–0) represents the 16 least significant bits of the register AC0.

Related Documentation From Texas Instruments

The following books describe the TMS320C55x devices and related support tools. To obtain a copy of any of these TI documents, call the Texas Instruments Literature Response Center at (800) 477-8924. When ordering, please identify the book by its title and literature number.

TMS320C55x Technical Overview (SPRU393). This overview is an introduction to the TMS320C55x digital signal processor (DSP). The TMS320C55x is the latest generation of fixed-point DSPs in the TMS320C5000 DSP platform. Like the previous generations, this processor is optimized for high performance and low-power operation. This book describes the CPU architecture, low-power enhancements, and embedded emulation features of the TMS320C55x.

TMS320C55x DSP CPU Reference Guide (literature number SPRU371) describes the architecture, registers, and operation of the CPU. This book also describes how to make individual portions of the DSP inactive to save power.

TMS320C55x DSP Mnemonic Instruction Set Reference Guide (literature number SPRU374) describes the mnemonic instructions individually. It also includes a summary of the instruction set, a list of the instruction opcodes, and a cross-reference to the algebraic instruction set.

TMS320C55x DSP Algebraic Instruction Set Reference Guide (literature number SPRU375) describes the algebraic instructions individually. It also includes a summary of the instruction set, a list of the instruction opcodes, and a cross-reference to the mnemonic instruction set.

TMS320C55x Optimizing C Compiler User's Guide (literature number SPRU281) describes the 'C55x C compiler. This C compiler accepts ANSI standard C source code and produces assembly language source code for TMS320C55x devices.

TMS320C55x Assembly Language Tools User's Guide (literature number SPRU280) describes the assembly language tools (assembler, linker, and other tools used to develop assembly language code), assembler directives, macros, common object file format, and symbolic debugging directives for TMS320C55x devices.

The CPU, the registers, and the instruction sets are also described in online documentation contained in Code Composer Studio™.

Trademarks

Code Composer Studio, TMS320C54x, C54x, TMS320C55x, and C55x are trademarks of Texas Instruments.

Contents

1	Introduction	1-1
	<i>Lists some key features of the TMS320C55x DSP architecture and recommends a process for code development.</i>	
1.1	TMS320C55x Architecture	1-2
1.2	Code Development Flow for Best Performance	1-3
2	Tutorial	2-1
	<i>Uses example code to walk you through the code development flow for the TMS320C55x DSP.</i>	
2.1	Introduction	2-2
2.2	Writing Assembly Code	2-3
2.2.1	Allocate Sections for Code, Constants, and Variables	2-5
2.2.2	Processor Mode Initialization	2-7
2.2.3	Setting up Addressing Modes	2-8
2.3	Understanding the Linking Process	2-10
2.4	Building Your Program	2-14
2.4.1	Creating a Project	2-14
2.4.2	Adding Files to the Work space	2-14
2.4.3	Modifying Build Options	2-15
2.4.4	Build the Program	2-15
2.5	Testing your code	2-16
2.6	Benchmarking your code	2-18
3	Optimizing C Code	3-1
	<i>Describes how to optimize your C code using compiler options, intrinsics, and code transformations.</i>	
3.1	The Compiler and Its Options	3-2
3.2	Using Program-level Optimization	3-5
3.3	Using Function Inlining	3-10
3.3.1	Using <code>-oi<size></code> option	3-10
3.4	Using Intrinsics	3-12
3.5	Using Long Data Accesses for 16-Bit Data	3-16

3.6	Generating Efficient Loop Code	3-17
3.6.1	Avoid function calls in repeated loops	3-17
3.6.2	Keep loops small to enable local repeat	3-17
3.6.3	Trip Count Issues	3-18
3.6.4	Using Unsigned Integer Types for Trip Counter	3-19
3.6.5	Use <code>_nassert</code> Intrinsic	3-21
3.6.6	Use <code>-o3</code> and <code>-pm</code> Compiler Options	3-24
3.7	Generating Efficient Control Code	3-27
3.8	Efficient Math Operations	3-28
3.8.1	Use 16-bit Data Types Whenever Possible	3-28
3.8.2	Special Considerations When Using MAC Constructs	3-28
3.8.3	Avoid Using Modulus Operator When Simulating Circular Addressing in C	3-29
3.9	Memory Management	3-34
3.9.1	Avoiding Holes Caused by Data Alignment	3-34
3.9.2	Local vs Global Symbol Declarations	3-35
3.9.3	Stack Configuration	3-35
3.9.4	Allocating Code and Data in the TMS320C55x Memory Map	3-36
3.10	Allocating Function Code to Different Sections	3-40
4	Optimizing Assembly Code	4-1
	<i>Describes some of the opportunities for optimizing TMS320C55x assembly code and provides corresponding code examples.</i>	
4.1	Efficient Use of the Dual-MAC Hardware	4-2
4.1.1	Data Memory Pointer Usage	4-3
4.1.2	Multi-Channel Applications	4-3
4.1.3	Multi-Algorithm Applications	4-4
4.1.4	Implicit Algorithm Symmetry	4-5
4.1.5	Loop Unrolling	4-9
4.2	Using Parallel Execution Features	4-20
4.2.1	Built-In Parallelism	4-20
4.2.2	User-Defined Parallelism	4-20
4.2.3	Architectural Features Supporting Parallelism	4-21
4.2.4	User-Defined Parallelism Rules	4-24
4.2.5	Process for Implementing User-Defined Parallelism	4-26
4.2.6	Parallelism Tips	4-28
4.2.7	Examples of Parallel Optimization Within CPU Functional Units	4-29
4.2.8	Example of Parallel Optimization Across the A-Unit, P-Unit, and D-Unit	4-39
4.3	Implementing Efficient Loops	4-46
4.3.1	Nesting of Loops	4-46
4.3.2	Efficient Use of <code>repeat(CSR)</code> Looping	4-50
4.3.3	Avoiding Pipeline Delays When Accessing Loop-Control Registers	4-52
4.4	Minimizing Instruction Pipeline Delays	4-53
4.4.1	Process to Resolve Pipeline Conflicts	4-55
4.4.2	Recommendations for Preventing Pipeline Delays	4-56
4.4.3	Memory Accesses and the Pipeline	4-77

5	Fixed-Point Arithmetic	5-1
	<i>Explains important considerations for doing fixed-point arithmetic with the TMS320C55x DSP. Includes code examples.</i>	
5.1	Fixed-Point Arithmetic	5-2
5.1.1	2s-Complement Numbers	5-2
5.1.2	Integers Versus Fractions	5-3
5.1.3	2s-Complement Arithmetic	5-5
5.1.4	Extended-Precision 2s-Complement Arithmetic	5-8
5.2	Extended-Precision Addition and Subtraction	5-10
5.3	Extended-Precision Multiplication	5-17
5.4	Division	5-21
5.4.1	Integer Division	5-21
5.4.2	Fractional Division	5-30
5.5	Methods of Handling Overflows	5-31
6	Bit-Reversed Addressing	6-1
	<i>Introduces bit-reverse addressing and its implementation on the TMS320C55x DSP. Includes code examples.</i>	
6.1	Introduction to Bit-Reverse Addressing	6-2
6.2	Using Bit-Reverse Addressing In FFT Algorithms	6-4
6.3	In-Place Versus Off-Place Bit-Reversing	6-6
6.4	Using the C55x DSPLIB for FFTs and Bit-Reversing	6-8
7	Application-Specific Instructions	7-1
	<i>Explains how to implement some common DSP algorithms using specialized TMS320C55x instructions. Includes code examples.</i>	
7.1	Symmetric and Asymmetric FIR Filtering (FIRS, FIRSN)	7-2
7.1.1	Symmetric FIR Filtering With the <code>firs</code> Instruction	7-3
7.1.2	Antisymmetric FIR Filtering With the <code>firsn</code> Instruction	7-4
7.1.3	Implementation of a Symmetric FIR Filter on the TMS320C55x DSP	7-4
7.2	Adaptive Filtering (LMS)	7-6
7.2.1	Delayed LMS Algorithm	7-7
7.3	Convolutional Encoding (BFXPA, BFXTR)	7-10
7.3.1	Bit-Stream Multiplexing and Demultiplexing	7-11
7.4	Viterbi Algorithm for Channel Decoding (ADDSUB, SUBADD, MAXDIFF)	7-16

8	TI C55x DSPLIB	8-1
	<i>Introduces the features and the C functions of the TI TMS320C55x DSP function library.</i>	
8.1	Features and Benefits	8-2
8.2	DSPLIB Data Types	8-2
8.3	DSPLIB Arguments	8-2
8.4	Calling a DSPLIB Function from C	8-3
8.5	Calling a DSPLIB Function from Assembly Language Source Code	8-4
8.6	Where to Find Sample Code	8-4
8.7	DSPLIB Functions	8-5
8.7.1	Description of Arguments	8-5
8.7.2	List of DSPLIB Functions	8-5

Figures

1–1	Code Development Flow	1-3
2–1	Section Allocation	2-6
2–2	Extended Auxiliary Registers Structure (XARn)	2-9
4–1	Data Bus Usage During a Dual-MAC Operation	4-2
4–2	Computation Groupings for a Block FIR (4-Tap Filter Shown)	4-10
4–3	Computation Groupings for a Single-Sample FIR With an Even Number of TAPS (4-Tap Filter Shown)	4-17
4–4	Computation Groupings for a Single-Sample FIR With an Odd Number of TAPS (5-Tap Filter Shown)	4-18
4–5	Matrix to Find Operators That Can Be Used in Parallel	4-22
4–6	CPU Operators and Buses	4-24
4–7	Process for Applying User-Defined Parallelism	4-27
4–8	Second Segment of the Pipeline (Execution Pipeline)	4-53
5–1	4-Bit 2s-Complement Integer Representation	5-4
5–2	8-Bit 2s-Complement Integer Representation	5-4
5–3	4-Bit 2s-Complement Fractional Representation	5-5
5–4	8-Bit 2s-Complement Fractional Representation	5-5
5–5	32-Bit Addition	5-11
5–6	32-Bit Subtraction	5-14
5–7	32-Bit Multiplication	5-17
6–1	FFT Flow Graph Showing Bit-Reversed Input and In-Order Output	6-4
7–1	Symmetric and Antisymmetric FIR Filters	7-3
7–2	Adaptive FIR Filter Implemented With the Least-Mean-Squares (LMS) Algorithm	7-6
7–3	Example of a Convolutional Encoder	7-10
7–4	Generation of an Output Stream G0	7-11
7–5	Bit Stream Multiplexing Concept	7-12
7–6	Butterfly Structure for K = 5, Rate 1/2 GSM Convolutional Encoder	7-17

Tables

3-1	Compiler Options Summary	3-2
3-2	TMS320C55x C Compiler Intrinsic	3-14
3-3	Section Descriptions	3-36
3-4	Possible Operand Combinations	3-38
4-1	CPU Data Buses and Constant Buses	4-23
4-2	Basic Parallelism Rules	4-25
4-3	Advanced Parallelism Rules	4-26
4-4	Steps in Process for Applying User-Defined Parallelism	4-27
4-5	Descriptions of the Execution Pipeline Stages	4-53
4-6	Recommendations for Preventing Pipeline Delays	4-56
4-7	Status Register Pipeline-Protection Granularity	4-64
4-8	MMR Pipeline-Protection Granularity	4-66
4-9	Half-Cycle Accesses to Dual-Access Memory (DARAM)	4-78
4-10	One-Cycle Accesses to Single-Access Memory (SARAM)	4-79
4-11	Cross-Reference Table Documented By Software Developers to Help Software Integrators Generate an Optional Application Mapping	4-81
6-1	Syntaxes for Bit-Reverse Addressing Modes	6-2
6-2	Bit-Reversed Addresses	6-3
6-3	Typical Bit-Reverse Initialization Requirements	6-5
7-1	Operands to the first or first instruction	7-4
8-1	DSPLIB Function Argument Descriptions	8-5
8-2	DSPLIB Functions	8-6

Examples

2-1	Final Assembly Code of test.asm	2-4
2-2	Partial Assembly Code of test.asm (Step 1)	2-6
2-3	Partial Assembly Code of test.asm (Step 2)	2-7
2-4	Partial Assembly Code of test.asm (Part3)	2-9
2-5	Linker command file (test.cmd)	2-11
2-6	Linker map file (test.map)	2-12
2-7	x Memory window	2-17
3-1	Main Function File	3-6
3-2	Sum Function File	3-6
3-3	Assembly Code Generated With -o3 and -pm Options	3-7
3-4	Assembly Generated Using -o3 -pm and -oi50	3-11
3-5	Implementing Saturated Addition in C	3-12
3-6	Assembly Code Generated by C Implementation of Saturated Addition	3-13
3-7	Single Call to _sadd Intrinsic	3-13
3-8	Assembly code Generated When Using Compiler Intrinsic for Saturated Add	3-14
3-9	Block Copy Using Long Data Access	3-16
3-10	Simple Loop that Allows Use of Local Repeat	3-17
3-11	Assembly Code for Local Repeat Generated by the Compiler	3-18
3-12	Inefficient Loop Code for Loop Variable and Constraints (C)	3-19
3-13	Inefficient Loop Code for Variable and Constraints(Assembly)	3-20
3-14	Using Unsigned Data Types	3-20
3-15	Assembly Code Generated for Unsigned Data Type	3-21
3-16	Using _nassert Directive	3-22
3-17	Assembly Code Generated With _nassert Directive	3-22
3-18	Main Function Calling sum	3-24
3-19	Assembly Code Generated With Main Calling sum	3-25
3-20	Assembly Source Output Using -o3 and -pm Options	3-26
3-21	Use Local Rather Than Global Summation Variables	3-28
3-22	Returning Q15 Result for Multiply Accumulate	3-29
3-23	Simulating Circular Addressing in C	3-30
3-24	Assembly Output for Circular Addressing C Code	3-31
3-25	Circular Addressing Using Modulus Operator	3-32
3-26	Assembly Output for Circular Addressing Using Modulo	3-33
3-27	Considerations for Long Data Objects in Structures	3-34
3-28	Declaration Using DATA_SECTION Pragma	3-37
3-29	Sample Linker Command File	3-39

3-30	Allocation of Functions Using CODE_SECTION Pragma	3-40
4-1	Complex Vector Multiplication Code	4-7
4-2	Block FIR Filter Code (Not Optimized)	4-12
4-3	Block FIR Filter Code (Optimized)	4-14
4-4	A-Unit Code With No User-Defined Parallelism	4-30
4-5	A-Unit Code in Example 4-4 Modified to Take Advantage of Parallelism	4-33
4-6	P-Unit Code With No User-Defined Parallelism	4-35
4-7	P-Unit Code in Example 4-6 Modified to Take Advantage of Parallelism	4-37
4-8	D-Unit Code With No User-Defined Parallelism	4-38
4-9	D-Unit Code in Example 4-8 Modified to Take Advantage of Parallelism	4-39
4-10	Code That Uses Multiple CPU Units But No User-Defined Parallelism	4-40
4-11	Code in Example 4-10 Modified to Take Advantage of Parallelism	4-43
4-12	Nested Loops	4-47
4-13	Branch-On-Auxiliary-Register-Not-Zero Construct (Shown in Complex FFT Loop Code)	4-48
4-14	Use of CSR (Shown in Real Block FIR Loop Code)	4-51
4-15	A-Unit Register Write/(Read in AD Stage) Sequence	4-59
4-16	A-Unit Register Read/(Write in AD Stage) Sequence	4-60
4-17	Register (Write in X Stage)/(Read in R Stage) Sequence	4-61
4-18	Good Use of MAR Instruction (Write/Read Sequence)	4-62
4-19	Bad Use of MAR Instruction (Read/Write Sequence)	4-63
4-20	Solution for Bad Use of MAR Instruction (Read/Write Sequence)	4-63
4-21	ST0_55 Course Granularity	4-65
4-22	ST0_55 Fine Granularity	4-65
4-23	MMR Coarse Granularity	4-67
4-24	MMR Fine Granularity	4-67
4-25	Protected BRC Read	4-69
4-26	Unprotected BRC Read	4-70
4-27	Unprotected BRC Write	4-70
4-28	BRC Initialization	4-71
4-29	CSR Initialization	4-72
4-30	Condition Evaluation Preceded by a X-stage Write to the Register Affecting the Condition	4-73
4-31	Making an Operation Conditional With execute(AD_unit)	4-74
4-32	Making an Operation Conditional With execute(D_unit)	4-75
4-33	Conditional Parallel Write Operation Followed by an AD-Stage Write to the Register Affecting the Condition	4-76
4-34	Write/Dual-Operand Read Sequence (Assumes Xmem, Ymem, and Smem Are in the Same SARAM block)	4-82
5-1	Signed 2s-Complement Binary Number Expanded to Decimal Equivalent	5-3
5-2	Computing the Negative of a 2s-Complement Number	5-3
5-3	Addition With 2s-Complement Binary Numbers	5-6
5-4	Subtraction With 2s-Complement Binary Numbers	5-7
5-5	Multiplication With 2s-Complement Binary Numbers	5-8
5-6	64-Bit Addition	5-12

5-7	64-Bit Subtraction	5-15
5-8	32-Bit Integer Multiplication	5-18
5-9	32-Bit Fractional Multiplication	5-20
5-10	Unsigned, 16-Bit By 16-Bit Integer Division	5-22
5-11	Unsigned, 32-Bit By 16-Bit Integer Division	5-23
5-12	Signed, 16-Bit By 16-Bit Integer Division	5-26
5-13	Signed, 32-Bit By 16-Bit Integer Division	5-28
6-1	Sequence of Auxiliary Registers Modifications in Bit-Reversed Addressing	6-2
6-2	Off-Place Bit Reversing of a Vector Array (in Assembly)	6-6
6-3	Using DSPLIB cbrev() Routine to Bit Reverse a Vector Array (in C)	6-8
7-1	Symmetric FIR Filter	7-5
7-2	Delayed LMS Implementation of an Adaptive Filter	7-9
7-3	Generation of Output Streams G0 and G1	7-11
7-4	Multiplexing Two Bit Streams With the Field Expand Instruction	7-13
7-5	Demultiplexing a Bit Stream With the Field Extract Instruction	7-15
7-6	Viterbi Butterflies for Channel Coding	7-19
7-7	Viterbi Butterflies Using Instruction Parallelism	7-21

Introduction

This chapter lists some of the key features of the TMS320C55x™ (C55x™) DSP architecture and shows a recommended process for creating code that runs efficiently.

Topic	Page
1.1 TMS320C55x Architecture	1-2
1.2 Code Development Flow for Best Performance	1-3

1.1 TMS320C55x Architecture

The TMS320C55x device is a fixed-point digital signal processor (DSP). The main block of the DSP is the central processing unit (CPU), which has the following characteristics:

- ☐ A unified program/data memory map. In program space, the map contains 16M bytes that are accessible at 24-bit addresses. In data space, the map contains 8M words that are accessible at 23-bit addresses.
- ☐ An input/output (I/O) space of 64K words for communication with peripherals.
- ☐ Software stacks that support 16-bit and 32-bit push and pop operations. You can use these stack for data storage and retrieval. The CPU uses these stacks for automatic context saving (in response to a call or interrupt) and restoring (when returning to the calling or interrupted code sequence).
- ☐ A large number of data and address buses, to provide a high level of parallelism. One 32-bit data bus and one 24-bit address bus support instruction fetching. Three 16-bit data buses and three 24-bit address buses are used to transport data to the CPU. Two 16-bit data buses and two 24-bit address buses are used to transport data from the CPU.
- ☐ An instruction buffer and a separate fetch mechanism, so that instruction fetching is decoupled from other CPU activities.
- ☐ The following computation blocks: one 40-bit arithmetic logic unit (ALU), one 16-bit ALU, one 40-bit shifter, and two multiply-and-accumulate units (MACs). In a single cycle, each MAC can perform a 17-bit by 17-bit multiplication (fractional or integer) and a 40-bit addition or subtraction with optional 32-/40-bit saturation.
- ☐ An instruction pipeline that is protected. The pipeline protection mechanism inserts delay cycles as necessary to prevent read operations and write operations from happening out of the intended order.
- ☐ Data address generation units that support linear, circular, and bit-reverse addressing.
- ☐ Interrupt-control logic that can block (or mask) certain interrupts known as the maskable interrupts.
- ☐ A TMS320C54x-compatible mode to support code originally written for a TMS320C54x™ DSP.

1.2 Code Development Flow for Best Performance

The following flow chart shows how to achieve the best performance and code-generation efficiency from your code. After the chart, there is a table that describes the phases of the flow.

Figure 1–1. Code Development Flow

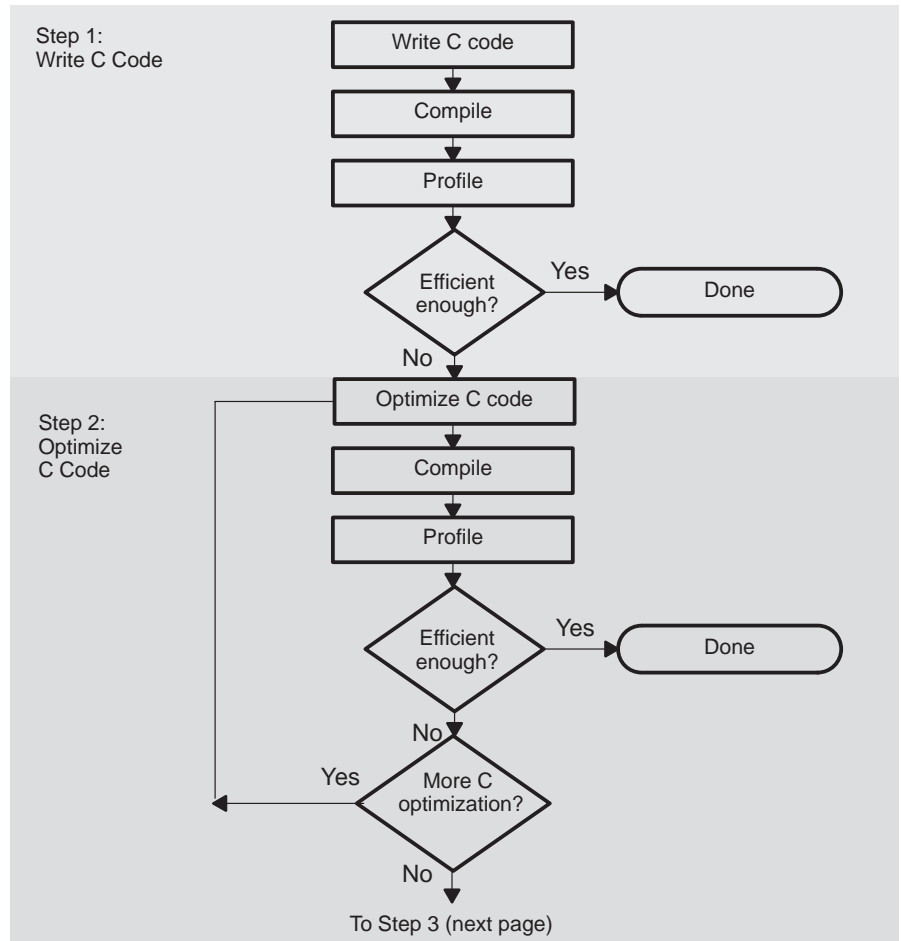
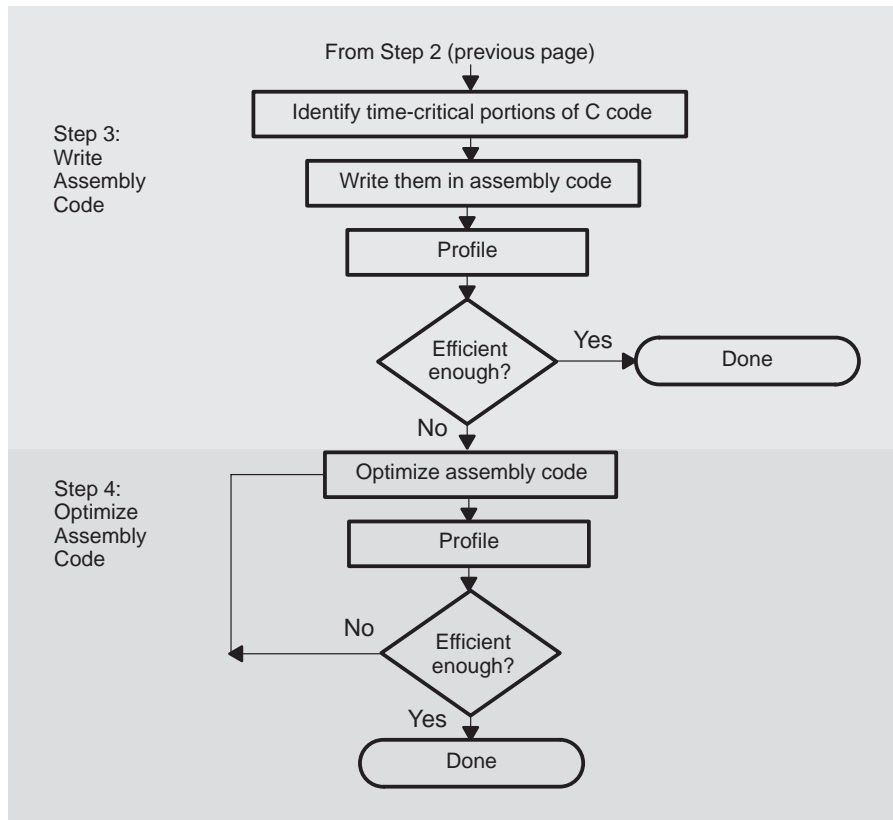



Figure 1–1. Code Development Flow (Continued)



Step	Goal
1	Write C Code: You can develop your code in C using the ANSI-compliant C55x C compiler without any knowledge of the C55x DSP. Use Code Composer Studio to identify any inefficient areas that you might have in your C code. After making your code functional, you can improve its performance by selecting higher-level optimization compiler options. If your code is still not as efficient as you would like it to be, proceed to step 2.
2	Optimize C Code: Explore potential modifications to your C code to achieve better performance. Some of the techniques you can apply include (see Chapter 3): <ul style="list-style-type: none"><input type="checkbox"/> Use specific types (register, volatile, const).<input type="checkbox"/> Modify the C code to better suit the C55x architecture.<input type="checkbox"/> Use an ETSI intrinsic when applicable.<input type="checkbox"/> Use C55x compiler intrinsics. After modifying your code, use the C55x profiling tools again, to check its performance. If your code is still not as efficient as you would like it to be, proceed to step 3.
3	Write Assembly Code: Identify the time-critical portions of your C code and rewrite them as C-callable assembly-language functions. Again, profile your code, and if it is still not as efficient as you would like it to be, proceed to step 4.
4	Optimize Assembly Code: After making your assembly code functional, try to optimize the assembly-language functions by using some of the techniques described in Chapter 4, <i>Optimizing Your Assembly Code</i> . The techniques include: <ul style="list-style-type: none"><input type="checkbox"/> Place instructions in parallel.<input type="checkbox"/> Rewrite or reorganize code to avoid pipeline protection delays.<input type="checkbox"/> Minimize stalls in instruction fetching.

Tutorial



This tutorial walks you through the code development flow introduced in Chapter 1, and introduces you to basic concepts of TMS320C55x™ (C55x™) DSP programming. It uses step-by-step instructions and code examples to show you how to use the software development tools integrated under Code Composer Studio (CCS).

Installing CCS before beginning the tutorial allows you to edit, build, and debug DSP target programs. For more information about CCS features, see the CCS Tutorial. You can access the CCS Tutorial within CCS by choosing Help→Tutorial→CCS Tutorial.

The examples in this tutorial were run on the CCS v1.2 software development tools, the most recent version available as of the publication of this book. Because the tools are being continuously improved, you may get different results if you are using a more recent version of the tools. The examples use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.

Topic	Page
2.1 Introduction	2-2
2.2 Writing Assembly Code	2-3
2.3 Understanding the Linking Process	2-10
2.4 Building Your Program	2-14
2.5 Testing Your Code	2-16
2.6 Benchmarking Your Code	2-18

2.1 Introduction

This tutorial presents a simple assembly code example that adds four numbers together ($y = x0 + x3 + x1 + x2$). This example helps you become familiar with the basics of C55x programming.

After completing the tutorial, you should know:

- ☐ The four common C55x addressing modes and when to use them.
- ☐ The basic C55x tools required to develop and test your software.

This tutorial does not replace the information presented in other C55x documentation and is not intended to cover all the topics required to program the C55x efficiently.

Refer to the related documentation listed in the preface of this book for more information about programming the C55x DSP. Much of this information has been consolidated as part of the C55x Code Composer Studio online help.

For your convenience, all the files required to run this example can be downloaded with the .PDF version of the *TMS320C55x Programmer's Guide (SPRU376)* from <http://www.ti.com/sc/docs/schome.htm>.

2.2 Writing Assembly Code

Writing your assembly code involves the following steps:

- ☐ Allocate sections for code, constants, and variables.
- ☐ Initialize the processor mode.
- ☐ Set up addressing modes and add the following values: $x0 + x1 + x2 + x3$.

The following rules should be considered when writing C55x assembly code:

- ☐ Labels

The first character of a label must be a letter or an underscore (`_`) followed by a letter, and must begin in the first column of the text file. Labels can contain up to 32 alphanumeric characters.

- ☐ Comments

When preceded by a semicolon (`;`), a comment may begin in any column. When preceded by an asterisk (`*`), a comment must begin in the first column.

The final assembly code product of this tutorial is displayed in Example 2–1, Final Assembly Code of test.asm. This code performs the addition of the elements in vector x. Sections of this code are highlighted in the three steps used to create this example.

For more information about assembly syntax, see the *TMS320C55x Assembly Language Tools User's Guide*.

Example 2–1. Final Assembly Code of test.asm

```
* Step 1: section allocation
* -----
        .def x,y,init
x        .usect "vars",4           ; reserve 4 uninitialized 16-bit locations for x
y        .usect "vars",1           ; reserve 1 uninitialized 16-bit location for y

        .sect "table"             ; create initialized section "table" to
init     .int 1,2,3,4             ; contain initialization values for x

        .text                     ; create code section (default is .text)
        .def start                ; define label to the start of the code
start

* Step 2: Processor mode initialization
* -----
        bit(ST1, #ST1_C54CM) = #0 ; set processor to '55x native mode instead of
                                   '54x compatibility mode (reset value)
        bit(ST2, #ST2_AR0LC) = #0 ; set AR0 register in linear mode
        bit(ST2, #ST2_AR6LC) = #0 ; set AR6 register in linear mode

* Step 3a: copy initialization values to vector x using indirect addressing
* -----
copy
        XAR0 = #x                  ; XAR0 pointing to variable x
        XAR6 = #init              ; XAR6 pointing to initialization table

        *AR0+ = *AR6+              ; copy starts from "init" to "x"
        *AR0+ = *AR6+
        *AR0+ = *AR6+
        *AR0 = *AR6

* Step 3b: add values of vector x elements using direct addressing
* -----
add
        XDP = #x                  ; XDP pointing to variable x
        .dp x                    ; and the assembler is notified

        AC0 = @x
        AC0 += @(x+3)
        AC0 += @(x+1)
        AC0 += @(x+2)

* Step 3c. write the result to y using absolute addressing
* -----
        *(#y) = AC0

end
        nop
        goto end
```


2.2.1 Allocate Sections for Code, Constants, and Variables

The first step in writing this assembly code is to allocate memory space for the different sections of your program.

Sections are modules consisting of code, constants, or variables needed to successfully run your application. These modules are defined in the source file using assembler directives. The following basic assembler directives are used to create sections and initialize values in the example code.

- ❑ `.sect "section_name"` creates initialized name section for code/data. Initialized sections are sections defining their initial values.
- ❑ `.usect "section_name", size` creates uninitialized named section for data. Uninitialized sections declare only their size in 16-bit words, but do not define their initial values.
- ❑ `.int value` reserves a 16-bit word in memory and defines the initialization value
- ❑ `.def symbol` makes a symbol global, known to external files, and indicates that the symbol is defined in the current file. External files can access the symbol by using the `.ref` directive. A symbol can be a label or a variable.

As shown in Example 2–2 and Figure 2–1, the example file `test.asm` contains three sections:

- ❑ `vars`, containing five uninitialized memory locations.
 - The first four are reserved for vector `x` (the input vector to add).
 - The last location, `y`, will be used to store the result of the addition.
- ❑ `table`, to hold the initialization values for `x`. The `init` label points to the beginning of section `table`.
- ❑ `text`, which contains the assembly code.


Example 2–2 shows the partial assembly code used for allocating sections.

Example 2–2. Partial Assembly Code of test.asm (Step 1)

```
* Step 1: section allocation
* -----
        .def x,y,init
x        .usect "vars",4           ; reserve 4 uninitialized locations for var x
y        .usect "vars",1           ; reserve 1 uninitialized location for result y

        .sect "table"              ; create initialized section "table" to
init     .int 1,2,3,4              ; contain initialization values for x

        .text                      ; create code section (default is .text)
        .def start                 ; make the start label global
start    ; define label to the start of the code
```

Figure 2–1. Section Allocation

2.2.2 Processor Mode Initialization

The second step is to make sure the status registers (ST0_55, ST1_55, ST2_55, and ST3_55) are set to configure your processor. You will either need to set these values or use the default values. Default values are placed in the registers after processor reset. You can locate the default register values after reset in the *TMS320C55x DSP CPU Reference Guide (SPRU371)*.

As shown in Example 2–3:

- ❑ The AR0 and AR6 registers are set to linear addressing (instead of circular addressing) using bit addressing mode to modify the status register bits. The syntax for bit addressing mode is:
 bit (register, #register_bitname)=#0 or #1
- ❑ The processor has been set in C55x native mode instead of C54x-compatible mode.

Example 2–3. Partial Assembly Code of test.asm (Step 2)

```
* Step 2: Processor mode initialization
* -----
    bit(ST1, #ST1_C54CM) = #0 ; set processor to '55x native mode instead of
                               '54x compatibility mode (reset value)
    bit(ST2, #ST2_AR0LC) = #0 ; set AR0 register in linear mode (reset value)
    bit(ST2, #ST2_AR6LC) = #0 ; set AR6 register in linear mode (reset value)
```

2.2.3 Setting up Addressing Modes

Four of the most common C55x addressing modes are used in this code:

- ☐ ARn Indirect addressing (identified by *), in which you use auxiliary registers (ARx) as pointers.
- ☐ DP direct addressing (identified by @), which provides a positive offset addressing from a base address specified by the DP register. The offset is calculated by the assembler and defined by a 7-bit value embedded in the instruction.
- ☐ k23 absolute addressing (identified by #), which allows you to specify the entire 23-bit data address with a label.
- ☐ Bit addressing (identified by the bit instruction), which allows you to modify a single bit of a memory location or MMR register.

For further details on these addressing modes, refer to the *TMS320C55x DSP CPU Reference Guide (SPRU371)*. Example 2–4 demonstrates the use of the addressing modes discussed in this section.

In Step 3a, initialization values from the *table* section are copied to vector *x* (the vector to perform the addition) using indirect addressing. Figure 2–2 illustrates the structure of the extended auxiliary registers (XARn). The XARn register is used only during register initialization. Subsequent operations use ARn because only the lower 16 bits are affected (ARn operations are restricted to a 64k main data page). AR6 is used to hold the address of *table*, and AR0 is used to hold the address of *x*.

In Step 3b, direct addressing is used to add the four values. Notice that the XDP register was initialized to point to variable *x*. The *.dp* assembler directive is used to define the value of XDP, so the correct offset can be computed by the assembler at compile time.

Finally, in Step 3c, the result was stored in the *y* vector using absolute addressing. Absolute addressing provides an easy way to access a memory location without having to make XDP changes, but at the expense of an increased code size.

Example 2–4. Partial Assembly Code of test.asm (Part3)

```

* Step 3a: copy initialization values to vector x using indirect addressing
* -----
copy
  XAR0 = #x                ; XAR0 pointing to startn of x array
  XAR6 = #init             ; XAR6 pointing to start of init array

  *AR0+ = *AR6+            ; copy from source "init" to destination "x"
  *AR0+ = *AR6+
  *AR0+ = *AR6+
  *AR0 = *AR6

* Step 3b: add values of vector x elements using direct addressing
* -----
add
  XDP = #x                ; XDP pointing to variable x
  .dp x                   ; and the assembler is notified

  AC0 = @x
  AC0 += @(x+3)
  AC0 += @(x+1)
  AC0 += @(x+2)

* Step 3c. write the result to y using absolute addressing
* -----
  *(#y) = AC0

end
  nop
  goto end

```

Figure 2–2. Extended Auxiliary Registers Structure (XARn)

Note: ARnH (upper 7 bits) specifies the 7-bit main data page. ARn (16-bit registers) specifies a 16-bit offset to the 7-bit main data page to form a 23-bit address.

2.3 Understanding the Linking Process

The linker (lnk55.exe) assigns the final addresses to your code and data sections. This is necessary for your code to execute.

The file that instructs the linker to assign the addresses is called the linker command file (test.cmd) and is shown in Example 2–5. The linker command file syntax is covered in detail in the *TMS320C55x Assembly Language Tools User's Guide (SPRU280)*.

- ❑ All addresses and lengths given in the linker command file uses byte addresses and byte lengths. This is in contrast to a TMS320C54x™ linker command file that uses 16-bit word addresses and word lengths.
- ❑ The MEMORY linker directive declares all the physical memory available in your system (For example, a DARAM memory block at location 0x100 of length 0x8000 bytes). Memory blocks cannot overlap.
- ❑ The SECTIONS linker directive lists all the sections contained in your input files and where you want the linker to allocate them.
- ❑ The following linker options are used in Example 2–5:
 - -o *filename*: names the executable file
 - -m: creates a map file
 - -e *entry_label*: provides the entry point for the code

When you build your project in Section 2.4, this code produces two files, test.out and a test.map. Review the test.map file, Example 2–6, to verify the addresses for x, y, and table. Notice that the linker reports byte addresses for program labels such as *start* and *text*, and 16-bit word addresses for data labels like *x*, *y*, and *table*. The C55x DSP uses byte addressing to access variable length instructions. Instructions can be 1-6 bytes long.

Example 2–5. Linker command file (test.cmd)

```
test.obj      /* input files */
-o test.out   /* output file */
-m test.map   /* map file */
-e start      /* entry point for the code */

MEMORY        /* byte address, byte len */
{
    DARAM: org= 000100h, len = 8000h
    SARAM: org= 010000h, len = 8000h
}

SECTIONS       /* byte address, byte len */
{
    vars :> DARAM
    table: > SARAM
    .text:> SARAM
}
```

Example 2–6. Linker map file (test.map)

```

*****
TMS320C55xx COFF Linker                      Version 1.03B
*****
>> Linked Mon Feb 14 14:52:21 2000

OUTPUT FILE NAME:    <test.out>
ENTRY POINT SYMBOL:  "start"  address: 00010008

MEMORY CONFIGURATION

      name      org (bytes)  len (bytes)  used (bytes)  attributes  fill
      ----      -
      DARAM      00000100    00000800    0000000a     RWIX
      SARAM      00010000    00000800    00000040     RWIX

SECTION ALLOCATION MAP

output
section  page  orgn(bytes)  orgn(words)  len(bytes)  len(words)  attributes/
-----  -
vars      0      00000080      00000080      00000005      00000005  UNINITIALIZED
                                     test.obj (vars)

table      0      00008000      00008000      00000004      00000004  test.obj
(table)

.text      0  00010008      00000038      00000037      00000037  test.obj
(.text)

          0001003f      00000001      00000001  --HOLE-- [fill
= 2020]

.data      0      00000000      00000000      00000000      00000000  UNINITIALIZED
          00000000      00000000      00000000  test.obj
(.data)

.bss      0      00000000      00000000      00000000      00000000  UNINITIALIZED
          00000000      00000000      00000000  test.obj (.bss)

```


Example 2–6. Linker map file (test.map), (Continued)

GLOBAL SYMBOLS: SORTED ALPHABETICALLY BY Name

abs. value/ byte addr	word addr	name
-----	-----	----
	00000000	.bss
	00000000	.data
00010008		.text
	00000000	__bss__
	00000000	__data__
	00000000	__edata__
	00000000	__end__
00010040		__etext__
00010008		__text__
	00000000	edata
	00000000	end
00010040		etext
	00008000	init
00010008		start
	00000080	x
	00000084	y

GLOBAL SYMBOLS: SORTED BY Symbol Address

abs. value/ byte addr	word addr	name
-----	-----	----
	00000000	__end__
	00000000	__edata__
	00000000	end
	00000000	edata
	00000000	__data__
	00000000	.data
	00000000	.bss
	00000000	__bss__
	00000080	x
	00000084	y
	00008000	init
00010008		start
00010008		.text
00010008		__text__
00010040		__etext__
00010040		etext

[16 symbols]

2.4 Building Your Program

At this point, you should have already successfully installed CCS and selected the C55x Simulator as the CCS configuration file to be used. You can select a configuration file to be used in CCS Setup.

Before building your program, you must set up your work environment and create a .mak file. Setting up your work environment involves the following tasks:

- ☐ Creating a project
- ☐ Adding files to the work space
- ☐ Modifying the build options
- ☐ Building your program

2.4.1 Creating a Project

First, create a new project called test.mak.

- 1) From the Project menu, choose New.
- 2) Navigate to c:\ti\myprojects\55x_examples\55x_prog.
- 3) Create the folder named 55x_prog if needed.
- 4) In the File Name field, Type tutor.mak and select Save.
- 5) You have now created a project named tutor.mak and saved it in c:\ti\myprojects\55x_examples\55x_prog

2.4.2 Adding Files to the Work space

Copy the tutorial files (test.asm and test.cmd) to the project directory.

- 1) Navigate to the directory where the tutorial files are located.
- 2) Copy the tutorial files and paste them into the project directory you created earlier: c:\ti\myprojects\55x_examples\55x_prog

As an alternative, you can create your own source files by choosing File→New→Source File and typing the source code from the examples in this book.

2.4.3 Modifying Build Options

Open the options dialog box and modify the Assembler and Linker options

- 1) From the Project menu, choose Options.
- 2) Select the Assembler tab.
- 3) Select Algebraic as the Assembler Type.
- 4) Select the Linker tab.
- 5) In the Map Filename field, type *test.map*.
- 6) In the Autoinitialization Model field, select No Autoinitialization.

2.4.4 Build the Program

From the Project menu, choose Rebuild All

When you build your project, CCS compiles, assembles, and links your code in one step. The assembler reads the assembly source file and converts C55x instructions to their corresponding binary encoding. The results of the assembly processes are an object file, *test.obj*, in industry standard COFF binary format. The object file contains all of your code and variables, but the addresses for the different sections of code are not assigned. This assignment takes place during the linking process.

Because there is no c code in your project, no compiler options were used.

The following basic assembler options were used to build the program:

- ☐ -as: Include symbols, make all symbols global
- ☐ -g: Enables assembly source debug
- ☐ -mg: Defines the assembler type as Algebraic

The following basic linker options were used to build the program:

- ☐ -m: Map filename
- ☐ -o: Output filename
- ☐ -x: Exhaustively Read Libraries

2.5 Testing your code

To test your code, inspect its execution using the C55x Simulator.

Load test.out

- 1) From the File menu, choose Load program.
- 2) Navigate to and select test.out, then choose Open.

CCS now displays the test.asm source code at the beginning of the start label because of the entry symbol defined in the linker command file (-e start). Otherwise, it would have shown the location pointed to by the reset vector. You can also define the start location under the Project Options Linker tab.

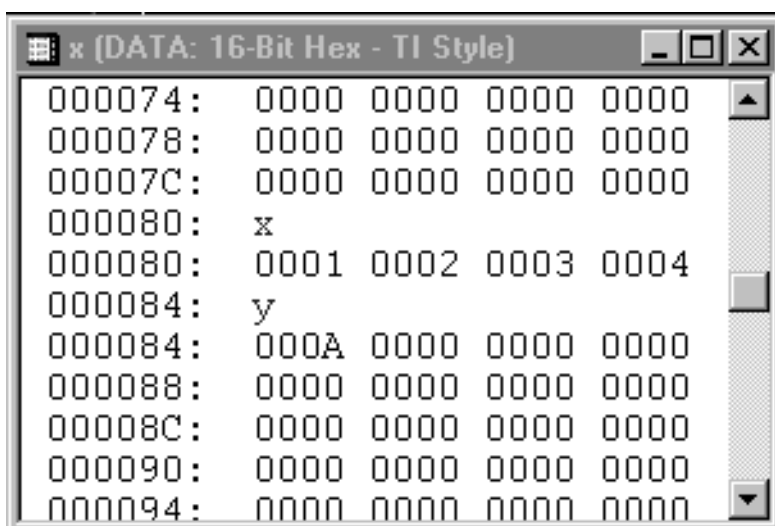
Display arrays *x*, *y*, and *init* by setting Memory Window options

- 1) From the View menu, chose Memory.
- 2) In the Title field, type *x*.
- 3) In the Address field, type *x*.
- 4) Repeat 1–3 for *y*.
- 5) Display the *init* array by selecting View→ Memory.
- 6) In the Title field, type *Table*.
- 7) In the Address field, type *init*.
- 8) Display AC0 by selecting View→CPU Registers→CPU Registers.

The labels *x*, *y*, and *init* are visible to the simulator (using View→ Memory) because they were exported as symbols (using the .def directive in test.asm). The -g option was used to enable assembly source debugging.

Now, single-step through the code to the *end* label by selecting Debug→Step Into. Examine the X Memory window to verify that the table values populate *x* and that *y* gets the value 0xa ($1 + 2 + 3 + 4 = 10 = 0xa$), as shown in Example 2–7.

Example 2-7. x Memory window



2.6 Benchmarking your code

After verifying the correct functional operation of your code, you can use CCS to calculate the number of cycles your code takes to execute.

Reload your code

From the File menu, choose *Reload Program*.

Enable clock for profiling

- 1) From the Profiler menu, choose *Enable clock*.
- 2) From the Profiler menu, choose *View Clock*.

Set breakpoints

- 1) Select the *test.asm* window.
- 2) Set one breakpoint at the beginning of the code you want to benchmark: Right-click on the instruction next to the *copy* label and choose *Toggle Breakpoint*.
- 3) Set one breakpoint marking the end: Right-click on the instruction next to the *end* label and choose *Toggle Breakpoint*.

Benchmark your code

- 1) Run to the first breakpoint by selecting Debug→ Run.
- 2) Double-click in the *Clock window* to clear the cycle count.
- 3) Run to the second breakpoint by selecting Debug→ Run.
- 4) The *Clock window* displays the number of cycles the code took to execute between the breakpoints, which was 15.

Optimizing C Code

You can maximize C performance by using compiler options, intrinsics, and code transformations. The assembly-language examples in this chapter use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.

Topic	Page
3.1 The Compiler and Its Options	3-2
3.2 Using Program-level Optimization	3-5
3.3 Using Function Inlining	3-10
3.4 Using Intrinsics	3-12
3.5 Using Long Data Accesses for 16 Bit Data	3-16
3.6 Generating Efficient Loop Code	3-17
3.7 Generating Efficient Control Code	3-27
3.8 Efficient Math Operations	3-28
3.9 Memory Management	3-34
3.10 Allocating Function Code to Different Sections	3-40

3.1 The Compiler and Its Options

The TMS320C55x™ (C55x™) C Compiler provides a wide range of options. Some of these options affect how your C source code is parsed (analyzed for syntactic and semantic conformance to ANSI C standards), others may select debug capabilities, such as generation of listing files (–pl), or insertion of symbolic debugging directives for runtime debug with Code Composer Studio (–g). This section focuses on that subset of compiler options which affect the compiler optimization phase. These options enable techniques that the compiler uses to generate more efficient assembly language than it does when no optimizations are selected.

Compiler Options

Options control the operation of the compiler. They can significantly affect the efficiency of the assembly source code generated by the compiler when translating your C source. The following subset of available compiler options provide the overall best performance boost in terms of cycle count and code size reduction in the translated assembly.

Table 3–1. Compiler Options Summary

Option		Description
–x2		Enable inlining of inline functions
–on	n = 0 (default)	Simplifies control flow Allocates variables to registers Performs loop rotation Eliminates unused code Simplifies expressions and statements Expands calls to inline functions
	n = 1	Same features as n=0 option PLUS: Performs local copy/constant propagation Removes unused assignments Eliminates local common expressions
	n = 2	Same features as n=1 option PLUS: <i>Performs loop optimizations</i> Eliminates global common sub–expressions Eliminates global unused assignments Performs loop unrolling

Table 3–1. Compiler Options Summary (Continued)

Option		Description
	n = 3	Same features as n=2 option PLUS: Removes all functions that are never called. Simplifies functions with return values that are never used. Inlines calls to small functions (regardless to declaration). Reorders functions so that the attributes of a called function are known when the caller is optimized. Identifies file-level variable characteristics.
-oimize		Enables inlining of functions based on a maximum size. Size here is internally determined by the optimizer and does not correspond to bytes or any known standard unit . Use -onx to check sizes of individual functions.
-onx	x= 0	Disables optimizer information file (default).
	x= 1	Produces optimizer information file.
	x= 2	Produces verbose optimizer information file.
-opn	n = 0	Specifies that the code contains functions and variables which may be accessed by code outside of the source provided to the compiler and that the compiler should not remove them.
	n = 1	Specifies that the source code contains variables that are modified outside the module, but does not call any functions from outside the current source. The compiler should not remove those variables.
	n = 2 (default)	Opposite to -op0. Specifies that the source code contains no functions or variables that are called or modified outside the source provided to the compiler. Therefore, the compiler is authorized to remove any "unused" functions and variables.
	n = 3	Opposite to -op1. Specifies that the module calls functions declared outside the current source, but does not use variables modified outside current source to the compiler. Therefore, the compiler is authorized to remove any "unused" variables.

Table 3–1. Compiler Options Summary (Continued)

Option	Description
–pm	Program mode option, tells the compiler to combine the individual C source programs included with this shell invocation. This option is best used in conjunction with –o3 option to take advantage of global and file level optimizations provided with level 3 optimization.
–mn	Re-enables optimizations disabled when using –g option (–g enables generation of symbolic debug information).
–ml	Selects large memory model.
–mr	Prevents generation of hardware block repeat, local repeat, and repeat instructions to reduce context save/restore for interrupts.
–ms	Optimizes for code space (default is to optimize for code speed).

3.2 Using Program-level Optimization

You can specify program-level optimization by using the `-pm` option with the `-o3` option. With program-level optimization, all your source files are compiled into one intermediate file called a module. The module moves to the optimization and code generation phases of the compiler. Because the compiler has access to the entire program, it performs several optimizations that are rarely applied during file-level optimization:

- ☐ If a particular argument in a function always has the same value, the compiler replaces the argument with the value and passes the value instead of the argument.
- ☐ If a return value of a function is never used, the compiler deletes the return code in the function.
- ☐ If a function is not called, directly or indirectly, the compiler removes the code in the function.

Also, using the `-pm` option can lead to better optimization for your loops. If the number of iterations of a loop is determined by a value passed into the function, and the compiler can determine what the value is from the caller, then the compiler will have more information about the minimum iterations on the loop, resulting in more efficient loop code.

Program-level optimization increases compilation time because the compiler performs more complex optimizations on a larger amount of code. For this reason you may want to optimize in stages. Using `-o3` option and `-oi` options first. And then in the last stage of code development, introduce the `-pm` option to further reduce code size and cycle time.

Example 3–1 and Example 3–2 show the content of two files. One file contains the source for the *main* function and the second file contains source for a small function called *sum*.

Example 3–1. Main Function File

```
extern int sum(const short *a, unsigned int n);
short a[10] = {1,2,3,4,5,6,7,8,9,10};
short b[10] = {11,12,13,14,15,16,17,18,19,20};
int sum1, sum2;
void main(void)
{
    sum1 = sum(a,9);
    sum2 = sum(b,9);
}
```

Example 3–2. Sum Function File

```
int sum(const short *a, unsigned int n)
{
    int sum = 0;
    unsigned int i;
    for(i=0; i<=n; i++)
    {
        sum += a[i];
    }
    return sum;
}
```

When compiled with `-o3` and `-pm` options the optimizer has enough information about the calls to *sum* to determine that the same loop count is used for both calls. It therefore eliminates the argument *n* from the call to the function and explicitly uses the count in the repeat code generated as is shown in Example 3–3:

Example 3–3. Assembly Code Generated With `-o3` and `-pm` Options

```

_sum:
    SP = SP - #1
                                           ; End Prolog Code

; **Parameter deleted: n == 9u;
    T0 = #0 ; |3|
    repeat(#9)
                                           ; loop starts ; |3|
    T0 = T0 + *AR0+ ; |8|
                                           ; loop ends
L1:
                                           ; Begin Epilog Code
    SP = SP + #1 ; |11|
    return ; |11|
                                           ; return occurs ; |11|

main:
    SP = SP - #1
                                           ; End Prolog Code
    AR0 = #_a ; |10|
    call #_sum ; |10|
                                           ; call occurs [#_sum] ; |10|
    *abs16(#_sum1) = T0 ; |10|
    AR0 = #_b ; |11|
    call #_sum ; |11|
                                           ; call occurs [#_sum] ; |11|
    *abs16(#_sum2) = T0 ; |11|
                                           ; Begin Epilog Code
    SP = SP + #1
    return
                                           ; return occurs

```

Considerations when mixing C and assembly

If you have any assembly functions in your program, exercise caution when using the `-pm` option. The compiler recognizes only the C source code and not any assembly code that might be present. Because the compiler does not recognize the assembly code calls and variable modifications to C functions, the `-pm` option optimizes out those C functions. To keep these functions, you can use 2 methods:

- ☐ place the `FUNC_EXT_CALLED` pragma before any declaration or reference to a function that you want to keep.
- ☐ Use the `-opn` option with the `-pm` and `-o3` options (see Controlling Program-Level Optimization).

In general, you achieve the best results through judicious use of the `FUNC_EXT_CALLED` pragma in combination with `-pm -o3` and `-op1` or `-op2`.

If any of the following situations apply to your application, use the suggested solution:

Situation 1.

Your application consists of C source code that calls assembly functions. Those assembly functions do not call any C functions or modify any C variables.

Solution

Compile with `-pm -o3 -op2` to tell the compiler that outside functions do not call C functions or modify C variables.

If you compile with the `-pm -o3` options only, the compiler reverts from the default optimization level (`-op2`) to `-op0`. The compiler uses `-op0` because it presumes that the calls to the assembly language functions that have a definition in C can call other C functions or modify C variables.

Situation 2.

Your application consists of C source code that calls assembly functions. The assembly language functions do not call C functions, but they modify C variables.

Solution

Try both of these solutions, and choose the one that works best with your code:

- ☐ Compile with `-pm -o3 -op1`.
- ☐ Add the `volatile` keyword to those variables that may be modified by the assembly functions, and compile with `-pm -o3 -op2`. The `volatile` keyword instructs the compiler that the variable may be changed by other code/processes that are not visible to it in the current compilation.

Situation 3.

Your application consists of C source code and assembly source code. The assembly functions are interrupt service routines that call C functions; the C functions that the assembly functions call are never called from C. These C functions act like main: they function as entry points into C.

Solution

You must add the volatile keyword to the C variables that can be modified by the interrupts. Then you can optimize your code in one of these ways:

- ☐ You achieve the best optimization by applying the `FUNC_EXT_CALLED` pragma to all of the entry-point functions called from the assembly language interrupts and then compiling with `-pm -o3 -op2`. Be sure that you use the pragma with all of the entry-point functions. If you do not, the compiler removes the entry-point functions that are not preceded by the `FUNC_EXT_CALL` pragma.
- ☐ Compile with `-pm -o3 -op3`. Because you do not use the `FUNC_EXT_CALL` pragma, you must use the `-op3` option, which is less aggressive than the `-op2` option, and your optimization may not be as effective.

Keep in mind that if you use `-pm -o3` without additional options, the compiler removes the C functions that the assembly functions call. Use the `FUNC_EXT_CALLED` pragma to keep these functions.

3.3 Using Function Inlining

You can have 2 modes of function inlining:

- ☐ Inlining controlled by the “inline” function declaration. To enable this mode use the `-x2` and `-o0` compiler options.
- ☐ Automatic inlining of small functions even if they are not declared as “inline”. To enable this mode use the `-o3` and `-oi<size>` compiler options.

3.3.1 Using `-oi<size>` option

The `-oi<size>` option may be used to specify automatic inlining of small functions even if they have not been declared with inline keyword and inlining has not been explicitly enabled using `-x2` option.

Example 3–4 shows the resulting assembly when the same code in Example 3–1 is compiled with `-o3 -pm` and `-oi50` options

In *main*, the function calls to *sum* have been inlined. However, code for the body of function *sum* has still been generated. The compiler must generate this code because it does not have enough information to eliminate the possibility that the function *sum* may be called by some other externally defined function or process.

3.4 Using Intrinsics

The C55x compiler provides intrinsics, special functions that map directly to inlined C55x instructions, to optimize your C code quickly. Intrinsics are specified with a leading underscore (`_`) and are accessed by calling them as you would call a function.

For example, saturated addition can only be expressed in C code by writing a multicycle function, such as the one in Example 3–5.

Example 3–5. Implementing Saturated Addition in C

```
int sadd(int a, int b)
{
    int result;

    result = a + b;

    // Check to see if 'a' and 'b' have the same sign

    if (((a^b) & 0x8000) == 0)
    {
        // If 'a' and 'b' have the same sign, check for underflow
        // or overflow

        if ((result ^ a) & 0x8000)
        {
            // If the result has a different sign than 'a'
            // then underflow or overflow has occurred.
            // if 'a' is negative, set result to max negative
            // If 'a' is positive, set result to max positive
            // value

            result = ( a < 0 ) ? 0x8000 : 0x7FFF;
        }
    }
    return result;
}
```

Example 3–6 shows the resultant assembly language code generated by the compiler.

Example 3–6. Assembly Code Generated by C Implementation of Saturated Addition

```

_sadd:
    SP = SP - #1
                                ; End Prolog Code
    AR1 = T1 ; |5|
    AR1 = AR1 + T0 ; |5|
    T1 = T1 ^ T0 ; |7|
    AR2 = T1 & #0x8000 ; |7|
    if (AR2!=#0) goto L2 ; |7|
                                ; branch occurs ; |7|
    AR2 = T0 ; |7|
    AR2 = AR2 ^ AR1 ; |7|
    AR2 = AR2 & #0x8000 ; |7|
    if (AR2==#0) goto L2 ; |7|
                                ; branch occurs ; |7|
    if (T0<#0) goto L1 ; |11|
                                ; branch occurs ; |11|
    T0 = #32767 ; |11|
    goto L3 ; |11|
                                ; branch occurs ; |11|
L1:
    AR1 = #-32768 ; |11|
L2:
    T0 = AR1 ; |14|
L3:
                                ; Begin Epilog Code
    SP = SP + #1 ; |14|
    return ; |14|
                                ; return occurs ; |14|

```

The code for the C simulated saturated addition can be replaced by a single call to the `_sadd` intrinsic as is shown in Example 3–7:

Example 3–7. Single Call to `_sadd` Intrinsic

```

int sadd(int a, int b)
{
    return _sadd(a,b);
}

```

The assembly code generated for this C source is shown in Example 3–8:

Example 3–8. Assembly code Generated When Using Compiler Intrinsic for Saturated Add

```

_sadd:
    SP = SP - #1
                                ; End Prolog Code

    bit(ST3, #ST3_SATA) = #1
    T0 = T0 + T1 ; |3|
                                ; Begin Epilog Code

    SP = SP + #1 ; |3|
    bit(ST3, #ST3_SATA) = #0
    return    ; |3|
                                ; return occurs ; |3|

```

Table 3–2 Lists the intrinsics supported by the C55x compiler. For more information on using intrinsics, please refer to the *TMS320C55x Optimizing C Compiler User's Guide*.

Table 3–2. TMS320C55x C Compiler Intrinsics

Intrinsic	C Compiler Intrinsic Description
int _sadd(int src1, int src2);	Adds two 16-bit integers, with SATA set, producing a saturated 16-bit result
long _lsadd(long src1, long src2);	Adds two 32-bit integers, with SATD set, producing saturated 32-bit result
int _ssub(int src1, int src2);	Subtracts src2 from src1 with SATA set, producing a saturated 16-bit result.
long _lssub(long src1, long src2);	Subtracts src2 from src1 with SATD set, producing a saturated 32-bit result.
int _smpy(int src1, int src2);	Multiplies src1 and src2, and shifts the result left by 1. Produces a saturated 16-bit result. (SATD and FRCT set)
long _lsmpy(int src1, int src2);	Multiplies src1 and src2, and shifts the result left by 1. Produces a saturated 32-bit result.(SATD and FRCT set)
long _smac(long src, int op1, int op2);	Multiplies op1 and op2, shifts the result left by 1, and adds it to src. Produces a saturated 32-bit result. (SATD , SMUL and FRCT set)
long _smas(long src, int op1, int op2);	Multiplies op1 and op2, shifts the result left by 1, and subtracts it from src. Produces a 32-bit result. (SATD, SMUL and FRCT set)
int _abss(int src);	Creates a saturated 16-bit absolute value. _abss(0x8000) => 0x7FFF (SATA set)

Table 3–2. TMS320C55x C Compiler Intrinsics (Continued)

Intrinsic	C Compiler Intrinsic Description
long _labss(long src);	Creates a saturated 32-bit absolute value. _labss(0x8000000) => 0x7FFFFFFF (SATD set)
int _sneg(int src);	Negates the 16-bit value with saturation. _sneg(0xffff8000) => 0x00007FFF
long _lsneg(long src);	Negates the 32-bit value with saturation. _lsneg(0x80000000) => 0x7FFFFFFF
int _smpyr(int src1, int src2);	Multiplies src1 and src2, shifts the result left by 1, and rounds by adding 2 ¹⁵ to the result. (SATD and FRCT set)
int _smacr(long src, int op1, int op2);	Multiplies op1 and op2, shifts the result left by 1, adds the result to src, and then rounds the result by adding 2 ¹⁵ . (SATD, SMUL and FRCT set)
int _smasr(long src, int op1, int op2);	Multiplies op1 and op2, shifts the result left by 1, subtracts the result from src, and then rounds the result by adding 2 ¹⁵ . (SATD, SMUL and FRCT set)
int _norm(int src);	Produces the number of left shifts needed to normalize src.
int _lnorm(long src);	Produces the number of left shifts needed to normalize src.
int _rnd(long src);	Rounds src by adding 2 ¹⁵ . Produces a 16-bit saturated result. (SATD set)
int _sshl(int src1, int src2);	Shifts src1 left by src2 and produces a 16-bit result. The result is saturated if src2 is less than or equal to 8. (SATD set)
long _lsshl(long src1, int src2);	Shifts src1 left by src2 and produces a 32-bit result. The result is saturated if src2 is less than or equal to 8. (SATD set)
int _shrs(int src1, int src2);	Shifts src1 right by src2 and produces a 16-bit result. Produces a saturated 16-bit result. (SATD set)
long _lshrs(long src1, int src2);	Shifts src1 right by src2 and produces a 32-bit result. Produces a saturated 32-bit result. (SATD set)
int _addc(int src1, int src2);	Adds src1, src2, and Carry bit and produces a 16-bit result.
long _laddc(long src1, int src2);	Adds src1, src2, and Carry bit and produces a 32-bit result.
int _subc(long src1, int src2);	Subtracts src2 and logical inverse of sign bit from src1, and produces a 16-bit result.
long _lsubc(long src1, int src2);	Subtracts src2 and logical inverse of sign bit from src1, and produces a 32-bit result.

3.5 Using Long Data Accesses for 16-Bit Data

The primary use of treating short data as long, is in the transfer of data from one memory location to another. Since long accesses also can occur in single cycle this could reduce the data movement time in half. The only limitation is that the data must be aligned on a long boundary (i.e. even word boundary). The code is even simpler if an additional requirement is that the number of items copied is a multiple of 2.

Example 3–9. Block Copy Using Long Data Access

```
void copy(const short *a, const short *b, unsigned short n)
{
    unsigned short i;
    unsigned short na;
    long *src, *dst;
    // This code assumes that the number of elements to transfer 'n'
    // is a multiple of 2. Divide the number of 1 word transfers
    // by 2 to get the number of double word transfers.
    na = (n>>1) -1;
    // Set beginning address of SRC and DST for long transfer.
    src = (long *)a;
    dst = (long *)b;
    for (i=0; i<= na; i++)
    {
        *dst++ = *src++;
    }
}
```

3.6 Generating Efficient Loop Code

You can realize substantial gains from the performance of your C loop code by refining your code in the following areas:

- ☐ Avoid function calls within the body of repeated loops. This enables compiler use of block and local repeat.
- ☐ Keep loop code small to enable compiler use of local repeat
- ☐ Analyzing trip count issues
- ☐ Using the `_nassert` intrinsic
- ☐ Use `-O3` and `-pm` Compiler Options

3.6.1 Avoid function calls in repeated loops

Whenever possible avoid using function calls within repeated loops. Because repeat labels and counts would have to be preserved across calls, the compiler opts to never generate block repeat or local repeat when function calls are present in a loop.

3.6.2 Keep loops small to enable local repeat

Keeping loop code small enables the compiler to make use of the native local repeat instruction. The compiler will generate local repeat for small loops that do not contain any control flow structures other than forward conditionals.

Example 3–10. Simple Loop that Allows Use of Local Repeat

```
void vecsum(const short *a, const short *b, short *c, unsigned int n)
{
    unsigned short i;

    for (i=0; i<= n-1; i++)
    {
        *c++ = *a++ + *b++;
    }
}
```

Example 3–11 displays the assembly code generated by the compiler.

Example 3–11. Assembly Code for Local Repeat Generated by the Compiler

```

_vecsum:
    SP = SP - #1
                                ; End Prolog Code

    AR3 = T0 - #1 ; |2|
    BRC0 = AR3 ; |2|
    localrepeat { ; |2|
                                ; loop starts ; |2|
L1:
    AR3 = *AR1+ ; |7|
    AR3 = AR3 + *AR0+ ; |7|
    *AR2+ = AR3 ; |7|
    }                            ; loop ends ; |8|
L2:
                                ; Begin Epilog Code

    SP = SP + #1
    return
                                ; return occurs

```

3.6.3 Trip Count Issues

A trip count is the number of times that a loop executes; the trip counter is the variable used to count each iteration. When the trip counter reaches the limit equal to the trip count, the loop terminates. Maximum performance for loop code is gained when the compiler can determine the exact minimum and maximum for the trip count. To this end, when invoking the compiler, use the following options to convey trip count information to the compiler:

- ☐ Use unsigned integer type for trip counter variable, whenever possible.
- ☐ Use `<=` comparison for loop termination, whenever possible.
- ☐ Use the `-o3 -pm` compiler options to allow optimizer to access the whole program or large parts of it and to characterize the behavior of loop trip counts.
- ☐ Use the `_nassert` intrinsic to help reduce code size by preventing the generation of a redundant loop or by allowing the compiler to software pipeline the innermost loop.

3.6.4 Using Unsigned Integer Types for Trip Counter

Using unsigned integer types for the trip counter, and trip maximum can reduce the amount of redundant code needed to check the loop entrance and termination conditions.

In Example 3–12, consider this simple for loop:

```
for(i = 0; i<n; i++)
```

If i has been declared an integer and the loop maximum, n , has been declared an integer, it is possible that the value of n could be negative. Even though the initial value for i is zero, this value could still be greater than the value of a negative n . Therefore the compiler is forced to generate code to test for this condition prior to entering the loop.

Example 3–12. Inefficient Loop Code for Loop Variable and Constraints (C)

```
int sum(const short *a, int n)
{
    int sum = 0;
    int i;
    for(i=0; i<n; i++)
    {
        sum += a[i];
    }
    return sum;
}
```

Example 3–13 displays the resulting assembly code.

Example 3–13. Inefficient Loop Code for Variable and Constraints(Assembly)

```

_sum:
    SP = SP - #1
                                ; End Prolog Code

    AR2 = #0 ; |3|
    if (T0<=#0) goto L1 ; |6|    ;== Note this compare
                                ; branch occurs ; |6|

    AR1 = T0 - #1 ; |6|          ;== Note effect of using <
    CSR = AR1
    repeat(CSR)

                                ; loop starts ; |6|

    AR2 = AR2 + *AR0+ ; |8|

                                ; loop ends

L1:
    T0 = AR2 ; |11|

                                ; Begin Epilog Code

    SP = SP + #1 ; |11|
    return ; |11|

                                ; return occurs ; |11|

```

Notice the comparison to skip around the loop code if n is negative. If the code uses simple counting loops, it is better to use unsigned types, as Example 3–14 and Example 3–15 illustrate:

Example 3–14. Using Unsigned Data Types

```

int sum(const short *a, unsigned int n)
{
    int sum = 0;
    unsigned int i;
    for(i=0; i<=n; i++)
    {
        sum += a[i];
    }
    return sum;
}

```

Example 3–15. Assembly Code Generated for Unsigned Data Type

```

_sum:
    SP = SP - #1
                                ; End Prolog Code

    AR1 = #0 ; |3|
    CSR = T0
    repeat(CSR)
                                ; loop starts ; |3|

    AR1 = AR1 + *AR0+ ; |8|
                                ; loop ends

L1:
    T0 = AR1 ; |11|
                                ; Begin Epilog Code

    SP = SP + #1 ; |11|
    return ; |11|
                                ; return occurs ; |11|

```

3.6.5 Use _nassert Intrinsic

The `_nassert` intrinsic may be used to convey the same kinds of information that were communicated by coding changes or compilation option selection:

For example, we could have used `_nassert` to specify that the loop counter variable was always positive rather than declare the type as unsigned (Example 3–12 through Example 3–15):

Example 3–16. Using _nassert Directive

```

int sum(const short *a, int n)
{
    int sum = 0;
    int i;
    _nassert(n > 0);
    for(i=0; i<=n; i++)
    {
        sum += a[i];
    }
    return sum;
}

```

Example 3–17 shows that the `_nassert` gives the compiler enough information about the loop maximum to eliminate the unnecessary check for loop maximum less than 0.

Example 3–17. Assembly Code Generated With _nassert Directive

```

_sum:
    SP = SP - #1
                                ; End Prolog Code

    AR1 = #0   ; |4|
    CSR = T0
    repeat(CSR)
                                ; loop starts ; |4|

    AR1 = AR1 + *AR0+ ; |11|
                                ; loop ends

L1:
    T0 = AR1   ; |14|
                                ; Begin Epilog Code

    SP = SP + #1 ; |14|
    return     ; |14|
                                ; return occurs ; |14|

```

`_nassert` may be used to communicate many things about the trip counter.

- ❑ It can convey that the trip count will be greater than some minimum value or smaller than some maximum value.

```
/* This loop will always execute at least 30 times */
_nassert( x  >= 30);
for(j=0; j<x; j++)
```

- ❑ It can convey that the trip count is always divisible by a value.

```
/* The trip count will execute some multiple of 4 times
*/
_nassert ( (x % 4) == 0);
for (j=0; j< x; j++)
```

- ❑ It can convey information about alignment of pointers and arrays.

```
void vecsum(short *a, short *b, const short *c)
{
    _nassert(((int) a & 0x3) == 0);
    _nassert(((int) b & 0x3) == 0);
    _nassert(((int) c & 0x3) == 0);
    ...
}
```

Several conditions may be combined within a single `_nassert` statement:

```
_nassert((x >= 8) && ( x<= 48) && (( x % 8) == 0));
for(j=0; j< x; j++)
```

The compiler knows that the loop will execute some multiple of 8 times (between 8 and 48) times. This information is useful in providing more information about unrolling a loop or the ability to perform word accesses on a loop.

3.6.6 Use `-o3` and `-pm` Compiler Options

The `-o3` and `-pm` options may be used to communicate to the optimizer the value of loop maximum/minimum.

The following code shows the effect of using the `-o3` and `-pm` options. The main program in Example 3–18 shows two calls to the `sum` function previously defined in Example 3–14.

Example 3–18. Main Function Calling `sum`

```
extern int sum(const short *a, unsigned int n);
short a[10] = {1,2,3,4,5,6,7,8,9,10};
short b[10] = {11,12,13,14,15,16,17,18,19,20};
int sum1, sum2;
void main(void)
{
    sum1 = sum(a,9);
    sum2 = sum(b,9);
}
```

Example 3–19 shows the assembly output by the compiler when the 'main' function is compiled with just `-o3` option.

Example 3–19. Assembly Code Generated With Main Calling sum

```

; End Prolog Code

AR0 = #_a ; |10|
T0 = #9   ; |10|
call #_sum ; |10|

; call occurs [#_sum] ; |10|

*abs16(#_sum1) = T0 ; |10|
AR0 = #_b ; |11|
T0 = #9   ; |11|
call #_sum ; |11|

; call occurs [#_sum] ; |11|

*abs16(#_sum2) = T0 ; |11|

; Begin Epilog Code

SP = SP + #1
return

; return occurs

_main:
SP = SP - #1

```

Example 3–20 shows the code generated when using both `-o3` and `-pm` options, because the source code is merged into a single compilation unit, the code for both functions appears in the same assembly source file. The loop in the `sum` function explicitly uses the loop count as an immediate value. This optimization is possible because the optimizer can directly determine the loop counter value. Note that the loop count is no longer passed as a parameter to the `sum` function when it is called in `main`.

Example 3–20. Assembly Source Output Using –o3 and –pm Options

```

_sum:
    SP = SP - #1
                                ; End Prolog Code
; **Parameter deleted: n == 9u;
    T0 = #0    ; |3|
    repeat(#9)                                ; Note explicit use of count
                                                ; loop starts ; |3|

    T0 = T0 + *AR0+ ; |8|
                                                ; loop ends

L1:
                                ; Begin Epilog Code

    SP = SP + #1 ; |11|
    return      ; |11|
                                ; return occurs ; |11|

_main:
    SP = SP - #1
                                ; End Prolog Code

    AR0 = #_a ; |10|
    call #_sum ; |10|
                                ; call occurs [#_sum] ; |10|

    *abs16(#_sum1) = T0 ; |10|
    AR0 = #_b ; |11|
    call #_sum ; |11|
                                ; call occurs [#_sum] ; |11|

    *abs16(#_sum2) = T0 ; |11|
                                ; Begin Epilog Code

    SP = SP + #1
    return
                                ; return occurs

```


3.7 Generating Efficient Control Code

The compiler generates similar constructs when implementing nested if-then-else and switch/case constructs when the number of case labels is less than eight. Because the first true condition is executed with the least amount of branching, it is best to allocate the most often executed conditional first. When the number of case labels exceeds eight, the compiler generates a .switch label section. In this case, it is still optimal to place the most often executed code at the first case label.

3.8 Efficient Math Operations

3.8.1 Use 16-bit Data Types Whenever Possible

Although memory accesses to long data types can occur in single cycle, arithmetic on long data types such as multiply, division, modulo, may require additional instructions and cycles to implement. Many long operations result in calls to run-time library routines that perform the indicated operation. Therefore, it is best to use a 16-bit data type whenever possible.

3.8.2 Special Considerations When Using MAC Constructs

The compiler can generate single repeat MAC operations. To facilitate the generation of single repeat MACs, use local rather than global variables for the summation. If a global variable is used, the compiler is obligated to perform an intervening store to the global object. This prevents it from generating a single repeat.

Example 3–21. Use Local Rather Than Global Summation Variables

Not Recommended:

```
int gsum = 0;

void dotp1(const short *x, const short *y, unsigned short n)
{
    unsigned int i;
    for(i=0; i<=n; i++)
        gsum += x[i] * y[i];
}
```

Recommended:

```
int dotp2(const short *x, const short *y, unsigned short n)
{
    unsigned int i;
    int lsum = 0;
    for(i=0; i<=n; i++)
        lsum += x[i] * y[i];
    return lsum;
}
```

In the case where Q15 arithmetic is being simulated, the result of the MAC operation may be accumulated into a long object. The result may then be shifted and truncated before return.

Example 3–22. Returning Q15 Result for Multiply Accumulate

```
int dotp (const short *x, const short *y, unsigned short n)
{
    unsigned int i;
    long sum = 0;
    for (i=0; i<=n; i++)
    {
        sum += x[i]*y[i];
    }
    return (int)((sum >> 15) & 0x0000FFFFu);
}
```

3.8.3 Avoid Using Modulus Operator When Simulating Circular Addressing in C

When simulating circular addressing in C avoid using the modulus operator. This can take several cycles to implement and may even result in a call to run-time library routines for modulus. Instead use code similar to that in Example 3–23.

Example 3–23. Simulating Circular Addressing in C

```
#define CIRC_UPDATE(arr,idx,inc,size)\
    (\
        idx = (idx+inc) & (size),\
        arr[idx]\
    )

long circ(const short *a, const short *h, short *r, short nh, short na)
{
    unsigned short i;
    unsigned long sum;
    short x = 0, nx = nh-1;

    sum = 0;
    for(i=0; i<na-1; i++)
    {
        sum += a[i] * CIRC_UPDATE(h,x,1,nx);
    }
    return sum;
}
```

Example 3–24 displays the resulting assembly generated by the compiler.

Example 3–24. Assembly Output for Circular Addressing C Code

```

_circ:
    SP = SP - #1
                                     ; End Prolog Code

    AR2 = T0   ; |9|
    T0 = AR1   ; |9|
    AR2 = AR2 - #1 ; |12|
    AR1 = T1 - #1 ; |15|
    AC0 = #0   ; |14|
    if (AR1==#0) goto L2 ; |15|
                                     ; branch occurs ;

|15|
    AR1 = AR1 - #1 ; |15|
    AR3 = #0   ; |12|
    BRC0 = AR1 ; |12|
    localrepeat { ; |12|
                                     ; loop starts ; |12|
L1:
    AR3 = AR3 + #1 ; |17|
    AR3 = AR3 & AR2 ; |17|
    T1 = *AR3(T0) ; |17|
    AC0 = AC0 + (T1 * *AR0+) ; |17|
    }                                     ; loop ends ; |18|
L2:
                                     ; Begin Epilog Code

    SP = SP + #1 ; |19|
    return      ; |19|
                                     ; return occurs ;

|19|

```

Example 3–25. Circular Addressing Using Modulus Operator

```
long modulo(const short *a, const short *h, short *r, short nh, short na)
{
    unsigned short i;
    unsigned long sum;
    short x = 0;

    sum = 0;
    for(i=0; i<na-1; i++)
    {
        ++x;
        sum += a[i] * h[x % nh];
    }

    return sum;
}
```

The resulting assembly code is displayed in Example 3–26:

Example 3–26. Assembly Output for Circular Addressing Using Modulo

```

_modulo:
    push(T2)
    push(AR5,T3) ;
    push(AR7,AR6) ;
    SP = SP - #2
                                                    ; End Prolog Code

    T2 = T0      ; |2|
    AR5 = AR0    ; |2|
    T3 = AR1     ; |2|
    AR6 = T1 - #1 ; |9|
    AC0 = #0     ; |7|
    dbl(*SP(#0)) = AC0 ; |7|
    if (AR6==#0) goto L2 ; |9|
                                                    ; branch occurs ; |9|

L1:
    AR7 = #0     ; |5|

    AR7 = AR7 + #1 ; |12|
    T0 = AR7     ; |12|
    T1 = T2      ; |12|
    call #I$$MOD ; |12|
                                                    ; call occurs [#I$$MOD] ; |12|

    AR3 = T0     ; |12|
    T0 = T3      ; |12|
    T1 = *AR3(T0) ; |12|
    AC0 = dbl(*SP(#0)) ; |12|
    AC0 = AC0 + (T1 * *AR5+) ; |12|
    dbl(*SP(#0)) = AC0 ; |12|
    AR6 = AR6 - #1 ; |13|
    if (AR6!=#0) goto L1 ; |13|
                                                    ; branch occurs ; |13|

L2:
    AC0 = dbl(*SP(#0))

                                                    ; Begin Epilog Code

    SP = SP + #2 ; |15|
    AR7 = pop() ; |15|
    AR6,AR5 = pop() ;
    T3,T2 = pop() ;
    return      ; |15|
                                                    ; return occurs ; |15|

```

3.9 Memory Management

This section provides a brief discussion on managing memory. Memory usage and subsequent code speed may be affected by a number of factors. The discussion in this chapter will focus on the following areas that affect memory usage:

- ☐ Avoiding holes caused by data alignment
- ☐ Local vs. global symbol declarations
- ☐ Stack configuration
- ☐ Allocating code and data in the C55x memory map

3.9.1 Avoiding Holes Caused by Data Alignment

The compiler requires that all long data be stored on an even word boundary. When declaring data objects (such as structures) that may contain a mixture of multi-word and single word elements, place the long data items in the structure definition first to avoid holes in memory. The compiler automatically aligns structure objects on an even word boundary. Placing these items first takes advantage of this default alignment.

Example 3–27. Considerations for Long Data Objects in Structures

```
Not recommended:
typedef struct _abc{
    int a;
    long b;
    int c;
} ABC;

Recommended:
typedef struct _abc{
    long a;
    int b,c;
}ABC;
```


3.9.2 Local vs Global Symbol Declarations

Locally declared symbols (symbols declared within a C function), are allocated space by the compiler on the software stack. Globally declared symbols (symbols declared at file level) are allocated space in the compiler generated .bss section by default. The C operating environment created by the C boot routine, `_c_int00`, places the C55x DSP in CPL mode. CPL mode enables stack based offset addressing and disables DP offset addressing. The compile accesses Global objects via absolute addressing modes. Because the full address of the global object is encoded as part of the instruction in absolute addressing modes, this can lead to larger code size and potentially slower code. CPL mode favors the use of locally declared objects, since it takes advantage of stack offset addressing. Therefore, if at all possible, it is better to declare and manipulate local objects rather than global objects. When function code requires multiple use of a non-volatile global object, it is better to declare a local object and assign it the appropriate value:

```
extern int Xflag;
int function(void)
{
    int lflag = Xflag;

    .
    x = lflag ? lflag & 0xfffe : lflag;
    .
    .
    return x;
}
```

3.9.3 Stack Configuration

The C55x has dual software stacks: the data stack pointer (SP) and the system stack pointer (SSP). These stacks can be indexed independently or simultaneously depending on the chosen operating mode. There are three possible operating modes for the stack:

- ☐ Dual 16-bit stack with fast return
- ☐ Dual 16-bit stack with slow return
- ☐ 32-bit stack with slow return.

The default mode is 32-bit stack with slow return. In this mode the SSP is incremented whenever SP is incremented. The primary use of SSP is to hold the upper 8 bits of the return address for context save. It is not used for data accesses. Because the C compiler allocates space on the data stack for all locally declared objects, operating in this mode doubles the space allocated for

each local object. This can rapidly increase memory usage. In dual 16-bit modes, the SSP is only incremented for context save (function calls, interrupt handling). Allocation of memory for local objects does not affect the system stack when either of the dual 16-bit modes is used.

Additionally, the selection of fast return mode enables use of the RETA and CFCT registers to effect return from functions. This potentially increases execution speed since it reduces the number of cycles required to return from a function. It is recommended to use dual 16-bit fast return mode to reduce memory space requirements and increase execution speed. The stack operating mode is selected by setting bits 28 and 29 of the reset vector address to the appropriate values. Dual 16-bit fast return mode may be selected by using the `.ivec` assembler directive when creating the address for the reset vector. For example:

```
.ivec      reset_isr_addr, USE_RETA
```

The assembler will automatically set the correct value for bits 28 and 29 when encoding the reset vector address.

3.9.4 Allocating Code and Data in the TMS320C55x Memory Map

The compiler groups generated code and data into logical units called sections. Sections are the building blocks of the common object file format (COFF) files created by the assembler. They are the logical units operated on by the linker when allocating space for code and data in the C55x memory map.

The compiler/assembler can create any of the following sections:

Table 3–3. Section Descriptions

Section	Description
<code>.cinit</code>	Initialization record table for global and static C variables
<code>.const</code>	Explicitly initialized global and static const symbols
<code>.text</code>	Executable code and constants
<code>.bss</code>	Global and static variables
<code>.ioport</code>	Uninitialized global and static variables of type <code>ioport</code>
<code>.stack</code>	Data stack (local variables, lower 16 bits of return address, etc.)
<code>.sysstack</code>	System stack (upper 8 bits of 24 bit return address)
<code>.sysmem</code>	Memory for dynamic allocation functions
<code>.switch</code>	Labels for switch/case
<code>.cio</code>	For CIO Strings and buffers

These sections are encoded in the COFF object file produced by the assembler. When linking the COFF objects, it is important to pay attention to where these sections are linked in memory to avoid as many memory conflicts as possible. Following are some recommendations:

- ❑ Allocate `.stack` and `.sysstack` in DARAM: the `.stack` and `.sysstack` sections are often accessed at the same time when a function call/return occurs. If these sections are allocated in the same SARAM block, then a memory conflict will occur adding additional cycles to the call/return operation. If they are allocated in DARAM or separate SARAM blocks, this will avoid such conflict.
- ❑ The location assigned to `.stack` and `.sysstack` sections are used to initialize the CPU data stack (SP) and system stack (SSP) registers, respectively. Because these two registers share a common data page pointer register (SPH) these sections must be allocated on the same 64k memory page.
- ❑ Similarly, local variable space is allocated on the stack, it is possible that there may be conflicts when global variables, whose allocation is in `.bss` section are accessed within the same instruction as a locally declared variable. Therefore, it may be best to allocate `.bss` and `.stack` in a single DARAM or separate SARAM memory spaces.
- ❑ Use the `DATA_SECTION` pragma: If an algorithm uses a set of coefficients that is applied to a known data array, use the `DATA_SECTION` pragma to place these variables in their own named section. Then explicitly allocate these sections in separate memory blocks to avoid conflicts. Example 3–28 shows sample C source for using the `DATA_SECTION` pragma to place variables in a user defined section.

Example 3–28. Declaration Using `DATA_SECTION` Pragma

```
#pragma DATA_SECTION(h, "coeffs")
short h[10];

#pragma DATA_SECTION(x, "mydata")
short x[10];
```

Most of the memory allocation recommendations are based on the assumption that the typical operation accesses at most two operands. Table 3–4 shows the possible operand combinations:

Table 3–4. Possible Operand Combinations

Operand 1	Operand 2	Comment
Local var (stack)	Local var (stack)	If stack is in DARAM then no memory conflict will occur
Local var(stack)	Global var(.bss)	If stack is in separate SARAM block or is in same DARAM block, then no conflict will occur
Local var(stack)	Const symbol (.const)	If .const is located in separate SARAM or same DARAM no conflict will occur
Global var(.bss)	Global var(.bss)	If .bss is allocated in DARAM, then no conflict will occur
Global var(.bss)	Const symbol(.const)	If .const and .bss are located in separate SARAM or same DARAM block, then no conflict will occur

When compiling with small memory model (compiler default) allocate all data sections, .bss, .stack, .sysmem, .sysstack, and .const on the first 64K page of memory (Page 0).

Example 3–29 contains a sample linker command file:

Example 3–29. Sample Linker Command File

```

/*****
    LINKER command file for LEAD3 memory map.
    Small memory model
*****/
MEMORY
{
    PAGE 0:
        MMR      : origin = 0000000h, length = 00000c0h
        SPRAM     : origin = 00000c0h, length = 0000040
        DARAM0    : origin = 0000100h, length = 0003F00h
        DARAM1    : origin = 0004000h, length = 0004000h
        DARAM2    : origin = 0008000h, length = 0004000h
        DARAM3    : origin = 000c000h, length = 0004000h
        SARAM0    : origin = 0010000h, length = 0002000h
        SARAM1    : origin = 0012000h, length = 0006000h
        SARAM2    : origin = 0018000h, length = 0004000h
        SARAM3    : origin = 001c000h, length = 0004000h
        SARAM4    : origin = 0020000h, length = 0004000h
        SARAM5    : origin = 0024000h, length = 0004000h
        SARAM6    : origin = 0028000h, length = 0004000h
        SARAM7    : origin = 002c000h, length = 0004000h
        SARAM8    : origin = 0030000h, length = 0004000h
        SARAM9    : origin = 0034000h, length = 0004000h
        SARAM10   : origin = 0038000h, length = 0004000h
        SARAM11   : origin = 003c000h, length = 0004000h
        SARAM12   : origin = 0040000h, length = 0004000h
        SARAM13   : origin = 0044000h, length = 0004000h
        SARAM14   : origin = 0048000h, length = 0004000h
        SARAM15   : origin = 004c000h, length = 0004000h
        CE0       : origin = 0050000h, length = 03b0000h
        CE1       : origin = 0400000h, length = 0400000h
        CE2       : origin = 0800000h, length = 0400000h
        CE3       : origin = 0c00000h, length = 03f8000h
        PDRAM     : origin = 0ff8000h, length = 07f00h
        VECS      : origin = 0ffff00h, length = 00100h /* reset vector */
    }
SECTIONS
{
    .vectors : { } > VECS PAGE 0 /* interrupt vector table */
    .cinit   : { } > SARAM0 PAGE 0 /* C initialization table */
    .text    : { } > SARAM1 PAGE 0 /* Code */
    .stack   : { } > DARAM0 PAGE 0 /* Data Stack */
    .sysstack : { } > DARAM0 PAGE 0 /* System Stack */
    .systemem : { } > DARAM1 PAGE 0 /* Dynamic Memory Allocation (heap) */
    .data    : { } > DARAM1 PAGE 0 /* Assembly data section */
    .bss     : { } > DARAM1 PAGE 0 /* C global and static variables */
    .const   : { } > DARAM1 PAGE 0 /* Explicitly declared C const symbols */
}

```

3.10 Allocating Function Code to Different Sections

The compiler provides a pragma to allow the placement of a function's code into a separate user defined section. The pragma is useful if it is necessary to have some granularity in the placement of code in memory.

Example 3–30. Allocation of Functions Using CODE_SECTION Pragma

```
#pragma CODE_SECTION(myfunction, "myfunc")  
void my function(void)  
{  
    .  
    .  
}
```

The pragma in Example 3–30, defines a new section called .myfunc. The code for the function myfunc will be placed by the compiler into this newly defined section. The section name can then be used within the SECTIONS directive of a linker command file to explicitly allocate memory for this function. For details on how to use the SECTIONS directive, see the *TMS320C55x Assembly Language Tools User's Guide*.

Optimizing Assembly Code

This chapter offers recommendations for producing TMS320C55x™ (C55x™) assembly code that:

- ❑ Makes good use of special architectural features, like the dual multiply-and-accumulate (MAC) hardware, parallelism, and looping hardware.
- ❑ Produces no pipeline conflicts, memory conflicts, or instruction-fetch stalls that would delay CPU operations.

This chapter shows ways you can optimize TMS320C55x assembly code, so that you have highly-efficient code in time-critical portions of your programs.


Topic	Page
4.1 Efficient Use of the Dual MAC Hardware	4-2
4.2 Using Parallel Execution Features	4-20
4.3 Implementing Efficient Loops	4-46
4.4 Minimizing Instruction Pipeline Delays	4-53

4.1 Efficient Use of the Dual-MAC Hardware

This section describes methods that help you develop assembly language code to efficiently use the dual multiply-and-accumulate (dual-MAC) hardware.

The two MAC units on the C55x DSP are economically fed data via three independent data buses: BB (the B bus), CB (the C bus), and DB (the D bus). During a dual-MAC operation, each MAC unit requires two data operands from memory (four operands total). However, the three data buses are capable of providing at most three independent operands. To obtain the required fourth operand, the data value on the B bus is used by both MAC units. This is illustrated in Figure 4–1. With this structure, the fourth data operand is not independent, but rather is dependent on one of the other three operands.

Figure 4–1. Data Bus Usage During a Dual-MAC Operation



In the most general case of two multiplications, one would expect a requirement of four fully independent data operands. While this is true on the surface, in most cases one can get by with only three independent operands and avoid degrading performance by specially structuring the DSP code at either the algorithm or application level. The special structuring can be categorized as follows:

- ☐ Multi-channel applications
- ☐ Multi-algorithm applications
- ☐ Implicit algorithm symmetry
- ☐ Loop unrolling

The specifics of each category are discussed in sections 4.1.2 through 4.1.5. Section 4.1.1 provides background information about the data-memory pointers used in dual-MAC operations.

4.1.1 Data Memory Pointer Usage

Nine different registers on the C55x DSP provide address sourcing for data memory operations using indirect addressing (pointer addressing). These registers are XAR0 through XAR7, and the XCDP register. During dual-MAC operations, addressing for the C bus and the D bus comes from two of XAR0 through XAR7, while the addressing for the B bus is provided only by the XCDP register. It should be noted that the B bus cannot be used to access external memory. Therefore, any data variable or array addressed using the XCDP register must be located in internal DSP memory. Refer to the specific device datasheet for memory map information.

4.1.2 Multi-Channel Applications

In multi-channel applications, the same signal processing is often done on two or more independent data streams. Depending on the specific type of processing being performed, it may be possible to process two channels of data in parallel, one channel in each MAC unit. In this way, a common set of constant coefficients can be shared.

An application example readily amenable to this approach is non-adaptive filtering where the same filter is being applied to two different data streams. Both channels are processed in parallel, one channel in each of the two MAC units. For example, the same FIR filter applied to two different data streams can be represented mathematically by the following expressions:

$$y_1(k) = \sum_{j=0}^{N-1} a_j x_1(k-j) \qquad y_2(k) = \sum_{j=0}^{N-1} a_j x_2(k-j)$$

where

N = Number of filter taps

a_j = Element in the coefficient array

$x_i()$ = Element in the i th vector of input values

$y_i()$ = Element in the i th vector of output values

k = Time index

The value a_j is common to both calculations. The two calculations can therefore be performed in parallel, with the common a_j delivered to the dual-MAC units using the B bus and using XCDP as the pointer. The C bus and the D bus are used along with two XARx registers to access the independent input elements $x_1(k-j)$ and $x_2(k-j)$.

A second example is the correlation computation between multiple incoming data streams and a fixed data vector. Suppose it is desired to compute the correlation between the vectors X_1 and Y , and also between the vectors X_2 and Y . One would need to compute the following for each element in the correlation vectors R_1 and R_2 :

$$r_{x_1y}(j) = \sum_k x_1(k+j)y(k) \quad r_{x_2y}(j) = \sum_k x_2(k+j)y(k)$$

The element $y(k)$ is common to both calculations. The two calculations can therefore be performed in parallel, with the common data $y(k)$ delivered to the dual-MAC units via the B bus with XCDP as the address pointer. The C bus and the D bus are used along with two XARx registers to access the independent elements $x_1(k+j)$ and $x_2(k+j)$.

4.1.3 Multi-Algorithm Applications

When two or more different processing algorithms are applied to the same data stream, it may be possible to process two such algorithms in parallel. For example, consider a statistical application that computes the autocorrelation of a vector X , and also the correlation between vector X and vector Y . One would need to compute the following for each element in the correlation vectors R_{xx} and R_{xy} :

$$r_{xx}(j) = \sum_k x(k+j)x(k) \quad r_{xy}(j) = \sum_k x(k+j)y(k)$$

The element $x(k+j)$ is common to both calculations. The two calculations can therefore be made in parallel, with the common data $x(k+j)$ delivered to the dual-MAC units using the B bus and using XCDP as the pointer. The C bus and the D bus are used along with two XARx registers to access the independent elements $x(k)$ and $y(k)$.

4.1.4 Implicit Algorithm Symmetry

When an algorithm has internal symmetry, it can sometimes be exploited for efficient dual-MAC implementation. One such example is a symmetric FIR filter. This filter has coefficients that are symmetrical with respect to delayed values of the input signal. The mathematical expression for a symmetric FIR filter can be described by the following discrete-time difference equation:

$$y(k) = \sum_{j=0}^{\frac{N}{2}-1} a_j [x(k-j) + x(k+j-N+1)]$$

where

N = Number of filter taps (even)

x() = Element in the vector of input values

y() = Element in the vector of output values

k = Time index

Similar in form to the symmetrical FIR filter is the anti-symmetrical FIR filter:

$$y(k) = \sum_{j=0}^{\frac{N}{2}-1} a_j [x(k-j) - x(k+j-N+1)]$$

Both the symmetrical and anti-symmetrical FIR filters can be implemented using a dual-MAC approach because only three data values need be fetched per inner loop cycle: a_j , $x(k-j)$, and $x(k+j-N+1)$. The coefficient a_j is delivered to the dual-MAC units using the B bus and using XCDP as the pointer. The C bus and the D bus are used along with two XARx registers to access the independent elements $x(k-j)$ and $x(k+j-N+1)$.

A second example of an algorithm with implicit symmetry is a complex vector multiplication. Let j be the imaginary unit value (that is, square root of -1). The mathematical expression for the algorithm is given by

$$\{C\} = \{A\}^T \times \{B\}$$

where $\{A\}$, $\{B\}$, and $\{C\}$ are complex vectors of length N and $\{A\}^T$ is the transpose of $\{A\}$. The components of $\{A\}$ and $\{B\}$ can be expressed as

$$a_i = a_i^{RE} + j a_i^{IM} \quad b_i = b_i^{RE} + j b_i^{IM} \quad (1 \leq i \leq N)$$

and the expression for each element in $\{C\}$ can be computed as


$$\begin{aligned} c_i &= (a_i^{RE} + j a_i^{IM}) (b_i^{RE} + j b_i^{IM}) \\ &= (a_i^{RE} b_i^{RE} - a_i^{IM} b_i^{IM}) + j (a_i^{RE} b_i^{IM} + a_i^{IM} b_i^{RE}) \end{aligned}$$

The required four multiplications in the above expression can be implemented with two dual-MAC instructions by grouping the multiplications as follows:

1st multiplication group: $a_i^{RE} b_i^{RE}$ and $a_i^{IM} b_i^{RE}$

2nd multiplication group: $a_i^{IM} b_i^{IM}$ and $a_i^{RE} b_i^{IM}$

Each dual-multiply grouping requires only three independent operands. An assembly code example for the complex vector multiply is given in Example 4–1 (part (a) shows algebraic instructions and part (b) shows mnemonic instructions). Note that this particular code assumes the following arrangement in memory for a complex vector:



In addition, the code stores both portions of the complex result to memory at the same time. This requires that the results vector be long-word aligned in memory. One way to achieve this is through use of the alignment flag option with the `.bss` directive, as was done with this code example. Alternatively, one could place the results array in a separate uninitialized named section using a `.usect` directive, and then use the linker command file to force long-word alignment of that section.

Example 4–1. Complex Vector Multiplication Code**(a) Mnemonic Instructions**

```

N      .set      3                      ; Length of each complex vector

      .data
A      .int      1,2,3,4,5,6           ; Complex input vector #1
B      .int      7,8,9,10,11,12        ; Complex input vector #2

;Results are: 0xffff7, 0x0016, 0xffff3, 0x0042, 0xffef, 0x007e

      .bss C, 2*N, ,1                  ; Results vector, long-word aligned

      .text
BCLR ARMS                      ; Clear ARMS bit (select DSP mode)
      .arms_off                      ; Tell assembler ARMS = 0

cplxmul:
      AMOV #A, XAR0                   ; Pointer to A vector
      AMOV #B, XCDP                   ; Pointer to B vector
      AMOV #C, XAR1                   ; Pointer to C vector
      MOV #(N-1), BRC0                ; Load loop counter
      MOV #1, T0                      ; Pointer offset
      MOV #2, T1                      ; Pointer increment

      RPTBLOCAL endloop               ; Start the loop

      MPY *AR0, *CDP+, AC0
      :: MPY *AR0(T0), *CDP+, AC1

      MAS *AR0(T0), *CDP+, AC0
      :: MAC *(AR0+T1), *CDP+, AC1

endloop:
      MOV pair(LO(AC0)), dbl(*AR1+) ; Store complex result
                                      ; End of loop

```

*Example 4–1. Complex Vector Multiplication Code (Continued)**(b) Algebraic Instructions*

```

N      .set      3                      ; Length of each complex vector

      .data
A      .int      1,2,3,4,5,6           ; Complex input vector #1
B      .int      7,8,9,10,11,12        ; Complex input vector #2

;Results are: 0xffff7, 0x0016, 0xffff3, 0x0042, 0xffef, 0x007e

      .bss C, 2*N, ,1                  ; Results vector, long-word aligned

      .text
bit(ST2,#ST2_ARMS) = #0                ; Clear ARMS bit (select DSP mode)
      .arms_off                        ; Tell assembler ARMS = 0

cplxmul:
XAR0 = #A                              ; Pointer to A vector
XCDP = #B                              ; Pointer to B vector
XAR1 = #C                              ; Pointer to C vector
BRC0 = #(N-1)                          ; Load loop counter
T0 = #1                                ; Pointer offset
T1 = #2                                ; Pointer increment

      localrepeat {                    ; Start the loop

AC0 = *AR0 * coef(*CDP+),
AC1 = *AR0(T0) * coef(*CDP+)

AC0 = AC0 - (*AR0(T0) * coef(*CDP+)),
AC1 = AC1 + (*(AR0+T1) * coef(*CDP+))

*AR1+ = pair(LO(AC0))                  ; Store complex result
      }                                ; End of loop

```

4.1.5 Loop Unrolling

Loop unrolling involves structuring computations to exploit the reuse of data among different time or geometric iterations of the algorithm. Many algorithms can be structured computationally to provide for such reuse and allow a dual-MAC implementation.

In filtering, input and/or output data is commonly stored in a delay chain buffer. Each time the filter is invoked on a new data point, the oldest value in the delay chain is discarded from the bottom of the chain, while the new data value is added to the top of the chain. A value in the chain will get reused (for example, multiplied by a coefficient) in the computations over and over again as successive time-step outputs are computed. The reuse will continue until such a time that the data value becomes the oldest value in the chain and is discarded. Dual-MAC implementation of filtering should therefore employ a time-based loop unrolling approach to exploit the reuse of the data. This scenario is presented in sections 4.1.5.1 and 4.1.5.2.

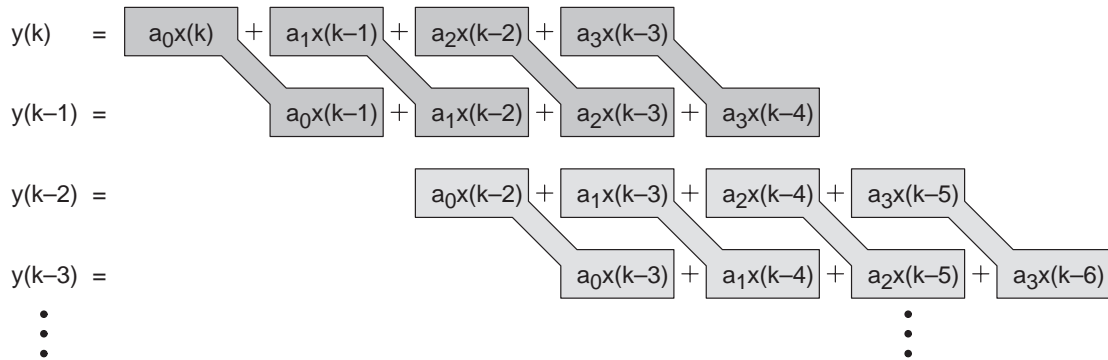
An application amenable to geometric based loop unrolling is matrix computations. In this application, successive rows in a matrix get multiplied and accumulated with the columns in another matrix. In order to obtain data reuse within the loop kernel, the computations using two different rows of data should be handled in parallel. This will be presented in section 4.1.5.3.

4.1.5.1 Temporal Loop Unrolling: Block FIR Filter

To efficiently implement a block FIR filter with the two MAC units, loop unrolling must be applied so that two time-based iterations of the algorithm are computed in parallel. This allows reuse of the coefficients.

Figure 2 illustrates the coefficient reuse for a 4-tap block FIR filter with constant, real-value coefficients. The implementation computes two sequential filter outputs in parallel so that only a single coefficient, a_i , is used by both MAC units. Consider, for example, the computation of outputs $y(k)$ and $y(k-1)$. For the first term in each of these two rows, one MAC unit computes $a_0x(k)$, while the second MAC unit computes $a_0x(k-1)$. These two computations combined require only three different values from memory: a_0 , $x(k)$, and $x(k-1)$. Proceeding to the second term in each row, $a_1x(k-1)$ and $a_1x(k-2)$ are computed similarly, and so on with the remaining terms. After fully computing the outputs $y(k)$ and $y(k-1)$, the next two outputs, $y(k-2)$ and $y(k-3)$, are computed in parallel. Again, the computation begins with the first two terms in each of these rows. In this way, DSP performance is maintained at two MAC operations per clock cycle.

Figure 4–2. Computation Groupings for a Block FIR (4-Tap Filter Shown)



Note that filters with either an even or odd number of taps are handled equally well by this method. However, this approach does require one to compute an even number of outputs $y()$. In cases where an odd number of outputs is desired, one can always zero-pad the input vector $x()$ with one additional zero element, and then discard the corresponding additional output.

Note also that not all of the input data must be available in advance. Rather, only two new input samples are required for each iteration through the algorithm, thereby producing two new output values.

A non-optimized assembly code example for the block FIR filter is shown in Example 4–2 (part (a) shows algebraic instructions and part (b) shows mnemonic instructions). An optimized version of the same code is found in Example 4–3 (part (a) shows algebraic instructions and part (b) shows mnemonic instructions). The following optimizations have been made in Example 4–3:

- ☐ The first filter tap was peeled out of the inner loop and implemented using a dual-multiply instruction (as opposed to a dual-multiply-and-accumulate instruction). This eliminated the need to clear AC0 and AC1 prior to entering the inner loop each time.
- ☐ The last filter tap was peeled out of the inner loop. This allows for the use of different pointer adjustments than in the inner loop, and eliminates the need to explicitly rewind the CDP, AR0, and AR1 pointers.

The combination of these first two optimizations results in a requirement that N_TAPS be a minimum of 3.

- ❑ Both results are now written to memory at the same time using a double store instruction. Note that this requires the results array (`OUT_DATA`) to be long-word aligned. One way to achieve this is through use of the alignment flag option with the `.bss` directive, as was done in with this code example. As an alternative, you could place the results array in a separate uninitialized named section using a `.usect` directive, and then use the linker command file to force long-word alignment of that section.
- ❑ The outer loop start instruction (`RPTBLOCAL` in mnemonic syntax, `localrepeat()` in algebraic syntax) has been put in parallel with the instruction that preceded it.

Example 4–2. Block FIR Filter Code (Not Optimized)**(a) Mnemonic Instructions**

```

N_TAPS    .set    4                      ; Number of filter taps
N_DATA    .set    11                    ; Number of input values

        .data
COEFFS    .int    1,2,3,4                ; Coefficients
IN_DATA    .int    1,2,3,4,5,6,7,8,9,10,11 ; Input vector

;Results are:    0x0014, 0x001E, 0x0028, 0x0032,
;                0x003C, 0x0046, 0x0050, 0x005A

        .bss    OUT_DATA, N_DATA - N_TAPS + 1 ; Output vector

        .text
BCLR ARMS                      ; Clear ARMS bit (select DSP mode)
.arms_off                      ; Tell assembler ARMS = 0

bfir:
        AMOV #COEFFS, XCDP          ; Pointer to coefficient array
        AMOV #(IN_DATA + N_TAPS - 1), XAR0 ; Pointer to input vector
        AMOV #(IN_DATA + N_TAPS), XAR1    ; 2nd pointer to input vector
        AMOV #OUT_DATA, XAR2           ; Pointer to output vector
        MOV  #((N_DATA - N_TAPS + 1)/2 - 1), BRC0
                                   ; Load outer loop counter
        MOV  #(N_TAPS - 1), CSR        ; Load inner loop counter

        RPTBLOCAL endloop           ; Start the outer loop

        MOV  #0, AC0                  ; Clear AC0
        MOV  #0, AC1                  ; Clear AC1

        RPT CSR                      ; Start the inner loop
        MAC  *AR0-, *CDP+, AC0        ; All taps
        :: MAC *AR1-, *CDP+, AC1

        MOV  AC0, *AR2+                ; Write 1st result
        MOV  AC1, *AR2+                ; Write 2nd result
        MOV  #COEFFS, CDP             ; Rewind coefficient pointer
        ADD  #(N_TAPS + 2), AR0        ; Adjust 1st input vector
                                   ; pointer
endloop:
        ADD  #(N_TAPS + 2), AR1        ; Adjust 2nd input vector
                                   ; pointer
                                   ; End of outer loop

```

Example 4–2. Block FIR Filter Code (Not Optimized) (Continued)**(b) Algebraic Instructions**

```

N_TAPS      .set      4                      ; Number of filter taps
N_DATA      .set      11                    ; Number of input values

      .data
COEFFS      .int      1,2,3,4              ; Coefficients
IN_DATA     .int      1,2,3,4,5,6,7,8,9,10,11 ; Input vector

;Results are:      0x0014, 0x001E, 0x0028, 0x0032,
;                  0x003C, 0x0046, 0x0050, 0x005A

      .bss      OUT_DATA, N_DATA - N_TAPS + 1 ; Output vector

      .text
bit(ST2,#ST2_ARMS) = #0                    ; Clear ARMS bit (select DSP mode)
      .arms_off                                ; Tell assembler ARMS = 0

bfir:
      XCDP = #COEFFS                          ; Pointer to coefficient array
      XAR0 = #(IN_DATA + N_TAPS - 1)          ; Pointer to input vector
      XAR1 = #(IN_DATA + N_TAPS)             ; 2nd pointer to input vector
      XAR2 = #OUT_DATA                       ; Pointer to output vector
      BRC0 = #((N_DATA - N_TAPS + 1)/2 - 1)  ; Load outer loop counter
      CSR = #(N_TAPS - 1)                   ; Load inner loop counter

      localrepeat {                          ; Start the outer loop

      AC0 = #0                               ; Clear AC0
      AC1 = #0                               ; Clear AC1

      repeat(CSR)                            ; Start the inner loop
      AC0 = AC0 + ( *AR0- * coef(*CDP+) ),    ; All taps
      AC1 = AC1 + ( *AR1- * coef(*CDP+) )

      *AR2+ = AC0                            ; Write 1st result
      *AR2+ = AC1                            ; Write 2nd result

      CDP = #COEFFS                          ; Rewind coefficient pointer
      AR0 = AR0 + #(N_TAPS + 2)              ; Adjust 1st input vector
      ; pointer
      AR1 = AR1 + #(N_TAPS + 2)              ; Adjust 2nd input vector
      ; pointer

      }                                       ; End of outer loop

```

Example 4–3. Block FIR Filter Code (Optimized)**(a) Mnemonic Instructions**

```

N_TAPS    .set    4                ; Number of filter taps
N_DATA    .set    11              ; Number of input values

        .data
COEFFS     .int    1,2,3,4        ; Coefficients
IN_DATA    .int    1,2,3,4,5,6,7,8,9,10,11 ; Input vector

;Results are:    0x0014, 0x001E, 0x0028, 0x0032,
;                0x003C, 0x0046, 0x0050, 0x005A

        .bss    OUT_DATA, N_DATA - N_TAPS + 1, ,1
                                ; Output vector, long word
aligned

        .text
BCLR ARMS                ; Clear ARMS bit (select DSP mode)
.arms_off                ; Tell assembler ARMS = 0

bfir:
    AMOV #COEFFS, XCDP        ; Pointer to coefficient array
    AMOV #(IN_DATA + N_TAPS - 1), XAR0 ; Pointer to input vector
    AMOV #(IN_DATA + N_TAPS), XAR1    ; 2nd pointer to input vector
    AMOV #OUT_DATA, XAR2          ; Pointer to output vector
    MOV  #((N_DATA - N_TAPS + 1)/2 - 1), BRC0
                                ; Load outer loop counter
    MOV  #(N_TAPS - 3), CSR       ; Load inner loop counter
    MOV  #(-(N_TAPS - 1)), T0     ; CDP rewind increment

    MOV  #(N_TAPS + 1), T1        ; ARx rewind increment
    ||RPTBLOCAL endloop          ; Start the outer loop

    MPY *AR0-, *CDP+, AC0         ; 1st tap
    :: MPY *AR1-, *CDP+, AC1

    RPT CSR                      ; Start the inner loop
    MAC *AR0-, *CDP+, AC0         ; Inner taps
    :: MAC *AR1-, *CDP+, AC1

    MAC *(AR0+T1), *(CDP+T0), AC0 ; Last tap
    :: MAC *(AR1+T1), *(CDP+T0), AC1

endloop:
    MOV pair(LO(AC0)), dbl(*AR2+) ; Store both results
                                ; End of outer loop

```

*Example 4–3. Block FIR Filter Code (Optimized) (Continued)**(b) Algebraic Instructions*

```

N_TAPS    .set    4                ; Number of filter taps
N_DATA    .set    11              ; Number of input values

        .data
COEFFS    .int    1,2,3,4          ; Coefficients
IN_DATA    .int    1,2,3,4,5,6,7,8,9,10,11 ; Input vector

;Results are:    0x0014, 0x001E, 0x0028, 0x0032,
;               0x003C, 0x0046, 0x0050, 0x005A

        .bss    OUT_DATA, N_DATA - N_TAPS + 1, ,1
                                ; Output vector, long-word
                                ;   aligned

        .text
bit(ST2,#ST2_ARMS) = #0          ; Clear ARMS bit (select DSP mode)
.arms_off                                ; Tell assembler ARMS = 0

bfir:
XCDP = #COEFFS                    ; Pointer to coefficient array
XAR0 = #(IN_DATA + N_TAPS - 1)    ; Pointer to input vector
XAR1 = #(IN_DATA + N_TAPS)        ; 2nd pointer to input vector
XAR2 = #OUT_DATA                  ; Pointer to output vector
BRC0 = #((N_DATA - N_TAPS + 1)/2 - 1)
                                ; Load outer loop counter
CSR = #(N_TAPS - 3)              ; Load inner loop counter
T0 = #(-(N_TAPS - 1))            ; CDP rewind increment

T1 = #(N_TAPS + 1)               ; ARx rewind increment
||localrepeat {
                                ; Start the outer loop

AC0 = *AR0- * coef(*CDP+),        ; 1st tap
AC1 = *AR1- * coef(*CDP+)

        repeat(CSR)            ; Start the inner loop
AC0 = AC0 + ( *AR0- * coef(*CDP+) ), ; Inner taps
AC1 = AC1 + ( *AR1- * coef(*CDP+) )

AC0 = AC0 + ( *(AR0+T1) * coef(*(CDP+T0)) ), ; Last tap
AC1 = AC1 + ( *(AR1+T1) * coef(*(CDP+T0)) )

*AR2+ = pair(LO(AC0))            ; Store both results
}                                ; End of outer loop

```

4.1.5.2 Temporal Loop Unrolling: Single-Sample FIR Filter

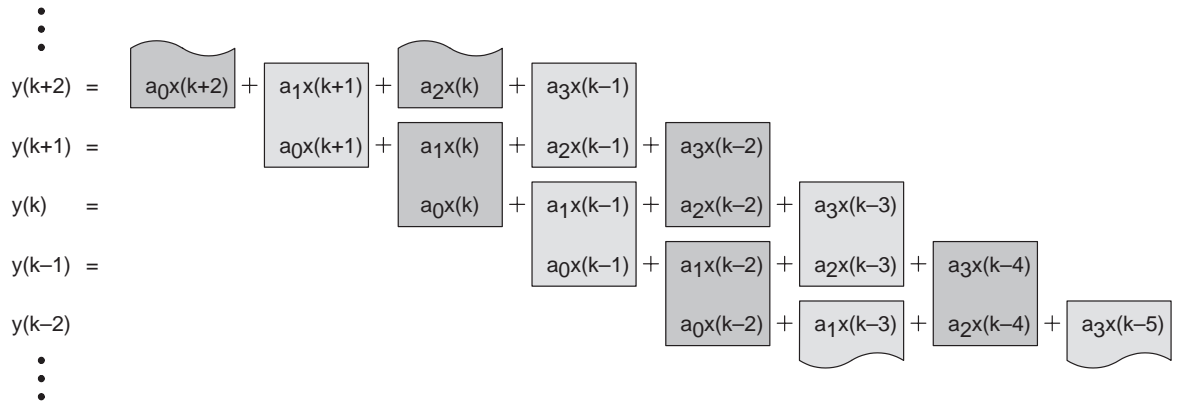
The temporally unrolled block FIR filter described in section 4.1.5.1 maintains dual-MAC throughput by sharing a common coefficient between the two MAC units. In some algorithms, the loop unrolling needs to be performed so that a common data variable is shared instead. The single-sample FIR filter is an example of such an algorithm. In the single-sample FIR filter, the calculations for the current sample period are interlaced with those of the next sample period in order to achieve a net performance of two MAC operations per cycle.

Figure 4–3 shows the needed computation groupings for a 4-tap FIR filter. At any given time step, one multiplies and accumulates every other partial product in the corresponding row, beginning with the first partial product in the row. In addition, one also multiplies and accumulates every other term in the next row (that is, the row above the current row) in advance of that time step, beginning with the second partial product in the next row. In this way, each row is fully computed over the course of two sample periods.

For example, at time step k , it is desired to compute $y(k)$. The first term in the $y(k)$ row is $a_0x(k)$, which is computed using one of the two MAC units. In addition, the second MAC unit is used to compute the second term in the $y(k+1)$ row, $a_1x(k)$, in advance of time step $k + 1$. These two computations combined require only three different values from memory: a_0 , a_1 , and $x(k)$. Note that the term $x(k)$ is not available until time k . This is why calculations at each time step must begin with the first term in the corresponding row.

The second term in the $y(k)$ row is $a_1x(k + 1)$. However, this would have been already computed during the first computation at time step $k - 1$ (similar to how $a_1x(k)$ was just pre-computed for time step $k+1$), so it can be skipped here. The third term in the $y(k)$ row, $a_2x(k - 2)$, is computed next, and at the same time, the term $a_3x(k - 2)$ is computed in the $y(k + 1)$ row in advance of time step $k+1$.

Figure 4–3. Computation Groupings for a Single-Sample FIR With an Even Number of TAPS (4-Tap Filter Shown)



Notice that two separate running sums are maintained, one with partial products for the current time step, the other with pre-computed terms for the next time step. At the next time step, the pre-computed running sum becomes the current running sum, and a new pre-computed running sum is started from zero. At the end of each sample period, the current running sum contains the current filter output, which can be dispatched as required by the application.

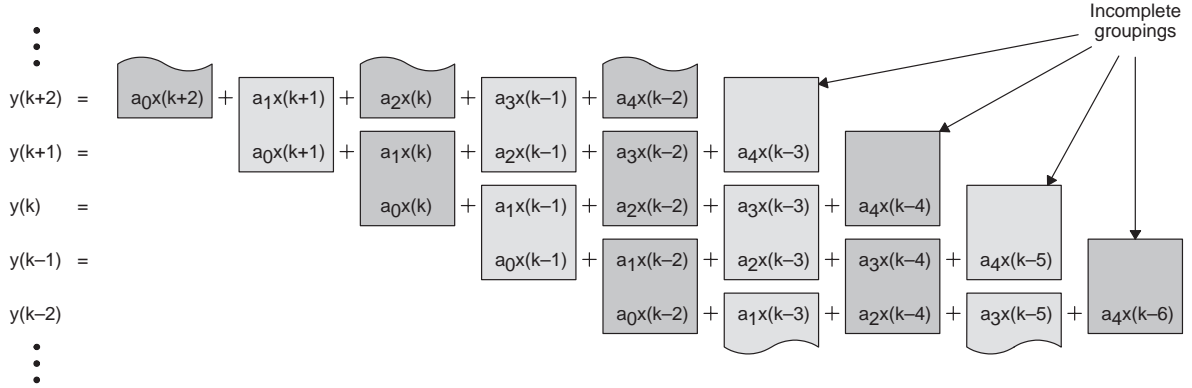
The above approach is not limited to the 4-tap filter illustrated in Figure 4–3. Any other filter with an even number of taps is a straightforward extension. For filters with an odd number of taps, the computation groupings become problematic, in that the last grouping in each row is missing the pre-calculation term in the row above it.

Figure 4–4 depicts this problem for a 5-tap filter. To overcome this problem, one should pad the filter to the next higher even number of taps by using a zero coefficient for the additional term. For example, the five tap filter is augmented to

$$y(k) = a_0x(k) + a_1x(k-1) + a_2x(k-2) + a_3x(k-3) + a_4x(k-4) + 0[x(k-5)]$$

In this way, any filter with an odd number of taps can be implemented as a filter with an even number of taps but retain the frequency response of the original odd-number-tap filter.

Figure 4–4. Computation Groupings for a Single-Sample FIR With an Odd Number of TAPS (5-Tap Filter Shown)



4.1.5.3 Geometric Loop Unrolling: Matrix Mathematics

Matrix mathematics typically involves considerable data reuse. Consider the general case of multiplying two matrices:

$$[C] = [A] \times [B]$$

where

$[A] = m \times n$ matrix

$[B] = n \times p$ matrix

$[C] = m \times p$ matrix

$m \geq 1, n \geq 1, p \geq 1$

The expression for each element in matrix C is given by:

$$c_{i,j} = \sum_{k=1}^n a_{i,k} b_{k,j} \quad (1 \leq i \leq m, 1 \leq j \leq p)$$

where the conventional notation $x_{i,j}$ is being used to represent the element of matrix X in the i th row and j th column. There are basically two different options for efficient dual-MAC implementation. First, one could compute $c_{i,j}$ and $c_{i,j+1}$ in parallel. The computations made are:

$$c_{i,j} = \sum_{k=1}^n a_{i,k} b_{k,j} \quad c_{i,j+1} = \sum_{k=1}^n a_{i,k} b_{k,j+1}$$

The element $a_{i,k}$ is common to both expressions. The computations can therefore be made in parallel, with the common data $a_{i,k}$ delivered to the dual-MAC units using the B bus and using XCDP as the pointer. The C bus and the D bus are used along with two XARx registers to access the independent elements $b_{k,j}$ and $b_{k,j+1}$.

Alternatively, one could compute $c_{i,j}$ and $c_{i+1,j}$ in parallel. The computations made are then:

$$c_{i,j} = \sum_{k=1}^n a_{i,k} b_{k,j} \qquad c_{i+1,j} = \sum_{k=1}^n a_{i+1,k} b_{k,j}$$

In this case, the element $b_{k,j}$ is common to both expressions. They can therefore be made in parallel, with the common data $b_{k,j}$ delivered to the dual-MAC units using the B bus and using XCDP as the pointer. The C bus and the D bus are used along with two XARx registers to access the independent elements $a_{i,k}$ and $a_{i+1,k}$.

The values of m and p determine which approach one should take. Because the inner loop will compute two elements in matrix C each iteration, clearly it is most efficient if an even number of elements can be computed. Therefore, if p is even, one should implement the first approach: compute $c_{i,j}$ and $c_{i,j+1}$ in parallel. Alternatively, if m is even, the second approach is more efficient: compute $c_{i,j}$ and $c_{i+1,j}$ in parallel. If both m and p are even, either approach is appropriate. Finally, if neither m nor p is even, there will be an extra element c that will need to be computed individually each time through the inner loop. One could add additional single-MAC code to handle the final element in the inner loop. Alternatively, one could pad either matrix A or matrix B with a row or column of zeros (as appropriate) to make either m or p even. The elements in matrix C computed using the pad row or column should then be discarded after computation.

4.2 Using Parallel Execution Features

The C55x architecture allows programmers to place two operations or instructions in parallel to reduce the total execution time. There are two types of parallelism: Built-in parallelism within a single instruction and user-defined parallelism between two instructions. Built-in parallelism (see section 4.2.1) is automatic; as soon as you write the instruction, it is put in place. User-defined parallelism is optional and requires decision-making. Sections 4.2.2 through 4.2.8 present the rules and restrictions associated with the use of user-defined parallelism, and give examples of using it.

4.2.1 Built-In Parallelism

Instructions that have built-in parallelism perform two different operations in parallel. In the algebraic syntax, they can be identified by the comma that separates the two operations, as in the following example:

```
AC0 = *AR0 * coef(*CDP),      ; The data referenced by AR0 is multiplied by
AC1 = *AR1 * coef(*CDP)      ; a coefficient referenced by CDP. At the same time
                               ; the data referenced by AR1 is multiplied by the
                               ; same coefficient.
```

In the mnemonic syntax, they can be identified by a double colon (::) that separates the two operations. The preceding example in the mnemonic syntax is:

```
MPY *AR0, *CDP, AC0           ; The data referenced by AR0 is multiplied by
:: MPY *AR1, *CDP, AC1        ; a coefficient referenced by CDP. At the same time
                               ; the data referenced by AR1 is multiplied by the
                               ; same coefficient.
```

4.2.2 User-Defined Parallelism

Two instructions may be placed in parallel to have them both execute in a single cycle. The two instructions are separated by the || separator. One of the two instructions may have built-in parallelism. The following algebraic code example shows a user-defined parallel instruction pair. One of the instructions in the pair also features built-in parallelism.

```
AC0 = AC0 + (*AR3+ * coef(*CDP+)), ; 1st instruction (has built-in parallelism)
AC1 = AC1 + (*AR4+ * coef(*CDP+))
|| repeat(CSR)                    ; 2nd instruction
```

The equivalent mnemonic code example is:

```
MPY *AR3+, *CDP+, AC0           ; 1st instruction (has built-in parallelism)
:: MPY *AR4+, *CDP+, AC1
|| RPT CSR                      ; 2nd instruction
```

4.2.3 Architectural Features Supporting Parallelism

The C55x architecture provides three main, independent computation units that are controlled by the instruction buffer unit (I unit):

- ☐ Program flow unit (P unit)
- ☐ Address-data flow Unit (A unit)
- ☐ Data computation unit (D unit)

The C55x instructions make use of dedicated operative resources (or operators) within each of the units. In total, there are 14 operators available across the three computation units, and the parallelism rules enable the use of two independent operators in parallel within the same cycle. If all other rules are observed, two instructions that independently use any two of the independent operators may be placed in parallel.

Figure 4–5 shows a matrix that reflects the 14 operators mentioned and the possible operator combinations that may be used in placing instructions in parallel. The operators are ordered from rows 1 through 14 as well as columns 1 through 14. A blank cell in any given position (row I, column J) in the matrix indicates that operator I may be placed in parallel with operator J, and an X in any given position indicates that the two operators cannot be placed in parallel. For example, a D-Unit MAC operation (row 7) may be placed in parallel with a P-Unit Load operation (column 13) but cannot be placed in parallel with a D-Unit ALU operation (column 5).

Figure 4–5. Matrix to Find Operators That Can Be Used in Parallel

		A-unit ALU	A-unit Swap	A-unit Load	A-unit Store	D-unit ALU	D-unit Shifter	D-unit MAC	D-unit Load	D-unit Store	D-unit Shift, Store	D-unit Swap	P-unit Control	P-unit Load	P-unit Store
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
A-unit ALU	1	X													
A-unit Swap	2		X												
A-unit Load	3														
A-unit Store	4														
D-unit ALU	5					X	X	X							
D-unit Shifter	6					X	X	X			X				
D-unit MAC	7					X	X	X							
D-unit Load	8														
D-unit Store	9														
D-unit Shift, Store	10						X				X				
D-unit Swap	11											X			
P-unit Control	12												X		
P-unit Load	13														
P-unit Store	14														

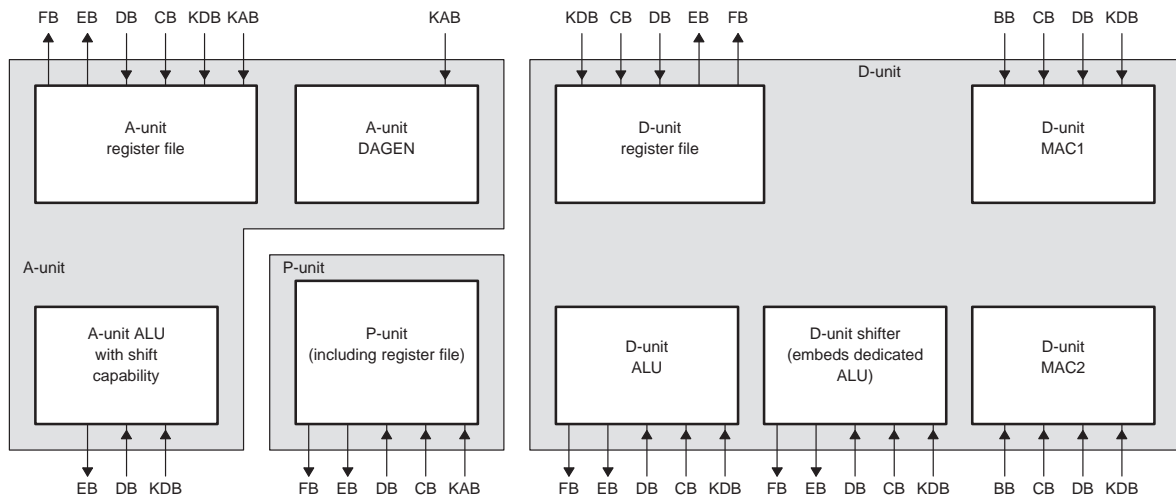
Note: X in a table cell indicates that the operator in that row and the operator in that column cannot be used in parallel with each other. A blank table cell indicates that the operators can be used in parallel.

Bus resources also play an important part in determining whether two instructions may be placed in parallel. Typically, programmers should be concerned with the data buses and the constant buses. Table 4–1 lists and describes the main CPU buses of interest and gives examples of instructions that use the different buses. These may also be seen pictorially in Figure 4–6. Figure 4–6 also shows all CPU buses and the registers/operators in each of the three functional units.

Table 4–1. CPU Data Buses and Constant Buses

Bus Type	Bus(es)	Description of Bus(es)	Example: Instruction That Uses The Bus(es)
Data	BB	Special coefficient read bus	$AC0 = (*AR1+) * coef(*CDP+),$ $AC1 = (*AR3+) * coef(*CDP+)$ The operand referenced by CDP is carried to the CPU on BB.
	CB, DB	Data-read buses	$AC0 = *AR3+$ The operand referenced by AR3 is carried to the CPU on DB.
	EB, FB	Data-write buses	$*AR3- = LO(AC0)$ The low half of AC0 is carried on EB to the location referenced by AR3.
Constant	KAB	Constant bus used in the address stage of the pipeline to carry addresses: <input type="checkbox"/> The P unit uses KAB to generate program-memory addresses. <input type="checkbox"/> The A unit uses KAB to generate data-memory addresses.	P-unit use: $Goto \#Routine2$ The constant Routine2 is carried on KPB to the P unit, where it is used for a program-memory address. A-unit use: $BRC0 = *SP(\#7)$ The immediate offset (7) is carried to the data-address generation unit (DAGEN) on KAB.
	KDB	Constant bus used by the A unit or D unit for computations. This bus is used in the execute stage of the instruction pipeline.	$AC0 = AC0 + \#1234h$ The constant 1234h is carried on KDB to the D-unit ALU, where it is used in the addition.

Figure 4–6. CPU Operators and Buses



4.2.4 User-Defined Parallelism Rules

This section describes the rules that a programmer must follow to place two instructions in parallel. It is essential to note here that all the rules must be observed for the parallelism to be valid. However, this section begins with a set of four basic rules (Table 4–2) that a programmer may use when implementing user-defined parallelism. If these are not sufficient, the set of advanced rules (Table 4–3) needs to be considered.

Table 4–2. Basic Parallelism Rules

Consideration	Rule
Hardware resource conflicts	Two instructions in parallel cannot compete for operators (see Figure 4–5, page 4-22) or buses (see Table 4–1, page 4-23).
Maximum instruction length	The combined length of the instruction pair cannot exceed 6 bytes.
Parallel enable bit OR Soft dual encoding	<p>If either of the following cases is true, the instructions can be placed in parallel:</p> <ul style="list-style-type: none"> <input type="checkbox"/> Parallel enable bit is present: At least one of two instructions in parallel must have a parallel enable bit in its instruction code. The instruction set reference guides (see <i>Related Documentation from Texas Instruments</i> in the preface) indicate whether a given instruction has a parallel enable bit. <input type="checkbox"/> Soft dual encoding is present: For parallel instructions that use Smem or Lmem operands, each instruction must use one of the indirect operands allowed for the dual AR indirect addressing mode: <ul style="list-style-type: none"> *ARn *ARn+ *ARn– *(ARn + T0) (Available if C54CM bit = 0) *(ARn + AR0) (Available if C54CM bit = 1) *(ARn – T0) (Available if C54CM bit = 0) *(ARn – AR0) (Available if C54CM bit = 1) *ARn(T0) (Available if C54CM bit = 0) *ARn(AR0) (Available if C54CM bit = 1) *(ARn + T1) *(ARn – T1)

Table 4–3. Advanced Parallelism Rules

Consideration	Rule
Byte extensions for constants	<p>An instruction that uses one of the following addressing-mode operands cannot be placed in parallel with another instruction. The constant (following the # symbol) adds 2 or 3 bytes to the instruction.</p> <p>*abs16(#k16) *port(#k16) (algebraic syntax) port(#k16) (mnemonic syntax) *(#k23) *ARn(#K16) *+ARn(#K16) *CDP(#K16) *+CDP(#K16)</p>
mmap() and port() qualifiers	<p>An instruction that uses the mmap() qualifier to indicate an access to a memory-mapped register or registers cannot be placed in parallel with another instruction. The use of the mmap() qualifier is a form of parallelism already.</p> <p>Likewise an instruction that uses a port() qualifier to indicate an access to I/O space cannot be placed in parallel with another instruction. The use of a port() qualifier is a form of parallelism already.</p>
Parallelism among A unit, D unit, and P unit	<p>Parallelism among the three computational units is allowed without restriction (see Figure 4–5).</p> <p>An operation executed within a single computational unit can be placed in parallel with a second operation executed in one of the other two computational units.</p>
Parallelism within the P unit	Two program-control instructions cannot be placed in parallel. However, other parallelism among the operators of the P unit is allowed.
Parallelism within the D unit	Certain restrictions apply to using operators of the D unit in parallel (see Figure 4–5).
Parallelism within the A unit	Two A-unit ALU operations or two A-unit swap operations cannot be performed in parallel. However, other parallelism among the operators of the A unit is allowed.

4.2.5 Process for Implementing User-Defined Parallelism

This section presents a process that may be used to simplify the process of using user-defined parallelism to produce optimized assembly language code. Figure 4–7 is a flow chart outlining this process, and the steps are also described in Table 4–4.

Figure 4–7. Process for Applying User-Defined Parallelism

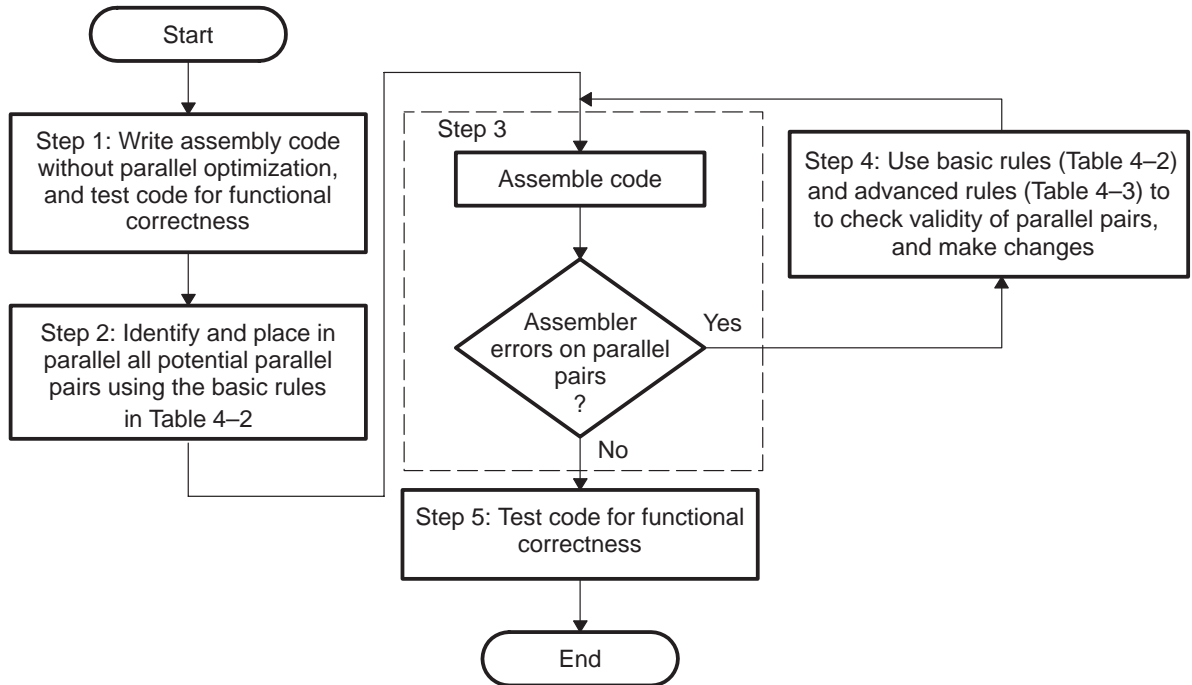


Table 4–4. Steps in Process for Applying User-Defined Parallelism

Step	Description
1	Write assembly code without the use of user-defined parallelism, and verify the functionality of the code. Note that in this step, you may take advantage of instructions with built-in parallelism.
2	Identify potential user-defined parallel instruction pairs in your code, and, using the basic rules outlined in Table 4–2 as guidelines, place instructions in parallel. Start by focusing on heavily used kernels of the code.
3	Run the optimized code through the assembler to see if the parallel instruction pairs are valid. The assembler will indicate any invalid parallel instruction pairs. If you have invalid pairs, go to step 4; otherwise go to step 5.
4	Refer to the set of parallelism rules in section 4.2.4 to determine why failing parallel pairs may be invalid. Make necessary changes and return to step 3.
5	Once all your parallel pairs are valid, make sure your code still functions correctly.

4.2.6 Parallelism Tips

As you try to optimize your code with user-defined parallelism, you might find the following tips helpful (these tips use algebraic instructions, but the concepts apply equally for the mnemonic instruction set):

- ❑ Place all load and store instructions in parallel. For example:

```
AC1 = *AR2      ; Load AC1
|| *AR3 = BRC0  ; Store BRC0
```

- ❑ The A-Unit ALU can handle (un)saturated 16-bit processing in parallel with the D-Unit ALU, MAC, and shift operators. For example:

```
; Modify AR1 in A unit, and perform an accumulator
; shift, saturate, and-store operation in the D unit.
AR1 = AR1 + T0
|| *AR2 = HI(saturate(uns(rnd(AC1 << #1))))
```

- ❑ Accumulator shift, saturate, and store operations can be placed in parallel with D-Unit ALU or MAC operations. For example:

```
; Shift, saturate, and store AC1 while
; modifying AC2.
*AR2 = HI(saturate(uns(rnd(AC1 << #1))))
|| AC2 = AC2 + AC1
```

- ❑ Control operations can be placed in parallel with DSP operations. For example:

```
; Switch control to a block-repeat loop, and
; Perform the first computation of the loop.
BLOCKREPEAT {
|| AC0 = rnd(AC0 + (*AR1+ * *AR3))
```

- ❑ Instructions with built-in parallelism increase the bandwidth of instructions paired by user-defined parallelism. For example:

```
; Place parallel accumulator load operations
; in parallel with an auxiliary register store
; operation.
HI(AC1) = HI(*AR2) + HI(AC0),
LO(AC1) = LO(*AR2) + LO(AC0)
|| dbl(*AR6)=AC2
```

- ❑ You can fill a buffer with a constant value efficiently. For example:

```
AC0 = #0          ; Clear AC0.
|| REPEAT(#9)     ; Switch control to repeat loop.
DBL(*AR1+) = AC0  ; Store 32-bit constant to buffer,
                  ; and increment pointer.
```

- ❑ Instructions to be executed conditionally can be placed in parallel with the *if* instruction. For example:

```
if (T0 == 0) execute(AD_unit) ; If T0 contains 0, ...
|| AR0 = #0                  ; ... Load AR0 with 0.
```

4.2.7 Examples of Parallel Optimization Within CPU Functional Units

This section provides examples to show how to make code more efficient by using parallelism within the A unit, the P unit, and D unit. (The examples in this section use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.)

4.2.7.1 A-Unit Example of Parallel Optimization

Example 4–4 shows a simple Host-DSP application in which a host sends a single command to tell the DSP which set of coefficients to use for a multiply-and-accumulate (MAC) operation. The DSP calls a COMPUTE function to perform the computation and returns the result to the host. The communication is based on a very simple handshaking, with the host and DSP exchanging flags (codes). The code in Example 4–4 does not use parallelism. Example 4–5 shows the code optimized through the use of parallel instruction pairs.

Example 4–4. A-Unit Code With No User-Defined Parallelism

```
; Variables
        .data

COEFF1   .word  0x0123           ; First set of coefficients
        .word  0x1234
        .word  0x2345
        .word  0x3456
        .word  0x4567

COEFF2   .word  0x7654           ; Second set of coefficients
        .word  0x6543
        .word  0x5432
        .word  0x4321
        .word  0x3210

HST_FLAG .set   0x2000           ; Host Flag Address
HST_DATA .set   0x2001           ; Host Data Address

CHANGE   .set   0x0000           ; "Change coefficients" command from host
READY    .set   0x0000           ; "READY" Flag from Host
BUSY     .set   0x1111           ; "BUSY" Flag set by DSP

        .global  start_al

        .text

start_al:
        AR0 = #HST_FLAG           ; AR0 points to Host Flag
        AR2 = #HST_DATA           ; AR2 points to Host Data
        AR1 = #COEFF1             ; AR1 points to COEFF1 buffer initially
        AR3 = #COEFF2             ; AR3 points to COEFF2 buffer initially
        CSR = #4                   ; Set CSR = 4 for repeat in COMPUTE
        BIT(ST1, #ST1_FRCT) = #1   ; Set fractional mode bit
        BIT(ST1, #ST1_SXMD) = #1   ; Set sign-extension mode bit
```

Example 4–4. A-Unit Code With No User-Defined Parallelism (Continued)

```

LOOP:
    T0 = *AR0                      ; T0 = Host Flag
    if (T0 == #READY) GOTO PROCESS ; If Host Flag is "READY", continue
    GOTO LOOP                      ; process - else poll Host Flag again

PROCESS:
    T0 = *AR2                      ; T0 = Host Data

    if (T0 == #CHANGE) EXECUTE(AD_UNIT)
                                ; The choice of either set of
                                ; coefficients is based on the value
                                ; of T0. COMPUTE uses AR3 for
                                ; computation, so we need to
                                ; load AR3 correctly here.

    SWAP(AR1, AR3)                ; Host message was "CHANGE", so we
                                ; need to swap the two coefficient
                                ; pointers.

    CALL COMPUTE                  ; Compute subroutine

    *AR2 = AR4                    ; Write result to Host Data
    *AR0 = #BUSY                  ; Set Host Flag to Busy
    GOTO LOOP                     ; Infinite loop continues

END

COMPUTE:
    AC1 = #0                      ; Initialize AC1 to 0
    REPEAT(CSR)                   ; CSR has a value of 4
    AC1 = AC1 + (*AR2 * *AR3+)    ; This MAC operation is performed
                                ; 5 times
    AR4 = AC1                     ; Result is in AR4

    RETURN

HALT:
    GOTO HALT

```

As mentioned, Example 4–5 shows the optimized code for Example 4–4. In Example 4–5, the parallel instruction pairs are highlighted. Notice the following points:

- ❑ The first four instructions (ARn loads) are immediate loads and cannot be placed in parallel due to constant bus conflicts and total instruction sizes.
- ❑ The first parallel pair shows an immediate load of CSR through the bus called KDB. This load is executed in parallel with the setting of the SXMD mode bit, which is handled by the A-unit ALU.
- ❑ The second parallel pair is a SWAP instruction in parallel with an *if* instruction. Despite the fact that the SWAP instruction is executed conditionally, it is valid to place it in parallel with the *if* instruction.
- ❑ The third parallel pair stores AR4 to memory via the D bus (DB), and stores a constant (BUSY) to memory via the bus called KDB.
- ❑ The fourth parallel pair loads AC1 with a constant that is carried on the bus called KDB and, in parallel, switches program control to a single-repeat loop.
- ❑ The last parallel pair stores AC1 to AR4 via a cross-unit bus and, in parallel, returns from the COMPUTE function.

Example 4–5. A-Unit Code in Example 4–4 Modified to Take Advantage of Parallelism

```

; Variables
        .data

COEFF1   .word  0x0123                ; First set of coefficients
        .word  0x1234
        .word  0x2345
        .word  0x3456
        .word  0x4567

COEFF2   .word  0x7654                ; Second set of coefficients
        .word  0x6543
        .word  0x5432
        .word  0x4321
        .word  0x3210

HST_FLAG .set   0x2000                ; Host Flag Address
HST_DATA .set   0x2001                ; Host Data Address

CHANGE   .set   0x0000                ; "Change coefficients" command from host
READY    .set   0x0000                ; "READY" Flag from Host
BUSY     .set   0x1111                ; "BUSY" Flag set by DSP

        .global  start_a2

        .text

start_a2:

        AR0 = #HST_FLAG              ; AR0 points to Host Flag
        AR2 = #HST_DATA              ; AR2 points to Host Data
        AR1 = #COEFF1                ; AR1 points to COEFF1 buffer initially
        AR3 = #COEFF2                ; AR3 points to COEFF2 buffer initially

        CSR = #4                     ; Set CSR = 4 for repeat in COMPUTE
        || BIT(ST1, #ST1_FRCT) = #1 ; Set fractional mode bit

        BIT(ST1, #ST1_SXMD) = #1     ; Set sign-extension mode bit

LOOP:
        T0 = *AR0                    ; T0 = Host Flag
        if (T0 == #READY) GOTO PROCESS ; If Host Flag is "READY", continue
        GOTO LOOP                    ; process - else poll Host Flag again

```

Example 4–5. A-Unit Code in Example 4–4 Modified to Take Advantage of Parallelism (Continued)

```

PROCESS:
    T0 = *AR2                                ; T0 = Host Data

    if (T0 == #CHANGE) EXECUTE(AD_UNIT)
                                ; The choice of either set of
                                ; coefficients is based on the value
                                ; of T0. COMPUTE uses AR3 for
                                ; computation, so we need to
                                ; load AR3 correctly here.

    || SWAP(AR1, AR3)                ; Host message was "CHANGE", so we
                                ; need to swap the two coefficient
                                ; pointers.

    CALL COMPUTE                    ; Compute subroutine

    *AR2 = AR4                        ; Write result to Host Data
    || *AR0 = #BUSY                ; Set Host Flag to Busy

    GOTO LOOP                        ; Infinite loop continues

END

COMPUTE:
    AC1 = #0                          ; Initialize AC1 to 0
    || REPEAT(CSR)                  ; CSR has a value of 4
    AC1 = AC1 + (*AR2 * *AR3+) ; This MAC operation is performed
                                ; 5 times
    AR4 = AC1                        ; Result is in AR4
    || RETURN

HALT:
    GOTO HALT

```


4.2.7.2 P-Unit Example of Parallel Optimization

Example 4–6 demonstrates a very simple nested loop and some simple control operations with the use of P-unit registers. This example shows the unoptimized code, and Example 4–7 shows the code optimized through the use of the P-unit parallel instruction pairs (parallel instruction pairs are highlighted).

Example 4–6. P-Unit Code With No User-Defined Parallelism

```

        .CPL_off                ; Tell assembler that CPL bit is 0
                                ; (Direct addressing is done with DP)

; Variables
        .data

var1     .word 0x0004
var2     .word 0x0000

        .global  start_p1

        .text

start_p1:
        XDP = #var1
        AR3 = #var2

        BRC0 = #0007h           ; BRC0 loaded using KPB
        BRC1 = *AR3             ; BRC1 loaded using DB

        AC2 = #0006h

        BLOCKREPEAT {
            AC1 = AC2
            AR1 = #8000h
            LOCALREPEAT {
                AC1 = AC1 - #1
                *AR1+ = AC1
            }
            AC2 = AC2 + #1
        }
        @(AC0_L) = BRC0          ; AC0_L loaded using EB
        @(AC1_L) = BRC1          ; AC1_L loaded using EB

        if (AC0 >= #0) goto start_p1:
        if (AC1 >= #0) goto start_p1:
end_p1

```

Notice the following about Example 4–7:

- ☐ The first three register loads are immediate loads using the same constant bus and, therefore, cannot be placed in parallel.
- ☐ The fourth register load (loading BRC0) can be placed in parallel with the next load (loading BRC1), which does not use the constant bus, but the data bus DB to perform the load.
- ☐ The next instruction, which is a control instruction (blockrepeat) is placed in parallel with the load of AC2. There are no conflicts and this pair is valid. The loading of AC2 is not part of the blockrepeat structure.
- ☐ The final parallel pair is a control instruction (localrepeat) and a load of AR1. The loading of AR1 is not part of the localrepeat structure but is part of the outer blockrepeat structure.

Example 4–7. P-Unit Code in Example 4–6 Modified to Take Advantage of Parallelism

```

        .CPL_off                ; Tell assembler that CPL bit is 0
                                ; (Direct addressing is done with DP)

; Variables
        .data

var1     .word 0x0004
var2     .word 0x0000

        .global  start_p2

        .text

start_p2:
        XDP = #var1
        AR3 = #var2

        BRC0 = #0007h           ; BRC0 loaded using KPB
        || BRC1 = *AR3          ; BRC1 loaded using DB

        AC2 = #0006h
        || BLOCKREPEAT {
            AC1 = AC2
            AR1 = #8000h
            || LOCALREPEAT {
                AC1 = AC1 - #1
                *AR1+ = AC1
            }
            AC2 = AC2 + #1
        }
        @(AC0_L) = BRC0          ; AC0_L loaded using EB
        @(AC1_L) = AR1           ; AC1_L loaded using EB

        if (AC0 >= #0) goto start_p2
        if (AC1 >= #0) goto start_p2
end_p2

```

4.2.7.3 D-Unit Example of Parallel Optimization

Example 4–8 demonstrates a very simple load, multiply, and store function with the use of D-unit registers. Example 4–9 shows this code modified to take advantage of user-defined parallelism (parallel instruction pairs are highlighted).

Example 4–8. D-Unit Code With No User-Defined Parallelism

```
; Variables
.data

var1 .word 0x8000
var2 .word 0x0004

.global start_d1

.text

start_d1:
    AR3 = #var1
    AR4 = #var2

    AC0 = #0004h          ; AC0 loaded using KDB
    AC2 = *AR3            ; AC2 loaded using DB

    T0 = #5A5Ah           ; T0 loaded with constant, 0x5A5A

    AC2 = AC2 + (AC0 * T0) ; MAC
    AC1 = AC1 + (AC2 * T0) ; MAC
    SWAP(AC0, AC2)        ; SWAP

    *AR3 = HI(AC1)        ; Store result in AC1
    *AR4 = HI(AC0)        ; Store result in AC0
end_d1
```

The following information determined the optimizations made in Example 4–9:

- ☐ As in the P-unit example on page 4-37, we cannot place the immediate register loads in parallel due to constant-bus conflicts.
- ☐ The instructions that load AC0 and AC2 have been placed in parallel because they are not both immediate loads and as such, there are no constant-bus conflicts.
- ☐ It is not possible to place the two single-MAC instructions in parallel since the same operator is required for both and as such a conflict arises. However, placing the second MAC instruction in parallel with the SWAP instruction is valid.
- ☐ The two 16-bit store operations at the end of the code are placed in parallel because there are two 16-bit write buses available (EB and FB).

Example 4–9. D-Unit Code in Example 4–8 Modified to Take Advantage of Parallelism

```

; Variables
.data

var1 .word 0x8000
var2 .word 0x0004

.global start_d2

.text

start_d2:
    AR3 = #var1
    AR4 = #var2

    AC0 = #0004h                ; AC0 loaded using KDB
    || AC2 = *AR3                ; AC2 loaded using DB

    T0 = #5A5Ah                ; T0 loaded with constant, 0x5A5A

    AC2 = AC2 + (AC0 * T0)      ; MAC
    AC1 = AC1 + (AC2 * T0)      ; MAC
    || SWAP(AC0, AC2)           ; SWAP

    *AR3 = HI(AC1)              ; Store result in AC1
    || *AR4 = HI(AC0)           ; Store result in AC0
end_d2

```

4.2.8 Example of Parallel Optimization Across the A-Unit, P-Unit, and D-Unit

Example 4–10 shows unoptimized code for an FIR (finite impulse response) filter. Example 4–11 is the result of applying user-defined parallelism to the same code. It is important to notice that the order of instructions has been altered in a number of cases to allow certain instruction pairs to be placed in parallel. The use of parallelism in this case has saved about 50% of the cycles outside the inner loop.

Example 4–10. Code That Uses Multiple CPU Units But No User-Defined Parallelism

```

.CPL_ON                ; Tell assembler that CPL bit is 1
                       ; (SP direct addressing like *SP(0) is enabled)

; Register usage
; -----

.asg  AR0, X_ptr        ; AR0 is pointer to input buffer - X_ptr
.asg  AR1, H_ptr        ; AR1 is pointer to coefficients - H_ptr
.asg  AR2, R_ptr        ; AR2 is pointer to result buffer - R_ptr
.asg  AR3, DB_ptr       ; AR3 is pointer to delay buffer - DB_ptr

FRAME_SZ .set 2

.global _fir

.text

;*****

_fir

; Create local frame for temp values
; -----

SP = SP - #FRAME_SZ

; Turn on fractional mode
; Turn on sign-extension mode
; -----

BIT(ST1, #ST1_FRCT) = #1      ; Set fractional mode bit
BIT(ST1, #ST1_SXMD) = #1     ; Set sign-extension mode bit

; Set outer loop count by subtracting 1 from nx and storing into
; block-repeat counter
; -----

AC1 = T1 - #1                ; AC1 = number of samples (nx) - 1
*SP(0) = AC1                 ; Top of stack = nx - 1
BRC0 = *SP(0)                ; BRC0 = nx - 1 (outer loop counter)

```

**Example 4–10. Code That Uses Multiple CPU Units But No User-Defined Parallelism
(Continued)**

```

; Store length of coefficient vector/delay buffer in BK register
; -----

BIT(ST2, #ST2_AR1LC) = #1      ; Enable AR1 circular configuration
BSA01 = *(#0011h)              ; Set buffer (filter) start address
                                ; AR1 used as filter pointer

BIT(ST2, #ST2_AR3LC) = #1      ; Enable AR3 circular configuration
*SP(1) = DB_ptr                ; Save pointer to delay buffer pointer
AC1 = *DB_ptr                  ; AC1 = delay buffer pointer
DB_ptr = AC1                   ; AR3 (DB_ptr) = delay buffer pointer
BSA23 = *(#0013h)              ; Set buffer (delay buffer) start address
                                ; AR3 used as filter pointer

*SP(0) = T0                    ; Save filter length, nh - used as buffer
                                ; size
BK03 = *SP(0)                  ; Set circular buffer size - size passed
                                ; in T0

AC1 = T0 - #3                  ; AC1 = nh - 3
*SP(0) = AC1
CSR = *SP(0)                   ; Set inner loop count to nh - 3

H_ptr = #0                     ; Initialize index of filter to 0
DB_ptr = #0                    ; Initialize index of delay buffer to 0

; Begin outer loop on nx samples
; -----

BLOCKREPEAT {
; Move next input sample into delay buffer
; -----

*DB_ptr = *X_ptr+

; Sum h * x for next y value
; -----

AC0 = *H_ptr+ * *DB_ptr+

REPEAT (CSR)
    AC0 = AC0 + (*H_ptr+ * *DB_ptr+)

AC0 = rnd(AC0 + (*H_ptr+ * *DB_ptr+))    ; Round result

```

**Example 4–10. Code That Uses Multiple CPU Units But No User-Defined Parallelism
(Continued)**

```

; Store result
; -----

    *R_ptr+ = HI(AC0)
    }

; Clear FRCT bit to restore normal C operating environment
; Return overflow condition of AC0 (shown in ACOV0) in T0
; Restore stack to previous value, FRAME, etc..
; Update current index of delay buffer pointer
; -----

END_FUNCTION:

    AR0 = *SP(1)           ; AR0 = pointer to delay buffer pointer
    SP = SP + #FRAME_SZ   ; Remove local stack frame
    *AR0 = DB_ptr         ; Update delay buffer pointer with current
                          ; index

    BIT(ST1, #ST1_FRCT) = #0      ; Clear fractional mode bit

    T0 = #0                 ; Make T0 = 0 for no overflow (return value)
    if(overflow(AC0)) execute(AD_unit)
    T0 = #1                 ; Make T0 = 1 for overflow (return value)

    RETURN
; *****

```

In Example 4–11, parallel pairs that were successful are shown in **bold type**; potential parallel pairs that failed are shown in *italic type*. The first failed due to a constant bus conflict, and the second failed due to the fact that the combined size is greater than 6 bytes. The third pair failed for the same reason, as well as being an invalid soft-dual encoding instruction. The last pair in italics failed because neither instruction has a parallel enable bit. Some of the load/store operations that are not in parallel were made parallel in the first pass optimization process; however, the parallelism failed due to bus conflicts and had to be removed.

Note:

Example 4–11 shows optimization *only* with the use of the parallelism features. Further optimization of this FIR function is possible by employing other optimizations.

Example 4–11. Code in Example 4–10 Modified to Take Advantage of Parallelism

```
.CPL_ON                ; Tell assembler that CPL bit is 1
                        ; (SP direct addressing like *SP(0) is enabled)

; Register usage
; -----

.asg  AR0, X_ptr        ; AR0 is pointer to input buffer - X_ptr
.asg  AR1, H_ptr        ; AR1 is pointer to coefficients - H_ptr
.asg  AR2, R_ptr        ; AR2 is pointer to result buffer - R_ptr
.asg  AR3, DB_ptr       ; AR3 is pointer to delay buffer - DB_ptr

FRAME_SZ .set 2

.global _fir

.text

;*****

_fir

; Create local frame for temp values
; -----

    SP = SP - #FRAME_SZ        ; (Attempt to put this in parallel with
                                ; the following AC1 modification failed)

; Set outer loop count by subtracting 1 from nx and storing into
; block-repeat counter
; Turn on fractional mode
; Turn on sign-extension mode
; -----

    AC1 = T1 - #1              ; AC1 = number of samples (nx) - 1

    BIT(ST1, #ST1_FRCT) = #1   ; Set fractional mode bit
    || *SP(0) = AC1            ; Top of stack = nx - 1

    BRC0 = *SP(0)              ; BRC0 = nx - 1 (outer loop counter)
    || BIT(ST1, #ST1_SXMD) = #1 ; Set sign-extension mode bit
```

*Example 4–11. Code in Example 4–10 Modified to Take Advantage of Parallelism
(Continued)*

```

; Store length of coefficient vector/ delay buffer in BK register
; -----

BIT(ST2, #ST2_AR1LC) = #1      ; Enable AR1 circular configuration
BSA01 = *(&0011h)              ; Set buffer (filter) start address
                                ; AR1 used as filter pointer

BIT(ST2, #ST2_AR3LC) = #1      ; Enable AR3 circular configuration
|| *SP(1) = DB_ptr              ; Save pointer to delay buffer pointer

AC1 = *DB_ptr                  ; AC1 = delay buffer pointer
DB_ptr = AC1                   ; AR3 (DB_ptr) = delay buffer pointer
|| *SP(0) = T0                  ; Save filter length, nh - used as buffer
                                ; size
BSA23 = *(&0013h)              ; Set buffer (delay buffer) start address
                                ; AR3 used as filter pointer

BK03 = *SP(0)                  ; Set circular buffer size - size passed
                                ; in T0

AC1 = T0 - #3                  ; AC1 = nh - 3
*SP(0) = AC1

CSR = *SP(0)                   ; Set inner loop count to nh - 3
|| H_ptr = #0                  ; Initialize index of filter to 0

DB_ptr = #0                    ; Initialize index of delay buffer to 0
                                ; (in parallel with BLOCKREPEAT below)

; Begin outer loop on nx samples
; -----

||BLOCKREPEAT {

; Move next input sample into delay buffer
; -----

*DB_ptr = *X_ptr+

```

*Example 4–11. Code in Example 4–10 Modified to Take Advantage of Parallelism
(Continued)*

```

; Sum h * x for next y value
; -----

AC0 = *H_ptr+ * *DB_ptr+
|| REPEAT (CSR)

AC0 = AC0 + (*H_ptr+ * *DB_ptr+)

AC0 = rnd(AC0 + (*H_ptr+ * *DB_ptr))    ; Round result

; Store result
; -----

*R_ptr+ = HI(AC0)
}

; Clear FRCT bit to restore normal C operating environment
; Return overflow condition of AC0 (shown in ACOV0) in T0
; Restore stack to previous value, FRAME, etc..
; Update current index of delay buffer pointer
; -----

END_FUNCTION:

AR0 = *SP(1)                ; AR0 = pointer to delay buffer pointer
|| SP = SP + #FRAME_SZ      ; Remove local stack frame

*AR0 = DB_ptr                ; Update delay buffer pointer with current
                             ; index

|| BIT(ST1, #ST1_FRCT) = #0 ; Clear fractional mode bit

T0 = #0                      ; Make T0 = 0 for no overflow (return value)
|| if(overflow(AC0)) execute(AD_unit)
T0 = #1                      ; Make T0 = 1 for overflow (return value)

|| RETURN
;*****

```

4.3 Implementing Efficient Loops

There are four common methods to implement instruction looping in the C55x DSP:

- ☐ Single repeat: *repeat(CSR/k8/k16),[CSR += TA/k4]*
- ☐ Local block repeat: *localrepeat{}*
- ☐ Block repeat: *blockrepeat{}*
- ☐ Branch on auxiliary register not zero: *if (ARn_mod != 0) goto #loop_start*

The selection of the looping method to use depends basically on the number of instructions that need to be repeated and in the way you need to control the loop counter parameter. The first three methods in the preceding list offer zero-overhead looping and the fourth one offers a 5-cycle loop overhead.

Overall, the most efficient looping mechanisms are the *repeat()* and the *localrepeat{}* mechanisms. The *repeat()* mechanism provides a way to repeat a single instruction or a parallel pair of instructions in an interruptible way. *repeat(CSR)*, in particular, allows you to compute the loop counter at runtime. Refer to section 4.3.2, *Efficient Use of repeat(CSR) Looping*.

Note:

If you are migrating code from a TMS320C54x™ DSP, be aware that a single-repeat instruction is interruptible on a TMS320C55x DSP. On a TMS320C54x DSP, a single-repeat instruction cannot be interrupted.

The *localrepeat{}* mechanism provides a way to repeat a block from the instruction buffer queue. Reusing code that has already been fetched and placed in the queue brings the following advantages:

- ☐ Fewer program-memory access pipeline conflicts
- ☐ Overall lower power consumption
- ☐ No repetition of wait-state and access penalties when executing loop code from external RAM

4.3.1 Nesting of Loops

You can create up to two levels of block-repeat loops without any cycle penalty. You can have one block-repeat loop nested inside another, creating an inner (level 1) loop and an outer (level0) loop. In addition, you can put any number of single-repeat loops inside each block-repeat loop. Example 4–12 shows a multi-level loop structure with two block-repeat loops and two single-repeat loops. (The examples in this section use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.)

Example 4–12. Nested Loops

```

    BRC0 = #(n0-1)
    BRC1 = #(n1-1)
;    ...
    localrepeat{                ; Level 0 looping (could instead be blockrepeat):
                                ; Loops n0 times
;    ...
        repeat( #(n2-1)
;    ...
        localrepeat{           ; Level 1 looping (could instead be blockrepeat):
                                ; Loops n1 times
;    ...
            repeat( #(n3-1) )
;    ...
        }
;    ...
    }

```

Example 4–12 shows one block-repeat loop nested inside another block-repeat loop. If you need more levels of multiple-instruction loops, use *branch on auxiliary register not zero* constructs to create the remaining outer loops. In Example 4–13 (page 4-48), a *branch on auxiliary register not zero* construct (see the last instruction in the example) forms the outermost loop of a Fast Fourier Transform algorithm. Inside that loop are two *localrepeat{}* loops. Notice that if you want the outermost loop to execute n times, you must initialize AR0 to $(n - 1)$ outside the loop.

**Example 4–13. Branch-On-Auxiliary-Register-Not-Zero Construct
(Shown in Complex FFT Loop Code)**

```

_cfft:

radix_2_stages:
; ...

outer_loop:
; ...
    BRC0 = DR1
; ...
    BRC1 = DR1
; ...
    AR4 = AR4 >> #1                ; outer loop counter
    || if (AR5 == #0) goto no_scale ; determine if scaling required
; ...

no_scale:

    localrepeat{

        AC0 = dbl(*AR3)                ; load ar,ai

        HI(AC2) = HI(*AR2) - HI(AC0),    ; tr = ar - br
        LO(AC2) = LO(*AR2) - LO(AC0)    ; ti = ai - bi

        localrepeat {

            HI(AC1) = HI(*AR2) + HI(AC0),    ; ar' = ar + br
            LO(AC1) = LO(*AR2) + LO(AC0)    ; ai' = ai + bi
            || dbl(*AR6)=AC2                ; store tr, ti

            AC2 = *AR6 * coef(*CDP+),        ; c*tr
            AC3 = *AR7 * coef(*CDP+)        ; c*ti

            dbl(*AR2+) = AC1                ; store ar, ai
            || AC0 = dbl(*AR3(DR0))          ; * load ar,ai

            AC3 = rnd(AC3 - (*AR6 * coef(*CDP-))), ; bi' = c*ti - s*tr
            AC2 = rnd(AC2 + (*AR7 * coef(*CDP-))) ; br' = c*tr + s*ti

            *AR3+ = pair(HI(AC2))            ; store br', bi'
            || HI(AC2) = HI(*AR2) - HI(AC0), ; * tr = ar - br
            LO(AC2) = LO(*AR2) - LO(AC0)    ; * ti = ai - bi

        }

    }

```

Note: This example shows portions of the file `cfft.asm` in the TI C55x DSPLIB (introduced in Chapter 8)

**Example 4–13. Branch-On-Auxiliary-Register-Not-Zero Construct
(Shown in Complex FFT Loop Code) (Continued)**

```

HI(AC1) = HI(*AR2) + HI(AC0),           ; ar' = ar + br
LO(AC1) = LO(*AR2) + LO(AC0)           ; ai' = ai + bi
|| dbl(*AR6)=AC2                        ; store tr, ti

AC2 = *AR6 * coef(*CDP+),               ; c*tr
AC3 = *AR7 * coef(*CDP+)                ; c*ti

dbl(*(AR2+DR1)) = AC1                   ; store ar, ai

AC3 = rnd(AC3 - (*AR6 * coef(*CDP+))),   ; bi' = c*ti - s*tr
AC2 = rnd(AC2 + (*AR7 * coef(*CDP+)))    ; br' = c*tr + s*ti

*(AR3+DR1) = pair(HI(AC2))              ; store br', bi'

}

AR3 = AR3 << #1
|| CDP = #0                             ; rewind coefficient pointer
DR3 = DR3 >> #1
|| if (AR4 != #0) goto outer_loop

```

Note: This example shows portions of the file `cfft.asm` in the TI C55x DSPLIB (introduced in Chapter 8)

To achieve an efficient nesting of loops, apply the following guidelines:

- ❑ Use a single-repeat instruction for the innermost loop if the loop contains only a single instruction (or a pair of instructions that have been placed in parallel).
- ❑ Use a local block-repeat instruction (*localrepeat* in the algebraic syntax) for a loop containing more than a single instruction or instruction pair—provided the loop contains no more than 56 bytes of code. If there are more than 56 bytes but you would still like to use the local block-repeat instruction, consider the following possibilities:
 - Split the existing loop into two smaller loops.
 - Reduce the number of bytes in the loop. For example, you can reduce the number of instructions that use embedded constants.
- ❑ Use a standard block-repeat instruction (*blockrepeat* in the algebraic syntax) in cases where a local block-repeat instruction cannot be used. The standard block-repeat mechanism always refetches the loop code from memory.

- ❑ When you nest a block-repeat loop inside another block-repeat loop, initialize the block-repeat counters (BRC0 and BRC1) in the code outside of both loops. This technique is shown in Example 4–12 (page 4-47).

Neither counter needs to be re-initialized inside its loop; placing such initializations inside the loops only adds extra cycles to the loops. The CPU uses BRC0 for the outer (level 0) loop and BRC1 for the inner (level 1) loop. BRC1 has a shadow register, BRS1, that preserves the initial value of BRC1. Each time the level 1 loop must begin again, the CPU automatically re-initializes BRC1 from BRS1.

4.3.2 Efficient Use of *repeat(CSR)* Looping

The single-repeat instruction syntaxes allow you to specify the repeat count as a constant (embedded in the repeat instruction) or as the content of the computed single-repeat register (CSR). When CSR is used, it is not decremented during each iteration of the single-repeat loop. Before the first execution of the instruction or instruction pair to be repeated, the content of CSR is copied into the single-repeat counter (RPTC). RPTC holds the active loop count and is decremented during each iteration. Therefore, CSR needs to be initialized only once. Initializing CSR outside the outer loop, rather than during every iteration of the outer loop, saves cycles. There are advantages to using CSR for the repeat count:

- ❑ The repeat count can be dynamically computed during runtime and stored to CSR. For example, CSR can be used when the number of times an instruction must be repeated depends on the iteration number of a higher loop structure.
- ❑ Using CSR saves outer loop cycles when the single-repeat loop is an inner loop.
- ❑ An optional syntax extension enables the repeat instruction to modify the CSR after copying the content of CSR to RPTC. When the single-repeat loop is repeated in an outer loop, CSR contains a new count.

Example 4–14 (page 4-51) uses CSR for a single-repeat loop that is nested inside a block-repeat loop. In the example, CSR is assigned the name `inner_cnt`.

Example 4–14. Use of CSR (Shown in Real Block FIR Loop Code)

```

;
; ...
.asg CSR, inner_cnt           ;inner loop count
.asg BRC0, outer_cnt         ;outer loop count
;
; ...
.asg AR0, x_ptr              ;linear pointer
.asg AR1, db_ptr1            ;circular pointer
.asg AR2, r_ptr              ;linear pointer
.asg AR3, db_ptr2            ;circular pointer
.asg CDP, h_ptr              ;circular pointer

; ...

_fir2:
; ...
    mar(*db_ptr2-)           ;index of 2nd oldest db entry
;
; Setup loop counts
;-----
    || T0 = T0 >> #1          ;T0 = nx/2

    T0 = T0 - #1              ;T0 = (nx/2 - 1)
    outer_cnt = T0            ;outer loop executes nx/2 times
    T0 = T1 - #3              ;T0 = nh-3
    inner_cnt = T0            ;inner loop executes nh-2 times
    T1 = T1 + #1              ;T1 = nh+1, adjustment for db_ptr1, db_ptr2
;
; Start of outer loop
;-----
    ||localrepeat {           ;start the outer loop

        *db_ptr1 = *x_ptr+     ;get 1st new input value
        *db_ptr2 = *x_ptr+     ;get 2nd new input value (newest)

;1st iteration
        AC0 = *db_ptr1+ * coef(*h_ptr+), ;part 1 of dual-MPY
        AC1 = *db_ptr2+ * coef(*h_ptr+), ;part 2 of dual-MPY

;inner loop
        ||repeat(inner_cnt)
        AC0 = AC0 + (*db_ptr1+ * coef(*h_ptr+)), ;part 1 of dual-MAC
        AC1 = AC1 + (*db_ptr2+ * coef(*h_ptr+)) ;part 2 of dual-MAC

```

Note: This example shows portions of the file fir2.asm in the TI C55x DSPLIB (introduced in Chapter 8)

Example 4–14. Use of CSR (Shown in Real Block FIR Loop Code) (Continued)

```
;last iteration has different pointer adjustment and rounding
    AC0 = rnd(AC0 + (*(db_ptr1-T1) * coef(*h_ptr+))), ;part 1 of dual-MAC
    AC1 = rnd(AC1 + (*(db_ptr2-T1) * coef(*h_ptr+))) ;part 2 of dual-MAC

;store result to memory
    *r_ptr+ = HI(AC0)                                ;store 1st Q15 result to memory
    *r_ptr+ = HI(AC1)                                ;store 2nd Q15 result to memory
}                                                    ;end of outer loop
; ...
```

Note: This example shows portions of the file `fir2.asm` in the TI C55x DSPLIB (introduced in Chapter 8)

4.3.3 Avoiding Pipeline Delays When Accessing Loop-Control Registers

Accesses to loop-control registers like CSR, BRC0, and BRC1 can cause delays in the instruction pipeline if nearby instructions make competing accesses. For recommendations on avoiding this type of pipeline delay, see the “Loop control” section of Table 4–6 (page 4-56) .

4.4 Minimizing Instruction Pipeline Delays

The C55x instruction pipeline is a protected pipeline that has two, decoupled segments. The first segment fetches instructions. The second segment, referred to as the *execution pipeline*, decodes instructions and performs data accesses and computations. The execution pipeline is illustrated in Figure 4–8 and described in Table 4–5. Because there are no potential data-access errors in the first segment, this section of the *Programmer’s Guide* focuses entirely on the second segment. For more details on the pipeline, refer to the CPU reference guide (see *Related Documentation From Texas Instruments* in the preface).

Figure 4–8. Second Segment of the Pipeline (Execution Pipeline)




Table 4–5. Descriptions of the Execution Pipeline Stages

Pipeline Stage	Description
Decode (D)	<ul style="list-style-type: none"> Read up to six bytes from the instruction buffer queue. Decode an instruction pair or a single instruction. Dispatch instructions to the appropriate CPU functional units.
Address (AD)	<ul style="list-style-type: none"> Compute data-space addresses in the data-address generation unit (DAGEN). Modify pointers and repeat counters as required. Compute the program-space addresses for PC-relative branching instructions.
Access 1 (AC1)	Send addresses for read operands as required on the following buses: BAB, CAB, and DAB.
Access 2 (AC2)	Allow one cycle for memories to respond to read accesses.

Table 4–5. Descriptions of the Execution Pipeline Stages (Continued)

Pipeline Stage	Description
Read (R)	<input type="checkbox"/> Transfer an operand or operands to the CPU via the data-read buses. <input type="checkbox"/> Generate an address for each operand write, and send the addresses on the data-write address buses. <input type="checkbox"/> Evaluate conditional operators
Execute (X)	<input type="checkbox"/> Execute data processing instructions that are executed in the A unit and the D unit. <input type="checkbox"/> Store results of computations into registers.
Write (W)	Send data to memory, to I/O space, or to registers accessed at their memory-mapped addresses.

Multiple instructions are executed simultaneously in the pipeline, and different instructions perform modifications to memory, I/O space, and register values during different stages of the pipeline. In an unprotected pipeline, this could lead to data-access errors—reads and writes at the same location happening out of the intended order. The pipeline-protection unit of the C55x DSP inserts extra cycles to prevent these errors. If an instruction (say, instruction 3) must access a location but a previous instruction (say, instruction 1) is not done with the location, instruction 3 is halted in the pipeline until instruction 1 is done. To minimize delays, you can take steps to prevent many of these pipeline-protection cycles.

The instruction set reference guides (see *Related Documentation From Texas Instruments* in the preface) show how many cycles an instruction takes to execute when the pipeline is full and experiencing no delays. Pipeline-protection cycles add to that best-case execution time. As it will be shown, most cases of pipeline conflict can be solved with instruction rescheduling.

This section provides examples to help you to better understand the impact of the pipeline structure on the way your code performs. It also provides you with recommendations for coding style and instruction usage to minimize conflicts or pipeline stalls. This section does not cover all of the pipeline potential conflicts, but some of the most common pipeline delays found when writing C55x code.

4.4.1 Process to Resolve Pipeline Conflicts

A pipeline conflict occurs when one instruction attempts to access a resource before a previous instruction is done with that resource. The pipeline-protection unit adds extra cycles to delay the later instruction. The following process is recommended for resolving pipeline conflicts that are causing delays. Try to focus your pipeline optimization effort on your key, inner code kernels first, to achieve a greater payback.

Step 1: Make your code functional.

Write your code first, without pipeline optimization in mind. In the C55x DSP, the pipeline is protected. Code is executed in the order in which it is written, and stalls are automatically inserted by the hardware to prevent incorrect operation. This makes programming the DSP easier and makes the code easier to debug than an open-pipeline device, in which the sequence of the code might not be the sequence of operation. There are only a few cases in which the C55x pipeline is not fully protected and they will be described in this documentation.

Step 2: Determine where the pipeline conflicts exist.

If you are using the C55x simulator, take advantage of its pipeline-conflict detection capabilities. Watch the *clock* variable when stepping through your code in the C55x simulator to view the intervention of the pipeline protection unit. If the clock increments by more than 1, there might be a pipeline or memory conflict in the instruction you just single stepped.

The C55x emulator/debugger does not support the *clock* variable, and setting breakpoints before and after may not give you accurate results for a single instruction due to the initial pipeline fill and the final pipeline flush during single stepping. In this case, you should try to benchmark small pieces of looped code.

Step 3: Apply the pipeline optimization coding recommendations summarized in Table 4–6.

After step 2 or if you are not using the simulator but the emulator, you should try to apply the recommendations directly.

Tip: When suspecting a pipeline conflict between two instructions, try to add NOP instructions in between. If the entire code cycle count does not increase by adding NOPs, then you can try to rearrange your code to replace those NOPs with useful instructions.

Software solutions to apply for pipeline and memory conflicts include:

- ☐ Reschedule instructions.
- ☐ Reduce memory accesses by using CPU registers to hold data.
- ☐ Reduce memory accesses by using a local repeat instruction, an instruction that enables the CPU to repeatedly execute a block of code from the instruction buffer queue.
- ☐ Relocate variables and data arrays in memory, or consider temporarily copying arrays to other nonconflicting memory banks at run time.

4.4.2 Recommendations for Preventing Pipeline Delays

Table 4–6 lists recommendations for avoiding common causes for pipeline delays. The rightmost column of the table directs you to the section that contains the details behind each recommendation. (The examples in this section use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.)

Table 4–6. Recommendations for Preventing Pipeline Delays

Recommendation Category	Recommendation	See ...
General	<input type="checkbox"/> In the case of a conflict, the front runner wins.	Section 4.4.2.1, page 4-57
	<input type="checkbox"/> Avoid write/read or read/write sequences to the same register.	Section 4.4.2.2, page 4-58
	<input type="checkbox"/> Use MAR instructions when possible, but avoid read/write sequences, and pay attention to instruction size.	Section 4.4.2.3, page 4-61
Pipeline-protection granularity	<input type="checkbox"/> Pay attention to pipeline-protection granularity when updating status register bit fields.	Section 4.4.2.4, page 4-64
	<input type="checkbox"/> Pay attention to pipeline-protection granularity when accessing MMRs in consecutive instructions.	Section 4.4.2.5, page 4-66
Loop control	<input type="checkbox"/> Understand when the loop-control registers are accessed in the pipeline	Section 4.4.2.6, page 4-68
	<input type="checkbox"/> Avoid accessing the BRC register in the last 4 instructions of a block-repeat or local block-repeat structure (to prevent a delay of the BRC decrement or an unprotected pipeline situation).	Section 4.4.2.7, page 4-68

Table 4–6. Recommendations for Preventing Pipeline Delays (Continued)

Recommendation Category	Recommendation	See ...
	<input type="checkbox"/> Do not access loop-control registers RSAX, REAX, and RPTC within the loop itself.	Section 4.4.2.8, page 4-71
	<input type="checkbox"/> Initialize the BRCx or CSR register at least 4 cycles before the repeat instruction, or initialize the register with an immediate value.	Section 4.4.2.9, page 4-71
Condition evaluation	<input type="checkbox"/> Try to set conditions well in advance of the time that the condition is tested.	Section 4.4.2.10, page 4-72
	<input type="checkbox"/> When making an instruction execute conditionally, use <i>execute(AD_unit)</i> instead of <i>execute(D_unit)</i> to create fully protected pipeline conditions, but be aware of potential pipeline delays.	Section 4.4.2.11, page 4-73
	<input type="checkbox"/> Understand the <i>execute (D_unit)</i> condition evaluation exception.	Section 4.4.2.12, page 4-75
Memory usage	<input type="checkbox"/> When working with dual MAC and FIR instructions, put the Cmem operand in a different memory bank.	Section 4.4.3.1, page 4-80
	<input type="checkbox"/> When working on data in the same memory bank, pay attention to instruction sequences that could generate memory conflicts.	Section 4.4.3.2, page 4-81
	<input type="checkbox"/> Map program code to a dedicated SARAM memory block to avoid conflicts with data accesses.	Section 4.4.3.3, page 4-82
	<input type="checkbox"/> For 32-bit accesses (using an Lmem operand), no performance hit is incurred if you use SARAM (there is no need to use DARAM).	Section 4.4.3.4, page 4-82

4.4.2.1 In the case of a conflict, the front runner wins.

A pipeline conflict arises when two instructions in different stages in the pipeline compete for the use of the same resource. The resource is granted to the instruction that is ahead in terms of pipeline execution, to increase overall instruction throughput.

4.4.2.2 Avoid write/read or read/write sequences to the same register.

All register accesses do not happen in the same pipeline stage. Registers are typically read in the read stage and written in the execute stage of the pipeline with the following notable exceptions:

- ☐ When using indirect addressing, AR0–7 and CDP registers are read and updated (for example, incremented according to the modifier *AR2+) during the AD stage.
- ☐ When using stack addressing or push, pop, call, and return instructions, SP is updated in the AD stage.
- ☐ When using the MAR () instruction, ARx and Tx registers are written during the AD stage.
- ☐ When using the swap () instruction, ARx and Tx registers are swapped during the AD stage.
- ☐ When you load one of the following registers with an immediate value (using a specific-CPU-register-load instruction), the CPU moves the immediate value into the register in the AD stage: BK03, BK47, BKC, BRC0, BRC1, CSR, DPH, PDP, BSA01, BSA23, BSA45, BSA67, BSAC, CDP, DP, SP, SSP.

Because registers are not accessed in the same pipeline stage, pipeline conflicts can occur in write/read or read/write sequences to the same register. Following are 4 common register pipeline conflict cases and how to resolve them.

Case 1: ARx write followed by an indirect addressing ARx read/update

Example 4–15 shows an AR write followed by an indirect-addressing AR read and update. In the example, I2 has a 4-cycle latency due to pipeline protection. I2 must wait in the AD stage until I1 finishes its X stage. I2 requires 5 cycles (minimum 1 cycle + 4 pipeline-protection cycles).

Example 4–15. A-Unit Register Write/(Read in AD Stage) Sequence

```

I1: AR1 = #y16      ; Load AR1 with constant
                      ; (AR1 modified in X stage of I1)
I2: AC0 = *AR1+      ; Load AC0 with value pointed to by AR1
                      ; (AR1 read in AD stage of I2)
                      ; Results: AC0 = content of memory at
                      ; location #y16, AR1 = #y16 + 1

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 4 cycles
	I2		I1				4	
	I2			I1			5	
	I2				I1		6	I1 writes to AR1
	I2					I1	7	I2 reads AR1
		I2					8	
			I2				9	
				I2			10	
					I2		11	
						I2	12	

One solution is instruction rescheduling. Between I1 and I2 you can place 4 cycles worth of instructions from elsewhere.

```

I1: AR1 = #y16
nop                      ; Replace NOPs with useful instructions
nop
nop
nop
I2: AC0 = *AR1+

```

Another solution is to use a MAR instruction to write to AR1:

```

I1: MAR(AR1 = #y16)
I2: AC0 = *AR1+

```

A MAR instruction modifies a register in the AD stage, preventing a pipeline conflict. This solution is covered in more detail in section 4.4.2.3 (page 4-61).

Case 2: ARx read followed by an indirect addressing ARx read/update

Example 4–16 shows an AR read followed by an indirect addressing AR update. In the example, I2 is delayed 2 cycles due to pipeline protection. I2 must

wait in the AD stage until I1 finishes its R stage. I2 executes in 3 cycles (minimum 1 cycle + 2 pipeline-protection cycles). Notice that AR1 is read by I1 and is incremented by I2 in the same cycle (cycle 5). This is enabled by an A-unit register prefetch mechanism activated in the R stage of the C55x DSP.

Example 4–16. A-Unit Register Read/(Write in AD Stage) Sequence

```
I1: AC1 = AR1      ; AR1 read in R stage
I2: AC0 = *AR1+    ; AR1 updated in AD stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 2 cycles
	I2		I1				4	
	I2			I1			5	I1 reads AR1; I2 increments AR1
		I2			I1		6	
			I2			I1	7	
				I2			8	
					I2		9	
						I2	10	

To prevent 1 to 3 of the pipeline-protection cycles, you can reschedule instructions. If possible, take up to 3 cycles worth of instructions from elsewhere in the program and place them between I1 and I2.

```
AC1 = AR1
nop          ; Replace NOPs with useful instructions
nop
AC0 = *AR1+
```

One could consider that using a MAR for the ARx register update could also be solution. However as Example 4–16 shows, using MAR can cause unnecessary pipeline conflicts. This is covered in detail in section 4.4.2.3.

Case 3: ACx write followed by an ACx read

Accumulators and registers not associated with address generation are read in the R stage and written at the end of the X stage.

The (write in X stage)/(read in R stage) sequence shown in Example 4–17 costs 1 cycle for pipeline protection. AC0 is updated in the X stage of I1. AC0 must be read in the R stage of I2, but I2 must wait until I1 has written to AC0. The 1 cycle can be regained if you move a 1-cycle instruction between I1 and I2:

```
I1: AC0 = AC0 + #1
      nop      ; Replace NOP with useful instruction
I2: AC2 = @AC0_L || mmap()
```

Notice that

```
AC0 = AC0 + #1
AC2 = AC0
```

will not cause pipeline conflicts because it uses only internal buses (no memory buses) to move data between registers. When AC0_L is accessed via the memory map (@AC0_L || mmap()), it is treated as a memory access.

Example 4–17. Register (Write in X Stage)/(Read in R Stage) Sequence

```
I1: AC0 = AC0 + #1      ; AC0 updated in X stage
I2: AC2 = @AC0_L || mmap() ; AC0_L read in R stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	I1 updates AC0; I2 delayed
				I2		I1	7	I2 reads AC0
					I2		8	
						I2	9	

4.4.2.3 Use MAR instructions when possible, but avoid read/write sequences, and pay attention to instruction size.

The MAR instruction uses independent hardware in the data-address generation unit (DAGEN) to update ARx and Tx registers in the AD stage of the pipeline. You can take advantage of this fact to avoid pipeline conflicts, as shown in Example 4–18. Because AR1 is updated by the MAR instruction prior to being used by I2 for addressing generation, no cycle penalty is incurred.

However, using a MAR instruction could increase instruction size. For example, `mar(AR1 = #y16)` requires 4 bytes, while `AR1 = #y16` instruction requires 3 bytes. You must consider the tradeoff between code size and speed.

Example 4–18. Good Use of MAR Instruction (Write/Read Sequence)

```
I1: mar(AR1 = #y16)    ; AR1 updated in AD stage
I2: AC0 = *AR1+        ; AR1 read in AD stage
                       ; (No cycle penalty)
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	I1 updates AR1
	I2	I1					3	I2 reads AR1
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	
					I2	I1	7	
						I2	8	

Example 4–19 shows that sometimes, using a MAR instruction can cause pipeline conflicts. The MAR instruction (I2) attempts to write to AR1 in the AD stage, but due to pipeline protection, I2 must wait for AR1 to be read in the R stage by I1. This causes a 2-cycle latency. Notice that AR1 is read by I1 and is updated by I2 in the same cycle (cycle 5). This is enabled by an A-unit register prefetch mechanism activated in the R stage of the C55x DSP.

One way to avoid the latency in Example 4–19 is to use the code in Example 4–20:

```
I1: AC1 = AR1 + T1 ; AR1 read in R stage and
I2: AR1 = AR1 + T1 ; AR1 updated in X stage
                       ; (No cycle penalty)
```

Example 4–19. Bad Use of MAR Instruction (Read/Write Sequence)

I1: AC1 = AR1 + T1 ; AR1 read in R stage and
 I2: mar(AR1 + T1) ; AR1 updated in AD stage

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 3 cycles
	I2		I1				4	
	I2			I1			5	I1 reads AR1; I2 updates AR1
		I2			I1		6	
			I2			I1	7	
				I2			8	
					I2		9	
						I2	10	

Example 4–20. Solution for Bad Use of MAR Instruction (Read/Write Sequence)

I1: AC1 = AR1 + T1 ; AR1 read in R stage and
 I2: AR1 = AR1 + T1 ; AR1 updated in X stage
 ; (No cycle penalty)

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	I1 reads AR1
				I2	I1		6	
					I2	I1	7	I2 updates AR1
						I2	8	

4.4.2.4 Pay attention to pipeline-protection granularity when updating status register bit fields.

The C55x pipeline-protection unit protects accesses to fields of the status registers (ST0_55–ST3_55) in different ways depending on the pipeline protection granularity allowed for a given instruction. The pipeline-protection unit can protect accesses to the content of a status register as a whole entity (coarse granularity) or as independent fields (fine granularity). The granularity options for the status registers are summarized in Table 4–7. Fine granularity is ideal, but its hardware implementation can be costly, and for this reason, fine granularity is not available for all registers.

Table 4–7. Status Register Pipeline-Protection Granularity

Register	Fine Granularity Used ...	Coarse Granularity Used ...
ST0_55	If an instruction accesses a bit field in ST0_55 but does not include a direct reference to ST0_55. Example: TC1 = *AR2 & #1h	If an instruction includes a direct reference to ST0_55. Examples: bit(ST0, #ST0_TC1) = #0 @ST0_L= 0x0 mmap()
ST1_55	If an instruction accesses a bit field in ST1_55 but does not include a direct reference to ST1_55. Exception: If the instruction uses one of the following syntaxes, ST1_55 will have fine granularity: bit(ST1, k4) = #0 bit(ST1, k4) = #1	If an instruction includes a direct reference to ST1_55. Example: push (@ST1_L) mmap()
ST2_55	Never	In all cases
ST3_55	Never	In all cases where protection exists. If you access one of these bits ... CAFRZ CAEN CACLR AVIS MPNMC ... you must use one of these syntaxes to have pipeline protection... bit(ST3,k4) = #0 bit(ST3,k4) = #1

In the case of coarse granularity, conflicts can occur when different instructions access two different register fields in a register. Example 4–21 shows a case where two register fields (TC1 and TC2) in ST0_55 cannot be accessed with

fine granularity. The TC1 access in I1 is protected as a write to all of ST0_55. I2 cannot test a bit in ST0_55 until I1 has updated ST0_55.

Example 4–22 shows an instruction sequence that allows for fine granularity during accesses to ST0_55. Because both instructions reference the bits directly rather than by a reference to ST0_55, the TC1 access and the TC2 access can occur in the same cycle.

Example 4–21. ST0_55 Course Granularity

```
I1: bit (ST0, #ST0_TC1) = #0      ; TC1 changed in X stage
I2: if (TC2) goto #label         ; TC2 read in R stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	I1 updates TC1; I2 delayed
				I2		I1	7	I2 tests TC2
					I2		8	
						I2	9	

Example 4–22. ST0_55 Fine Granularity

```
I1: TC1 = AR1 > AC1              ; TC1 changed in X stage
I2: if (TC2) goto #1 2           ; TC2 read in R
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	I1 updates TC1; I2 tests TC2
					I2	I1	7	
						I2	8	

4.4.2.5 Pay attention to pipeline-protection granularity when accessing MMRs in consecutive instructions.

Memory-mapped registers (MMRs) can be viewed by the pipeline-protection unit as individual registers or as part of the same MMR group.

A conflict can occur if MMRs belonging to the same group are accessed at the same time. No conflict will occur if

- ☐ At least one of the MMRs is protected individually
- ☐ Or both MMRs belong to different group

Basically, we have coarse granularity within the groups and fine granularity (or individual protection) among groups. This is shown in Table 4–8.

Table 4–8. MMR Pipeline-Protection Granularity

Pipeline-Protection Granularity	Registers
Fine (individual protection)	ST0_55, ST1_55, ST2_55, ST3_55, AR0, AR1, AR2, AR3, AR4, AR5, AR6, AR7, T0, T1, T2, T3, CDP, SSP, SP, CSR, RPTC, IER0, IER1, DBIER0, DBIER1, IFR0, IFR1, IVPD, IVPH
Coarse (group protection)	Group 1: AC0L, AC0H, AC0G Group 2: AC1L, AC0H, AC0G Group 3: AC2L, AC2H, AC2G Group 4: AC3L, AC3H, AC3G Group 5: TRN0, TRN1 Group 6: BK03, BK47, BKC, BSA01, BSA23, BSA45, BSA67, BSAC Group 7: BRC0, BRC1, BRS1, RSA0, RSA1, REA0, REA1

Example 4–23 shows a case of MMR coarse granularity. Even though BSA23 is not the same register as BSA01, they both belong to the same MMR group (as shown in Table 4–8). For this reason, I2 waits for I1 to complete the write to BSA23 (in the W stage) before proceeding with the read from BSA01 (in the R stage). The result is a 2-cycle delay of I2. The solution is to move instructions from elsewhere in the program such that sufficient space is between the I1 and I2.

Example 4–24 shows MMR fine granularity. Because the pipeline-protection unit has fine granularity (individual protection) for CDP and AR1, it does not delay I2.

Example 4–23. MMR Coarse Granularity

```

I1: @BSA23_L = AC0 || mmap(); BSA23 written to in W stage
I2: AC1 = @BSA01_L || mmap(); BSA01 read in R stage

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 4 cycles
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	
				I2		I1	7	I1 updates BSA23
				I2			8	I2 reads BSA01
					I2		9	
						I2	10	

Example 4–24. MMR Fine Granularity

```

I1: @CDP = AC1 || mmap() ; CDP written to in X stage
I2: AC1 = *AR1 ; AR1 read in AD stage

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 reads AR1
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	I1 updates CDP
					I2	I1	7	
						I2	8	

4.4.2.6 Understand when the loop-control registers are accessed in the pipeline

As in any register, the h/w loop controller registers (CSR, RPTC, RSA0, REA0, RSA1, REA1, BRC0, BRC1, and BRS1) are read in R stage and written in X stage. The exception is when being modified by the repeat instruction themselves. For example:

- During *repeat()* loop execution:
 - CSR or an instruction constant is loaded into RPTC in the address stage of the single *repeat()* instruction.
 - RPTC is tested and decremented in the address stage of each repeated instruction.
- During *blockrepeat{}* and *localrepeat{}* loop execution:
 - BRS1 is loaded into BRC1 in the address stage of the *blockrepeat* or *localrepeat* instructions.
 - RSAx and REAx are loaded in the address stage of the *blockrepeat* or *localrepeat* instructions.
 - BRCx is decremented in the decode stage of the last instruction of the loop.
 - RSAx and REAx are constantly read in the address stage of each loop instruction.

4.4.2.7 Avoid accessing the BRC register in the last 4 instructions of a block-repeat or local block-repeat structure (to prevent a delay of the BRC decrement or an unprotected pipeline situation).

Reading BRC in one of the last three instructions of a block-repeat structure could cause delay of the BRC automatic decrement inside the loop or cause an unprotected pipeline situation. Similarly, writing to BRC in one of the last four instructions of a block-repeat structure could cause an unprotected pipeline situation.

Example 4–25 shows the effect of reading BRC0 within the last three instructions of a block-repeat structure. $\mathbb{I}1$ reads BRC0 in the R stage of the pipeline (cycle 5), and the CPU must decrement BRC0 in the D stage of $\mathbb{I}2$ (the last instruction of a block-repeat structure). The pipeline-protection unit keeps the proper sequence of these operations (read BRC0 and then decrement BRC0) by delaying $\mathbb{I}2$ in the D stage for 4 cycles.

Example 4–25. Protected BRC Read

```

I1: T1 = BRC0                ; BRC0 read in R stage
I2: Inst 2 (last instruction of a
    block repeat structure)   ; BRC0 decremented in D
                                ; stage

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	I2 delayed 4 cycles
I2		I1					3	
I2			I1				4	
I2				I1			5	I1 reads BRC0
I2					I1		6	I2 decrements BRC0
	I2					I1	7	
		I2					8	
			I2				9	
				I2			10	
					I2		11	
						I2	12	

Example 4–26 shows the effect of reading the BRC register in the last instruction of a block-repeat structure. BRC0 is to be read by I1 in the R stage of the pipeline (cycle 5) while BRC0 is to be decremented in the AD stage of I2 (the last instruction of a block-repeat structure). The pipeline-protection unit *cannot* guarantee the sequence of these operations (read BRC0 and then write to BRC0). BRC0 is decremented first, and then the decremented content of BRC0 is read by I1.

Example 4–26. Unprotected BRC Read

```

I1: T2 = BRC0    ; Last instruction of block-repeat
    }           ; structure
I2: Inst 2

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
	I1						2	CPU decrements BRC0
		I1					3	
			I1				4	
				I1			5	I1 reads decremented BRC0
					I1		6	
						I1	7	

BRC write accesses are not protected in the last 4 cycles of a block-repeat structure. Do not write to BRC0 or BRC1 within those cycles.

In Example 4–27 BRC0 is to be written to by I1 in the W stage (cycle 7), while BRC0 is decremented in the D stage of I2. The pipeline-protection unit *cannot* guarantee the proper sequence of these operations (write to BRC0 and then decrement BRC0). BRC0 is decremented by I2 before BRC0 changed by I1.

Example 4–27. Unprotected BRC Write

```

I1: @BRC0_L = BRC0_L + #1    ;BRC0 written in W stage
    || mmap()
I2: Inst2    ; Last instruction of a block-repeat loop
           ; BRC0 decremented in D stage of last
           ; instruction

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	CPU decrements BRC0
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	
					I2	I1	7	I1 changes BRC0 late
						I2	8	

4.4.2.8 Do not access loop-control registers RSAx, REAx, and RPTC within the loop itself.

Accesses to the register RSA0, RSA1, REA0, and REA1 are not protected when the read or the write is performed within block repeat structures. Similarly, RPTC register accesses are not protected when the accesses are performed within a (conditional) single-instruction repeat loop. Do not access any of these loop-control registers within the loop itself.

4.4.2.9 Initialize the BRCx or CSR register at least 4 cycles before the repeat instruction, or initialize the register with an immediate value.

Whenever BRC1 is loaded, BRS1 is loaded with the same value. In Example 4–28, BRC1 and BRS1 are to be loaded by I1 in X stage of the pipeline (cycle 6), while BRS1 is to be read by I2 in the AD stage (cycle 3) to initialize BRC1. The pipeline-protection unit keeps the proper sequence of these operations (write to BRS1 and then read BRS1) by delaying the completion of I2 by 4 cycles. Example 4–29 shows a similar situation with CSR.

Instruction rescheduling or initialization of BRC1/CSR with an immediate value will remove the pipeline conflict in Example 4–28. An instruction that loads BRC1/CSR with an immediate value will do so in the AD stage of the instruction.

Example 4–28. BRC Initialization

```
I1: BRC1 = *AR1          ; BRC1 and BRS1 loaded in X stage
I2: blockrepeat {        ; BRS1 is copied to BRC1 in AD stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 4 cycles
	I2		I1				4	
	I2			I1			5	
	I2				I1		6	I1 loads BRS1
	I2					I1	7	I2 copies BRS1 into BRC1
		I2					8	
			I2				9	
				I2			10	
					I2		11	
						I2	12	

Example 4–29. CSR Initialization

```

I1: CSR = * AR1      ; CSR written in X stage
I2: repeat(CSR)     ; CSR read in AD stage

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	I2 delayed 4 cycles
	I2		I1				4	
	I2			I1			5	
	I2				I1		6	I1 loads CSR
	I2					I1	7	I2 reads CSR
		I2					8	
			I2				9	
				I2			10	
					I2		11	
						I2	12	

4.4.2.10 Try to set conditions well in advance of the time that the condition is tested.

Conditions are typically evaluated in the R stage of the pipeline with the following exceptions:

- ☐ When the (AD_unit) keyword is used, the condition is evaluated in the AD stage; for example:

```

if (cond) execute (AD_unit)

```
- ☐ When the (D_unit) keyword is used, the condition is evaluated in the X stage; for example:

```

if (cond) execute (D_unit)

```

The exception is when the D_unit instruction conditions a write to memory. In this case the condition is evaluated in the read stage; for example:

```

if (cond) execute (D_unit)
|| *ar1 += ac1

```
- ☐ When a condition is used to make a “forward” branch within a local block repeat structure, the condition is evaluated in the AD stage; for example:

```

if (cond) goto #L16

```

Example 4–30 involves a condition evaluation preceded too closely by a write to the register affecting the condition. AC1 is updated by I1 in the X stage,

while I2 must read AC1 in the R stage to evaluate the condition (AC1 == #0). The pipeline-protection unit ensures the proper sequence of these operations (write to AC1 and then test AC1) by delaying the completion of I2 by 1 cycle. The solution is to move I1 within the program, such that it updates AC1 at least 1 cycle sooner.

Example 4–30. Condition Evaluation Preceded by a X-stage Write to the Register Affecting the Condition

```
I1: AC1 = AC1 + #1           ; AC1 update in X stage
I2: If (AC1 == #0) goto #subroutine ; AC1 test in R stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	I1 updates AC1; I2 delayed
				I2		I1	7	I2 tests AC1
					I2		8	
						I2	9	

4.4.2.11 When making an instruction execute conditionally, use `execute(AD_unit)` instead of `execute(D_unit)` to create fully protected pipeline conditions, but be aware of potential pipeline delays.

Example 4–31 shows a case where a register load operation (AC1 = *AR3+) is made conditional with the *if (condition) execute(AD-unit)* instruction. When the *execute(AD_unit)* keyword is used, the condition is evaluated in the AD stage, and if the condition is true, the conditional instruction performs its operations in the AD through W stages. In Example 4–31, the AR3+ update in I3 depends on whether AC0 > 0; therefore, the update is postponed until I1 updates AC0. As a result, I3 is delayed by 4 cycles.

One solution to the latency problem in Example 4–31 is to move I1 such that AC0 is updated 4 cycles earlier. Another solution is to replace the *execute(AD_unit)* keyword with the *execute(D_unit)* keyword, as shown in Example 4–32. When the *execute(D_unit)* keyword is used, the condition is evaluated in the X stage, and if the condition is true, the conditional instruction performs its X stage operation. Operations in the AD through R stages will happen unconditionally. In Example 4–32, AR3 is updated in the AD stage regard-

less of the condition. Also, the memory read operation in `I2` will always happen. However, the memory value will be written `AC1` only if the condition is true. Otherwise, the memory value is discarded. Overall, a zero-latency execution is achieved.

Notice that the advantage of using `execute(AD_unit)` is that it provides a 100% pipeline protection environment. However, this comes at the expense of added latency cycles.

Example 4–31. Making an Operation Conditional With `execute(AD_unit)`

```
I1: AC0 = AC0 + #1 ; AC0 updated in X stage
I2: If (AC0 > 0) execute (AD_Unit) ; AC0 tested in AD stage
I3: AC1 = *AR3+
    ; If AC0 > 0, increment AR3 in AD stage and load AC1
    ; in X stage
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
I3	I2	I1					3	I2 and I3 delayed 4 cycles
I3	I2		I1				4	
I3	I2			I1			5	
I3	I2				I1		6	I1 updates AC0
I3	I2					I1	7	I2 tests AC0
	I3	I2					8	If AC0 > 0, I3 updates AR3
		I3	I2				9	
			I3	I2			10	
				I3	I2		11	
					I3	I2	12	If AC0 > 0, I3 updates AC1
						I3	13	

Example 4–32. Making an Operation Conditional With `execute(D_unit)`

```

I1: AC0 = AC0 + #1           ; AC0 updated in X stage
I2: If (AC0 > 0) execute (D_Unit) ; AC0 tested in X stage
I3:  AC1 = *AR3+
    ; If AC0 > 0, load AC1 in X stage. Update AR3 in AD
    ; stage regardless of the condition.

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1								
I2	I1							
I3	I2	I1						
	I3	I2	I1					I3 updates AR3 unconditionally
		I3	I2	I1				
			I3	I2	I1			I1 updates AC0
				I3	I2	I1		I2 tests AC0
					I3	I2		If AC0 > 0, I3 updates AC1
						I3		

4.4.2.12 Understand the `execute (D_unit)` condition evaluation exception.

Typically, the *execute (D_unit)* keyword causes the CPU to evaluate the condition in the execute (X) stage. The exception is when you make a memory write operation dependent on a condition, in which case the condition is evaluated in the read (R) stage.

In Example 4–33, AR3 is to be read by I1 in R stage to evaluate the condition, while AR3 is to be modified by I3 in the AD stage. The pipeline-protection unit keeps the proper sequence of these operations (read AR3 and then write to AR3) by delaying the completion of I3 by 2 cycles. Notice that AR3 is tested by I1 and is modified by I3 in the same cycle (cycle 5). This is enabled by an A-unit register prefetch mechanism activated in the R stage of the C55x DSP.

To prevent the 2-cycle delay in Example 4–33, you can insert two other, non-conflicting instructions between I2 and I3.

Example 4–33. Conditional Parallel Write Operation Followed by an AD-Stage Write to the Register Affecting the Condition

```

I1: If (AR3 == #0) execute(D-Unit) ; AR3 tested in R stage
I2: || *AR1+ = AC1
    ; If AR3 contains 0, write to memory in W stage.
    ; Update AR1 regardless of condition.
I3: AC1 = *(AR3 + T0)          ; AR3 updated in AD stage

```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1 I2							1	
I3	I1 I2						2	I2 updates AR1 unconditionally
	I3	I1 I2					3	I3 delayed 2 cycles
	I3		I1 I2				4	
	I3			I1 I2			5	I1 tests AR3; I3 modifies AR3
		I3			I1 I2		6	
			I3			I1 I2	7	
				I3			8	
					I3		9	
						I3	10	

4.4.3 Memory Accesses and the Pipeline

This section discusses the impact of the memory mapping of variables/array on the execution of C55x instructions. The two factors to consider are the type of memory (DARAM vs SARAM) and the memory bank allocation. DARAM supports 2 accesses/cycle to the same memory bank. SARAM supports 1 access/cycle to one memory bank.

Even though the same buses are used to access SARAM and DARAM memory banks, the specifics of the accesses are different in the two cases. Table 4–9 and Table 4–10 show bus activity by access type for DARAM and SARAM, respectively. A shaded cell in the table indicates bus activity. Within the shaded cells are references to the CPU memory buses being used:

Symbol	Bus
B	BB. This data-read data bus carries a 16-bit coefficient data value (Cmem) from data space to the CPU.
C, D	CB, DB. Each of these data-read data buses carries a 16-bit data value to the CPU. DB carries a value from data space or from I/O-space. CB carries a second value in certain instances when the CPU reads two data-space values at the same time.
E, F	EB, FB. Each of these data-write data buses carries a 16-bit data value from the CPU. EB carries a value to data space or to I/O-space. FB carries a second value when the CPU writes two values to data space at the same time.

Table 4–9. Half-Cycle Accesses to Dual-Access Memory (DARAM)

Access Type	D	AD	AC1	AC2	R	X	W	
Smem read				D				
Smem write								E
Smem read/ modify/write				D				E
Lmem read				C D				
Lmem write								E F
Xmem read Ymem read				D	C			
Xmem write Ymem write							F	E
Xmem read Ymem read Cmem read				D	C B			
Xmem read Ymem write				D				E
Lmem read Lmem write				C D				E F

Table 4–10. One-Cycle Accesses to Single-Access Memory (SARAM)

Access Type	D	AD	AC1	AC2	R	X	W	
Smem read					D			
Smem write								E
Smem read/ modify/write					D			E
Lmem read					C D			
Lmem write								E F
Xmem read Ymem read					D C			
Xmem write Ymem write								E F
Xmem read Ymem read Cmem read					D C B			
Xmem read Ymem write					D			E
Lmem read Lmem write					C D			E F

Ideally, we should allocate all data into DARAM due to its higher memory bandwidth (2 cycles/access). However, DARAM is a limited resource and should be used only when it is advantageous. Following are recommendations to guide your memory mapping decisions.

- ☐ Reschedule instructions.
- ☐ Reduce memory accesses by using CPU registers to hold data.
- ☐ Reduce memory accesses by using a local repeat instruction, an instruction that enables the CPU to repeatedly execute a block of code from the instruction buffer queue.
- ☐ Relocate variables and data arrays in memory, or consider temporarily copying arrays to other nonconflicting memory banks at run time.

4.4.3.1 When working with dual MAC and FIR instructions, put the Cmem operand in a different memory bank.

The only memory access type which can generate a conflict in a DARAM is the execution of instructions requiring three data operands in 1 cycle: Xmem, Ymem, and Cmem (coefficient operand). The instructions that use three data operands are:

- ☐ Dual multiply-and-accumulate (MAC) instruction:
$$ACx = M40(rnd(ACx + uns(Xmem) * uns(Cmem))) ,$$
$$ACy = M40(rnd(ACy + uns(Ymem) * uns(Cmem)))$$
- ☐ Finite impulse response filter instructions:
$$firs(Xmem, Ymem, Cmem, ACx, ACy)$$
$$firsn(Xmem, Ymem, Cmem, ACx, ACy)$$

This memory conflict can be solved by maintaining the Ymem and Xmem operands in the same DARAM memory bank but putting the Cmem operand into a different memory bank (SARAM or DARAM).

When cycle intensive DSP kernels are developed, it's extremely important to identify and document software integration recommendations stating which variables/arrays must not be mapped in the same DARAM memory block.

The software developer should also document the associated cycle cost when the proposed optimized mapping is not performed. That information will provide the software integrator enough insight to make trade-offs.

When cycle intensive DSP kernels are developed, it is extremely important to identify and document s/w integration recommendations stating which variables/arrays must not be mapped in the same dual access memory. The software developer should also document the associated cycle cost when the proposed optimized mapping is not performed. That information will provide the software integrator enough insight to make trade-offs. Table 4–11 provides an example of such table: if the 3 arrays named “input,” “output,” and “coefficient” are in the same DARAM, the subroutine named “filter” will have 200 cycle overhead.

Table 4–11. *Cross-Reference Table Documented By Software Developers to Help Software Integrators Generate an Optional Application Mapping*

Routine	Cycle Weight	Array 1 Name (Xmem)	Size	Array 2 Name (Ymem)	Size	Array 3 Name (Cmem)	Size	Cycle Cost per Routine
filter	10%	input	40	output	40	coefficient	10	10*20
...

4.4.3.2 When working on data in the same memory bank, pay attention to instruction sequences that could generate memory conflicts.

Example 4–34 shows how instruction sequences can generate memory conflicts (2 accesses per cycle per bank) for a SARAM block. The code in the example performs a memory write operation followed by a dual-operand read operation within the next 2 cycles.

If an instruction performing a dual operand read access is 2 cycle away from a single operand write instruction, a memory bank access conflict will occur if the 3 accesses are performed in the same memory bank. This is shown in Example 4–34. The solution is to use DARAM (0 cycle delay).

Example 4–34. Write/Dual-Operand Read Sequence
(Assumes Xmem, Ymem, and Smem Are in the Same SARAM block)

```
I1: Smem = ACx
I2: NOP
I3: ACx = (Xmem << #16) + (Ymem << #16)
```

D	AD	AC1	AC2	R	X	W	Cycle	Comment
I1							1	
I2	I1						2	
	I2	I1					3	
		I2	I1				4	
			I2	I1			5	
				I2	I1		6	
				I3	I2	I1	7	I1 writes to Smem; I3 delayed
				I3		I2	8	I3 reads Xmem
				I3			9	I3 reads Ymem
					I3		10	
						I3	11	

4.4.3.3 Map program code to a dedicated SARAM memory block to avoid conflicts with data accesses

If a DARAM block maps both program and data spaces of the same routine, a program code fetch will conflict with a dual(or triple) data operand read (or write) access if they are performed in the same memory block. C55x DSP resolves the conflict by delaying the program code fetch by one cycle. It is therefore recommended to map the program code in a dedicated program memory bank: generally a SARAM memory bank is preferred. This enables to avoid conflicts with data variables mapped in the high bandwidth DARAM banks.

Another way to avoid memory conflicts is to use the 64-byte instruction buffer to execute blocks of instructions without refetching code after the 1st iteration (see *localrepeat{}* instruction). Conflict will only occur in the first loop iteration.

4.4.3.4 For 32-bit accesses (using an Lmem operand), no performance hit is incurred if you use SARAM (there is no need to use DARAM).

When a 32-bit memory access is performed with Lmem, only one *address* bus (DAB or EAB) is used to specify the most and least significant words of the 32-bit value. Therefore, reading from or writing to a 32-bit memory location in an SARAM bank occurs in 1 cycle.

Fixed-Point Arithmetic



The TMS320C55x™ (C55x™) DSP is a 16-bit, fixed-point processor. This chapter explains important considerations for performing standard- and extended-precision fixed-point arithmetic with the DSP. Assembly-language code examples are provided to demonstrate the concepts.

Topic	Page
5.1 Fixed-Point Arithmetic	5-2
5.2 Extended-Precision Addition and Subtraction	5-10
5.3 Extended-Precision Multiplication	5-17
5.4 Division	5-21
5.5 Methods of Handling Overflows	5-31

5.1 Fixed-Point Arithmetic

Digital signal processors (DSPs) have been developed to process complex algorithms that require heavy computations. DSPs can be divided into two groups: floating-point DSPs and fixed-point DSPs.

Typically, floating-point DSPs use 32-bit words composed of a 24-bit mantissa and an 8-bit exponent, which together provide a dynamic range from 2^{-127} to $2^{128}(1 - 2^{-23})$. This vast dynamic range in floating-point devices means that dynamic range limitations may be virtually ignored in a design. Floating-point devices are usually more expensive and consume more power than fixed-point devices.

Fixed-point DSPs, like the TMS320C55x DSP, typically use 16-bit words. They use less silicon area than their floating-point counterparts, which translates into cheaper prices and less power consumption. Due to the limited dynamic range and the rules of fixed-point arithmetic, a designer must play a more active role in the development of a fixed-point DSP system. The designer has to decide whether the 16-bit words will be interpreted as integers or fractions, apply scale factors if required, and protect against possible register overflows.

5.1.1 2s-Complement Numbers

In binary form, a number can be represented as a signed magnitude, where the left-most bit represents the sign and the remaining bits represent the magnitude.

$$+52 = 0\ 011\ 0100b$$

$$-52 = 1\ 011\ 0100b$$

This representation is not used in a DSP architecture because the addition algorithm would be different for numbers that have the same signs and for numbers that have different signs. The DSP uses the 2s-complement format, in which a positive number is represented as a simple binary value and a negative value is represented by inverting all the bits of the corresponding positive value and then adding 1.

Example 5–1 shows the decimal number 353 as a 16-bit signed binary number. Each bit position represents a power of 2, with 2^0 at the position of the least significant bit and 2^{15} at the position of the most significant bit. The 0s and 1s of the binary number determine how these powers of 2 are weighted (times 0 or times 1) when summed to form 353. Because the number is signed, 2^{15} is given a negative sign. Example 5–2 shows how to compute the negative of a 2s-complement number.

Example 5–1. Signed 2s-Complement Binary Number Expanded to Decimal Equivalent

2^{15}	2^{14}	2^{13}	2^{12}	2^{11}	2^{10}	2^9	2^8	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
0	0	0	0	0	0	0	1	0	1	1	0	0	0	0	1

$$\begin{aligned}
&= (0 \times (-2^{15})) + (0 \times 2^{14}) + (0 \times 2^{13}) + (0 \times 2^{12}) + (0 \times 2^{11}) + (0 \times 2^{10}) + (0 \times 2^9) + (1 \times 2^8) + (0 \times 2^7) \\
&\quad + (1 \times 2^6) + (1 \times 2^5) + (0 \times 2^4) + (0 \times 2^3) + (0 \times 2^2) + (0 \times 2^1) + (1 \times 2^0) \\
&= (1 \times 2^8) + (1 \times 2^6) + (1 \times 2^5) + (1 \times 2^0) \\
&= 256 + 64 + 32 + 1 = 353
\end{aligned}$$

Example 5–2. Computing the Negative of a 2s-Complement Number

Begin with a positive binary number (353 decimal):	0000 0001 0110 0001
Invert all bits to get the 1s complement:	1111 1110 1001 1110
Add 1 to get the 2s complement:	$ \begin{array}{r} 1111\ 1110\ 1001\ 1110 \\ + 1 \\ \hline 1111\ 1110\ 1001\ 1111 \end{array} $
Result: negative binary number (–353 decimal):	1111 1110 1001 1111

5.1.2 Integers Versus Fractions

The most common formats used in DSP programming are integers and fractions. In signal processing, fractional representation is more common. A fraction is defined as a ratio of two integers such that the absolute value of the ratio is less than or equal to 1. When two fractions are multiplied together, the result is also a fraction. Multiplicative overflow, therefore, never occurs. Note, however, that additive overflow can occur when fractions are added. Overflows are discussed in section 5.5, beginning on page 5-31.

Figure 5–1 shows how you can interpret 2s-complement numbers as integers. The most significant bit (MSB) is given a negative weight, and the integer is the sum of all the applicable bit weights. If a bit is 1, its weight is included in the sum; if the bit is 0, its weight is not applicable (the effective weight is 0). For simplicity, the figure shows 4-bit binary values; however, the concept is easily extended for larger binary values. Compare the 4-bit format in Figure 5–1 with the 8-bit format in Figure 5–2. The LSB of a binary integer always has a bit weight of 1, and the absolute values of the bit weights increase toward the MSB. Adding bits to the left of a binary integer does not change the absolute bit weights of the original bits.

Figure 5–1. 4-Bit 2s-Complement Integer Representation

4-bit 2s-complement integer	<table><tr><td>MSB</td><td></td><td></td><td>LSB</td></tr></table>				MSB			LSB
MSB			LSB					
Bit weights	$-2^3 = -8$	$2^2 = 4$	$2^1 = 2$	$2^0 = 1$				
<hr/>								
Least positive value	0	0	0	1	$= 0 + 0 + 0 + 1 = 1$			
Most positive value	0	1	1	1	$= 0 + 4 + 2 + 1 = 7$			
Least negative value	1	1	1	1	$= -8 + 4 + 2 + 1 = -1$			
Most negative value	1	0	0	0	$= -8 + 0 + 0 + 0 = -8$			
Other examples:	0	1	0	1	$= 0 + 4 + 0 + 1 = 5$			
	1	1	0	1	$= -8 + 4 + 0 + 1 = -3$			

Figure 5–2. 8-Bit 2s-Complement Integer Representation

MSB							LSB
$-2^7 = -128$	$2^6 = 64$	$2^5 = 32$	$2^4 = 16$	$2^3 = 8$	$2^2 = 4$	$2^1 = 2$	$2^0 = 1$

Figure 5–3 shows how 2s-complement numbers can be interpreted as a fractions. The concept is much the same as that in Figure 5–1, but the bit weights are fractional, meaning that the number cannot have an absolute value larger than 1. Compare the 4-bit format in Figure 5–3 with the 8-bit format in Figure 5–4. The MSB of a binary fraction always has a bit weight of -1 , and the absolute values of the bit weights decrease toward the LSB. Unlike adding bits to the left of a binary integer, adding bits to the left of a binary fraction changes the bit weights of the original bits.

Figure 5–3. 4-Bit 2s-Complement Fractional Representation

4-bit binary fraction	MSB			LSB	
Bit weights	$-2^0 = -1$	$2^{-1} = 1/2$	$2^{-2} = 1/4$	$2^{-3} = 1/8$	
Least positive value	0	0	0	1	$= 0 + 0 + 0 + 1/8$ $= 1/8$
Most positive value	0	1	1	1	$= 0 + 1/2 + 1/4 + 1/8$ $= 7/8$
Least negative value	1	1	1	1	$= -1 + 1/2 + 1/4 + 1/8$ $= -1/8$
Most negative value	1	0	0	0	$= -1 + 0 + 0 + 0$ $= -1$
Other examples:	0	1	0	1	$= 0 + 1/2 + 0 + 1/8$ $= 5/8$
	1	1	0	1	$= -1 + 1/2 + 0 + 1/8$ $= -3/8$

Figure 5–4. 8-Bit 2s-Complement Fractional Representation

MSB							LSB
$-2^0 = -1$	$2^{-1} = 1/2$	$2^{-2} = 1/4$	$2^{-3} = 1/8$	$2^{-4} = 1/16$	$2^{-5} = 1/32$	$2^{-6} = 1/64$	$2^{-7} = 1/128$

5.1.3 2s-Complement Arithmetic

An important advantage of the 2s-complement format is that addition is performed with the same algorithm for all numbers. To become more familiar with 2s-complement binary arithmetic, refer to the examples in this section. You may want to try a few examples yourself. It is important to understand how 2s-complement arithmetic is performed by the DSP instructions, in order to efficiently debug your program code.

Example 5–3 shows two 2s-complement additions. These binary operations are completely independent of the convention the programmer uses to convert

them into decimal numbers. To highlight this fact, an integer interpretation and a fractional interpretation are shown for each addition. For simplicity, the examples use 8-bit binary values; however, the concept is easily extended for larger binary values. For a better understanding of how the integer and fractional interpretations were derived for the 8-bit binary numbers, see Figure 5–2 (page 5-4) and Figure 5–4 (page 5-5), respectively.

Example 5–3. Addition With 2s-Complement Binary Numbers

2s-Complement Addition	Integer Interpretation	Fractional Interpretation
<div>1 (carry)</div> <div>0000 0101</div> <div>+ 0000 0100</div> <div>-----</div> <div>0000 1001</div>	<div>5</div> <div>+ 4</div> <div>-----</div> <div>9</div>	<div>5/128</div> <div>+ 4/128</div> <div>-----</div> <div>9/128</div>
<div>1 1 1 (carries)</div> <div>0000 0101</div> <div>+ 0000 1101</div> <div>-----</div> <div>0001 0010</div>	<div>5</div> <div>+ 13</div> <div>-----</div> <div>18</div>	<div>5/128</div> <div>+ 13/128</div> <div>-----</div> <div>18/128</div>

Example 5–4 shows subtraction. As with the additions in Example 5–3, an integer interpretation and a fractional interpretation are shown for each computation. It is important to notice that 2s-complement subtraction is the same as the addition of a positive number and a negative number. The first step is to find the 2s-complement of the number to be subtracted. The second step is to perform an addition using this negative number.

Example 5–4. Subtraction With 2s-Complement Binary Numbers

2s-Complement Subtraction	Integer Interpretation	Fractional Interpretation
Original form: $\begin{array}{r} 0000\ 0101 \\ -\ 0000\ 0100 \\ \hline \end{array}$	$\begin{array}{r} 5 \\ -\ 4 \\ \hline \end{array}$	$\begin{array}{r} 5/128 \\ -\ 4/128 \\ \hline \end{array}$
2s complement of subtracted term: $\begin{array}{r} 11\ (\text{carries}) \\ 1111\ 1011 \\ +\quad\quad 1 \\ \hline 1111\ 1100 \end{array}$		
Addition form: $\begin{array}{r} 11111\ 1\ (\text{carries}) \\ 0000\ 0101 \\ +\ 1111\ 1100 \\ \hline 0000\ 0001 \\ (\text{final carry ignored}) \end{array}$	$\begin{array}{r} 5 \\ +\ (-4) \\ \hline 1 \end{array}$	$\begin{array}{r} 5/128 \\ +\ (-4/128) \\ \hline 1/128 \end{array}$
Original form: $\begin{array}{r} 0000\ 0101 \\ -\ 0000\ 1101 \\ \hline \end{array}$	$\begin{array}{r} 5 \\ -\ 13 \\ \hline \end{array}$	$\begin{array}{r} 5/128 \\ -\ 13/128 \\ \hline \end{array}$
2s complement of subtracted term: $\begin{array}{r} 1111\ 0010 \\ +\quad\quad 1 \\ \hline 1111\ 0011 \end{array}$		
Addition form: $\begin{array}{r} 111\ (\text{carries}) \\ 0000\ 0101 \\ +\ 1111\ 0011 \\ \hline 1111\ 1000 \end{array}$	$\begin{array}{r} 5 \\ +\ (-13) \\ \hline -8 \end{array}$	$\begin{array}{r} 5/128 \\ +\ (-13/128) \\ \hline -8/128 \end{array}$

Example 5–5 shows 2s-complement multiplication. For simplicity, the example uses a 4-bit by 4-bit multiplication and assumes an 8-bit accumulator for the result. Notice that the 7-bit mathematical result is sign extended to fill the accumulator. The C55x DSP sign extends multiplication results in this way, extending the result to either 32 bits or 40 bits. The effects of this type of sign extension can be seen in the integer and fractional interpretations of Example 5–5. The integer is not changed by sign extension, but the fraction can be misinterpreted. Sign extension adds an extra sign bit. If your program assumes that the MSB of the result is the only sign bit, you must shift the result left by 1 bit to remove the extra sign bit. In the C55x DSP, there is a control bit called FRCT to automate this shift operation. When FRCT = 1, the DSP automatically performs a left shift by 1 bit after a multiplication. You can clear and set FRCT with the following instructions:

```
bit(ST1,#ST1_FRCT) = #0 ; Clear FRCT
bit(ST1,#ST1_FRCT) = #1 ; Set FRCT
```

Example 5–5. Multiplication With 2s-Complement Binary Numbers

2s-Complement Multiplication		Integer Interpretation	Fractional Interpretation
0100	Multiplicand	4	4/8
x 1101	Multiplier	x (–3)	x (–3/8)
<hr/>		<hr/>	<hr/>
0000100			
000000			
00100			
1100	(see Note)		
<hr/>			
1110100	7-bit mathematical result	–12	–12/64 (The MSB of the result is the only sign bit)
11110100	8-bit sign-extended result in accumulator	–12	–12/64 if properly interpreted; –12/128 if incorrectly interpreted. To remove extra sign bit in MSB position, shift result left by 1 bit.
Note: Because the MSB is a sign bit, the final partial product is the 2s complement negative of the multiplicand.			

5.1.4 Extended-Precision 2s-Complement Arithmetic

Numerical analysis, floating-point computations, and other operations may require arithmetic operations with more than 32 bits of precision. Because the C55x device is a 16-/32-bit fixed-point processor, software is required for arithmetic operations with extended precision. These arithmetic functions are performed in parts, similar to the way in which longhand arithmetic is done.

The DSP has several features that help make extended-precision calculations more efficient. One of the features is the CARRY status bit, which is affected by most arithmetic D-unit ALU instructions, as well as the rotate and shift operations. CARRY depends on the M40 status bit. When $M40 = 0$, the carry/borrow is detected at bit position 31. When $M40 = 1$, the carry/borrow reflected in CARRY is detected at bit position 39. Your code can also explicitly modify CARRY by loading ST0_55 or by using a status bit clear/set instruction. For proper extended-precision arithmetic, the saturation mode bit should be cleared ($SATD = 0$) to prevent the accumulator from saturating during the computations.

Two C55x data buses, CB and DB, allow some instructions to handle 32-bit operands in a single cycle. The long-word load and double-precision add/subtract instructions use 32-bit operands and can efficiently implement extended-precision arithmetic.

The hardware multiplier can multiply signed/unsigned numbers, as well as multiply two signed numbers and two unsigned numbers. This makes 32-bit multiplication operations efficient.

5.2 Extended-Precision Addition and Subtraction

The CARRY bit is set in status register 0 (ST0_55) if a carry is generated when an accumulator value is added to:

- ☐ Another accumulator value
- ☐ A data-memory operand
- ☐ An immediate operand (embedded in the instruction)

A carry can also be generated when two data-memory operands are added or when a data-memory operand is added to an immediate operand. If a computation does not generate a carry, the CARRY bit is cleared.

The ADD instruction with a 16-bit shift (shown following this paragraph) is an exception because it can only set the CARRY bit. If this instruction does not generate a carry, the CARRY bit is left unchanged. This allows the D-unit ALU to generate the appropriate carry when adding to the lower or upper half of the accumulator causes a carry.

Mnemonic instruction: ADD Smem << #16, ACx, ACy

Algebraic instruction: $ACy = ACx + (Smem \ll \#16)$

Figure 5–5 shows several 32-bit additions and their effects on the CARRY bit, which is referred to as C in the figure.

Figure 5–5. 32-Bit Addition

C	MSB												LSB	
X	F	F	F	F	F	F	F	F	F	F	F	F	F	ACx
	+												1	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	

C	MSB												LSB	
X	F	F	F	F	F	F	F	F	F	F	F	F	F	ACx
	+F												F	
1	F	F	F	F	F	F	F	F	F	F	F	F	F	E

C	MSB												LSB	
X	0	0	7	F	F	F	F	F	F	F	F	F	F	ACx
	+												1	
0	0	0	8	0	0	0	0	0	0	0	0	0	0	

C	MSB												LSB	
X	0	0	7	F	F	F	F	F	F	F	F	F	F	ACx
	+F												F	
1	0	0	7	F	F	F	F	F	F	F	F	F	F	E

C	MSB												LSB	
X	F	F	8	0	0	0	0	0	0	0	0	0	0	ACx
	+												1	
0	F	F	8	0	0	0	0	0	0	0	0	0	1	

C	MSB												LSB	
X	F	F	8	0	0	0	0	0	0	0	0	0	0	ACx
	+F												F	
1	F	F	7	F	F	F	F	F	F	F	F	F	F	

$$ACy = ACx + Smem + CARRY$$

C	MSB												LSB	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	ACx
	+												0	
0	0	0	0	0	0	0	0	0	0	0	0	0	1	ACy

C	MSB												LSB	
1	F	F	F	F	F	F	F	F	F	F	F	F	F	ACx
	+												0	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	

$$ACy = ACx + (Smem \ll 16)$$

C	MSB												LSB	
1	F	F	8	0	0	0	F	F	F	F	F	F	F	ACx
	+0												0	
1	F	F	8	0	0	1	F	F	F	F	F	F	F	ACy

C	MSB												LSB	
1	F	F	8	0	0	0	F	F	F	F	F	F	F	ACx
	+0												0	
1	F	F	F	F	F	F	F	F	F	F	F	F	F	

The code in Example 5–6 adds two 64-bit numbers to obtain a 64-bit result. The partial sum of the 64-bit addition is efficiently performed by the following instructions, which handle 32-bit operands in a single cycle.

Mnemonic instructions: MOV40 dbl(Lmem), ACx

ADD dbl(Lmem), ACx

Algebraic instructions: ACx = dbl(Lmem)

ACx = ACx + dbl(Lmem)

For the upper half of a partial sum, the instruction that follows this paragraph uses the carry bit generated in the lower 32-bit partial sum. Each partial sum is stored in two memory locations using MOV ACx, dbl(Lmem) or dbl(Lmem) = ACx.

Mnemonic instruction: ADD uns(Smem), CARRY, ACx

Algebraic instruction: ACx = ACx + uns(Smem) + CARRY

Example 5–6. 64-Bit Addition**(a) Mnemonic Instructions**

```

;*****
; 64-Bit Addition      Pointer assignments:
;
;   X3 X2 X1 X0      AR1 -> X3 (even address)
; + Y3 Y2 Y1 Y0      X2
; -----          X1
;   W3 W2 W1 W0      X0
;
;                   AR2 -> Y3 (even address)
;                   Y2
;                   Y1
;                   Y0
;
;                   AR3 -> W3 (even address)
;                   W2
;                   W1
;                   W0
;
;*****

MOV40  dbl(*AR1(#2)), AC0          ; AC0 = X1 X0
ADD    dbl(*AR2(#2)), AC0          ; AC0 = X1 X0 + Y1 Y0
MOV    AC0,dbl(*AR3(#2))           ; Store W1 W0.
MOV40  dbl(*AR1), AC0              ; AC0 = X3 X2
ADD    uns(*AR2(#1)),CARRY,AC0     ; AC0 = X3 X2 + 00 Y2 + CARRY
ADD    *AR2<< #16, AC0             ; AC0 = X3 X2 + Y3 Y2 + CARRY
MOV    AC0, dbl(*AR3)              ; Store W3 W2.

```

Example 5–6. 64-Bit Addition (Continued)**(b) Algebraic Instructions**

```

;*****
; 64-Bit Addition      Pointer assignments:
;
;   X3 X2 X1 X0      AR1 -> X3 (even address)
; + Y3 Y2 Y1 Y0      X2
; -----           X1
;   W3 W2 W1 W0      X0
;
;                   AR2 -> Y3 (even address)
;                   Y2
;                   Y1
;                   Y0
;
;                   AR3 -> W3 (even address)
;                   W2
;                   W1
;                   W0
;
;*****

AC0 = dbl(*AR1(#2))           ; AC0 = X1 X0
AC0 = AC0 + dbl(*AR2(#2))     ; AC0 = X1 X0 + Y1 Y0
dbl(*AR3(#2)) = AC0           ; Store W1 W0.
AC0 = dbl(*AR1)               ; AC0 = X3 X2
AC0 = AC0 + uns(*AR2(#1)) + CARRY ; AC0 = X3 X2 + 00 Y2 + CARRY
AC0 = AC0 + (*AR2<< #16)      ; AC0 = X3 X2 + Y3 Y2 + CARRY
dbl(*AR3) = AC0               ; Store W3 W2.

```

During subtraction, the effect on the CARRY bit is similar to that during addition. CARRY is cleared if a borrow is generated when an accumulator value is subtracted from:

- ☐ Another accumulator
- ☐ A data-memory operand
- ☐ An immediate operand

A borrow can also be generated when two data-memory operands are subtracted or when a data-memory operand is subtracted from a data-memory operand. If a borrow is not generated, the CARRY is set.

The SUB instruction with a 16-bit shift (shown following this paragraph) is an exception because it only resets the carry bit. This allows the D-unit ALU to generate the appropriate carry when subtracting to the lower or upper half of the accumulator causes a borrow.

Mnemonic instruction: SUB Smem << #16, ACx, ACy

Algebraic instruction: $ACy = ACx - (Smem \ll \#16)$

Figure 5–6 shows several 32-bit subtractions and their effects on the carry bit.

Figure 5–6. 32-Bit Subtraction

C	MSB											LSB
X	0	0	0	0	0	0	0	0	0	0	0	ACx
												1
0	F	F	F	F	F	F	F	F	F	F	F	

C	MSB											LSB
X	0	0	7	F	F	F	F	F	F	F	F	ACx
												1
1	0	0	7	F	F	F	F	F	F	F	F	E

C	MSB											LSB
X	F	F	8	0	0	0	0	0	0	0	0	ACx
												1
1	F	F	7	F	F	F	F	F	F	F	F	

C	MSB											LSB
X	F	F	8	0	0	0	0	0	0	0	0	ACx
												1
0	F	F	8	0	0	0	0	0	0	0	0	

$$ACy = ACx - Smem - BORROW$$

C	MSB											LSB
0	0	0	0	0	0	0	0	0	0	0	0	ACx
												0
0	F	F	F	F	F	F	F	F	F	F	F	ACy

C	MSB											LSB
0	F	F	F	F	F	F	F	F	F	F	F	ACx
												0
1	F	F	F	F	F	F	F	F	F	F	F	E

$$ACy = ACx - (Smem \ll 16)$$

C	MSB											LSB
1	F	F	8	0	0	0	F	F	F	F	F	ACx
												0
0	0	0	7	F	F	F	F	F	F	F	F	ACy

C	MSB											LSB
0	F	F	8	0	0	0	F	F	F	F	F	ACx
												0
0	F	F	8	0	0	1	F	F	F	F	F	

Example 5–7 subtracts two 64-bit numbers to obtain a 64-bit result. The partial remainder of the 64-bit subtraction is efficiently performed by the following instructions, which handle 32-bit operands in a single cycle.

Mnemonic instructions: MOV40 dbl(Lmem), ACx
SUB dbl(Lmem), ACx
Algebraic instructions: ACx = dbl(Lmem)
ACx = ACx – dbl(Lmem)

For the upper half of the partial remainder, the instruction that follows this paragraph uses the borrow generated in the lower 32-bit partial remainder. The borrow is not a physical bit in a status register; it is the logical inverse of CARRY. Each partial sum is stored in two memory locations using MOV ACx, dbl(Lmem) or dbl(Lmem) = ACx.

Mnemonic instruction: SUB uns(Smem), BORROW, ACx
Algebraic instruction: ACx = ACx – uns(Smem) – BORROW

Example 5–7. 64-Bit Subtraction**(a) Mnemonic Instructions**

```

;*****
; 64-Bit Subtraction      Pointer assignments:
;
;   X3 X2 X1 X0          AR1 -> X3 (even address)
; - Y3 Y2 Y1 Y0          X2
; -----              X1
;   W3 W2 W1 W0          X0
;                        AR2 -> Y3 (even address)
;                        Y2
;                        Y1
;                        Y0
;                        AR3 -> W3 (even address)
;                        W2
;                        W1
;                        W0
;
;*****

MOV40 db1(*AR1(#2)), AC0          ; AC0 = X1 X0
SUB db1(*AR2(#2)), AC0           ; AC0 = X1 X0 - Y1 Y0
MOV AC0, db1(*AR3(#2))           ; Store W1 W0.
MOV40 db1 (*AR1), AC0            ; AC0 = X3 X2
SUB uns(*AR2(#1)), BORROW, AC0   ; AC0 = X3 X2 - 00 Y2 - BORROW
SUB *AR2 << #16, AC0            ; AC0 = X3 X2 - Y3 Y2 - BORROW
MOV AC0, db1(*AR3)              ; Store W3 W2.

```

Example 5–7. 64-Bit Subtraction (Continued)*(b) Algebraic Instructions*

```

;*****
; 64-Bit Subtraction      Pointer assignments:
;
;   X3 X2 X1 X0          AR1 -> X3 (even address)
; - Y3 Y2 Y1 Y0          X2
; -----              X1
;   W3 W2 W1 W0          X0
;                        AR2 -> Y3 (even address)
;                        Y2
;                        Y1
;                        Y0
;                        AR3 -> W3 (even address)
;                        W2
;                        W1
;                        W0
;
;*****

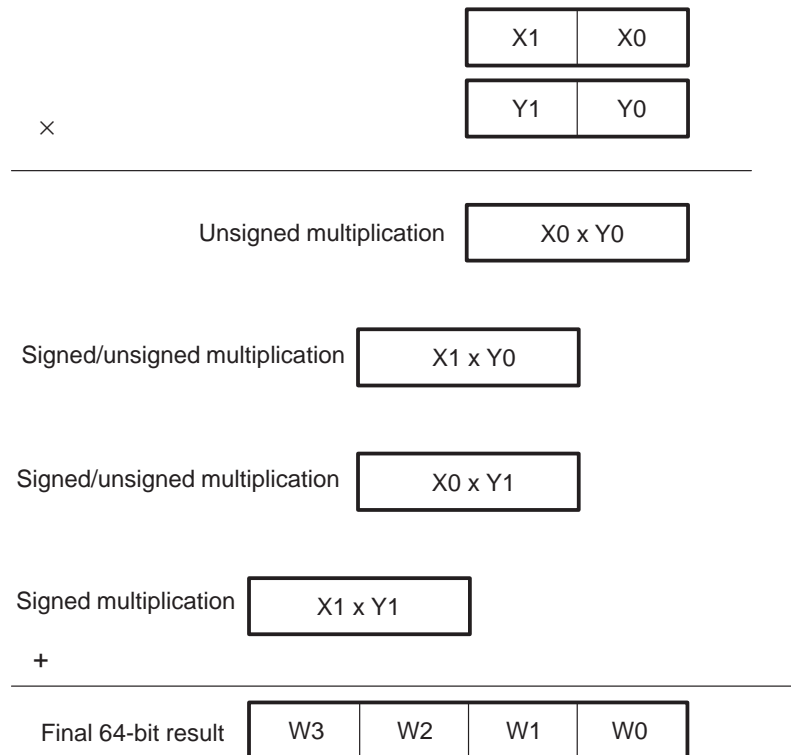
AC0 = dbl(*AR1(#2))          ; AC0 = X1 X0
AC0 = AC0 - dbl(*AR2(#2))    ; AC0 = X1 X0 - Y1 Y0
dbl(*AR3(#2)) = AC0          ; Store W1 W0.
AC0 = dbl (*AR1)             ; AC0 = X3 X2
AC0 = AC0 - uns(*AR2(#1)) - BORROW ; AC0 = X3 X2 - 00 Y2 - BORROW
AC0 = AC0 - (*AR2<< #16)    ; AC0 = X3 X2 - Y3 Y2 - BORROW
dbl(*AR3) = AC0              ; Store W3 W2.

```


5.3 Extended-Precision Multiplication

Extended precision multiplication can be performed using basic C55x instructions. The C55x instruction set provides the user with a very flexible set of 16-bit multiply instructions that accept signed and unsigned operands and with a very efficient set of multiply-and-accumulate instructions that shift the value of the accumulator before adding it to the multiplication result. Figure 5–5 shows how two 32-bit numbers yield a 64-bit product.

Figure 5–7. 32-Bit Multiplication



Example 5–8 shows that a multiplication of two 32-bit integer numbers requires one multiplication, two multiply/accumulate/shift operations, and a multiply/accumulate operation. The product is a 64-bit integer number. Example 5–9 shows a fractional multiplication. The operands are in Q31 format, while the product is in Q31 format.

Example 5–8. 32-Bit Integer Multiplication**(a) Mnemonic Instructions**

```

;*****
; This routine multiplies two 32-bit signed integers, giving a
; 64-bit result. The operands are fetched from data memory and the
; result is written back to data memory.
;
; Data Storage:                      Pointer Assignments:
; X1 X0          32-bit operand      AR0 -> X1
; Y1 Y0          32-bit operand      X0
; W3 W2 W1 W0    64-bit product      AR1 -> Y1
;                                     Y0
; Entry Conditions:                  AR2 -> W0
; SXMD = 1 (sign extension on)       W1
; SATD = 0 (no saturation)           W2
; FRCT = 0 (fractional mode off)     W3
;
; RESTRICTION: The delay chain and input array must be
; long-word aligned.
;*****

AMAR *AR0+                      ; AR0 points to X0
|| AMAR *AR1+                   ; AR1 points to Y0
MPYM uns(*AR0), uns(*AR1), AC0   ; AC0 = X0*Y0
MOV AC0,*AR2+                   ; Save W0
MACM *AR0+, uns(*AR1-), AC0 >> #16, AC0 ; AC0 = X0*Y0>>16 + X1*Y0
MACM uns(*AR0-), *AR1, AC0      ; AC0 = X0*Y0>>16 + X1*Y0 + X0*Y1
MOV AC0, *AR2+                 ; Save W1
MACM *AR0, *AR1, AC0 >> #16, AC0 ; AC0 = AC0>>16 + X1*Y1
MOV AC0, *AR2+                 ; Save W2
MOV HI(AC0), *AR2              ; Save W3

```

*Example 5–8. 32-Bit Integer Multiplication (Continued)**(b) Algebraic Instructions*

```

;*****
; This routine multiplies two 32-bit signed integers, giving a
; 64-bit result. The operands are fetched from data memory and the
; result is written back to data memory.
;
; Data Storage:                      Pointer Assignments:
; X1 X0          32-bit operand      AR0 -> X1
; Y1 Y0          32-bit operand      X0
; W3 W2 W1 W0    64-bit product      AR1 -> Y1
;                                     Y0
; Entry Conditions:                  AR2 -> W0
; SXMD = 1 (sign extension on)      W1
; SATD = 0 (no saturation)          W2
; FRCT = 0 (fractional mode off)    W3
;
; RESTRICTION: The delay chain and input array must be
; long-word aligned.
;*****

mar(*AR0+)                      ; AR0 points to X0
|| mar(*AR1+)                   ; AR1 points to Y0
AC0 = uns(*AR0-)*uns(*AR1)      ; ACO = X0*Y0
*AR2+ = AC0                     ; Save W0
AC0 = (AC0 >> #16) + ((*AR0+)*uns(*AR1-)) ; AC0 = X0*Y0>>16 + X1*Y0
AC0 = AC0 + (uns(*AR0-)* (*AR1)) ; AC0 = X0*Y0>>16 + X1*Y0 + X0*Y1
*AR2+ = AC0                     ; Save W1
AC0 = (AC0 >> #16) + ((*AR0)*(*AR1)) ; AC0 = AC0>>16 + X1*Y1
*AR2+ = AC0                     ; Save W2
*AR2 = HI(AC0)                 ; Save W3

```

Example 5–9. 32-Bit Fractional Multiplication**(a) Mnemonic Instructions**

```

;*****
; This routine multiplies two Q31 signed integers, resulting in a
; Q31 result. The operands are fetched from data memory and the
; result is written back to data memory.
;
; Data Storage:                                Pointer Assignments:
; X1 X0      Q31 operand                        AR0 -> X1
; Y1 Y0      Q31 operand                        X0
; W1 W0      Q31 product                       AR1 -> Y1
;                                                  Y0
; Entry Conditions:                            AR2 -> W1 (even address)
; SXMD = 1 (sign extension on)                  W0
; SATD = 0 (no saturation)
; FRCT = 1 (shift result left by 1 bit)
;
; RESTRICTION: W1 W0 is aligned such that W1 is at an even address.
;*****
AMAR *AR0+                                ; AR0 points to X0
MPYM uns(*AR0), *AR1+, AC0                ; AC0 = X0*Y1
MACM *AR0, uns(*AR1-), AC0                ; AC0 =X0*Y1 + X1*Y0
MACM *AR0, *AR1, AC0 >> #16, AC0          ; AC0 = AC0>>16 + X1*Y1
MOV AC0, dbl(*AR2)                        ; Save W1 W0

```

(b) Algebraic Instructions

```

;*****
; This routine multiplies two Q31 signed integers, resulting in a
; Q31 result. The operands are fetched from data memory and the
; result is written back to data memory.
;
; Data Storage:                                Pointer Assignments:
; X1 X0      Q31 operand                        AR0 -> X1
; Y1 Y0      Q31 operand                        X0
; W1 W0      Q31 product                       AR1 -> Y1
;                                                  Y0
; Entry Conditions:                            AR2 -> W1 (even address)
; SXMD = 1 (sign extension on)                  W0
; SATD = 0 (no saturation)
; FRCT = 1 (shift result left by 1 bit)
;
; RESTRICTION: W1 W0 is aligned such that W1 is at an even address.
;*****
mar(*AR0+)                                ; AR0 points to X0
AC0 = uns(*AR0-)*(*AR1+)                  ; AC0 = X0*Y1
AC0 = AC0 + ((*AR0)* uns(*AR1-))          ; AC0 =X0*Y1 + X1*Y0
AC0 = (AC0 >> #16) + ((*AR0)*(*AR1))      ; AC0 = AC0>>16 + X1*Y1
dbl(*AR2) = AC0                            ; Save W1 W0

```

5.4 Division

Binary division is the inverse of multiplication. Multiplication consists of a series of shift and add operations, while division can be broken into a series of subtract and shift operations. On the C55x DSP you can implement this kind of division by repeating a form of the conditional subtract (SUBC) instruction.

In a fixed-point processor, the range of the numbers we can use is limited by the number of bits and the convention we use to represent these numbers. For example, with a 16-bit unsigned representation, it is not possible to represent a number larger than $2^{16} - 1$ (that is, 65 535). There can be problems with division operations that require computing the inverse of a very small number. Some digital signal processing algorithms may require integer or fractional division operations that support a large range of numbers. This kind of division can be implemented with the conditional subtract (SUBC) instruction.

The difference between integers and fractional numbers is so great in a fixed point architecture that it requires different algorithms to perform the division operation. Section 5.4.1 shows how to implement signed and unsigned integer division. Section 5.4.2 (page 5-30) describes fractional division.

5.4.1 Integer Division

To prepare for a SUBC integer division, place a 16-bit positive dividend in an accumulator. Place a 16-bit positive divisor in memory. When you write the SUBC instruction, make sure that the result will be in the same accumulator that supplies the dividend; this creates a cumulative result in the accumulator when the SUBC instruction is repeated. Repeating the SUBC instruction 16 times produces a 16-bit quotient in the low part of the accumulator (bits 15–0) and a remainder in the high part of the accumulator (bits 31–16). During each execution of the conditional subtract instruction:

- 1) The 16-bit divisor is shifted left by 15 bits and is subtracted from the value in the accumulator.
- 2) If the result of the subtraction is greater than or equal to 0, the result is shifted left by 1 bit, added to 1, and stored in the accumulator. If the result of the subtraction is less than 0, the result is discarded and the value in the accumulator is shifted left by 1 bit.

The following examples show the implementation of the signed/unsigned integer division using the SUBC instruction.

5.4.1.1 Examples of Unsigned Integer Division

Example 5–10 shows how to use the SUBC instruction to implement unsigned division with a 16-bit dividend and a 16-bit divisor.

Example 5–10. Unsigned, 16-Bit By 16-Bit Integer Division

(a) Mnemonic Instructions

```

;*****
; Pointer assignments:
;   AR0 -> Dividend      Divisor ) Dividend
;   AR1 -> Divisor
;   AR2 -> Quotient
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Unsigned division, 16-bit dividend, 16-bit divisor
;   - Sign extension turned off. Dividend & divisor are positive numbers.
;   - After division, quotient in AC0(15-0), remainder in AC0(31-16)
;*****

BCLR SXMD                ; Clear SXMD (sign extension off)
MOV *AR0, AC0             ; Put Dividend into AC0
RPT #(16 - 1)            ; Execute subc 16 times
    SUBC *AR1, AC0, AC0   ; AR1 points to Divisor
MOV AC0, *AR2             ; Store Quotient
MOV HI(AC0), *AR3        ; Store Remainder

```

(b) Algebraic Instructions

```

;*****
; Pointer assignments:
;   AR0 -> Dividend      Divisor ) Dividend
;   AR1 -> Divisor
;   AR2 -> Quotient
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Unsigned division, 16-bit dividend, 16-bit divisor
;   - Sign extension turned off. Dividend & divisor are positive numbers.
;   - After division, quotient in AC0(15-0), remainder in AC0(31-16)
;*****

bit(ST1,#ST1_SXMD) = #0   ; Clear SXMD (sign extension off)
AC0 = *AR0                ; Put Dividend into AC0
repeat( #(16 - 1) )      ; Execute subc 16 times
    subc( *AR1, AC0, AC0 ) ; AR1 points to Divisor
*AR2 = AC0                ; Store Quotient
*AR3 = HI(AC0)            ; Store Remainder

```

Example 5–11 shows how to implement unsigned division with a 32-bit dividend and a 16-bit divisor. The code uses two stages of 16-bit by 16-bit integer division. The first stage takes as inputs the high 16 bits of the 32-bit dividend and the 16-bit divisor. The result in the low half of the accumulator is the high 16 bits of the quotient. The result in the high half of the accumulator is shifted left by 16 bits and added to the lower 16 bits of the dividend. This sum and the 16-bit divisor are the inputs to the second stage of the division. The lower 16 bits of the resulting quotient is the final quotient and the resulting remainder is the final remainder.

Example 5–11. Unsigned, 32-Bit By 16-Bit Integer Division

(a) Mnemonic Instructions

```

;*****
; Pointer assignments:
;   AR0 -> Dividend high half          Divisor ) Dividend
;           Dividend low half
;   ...
;   AR1 -> Divisor
;   ...
;   AR2 -> Quotient high half
;           Quotient low half
;   ...
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Unsigned division, 32-bit dividend, 16-bit divisor
;   - Sign extension turned off. Dividend & divisor are positive numbers.
;   - Before 1st division: Put high half of dividend in AC0
;   - After 1st division:  High half of quotient in AC0(15-0)
;   - Before 2nd division: Put low part of dividend in AC0
;   - After 2nd division:  Low half of quotient in AC0(15-0) and
;                           Remainder in AC0(31-16)
;*****

BCLR SXMD                ; Clear SXMD (sign extension off)
MOV *AR0+, AC0            ; Put high half of Dividend in AC0
|| RPT #(15 - 1)          ; Execute subc 15 times
    SUBC *AR1, AC0, AC0    ; AR1 points to Divisor
SUBC *AR1, AC0, AC0        ; Execute subc final time
|| MOV #8, AR4            ; Load AR4 with AC0_L memory address
MOV AC0, *AR2+            ; Store high half of Quotient
MOV *AR0+, *AR4           ; Put low half of Dividend in AC0_L
RPT #(16 - 1)            ; Execute subc 16 times
    SUBC *AR1, AC0, AC0
MOV AC0, *AR2+            ; Store low half of Quotient
MOV HI(AC0), *AR3        ; Store Remainder
BSET SXMD                ; Set SXMD (sign extension on)

```

Example 5–11. Unsigned, 32-Bit By 16-Bit Integer Division (Continued)**(b) Algebraic Instructions**

```

;*****
; Pointer assignments:
;   AR0 -> Dividend high half      Divisor ) Dividend
;           Dividend low half
;   ...
;   AR1 -> Divisor
;   ...
;   AR2 -> Quotient high half
;           Quotient low half
;   ...
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Unsigned division, 32-bit dividend, 16-bit divisor
;   - Sign extension turned off. Dividend & divisor are positive numbers.
;   - Before 1st division: Put high half of dividend in AC0
;   - After 1st division:  High half of quotient in AC0(15-0)
;   - Before 2nd division: Put low part of dividend in AC0
;   - After 2nd division:  Low half of quotient in AC0(15-0) and
;                           Remainder in AC0(31-16)
;*****

bit(ST1,#ST1_SXMD) = #0      ; Clear SXMD (sign extension off)
AC0 = *AR0+                  ; Put high half of Dividend in AC0
|| repeat( #(15 - 1) )      ; Execute subc 15 times
    subc( *AR1, AC0, AC0)    ; AR1 points to Divisor
subc( *AR1, AC0, AC0)        ; Execute subc final time
|| AR4 = #8                  ; Load AR4 with AC0_L memory address
*AR2+ = AC0                  ; Store high half of Quotient
*AR4 = *AR0+                  ; Put low half of Dividend in AC0_L
repeat( #(16 - 1) )          ; Execute subc 16 times
    subc( *AR1, AC0, AC0)
*AR2+ = AC0                  ; Store low half of Quotient
*AR3 = HI(AC0)               ; Store Remainder
bit(ST1,#ST1_SXMD) = #1      ; Set SXMD (sign extension on)

```


5.4.1.2 Examples of Signed Integer Division

Some applications might require doing division with signed numbers instead of unsigned numbers. The conditional subtract instruction works only with positive integers. The signed integer division algorithm computes the quotient as follows:

- 1) The sign of the quotient is determined and preserved in AC0
- 2) The quotient of the absolute values of the dividend and the divisor is determined using repeated conditional subtract instructions
- 3) The negative of the result is computed if required, according to the sign of AC0

Example 5–12 shows the implementation of division with a signed 16-bit dividend and a 16-bit signed divisor, and Example 5–13 extends this algorithm to handle a 32-bit dividend.

Example 5–12. Signed, 16-Bit By 16-Bit Integer Division**(a) Mnemonic Instructions**

```

;*****
; Pointer assignments:
;   AR0 -> Dividend      Divisor ) Dividend
;   AR1 -> Divisor
;   AR2 -> Quotient
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Signed division, 16-bit dividend, 16-bit divisor
;   - Sign extension turned on. Dividend and divisor can be negative.
;   - Expected quotient sign saved in AC0 before division
;   - After division, quotient in AC1(15-0), remainder in AC1(31-16)
;*****

      BSET SXMD                ; Set SXMD (sign extension on)
      MPYM *AR0, *AR1, AC0     ; Sign of (Dividend x Divisor) should be
                                ;   sign of Quotient
      MOV *AR1, AC1            ; Put Divisor in AC1
      ABS AC1, AC1             ; Find absolute value, |Divisor|
      MOV AC1, *AR2            ; Store |Divisor| temporarily
      MOV *AR0, AC1            ; Put Dividend in AC1
      ABS AC1, AC1             ; Find absolute value, |Dividend|
      RPT #(16 - 1)            ; Execute subc 16 times
      SUBC *AR2, AC1, AC1      ; AR2 -> |Divisor|
      MOV HI(AC1), *AR3        ; Save Remainder
      MOV AC1, *AR2            ; Save Quotient
      SFTS AC1, #16            ; Shift quotient: Put MSB in sign position
      NEG AC1, AC1             ; Negate quotient
      XCCPART label, AC0 < #0  ; If sign of Quotient should be negative,
      MOV HI(AC1), *AR2        ;   replace Quotient with negative version
label:

```

*Example 5–12. Signed, 16-Bit By 16-Bit Integer Division (Continued)**(b) Algebraic Instructions*

```

;*****
; Pointer assignments:
;   AR0 -> Dividend      Divisor ) Dividend
;   AR1 -> Divisor
;   AR2 -> Quotient
;   AR3 -> Remainder
;
; Algorithm notes:
;   - Signed division, 16-bit dividend, 16-bit divisor
;   - Sign extension turned on. Dividend and divisor can be negative.
;   - Expected quotient sign saved in AC0 before division
;   - After division, quotient in AC1(15-0), remainder in AC1(31-16)
;*****

bit(ST1,#ST1_SXMD) = #1      ; Set SXMD (sign extension on)
AC0 = (*AR0) * (*AR1)        ; Sign of (Dividend x Divisor) should be
                               ;   sign of Quotient
AC1 = *AR1                   ; Put Divisor in AC1
*AR2 = AC1                   ; Find absolute value, |Divisor|
AC1 = *AR0                   ; Store |Divisor| temporarily
AC1 = |AC1|                  ; Put Dividend in AC1
repeat( #(16 - 1) )         ; Find absolute value, |Dividend|
    subc( *AR2, AC1, AC1)    ; Execute subc 16 times
*AR3 = HI(AC1)              ; AR2 -> |Divisor|
*AR2 = AC1                  ; Save Remainder
AC1 = AC1 << #16            ; Save Quotient
AC1 = - AC1                 ; Shift quotient: Put MSB in sign position
if(AC0 < #0) execute (D_unit) ; Negate quotient
*AR2 = HI(AC1)              ; If sign of Quotient should be negative,
                               ;   replace Quotient with negative version

```

Example 5–13. Signed, 32-Bit By 16-Bit Integer Division**(a) Mnemonic Instructions**

```

;*****
;  Pointer assignments:  (Dividend and Quotient are long-word aligned)
;    AR0 -> Dividend high half (NumH) (even address)
;           Dividend low half (NumL)
;    AR1 -> Divisor (Den)
;    AR2 -> Quotient high half (QuotH) (even address)
;           Quotient low half (QuotL)
;    AR3 -> Remainder (Rem)
;
;  Algorithm notes:
;    - Signed division, 32-bit dividend, 16-bit divisor
;    - Sign extension turned on. Dividend and divisor can be negative.
;    - Expected quotient sign saved in AC0 before division
;    - Before 1st division: Put high half of dividend in AC1
;    - After 1st division:  High half of quotient in AC1(15-0)
;    - Before 2nd division: Put low part of dividend in AC1
;    - After 2nd division:  Low half of quotient in AC1(15-0) and
;                           Remainder in AC1(31-16)
;*****

BSET SXMD                ; Set SXMD (sign extension on)
MPYM  *AR0, *AR1, AC0     ; Sign( NumH x Den ) is sign of actual result
MOV  *AR1, AC1            ; AC1 = Den
ABS  AC1, AC1             ; AC1 = abs(Den)
MOV  AC1, *AR3            ; Rem = abs(Den) temporarily
MOV40 dbl(*AR0), AC1      ; AC1 = NumH NumL
ABS  AC1, AC1             ; AC1 = abs(Num)
MOV  AC1, dbl(*AR2)       ; QuotH = abs(NumH) temporarily
                               ; QuotL = abs(NumL) temporarily

MOV  *AR2, AC1            ; AC1 = QuotH
RPT  #(15 - 1)            ; Execute subc 15 times
    SUBC *AR3, AC1, AC1
SUBC *AR3, AC1, AC1       ; Execute subc final time
|| MOV #11, AR4           ; Load AR4 with AC1_L memory address
MOV  AC1, *AR2+           ; Save QuotH
MOV  *AR2, *AR4           ; AC1_L = QuotH
RPT  #(16 - 1)            ; Execute subc 16 times
    SUBC *AR3, AC1, AC1
MOV  AC1, *AR2-           ; Save QuotL
MOV  HI(AC1), *AR3        ; Save Rem

BCC skip, AC0 >= #0       ; If actual result should be positive, goto skip.
MOV40 dbl(*AR2), AC1      ; Otherwise, negate Quotient.
NEG  AC1, AC1
MOV  AC1, dbl(*AR2)

skip:
RET

```

Example 5–13. Signed, 32-Bit By 16-Bit Integer Division (Continued)**(b) Algebraic Instructions**

```

;*****
;  Pointer assignments:  (Dividend and Quotient are long-word aligned)
;    AR0 -> Dividend high half (NumH) (even address)
;           Dividend low half (NumL)
;    AR1 -> Divisor (Den)
;    AR2 -> Quotient high half (QuotH) (even address)
;           Quotient low half (QuotL)
;    AR3 -> Remainder (Rem)
;
;  Algorithm notes:
;    - Signed division, 32-bit dividend, 16-bit divisor
;    - Sign extension turned on. Dividend and divisor can be negative.
;    - Expected quotient sign saved in AC0 before division
;    - Before 1st division: Put high half of dividend in AC1
;    - After 1st division:  High half of quotient in AC1(15-0)
;    - Before 2nd division: Put low part of dividend in AC1
;    - After 2nd division:  Low half of quotient in AC1(15-0) and
;                           Remainder in AC1(31-16)
;*****

bit(ST1,#ST1_SXMD) = #1      ; Set SXMD (sign extension on)
AC0 = (*AR0)* (*AR1)          ; Sign( NumH x Den ) is sign of actual result
AC1 = *AR1                    ; AC1 = Den
AC1 = |AC1|                   ; AC1 = abs(Den)
*AR3 = AC1                    ; Rem = abs(Den) temporarily
AC1 = dbl(*AR0)               ; AC1 = NumH NumL
AC1 = |AC1|                   ; AC1 = abs(Num)
dbl(*AR2) = AC1               ; QuotH = abs(NumH) temporarily
                               ; QuotL = abs(NumL) temporarily

AC1 = *AR2                    ; AC1 = QuotH
repeat( #(15 - 1) )          ; Execute subc 15 times
    subc( *AR3, AC1, AC1)
subc( *AR3, AC1, AC1)         ; Execute subc final time
||AR4 = #11                  ; Load AR4 with AC1_L memory address
*AR2+ = AC1                   ; Save QuotH
*AR4 = *AR2                   ; AC1_L = QuotH
repeat( #(16 - 1) )          ; Execute subc 16 times
    subc( *AR3, AC1, AC1)
*AR2- = AC1                   ; Save QuotL
*AR3 = HI(AC1)                ; Save Rem

if (AC0 >= #0) goto skip      ; If actual result should be positive, goto skip.
AC1 = dbl(*AR2)               ; Otherwise, negate Quotient.
AC1 = - AC1
dbl(*AR2) = AC1

skip:
return

```

5.4.2 Fractional Division

The algorithms that implement fractional division compute first an approximation of the inverse of the divisor (denominator) using different techniques such as Taylor Series expansion, line of best fit, and successive approximation. The result is then multiplied by the dividend (numerator). The C55x DSP function library (see Chapter 8) implements this function under the name `ldiv16`.

To calculate the value of Y_m this algorithm uses the successive approximation method. The approximations are performed using the following equation:

$$Y_{m+1} = 2 Y_m - Y_m^2 X_{norm}$$

If we start with an initial estimate of Y_m , then the equation will converge to a solution very rapidly (typically in three iterations for 16-bit resolution). The initial estimate can either be obtained from a look-up table, from choosing a midpoint, or simply from linear interpolation. The `ldiv16` algorithm uses linear interpolation. This is accomplished by taking the complement of the least significant bits of the X_{norm} value.

5.5 Methods of Handling Overflows

An overflow occurs when the result of an arithmetical operation is larger than the largest number that can be represented in the register that must hold the result. Due to the 16-bit format, fixed-point DSPs provide a limited dynamic range. You must manage the dynamic range of your application to avoid possible overflows. The overflow depends on the nature of the input signal and of the algorithm in question.

5.5.1 Hardware features for overflow handling

The C55x DSP offers several hardware features for overflow handling:

☐ Guard bits:

Each of the C55x accumulators (AC0, AC1, AC2, and AC3) has eight guard bits (bits 39–32), which allow up to 256 consecutive multiply-and-accumulate operations before an accumulator overflow.

☐ Overflow flags:

Each C55x accumulator has an associated overflow flag (see the following table). When an operation on an accumulator results in an overflow, the corresponding overflow flag is set.

☐ Saturation mode bits:

The DSP has two saturation mode bits: SATD for operations in the D unit of the CPU and SATA for operations in the A unit of the CPU. When the SATD bit is set and an overflow occurs in the D unit, the CPU saturates the result. Regardless of the value of SATD, the appropriate accumulator overflow flag is set. Although no flags track overflows in the A unit, overflowing results in the A unit are saturated when the SATA bit is set.

Saturation replaces the overflowing result with the nearest range boundary. Consider a 16-bit register which has range boundaries of 8000h (largest negative number) and 7FFFh (largest positive number). If an operation generates a result greater than 7FFFh, saturation can replace the result with 7FFFh. If a result is less than 8000h, saturation can replace the result with 8000h.

5.5.1.1 Overview of overflow handling techniques

There are a number of general methodologies to handle overflows. Among the methodologies are saturation, input scaling, fixed scaling, and dynamic scaling. We will give an overview of these methodologies and will see some examples illustrating their application.

☐ Saturation:

One possible way to handle overflows is to use the hardware saturation modes mentioned in section 5.5.1. However, saturation has the effect of clipping the output signal, potentially causing data distortion and non-linear behavior in the system.

☐ Input scaling:

You can analyze the system that you want to implement and scale the input signal, assuming worst conditions, to avoid overflow. However, this approach can greatly reduce the precision of the output.

☐ Fixed scaling:

You can scale the intermediate results, assuming worst conditions. This method prevents overflow but also increases the system's signal-to-noise ratio.

☐ Dynamic scaling:

The intermediate results can be scaled only when needed. You can accomplish this by monitoring the range of the intermediate results. This method prevents overflow but increases the computational requirements.

The next sections demonstrate these methodologies applied to FIR (finite impulse response) filters, IIR (infinite impulse response) filters and FFTs (fast Fourier transforms).

5.5.1.2 Scaling methods for FIR filters

The best way to handle overflow problems in FIR (finite impulse response) filters is to design the filters with a gain less than 1 to avoid having to scale the input data. This method, combined with the guard bits available in each of the accumulators, provides a robust way to handle overflows in the filters.

Fixed scaling and input scaling are not used due to their negative impact on signal resolution (basically one bit per multiply-and-accumulate operation). Dynamic scaling can be used for an FIR filter if the resulting increase in cycles is not a concern. Saturation is also a common option for certain types of audio signals.

5.5.1.3 Scaling methods for IIR filters

Fixed-point realization of an IIR (infinite impulse response) filter in cascaded second-order stages is recommended to minimize the frequency response sensitivity of high-order filters. In addition to round-off error due to filter coefficient quantization, overflow avoidance is critical due to the recursive nature of the IIR filter.

Overflow between stages can be avoided by maintaining the intermediate value in the processor accumulator. However, overflow can happen at the internal filter state (delay buffer) inside each stage. To prevent overflow at stage k , the filter unit-pulse response $f(n)$ must be scaled (feed forward path) by a gain factor G_k given by

$$\text{Option 1 : } G_k = \sum_n \text{abs}(f(n))$$

or

$$\text{Option 2 : } G_k = \sum_n \left(\text{abs}(f(n))^2 \right)^{1/2}$$

Option 1 prevents overflows, but at the expense of precision. Option 2 allows occasional overflows but offers an improved precision. In general, these techniques work well if the input signal does not have a large dynamic range.

Another method to handle overflow in IIR filters is to use dynamic scaling. In this approach, the internal filter states are scaled down by half only if an overflow is detected at each stage. The result is a higher precision but at the expense of increased MIPS.

5.5.1.4 Scaling methods for FFTs

In FFT (Fast Fourier Transform) operations, data will grow an average of one bit on the output of each butterfly. Input scaling will require shifting the data input by $\log n$ (n = size of FFT) that will cause a $6(\log n)$ dB loss even before computing the FFT. In fixed scaling, the output of the butterfly will be scaled by 2 at each stage. This is probably the most common scaling approach for FFTs because it is simple and has a better sound-to-noise ratio (SNR). However, for larger FFTs this scaling may cause information loss.

Another option is to implement a dynamic scaling approach in which scaling by 2 at each stage occurs only when bit growth occurs. In this case, an exponent is assigned to the entire stage block (block floating-point method). When scaling by 2 happens, the exponent is incremented by 1. At the end of the FFT, the data are scaled up by the resulting exponent. Dynamic scaling provides the best SNR but increases the FFT cycle count because you have to detect bit growth and update the exponent accordingly.

Bit-Reversed Addressing

This chapter introduces the concept and the syntax of bit-reverse addressing. It then explains how bit-reverse addressing can help to speed up a Fast Fourier Transform (FFT) algorithm. To find code that performs complex and real FFTs (forward and reverse) and bit-reversing of FFT vectors, see Chapter 8, *TIC55x DSPLIB*.

Topic	Page
6.1 Introduction to Bit-Reverse Addressing	6-2
6.2 Using Bit-Reverse Addressing in FFT Algorithms	6-4
6.3 In-Place Versus Off-Place Bit-Reversing	6-6
6.4 Using the C55x DSPLIB for FFTs and Bit-Reversing	6-8

6.1 Introduction to Bit-Reverse Addressing

Bit-reverse addressing is a special type of indirect addressing. It uses one of the auxiliary registers (AR0–AR7) as a base pointer of an array and uses temporary register 0 (T0) as an index register. When you add T0 to the auxiliary register using bit-reverse addressing, the address is generated in a bit-reversed fashion, with the carry propagating from left to right instead of from right to left.

Table 6–1 shows the syntaxes for each of the two bit-reversed addressing modes supported by the TMS320C55x™ (C55x™) DSP.

Table 6–1. Syntaxes for Bit-Reverse Addressing Modes

Operand Syntax	Function	Description
*(ARx–T0B)	address = ARx ARx = (ARx – T0)	After access, T0 is subtracted from ARx with reverse carry (rc) propagation.
*(ARx+T0B)	address = ARx ARx = (ARx + T0)	After access, T0 is added to ARx with reverse carry (rc) propagation.

Assume that the auxiliary registers are 8 bits long, that AR2 represents the base address of the data in memory (01100000b), and that T0 contains the value 00001000b (decimal 8). Example 6–1 shows a sequence of modifications of AR2 and the resulting values of AR2.

Table 6–2 shows the relationship between a standard bit pattern that is repeatedly incremented by 1 and a bit-reversed pattern that is repeatedly incremented by 1000b with reverse carry propagation. Compare the bit-reversed pattern to the 4 LSBs of AR2 in Example 6–1.

Example 6–1. Sequence of Auxiliary Registers Modifications in Bit-Reversed Addressing

*(AR2+T0B)	;AR2	=	0110 0000	(0th value)
*(AR2+T0B)	,AR2	=	0110 1000	(1st value)
*(AR2+T0B)	;AR2	=	0110 0100	(2nd value)
*(AR2+T0B)	;AR2	=	0110 1100	(3rd value)
*(AR2+T0B)	;AR2	=	0110 0010	(4th value)
*(AR2+T0B)	;AR2	=	0110 1010	(5th value)
*(AR2+T0B)	;AR2	=	0110 0110	(6th value)
*(AR2+T0B)	;AR2	=	0110 1110	(7th value)


Table 6–2. Bit-Reversed Addresses

Step	Bit Pattern	Bit-Reversed Pattern	Bit-Reversed Step
0	0000	0000	0
1	0001	1000	8
2	0010	0100	4
3	0011	1100	12
4	0100	0010	2
5	0101	1010	10
6	0110	0110	6
7	0111	1110	14
8	1000	0001	1
9	1001	1001	9
10	1010	0101	5
11	1011	1101	13
12	1100	0011	3
13	1101	1011	11
14	1110	0111	7
15	1111	1111	15

6.2 Using Bit-Reverse Addressing In FFT Algorithms

Bit-reversed addressing enhances execution speed for Fast-Fourier Transform (FFT) algorithms. Typical FFT algorithms either take an in-order vector input and produce a bit-reversed vector output or take a bit-reversed vector input and produce an in-order vector output. In either case, bit-reverse addressing can be used to resequence the vectors. Figure 6–1 shows a flow graph of an 8-point decimation-in-frequency FFT algorithm with a bit-reversed input and an in-order output.

Figure 6–1. FFT Flow Graph Showing Bit-Reversed Input and In-Order Output



Consider a complex FFT of size N (that is, an FFT with an input vector that contains N complex numbers). You can bit-reverse either the input or the output vectors by executing the following steps:

- 1) Write 0 to the ARMS bit of status register 2 to select the DSP mode for AR indirect addressing. (Bit-reverse addressing is not available in the control mode of AR indirect addressing.) Then use the `.arms_off` directive to notify the assembler of this selection.
- 2) Use Table 6–3 to determine how the base pointer of the input array must be aligned to match the given vector format. Then load an auxiliary register with the proper base address.
- 3) Consult Table 6–3 to properly load the index register, T0.
- 4) Ensure that the entire array fits within a 64K boundary (the largest possible array addressable by the 16-bit auxiliary register).

As an example of how to use Table 6–3, suppose you need to bit-reverse a vector with $N = 64$ complex elements in the Re–Im–Re–Im format. The $(n + 1)$ least significant bits (LSBs) of the base address must be 0s, where

$$(n + 1) = (\log_2 N + 1) = (\log_2 64 + 1) = (6 + 1) = 7 \text{ LSBs}$$

Therefore, AR0 must be loaded with a base address of this form (Xs are don't cares):

$$\text{AR0} = \text{XXXX XXXX X000 0000b}$$

The index loaded into T0 is equal to the number of elements:

$$\text{T0} = 2^n = 2^6 = 64$$

Table 6–3. Typical Bit-Reverse Initialization Requirements

Vector Format in Memory	T0 Initialization Value	Alignment of Vector Base Address
Re–Im–Re–Im. The real and imaginary parts of each complex element are stored at consecutive memory locations. For example: Real Imaginary ... Real Imaginary	2^n where $n = \log_2 N$	$(n+1)$ LSBs must be 0s, where $n = \log_2 N$
Re–Re...Im–Im. The real and imaginary data are stored in separate arrays. For example: Real Real ... Imaginary Imaginary	$2^{(n-1)}$ where $n = \log_2 N$	n LSBs must be 0s, where $n = \log_2 N$

6.3 In-Place Versus Off-Place Bit-Reversing

You can have a bit-reversed array write over the original array (in-place bit-reversing) or you can place the bit-reversed array in a separate place in memory (off-place bit-reversing). Example 6–2 shows assembly language code for an off-place bit-reverse operation. An input array of N complex elements pointed to by AR0 is bit-reversed into an output array pointed to by AR1. The vector format for input and output vectors is assumed to be Re–Im–Re–Im. Each element (Re–Im) of the input array is loaded into AC0. Then the element is transferred to the output array using bit-reverse addressing. Each time the index in T0 is added to the address in AR1, the addition is done with carries propagating from left to right, instead of from right to left.

Although it requires twice the memory, off-place bit-reversing is faster than in-place bit-reversing. Off-place bit-reversing requires 2 cycles per complex data point, while in-place bit-reversing requires approximately 4 cycles per complex data point.

The code shown in Example 6–2 (b) is from the cbrev() C-callable optimized assembly function available in the C55x DSPLIB (see Chapter 8). The DSPLIB cbrev() function supports both in-place and off-place bit-reversing and DSPLIB is provided in source format to show how efficient in-place bit-reversing can also be achieved.

Example 6–2. Off-Place Bit Reversing of a Vector Array (in Assembly)

(a) Mnemonic Instructions

```

; ...
BCLR ARMS                      ; reset ARMS bit to allow bit-reverse addressing
.arms_off                      ; notify the assembler of ARMS bit = 0
; ...
off_place:
    RPTBLOCAL{
        MOV db1(*AR0+), AC0      ; AR0 points to input array
        MOV AC0, db1(*(AR1+T0B)) ; AR1 points to output array
                                ; T0 = NX = number of complex elements in
                                ; array pointed to by AR0
    }

```

(b) Algebraic Instructions

```
;...
bit (ST2, #ST2_ARMS) = #0 ; reset ARMS bit to allow bit-reverse addressing
.arms_off                ; notify the assembler of ARMS bit = 0
;...
off_place:
localrepeat{
    AC0 = dbl(*AR0+)      ; AR0 points to input array
    dbl(*(AR1+T0B)) = AC0 ; AR1 points to output array
                          ; T0 = NX = number of complex elements in
                          ; array pointed to by AR0
}
```

Note: This example shows portions of the file cbrev.asm in the TI C55x DSPLIB (introduced in Chapter 8)

6.4 Using the C55x DSPLIB for FFTs and Bit-Reversing

The C55x DSP function library (DSPLIB) offers C-callable DSP assembly-optimized routines. Among these are a bit-reversing routine (`cbrev()`) and complex and real FFT routines.

Example 6–3 shows how you can invoke the `cbrev()` DSPLIB function from C to do in-place bit-reversing. The function bit-reverses the position of in-order elements in a complex vector `x` and then computes a complex FFT of the bit-reversed vector. The function uses in-place bit-reversing. See Chapter 8 for an introduction to the C55x DSPLIB.

Example 6–3. Using DSPLIB `cbrev()` Routine to Bit Reverse a Vector Array (in C)

```
#define NX 64
short x[2*NX]      ;
short scale = 1    ;

void main(void)
{
    ;...
    cbrev(x,x,NX)    // in-place bit-reversing on input data (Re-Im format)
    cfft(x,NX,scale) // 64-point complex FFT on bit-reversed input data with
                    // scaling by 2 at each stage enabled
    ;...
}
```

Note: This example shows portions of the file `cfft_t.c` in the TI C55x DSPLIB (introduced in Chapter 8)

Application-Specific Instructions

This chapter presents examples of efficient implementations of some common signal processing and telecommunications functions. These examples illustrate the use of some application-specific instructions on the TMS320C55x™ (C55x™) DSP. (Most of the examples in this chapter use instructions from the algebraic instruction set, but the concepts apply equally for the mnemonic instruction set.)

Topic	Page
7.1 Symmetric and Asymmetric FIR Filtering (FIRS, FIRSN)	7-2
7.2 Adaptive Filtering (LMS)	7-6
7.3 Convolutional Encoding (BFXPA, BFXTR)	7-10
7.4 Viterbi Algorithm for Channel Decoding (ADDSUB, SUBADD, MAXDIFF)	7-16

7.1 Symmetric and Asymmetric FIR Filtering (FIRS, FIRSN)

FIR (finite impulse response) filters are often used in telecommunications applications because they are unconditionally stable and they may be designed to preserve important phase information in the processed signal. A *linear phase* FIR provides a phase shift that varies in proportion to the input frequency and requires that the impulse response be **symmetric**: $h(n) = h(N-n)$.

Another class of FIR filter is the **antisymmetric** FIR: $h(n) = -h(N-n)$. A common example is the Hilbert transformer, which shifts positive frequencies by +90 degrees and negative frequencies by -90 degrees. Hilbert transformers may be used in applications, such as modems, in which it is desired to cancel lower sidebands of modulated signals.

Figure 7–1 gives examples of symmetric and antisymmetric filters, each with eight coefficients (a_0 through a_7). Both symmetric and antisymmetric filters may be of even or odd length. However, even-length symmetric filters lend themselves to computational shortcuts which will be described in this section. It is sometimes possible to reformulate an odd-length filter as a filter with one more tap, to take advantage of these constructs.

Because (anti)symmetric filters have only $N/2$ distinct coefficients, they may be folded and performed with $N/2$ additions (subtractions) and $N/2$ multiply-and-accumulate operations. Folding means that pairs of elements in the delay buffer which correspond to the same coefficient are pre-added(subtracted) prior to multiplying and accumulating.

The C55x DSP offers two different ways to implement symmetric and asymmetric filters. This section shows how to implement these filters using specific instructions, FIRS and FIRSN. To see how to implement symmetric and asymmetric filters using the dual-MAC hardware, see section 4.1.4, *Implicit Algorithm Symmetry*, which begins on page 4-5. The *firs*/*firsn* implementation and the dual-MAC implementation are equivalent from a throughput standpoint.

Figure 7–1. Symmetric and Antisymmetric FIR Filters



Symmetric FIR Filter $a[n]$:

$$a[n] = a[N - n] \text{ where } 0 < n < \frac{N}{2}$$

Symmetric FIR filter output:

$$Y(n) = a[0](x[n - 7] + x[0]) + a[1](x[n - 6] + x[1]) + a[2](x[n - 5] + x[2]) + a[3](x[n - 4] + x[3])$$

Antisymmetric FIR filter output:

$$Y(n) = a[0](x[n - 7] - x[0]) + a[1](x[n - 6] - x[1]) + a[2](x[n - 5] - x[2]) + a[3](x[n - 4] - x[3])$$

Definitions:

a = Filter coefficient

Y = Filter output

n = Sample index

x = Filter input data value

N = Number of filter taps

7.1.1 Symmetric FIR Filtering With the *firs* Instruction

The C55x instruction for symmetric FIR filtering is:

`firs(Xmem,Ymem,Cmem,ACx,ACy)`

This instruction performs two parallel operations: a multiply-and-accumulate (MAC) operation, and an addition. The `firs()` instruction performs the following parallel operations:

$ACy = ACy + (ACx * Cmem),$

$ACx = (Xmem \ll \#16) + (Ymem \ll \#16)$

The first operation performs a multiplication and an accumulation in a MAC unit of the CPU. The input operands of the multiplier are the content of $ACx(32-16)$ and a data memory operand, which is addressed using the coefficient addressing mode and is sign extended to 17 bits. Table 7–1 explains the operands necessary for the operation.

Table 7–1. Operands to the *firs* or *firsn* instruction

Operand(s)	Description
Xmem and Ymem	One of these operands points to the newest value in the delay buffer. The other points to the oldest value in the delay buffer.
Cmem	This operand points to the filter coefficient.
ACx	ACx is one of the four accumulators (AC0–AC3). It holds the sum of the two delayed input values referenced by Xmem and Ymem.
ACy	ACy is one of the four accumulators (AC0–AC3) but is not the same accumulator as ACx. ACy holds the output of each filter tap. After all the filter taps have been performed, ACy holds the final result.

7.1.2 Antisymmetric FIR Filtering With the *firsn* Instruction

The antisymmetric FIR is the same as the symmetric FIR except that the pre-addition of sample pairs is replaced with a pre-subtraction. The C55x instruction for antisymmetric FIR filtering is:

```
firsn(Xmem,Ymem,Cmem,ACx,ACy)
```

This instruction performs two parallel operations: a multiply-and-accumulate (MAC) operation, and a subtraction. The *firsn*() instruction performs the following parallel operations:

$$\begin{aligned} \text{ACy} &= \text{ACy} + (\text{ACx} * \text{Cmem}), \\ \text{ACx} &= (\text{Xmem} \ll \#16) - (\text{Ymem} \ll \#16) \end{aligned}$$

The first operation performs a multiplication and an accumulation in a MAC unit of the CPU. The input operands of the multiplier are the content of ACx(32–16) and a data memory operand, which is addressed using the coefficient addressing mode and is sign extended to 17 bits. Table 7–1 (page 7-4) explains the operands necessary for the operation.

7.1.3 Implementation of a Symmetric FIR Filter on the TMS320C55x DSP

The C55x DSPLIB features an efficient implementation of the Symmetric FIR on the C55x device. Example 7–1 presents the kernel of that implementation to illustrate the usage of the *firs* instruction.

Example 7–1. Symmetric FIR Filter

```

;
; Start of outer loop
;-----
    localrepeat {           ; Start the outer loop

; Get next input value

    *db_ptr1 = *x_ptr+      ; x_ptr: pointer to input data buffer
                          ; db_ptr1: pointer to newest input value

; Clear AC0 and pre-load AC1 with the sum of the 1st and last inputs
    ||AC0 = #0;

; 1st and last inputs
    AC1 = (*db_ptr1+ << #16) + (*db_ptr2- << #16)
; Inner loop
    ||repeat(inner_cnt)
    firs(*db_ptr1+, *db_ptr2-, *h_ptr+, AC1, AC0)

; 2nd to last iteration has different pointer adjustment
    firs(*(db_ptr1-T0), *(db_ptr2+T1), coef(*h_ptr+), AC1, AC0)

; Last iteration is a MAC with rounding
    AC0 = rnd(AC0 + (*h_ptr+ * AC1))

; Store result to memory
    *r_ptr+ = HI(AC0)      ;store Q15 value to memory

    }                     ;end of outer loop

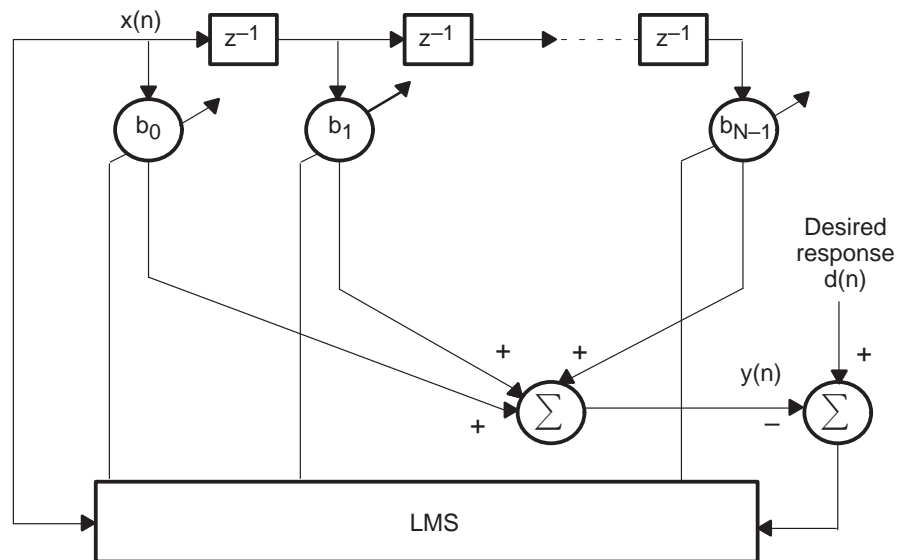
```

Note: This example shows portions of the file `firs.asm` in the TI C55x DSPLIB (introduced in Chapter 8)

7.2 Adaptive Filtering (LMS)

Some applications for adaptive FIR (finite impulse response) and IIR (infinite impulse response) filtering include echo and acoustic noise cancellation. In these applications, an adaptive filter tracks changing conditions in the environment. Although in theory, both FIR and IIR structures can be used as adaptive filters, stability problems and the local optimum points of IIR filters makes them less attractive for this use. Therefore, FIR filters are typically used for practical adaptive filter applications. The least mean square (LMS), local block-repeat, and parallel instructions on the C55x DSP can be used to efficiently implement adaptive filters. The block diagram of an adaptive FIR filter is shown in Figure 7–2.

Figure 7–2. Adaptive FIR Filter Implemented With the Least-Mean-Squares (LMS) Algorithm



Two common algorithms employed for least mean squares adaptation are the non-delayed LMS and the delayed LMS algorithm. When compared to non-delayed LMS, the more widely used delayed LMS algorithm has the advantage of greater computational efficiency at the expense of slightly relaxed convergence properties. Therefore, section 7.2.1 describes only the delayed LMS algorithm.

7.2.1 Delayed LMS Algorithm

In the delayed LMS, the convolution is performed to compute the output of the adaptive filter:

$$y(n) = \sum_{k=0}^{N-1} b_k x(n - k)$$

where

y = Filter output

n = Sample index

k = Delay index

N = Number of filter taps

b_k = Adaptive coefficient

x = Filter input data value

The value of the error is computed and stored to be used in the next invocation:

$$e(n) = d(n) - y(n)$$

where

e = Error

d = Desired response

y = Actual response (filter output)

The coefficients are updated based on an error value computed in the previous invocation of the algorithm (β is the conversion constant):

$$b_k(n + 1) = b_k(n) + 2\beta e(n - 1)x(n - k - 1)$$

The delayed LMS algorithm can be implemented with the LMS instruction—`lms(Xmem, Ymem, ACx, ACy)`—which performs a multiply-and-accumulate (MAC) operation and, in parallel, an addition with rounding:

$ACy = ACy + (Xmem * Ymem),$

$ACx = rnd(ACx + (Xmem << \#16))$

The input operands of the multiplier are the content of data memory operand Xmem, sign extended to 17 bits, and the content of data memory operand Ymem, sign extended to 17 bits. One possible implementation would assign the following roles to the operands of the LMS instruction:

Operand(s)	Description
Xmem	This operand points to the coefficient array.
Ymem	This operand points to the data array.
ACx	ACx is one of the four accumulators (AC0–AC3). ACx is used to update the coefficients.
ACy	ACy is one of the four accumulators (AC0–AC3) but is not the same accumulator as ACx. ACy holds the output of the FIR filter.

An efficient implementation of the delayed LMS algorithm is available in the C55x DSP function library (see Chapter 8). Example 7–2 shows the kernel of this implementation.

Example 7–2. Delayed LMS Implementation of an Adaptive Filter

```

; ar_data: index in the delay buffer
; ar_input: pointer to input vector
; ar_coef: pointer to coefficient vector

StartSample:

; Clear AC0 for initial error term
AC1 = #0
|| localrepeat {
    *ar_data+ = *ar_input+      ;copy input -> state(0)

; Place error term in T3
    T3 = HI(AC1)

;place first update term in AC0
;...while clearing FIR value
    AC0 = T3 * *ar_data+
    || AC1 = #0

;AC0 = update coef
;AC1 = start of FIR output
    LMS(*ar_coef, *ar_data, AC0, AC1)
    || localrepeat {
        *ar_coef+ = HI(AC0)
        || AC0 = T3 * *ar_data+

;AC0 = update coef
;AC1 = update of FIR output
        LMS(*ar_coef, *ar_data, AC0, AC1)
    }

; Store Calculated Output
    *ar_coef+ = HI(AC0)
    || *ar_output+ = HI(rnd(AC1))

; AC2 is error amount
; Point to oldest data sample
    AC2 = (*ar_des+ << #16) - AC1
    || mar(*ar_data+)

; Place updated mu_error term in AC1
    AC1 = rnd(T_step*AC2)
}

```

Note: This example shows portions of the file `dlms.asm` in the TI C55x DSPLIB (introduced in Chapter 8)

7.3 Convolutional Encoding (BFXPA, BFXTR)

The goal in every telecommunication system is to achieve maximum data transfer, using a minimum bandwidth, while maintaining an acceptable quality of transmission. Convolutional codes are a forward error control (FEC) technique in which extra binary digits are added to the original information binary digits prior to transmission to create a code structure which is resistant to errors that may occur within the channel. A decoder at the receiver exploits the code structure to correct any errors that may have occurred. The redundant bits are formed by XORing the current bit with time-delayed bits within the past K input sample history. This is effectively a 1-bit convolution sum; hence the term convolutional encoder. The coefficients of the convolution sum are described using polynomial notation. A convolutional code is defined by the following parameters:

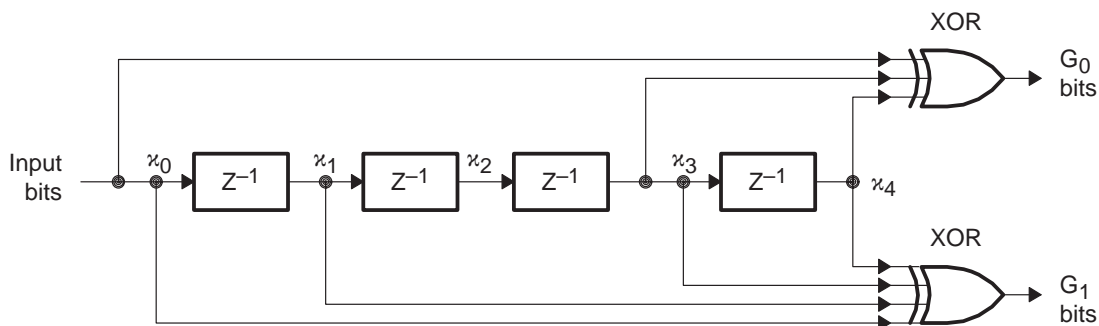
n = Number of function generators

G_0, G_1, \dots, G_n = Polynomials that define the convolutions of bit streams


K = Constraint length (number of delays plus 1)

The rate of the convolutional encoder is defined as $R = 1/n$. Figure 7–3 gives an example of a convolutional encoder with $K=5$ and $R = 1/2$.

Figure 7–3. Example of a Convolutional Encoder



The C55x DSP creates the output streams (G0 and G1) by XORing the shifted input stream (see Figure 7–4).

Figure 7–4. Generation of an Output Stream G_0 

Example 7–3 shows an implementation of the output bit streams for the convolutional encoder of Figure 7–3.

Example 7–3. Generation of Output Streams G_0 and G_1

```

AR3 = #09H                                ; AC0_H
AR1 = #in_bit_stream

; Load 32 bits into the accumulators
AC0 = *AR1+
AC0 = AC0 + (*AR1 << #16)
AC1 = AC0

; Generate G0
AC0 = AC0 ^ (AC1 <<< #-1)                ; A = A XOR B>>3
AC0 = AC0 ^ (AC1 <<< #-3)                ; A = A XOR B>>4
T0 = AC0                                ; Save G0

; Generate G1
AC0 = AC0 ^ (AC1 <<< #-1)                ; A = A XOR B>>1
AC0 = AC0 ^ (AC1 <<< #-3)                ; A = A XOR B>>3
AC0 = AC0 ^ (AC1 <<< #-4)                ; A = A XOR B>>4 ----> AC0_L = G1

*AR3 = T0                                ; AC0_H = G0 -----> AC0 = G0G1

```

7.3.1 Bit-Stream Multiplexing and Demultiplexing

After a bit stream is convolved by the various generating polynomials, the redundant bits are typically multiplexed back into a single higher-rate bit stream. The C55x DSP has dedicated instructions that allow the extraction or insertion of a group of bits anywhere within a 32-bit accumulator. These instructions can be used to greatly expedite the bit-stream multiplexing operation on convolutional encoders.

Figure 7–5 illustrates the concept of bit multiplexing.

Figure 7–5. Bit Stream Multiplexing Concept



The C55x DSP has a dedicated instruction to perform the multiplexing of the bit streams:

`dst = field_expand(ACx,k16)`

This instruction executes in 1 cycle according to the following algorithm:

- 1) Clear the destination register.
- 2) Reset to 0 the bit index pointing within the destination register:
`index_in_dst`.
- 3) Reset to 0 the bit index pointing within the source accumulator:
`index_in_ACx`.
- 4) Scan the bit field mask k16 from bit 0 to bit 15, testing each bit. For each tested mask bit:
 - If the tested bit is 1:
 - a) Copy the bit pointed to by `index_in_ACx` to the bit pointed to by `index_in_dst`.
 - b) Increment `index_in_ACx`.
 - c) Increment `index_in_dst`, and test the next mask bit.

If the tested bit is 0:

Increment `index_in_dst`, and test the next mask bit.

Example 7–4 demonstrates the use of the `field_expand()` instruction.

Example 7–4. Multiplexing Two Bit Streams With the Field Expand Instruction

```

.asg  AC0, G0                ; Assign G0 to register AC0.
.asg  AC1, G1                ; Assign G1 to register AC1.
.asg  AC2, Temp              ; Assign Temp to register AC2.
.asg  AC3, G0G1              ; Assign G0G1 to register AC3.

G0 = *get_G0                 ; Load G0 stream.
G1 = *get_G1                 ; Load G1 stream.
G0G1 = field_expand(G0, #5555h) ; Expand G0 stream.
Temp = G0G1                  ; Temporarily store expanded G0 stream.
G0G1 = field_expand(G1, #AAAAh) ; Expand G1 stream
G0G1 = G0G1 | Temp            ; Interleave expanded streams

```

G0G1 = field_expand(G0,#5555h)

5555h

0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

G0(15–0)

X	X	X	X	X	X	X	X	1	1	0	0	1	1	1	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

0	1	0	1	0	0	0	0	0	1	0	0	0	1	0	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

G0G1 = field_expand(G1,#AAAAh)

AAAAh

1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

G1(15–0)

X	X	X	X	X	X	X	X	0	1	1	1	0	0	1	0
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

0	0	1	0	1	0	1	0	0	0	0	0	1	0	0	0
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

G0G1 = G0G1 | Temp

0	1	1	1	1	0	1	0	0	1	0	0	1	1	0	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

At the receiver G0 and G1 must be extracted by de-interleaving the G0 G1 stream. The DSP has dedicated instructions to perform the de-interleaving of the bit streams:

```
dst = field_extract(ACx,k16)
```

This instruction executes in 1 cycle according to the following algorithm:

- 1) Clear the destination register.
- 2) Reset to 0 the bit index pointing within the destination register:
index_in_dst.
- 3) Reset to 0 the bit index pointing within the source accumulator:
index_in_ACx.
- 4) Scan the bit field mask k16 from bit 0 to bit 15, testing each bit. For each tested mask bit:

If the tested bit is 1:

- a) Copy the bit pointed to by index_in_ACx to the bit pointed to by index_in_dst.
- b) Increment index_in_dst.
- c) Increment index_in_ACx, and test the next mask bit.

If the tested bit is 0:

Increment index_in_ACx, and test the next mask bit.

Example 7–5 demonstrates the use of the field_extract() instruction. The example shows how to de-multiplex the signal created in Example 7–4.

Example 7–5. Demultiplexing a Bit Stream With the Field Extract Instruction


```

.asg  T2, G0                ; Assign G0 to register T2.
.asg  T3, G1                ; Assign G1 to register T3.
.asg  AC0, G0G1             ; Assign G0G1 to register AC0.
.asg  AR1, receive           ; Assign receive to register AR1.


G0G1 = *receive              ; Get bit stream.
G0 = field_extract(G0G1, #05555h) ; Extract G0 from bit stream.
G1 = field_extract(G0G1, #0AAAAh) ; Extract G1 from bit stream.

```

G0 = field_extract(G0G1, #5555h)



G1 = field_extract(G0G1, #AAAAh)



7.4 Viterbi Algorithm for Channel Decoding (ADDSUB, SUBADD, MAXDIFF)

The Viterbi algorithm is widely used in communications systems for decoding information that has been convolutionally encoded. The most computationally intensive part of the routine is comprised of many add-compare-select (ACS) iterations. Minimizing the time for each ACS calculation is important. For a given system, the number of ACS calculations depends on the constraint length K and is equal to $2^{(K-2)}$. Thus, as K increases, the number of ACS calculations increases exponentially. The C55x DSP can perform the ACS operation in 1 cycle, due to dedicated instructions that support the Viterbi algorithm.


The convolutional encoder depicted in Figure 7–3 (page 7-10) is used in the global system for mobile communications (GSM) and is described by the following polynomials ($K=5$):

$$G_0(x) = 1 + x^3 + x^4 \qquad G_1(x) = 1 + x + x^3 + x^4$$

The convolutionally encoded outputs are dependent on past data inputs. Moreover, the contents of the encoder can be viewed as a finite state machine. A trellis diagram can represent the allowable state transitions, along with their corresponding path states. Decoding the data involves finding the optimal path through the trellis, by iteratively selecting possible paths to each delay state, for a given number of symbol time intervals. Two path metrics are calculated by adding a local distance to two old metrics. A comparison is made and a new path metric is selected from the two.

In the case of the GSM encoder, there are 16 possible states for every symbol time interval. For rate $1/n$ systems, there is some inherent symmetry in the trellis structure, which simplifies the calculations. The path states leading to a delay state are complementary. That is, if one path has $G_0G_1 = 00$, the other path has $G_0G_1 = 11$. This symmetry is based on the encoder polynomials and is true for most systems. Two starting and ending complementary states can be paired together, including all the paths between them, to form a butterfly structure (see Figure 7–6). Hence, only one local distance is needed for each butterfly; it is added and subtracted for each new state. Additionally, the old metric values are the same for both updates, so address manipulation is minimized.

Figure 7–6. Butterfly Structure for $K = 5$, Rate 1/2 GSM Convolutional Encoder



The following equation defines a local distance for the rate 1/2 GSM system:

$$LD = SD_0 G_0(j) + SD_1 G_1(j)$$

where

SD_x = Soft-decision input into the decoder

$G_x(j)$ = Expected encoder output for the symbol interval j

Usually, the $G_x(j)$ s are coded as signed antipodal numbers, meaning that “0” corresponds to +1 and “1” corresponds to –1. This coding reduces the local distance calculation to simple addition and subtraction.

As shown in Example 7–6, the DSP can calculate a butterfly quickly by using its accumulators in a dual 16-bit computation mode. To determine the new path metric j , two possible path metrics, $2j$ and $2j+1$, are calculated in parallel with local distances (LD and –LD) using the add-subtract (ADDSUB) instruction and an accumulator. To determine the new path metric $(j+2^{(K-2)})$, the subtract-add (SUBADD) instruction is also used, using the old path metrics plus local distances stored in a separate accumulator. The MAXDIFF instruction is then used on both accumulators to determine the new path metrics. The MAXDIFF instruction compares the upper and lower 16-bit values for two given accumulators, and stores the larger values in a third accumulator.

Example 7–6 shows two macros for Viterbi butterfly calculations. The add-subtract (ADDSUB) and subtract-add (SUBADD) computations in the two macros are performed in alternating order, which is based on the expected encoder state. A local distance is stored in the register T3 beforehand. The MAXDIFF instruction performs the add-compare-select function in 1 cycle. The updated path metrics are saved to memory by the next two lines of code.

Two 16-bit transition registers (TRN0 and TRN1) are updated with every comparison done by the MAXDIFF instruction, so that the selected path metric can be tracked. TRN0 tracks the results from the high part data path, and TRN1 tracks the low part data path. These bits are later used in traceback, to determine the original uncoded data. Using separate transition registers allows for storing the selection bits linearly, which simplifies traceback. In contrast, the TMS320C54x™ (C54x™) DSP has only one transition register, storing the selection bits as 0, 8, 1, 9, etc. As a result, on the C54x DSP, additional lines of code are needed to process these bits during traceback.

You can make the Viterbi butterfly calculations faster by implementing user-defined instruction parallelism (see section 4.2, page 4-20) and software pipelining. Example 7–7 (page 7-21) shows the inner loop of a Viterbi butterfly algorithm. The algorithm places some instructions in parallel (||) in the CPU, and the algorithm implements software pipelining by saving previous results at the same time it performs new calculations. Other operations, such as loading the appropriate local distances, are coded with the butterfly algorithm.

Example 7–6. Viterbi Butterflies for Channel Coding**(a) Mnemonic Instructions**

```

BFLY_DIR_MNEM .MACRO
;new_metric(j)&(j+2^(K-2))

    ADDSUB T3, *AR5+, AC0                ; AC0(39-16) = Old_Met(2*j)+LD
                                          ; AC0(15-0) = Old_met(2*j+1)-LD

    SUBADD T3, *AR5+, AC1                ; AC1(39-16) = Old_Met(2*j)-LD
                                          ; AC1(15-0) = Old_met(2*j+1)+LD

    MAXDIFF AC0, AC1, AC2, AC3           ; Compare AC0 and AC1
    MOV AC2, *AR3+, *AR4+               ; Store the lower maxima (with AR3)
                                          ; and upper maxima (with AR4)

.ENDM

BFLY_REV_MNEM .MACRO
;new_metric(j)&(j+2^(K-2))

    SUBADD T3, *AR5+, AC0                ; AC0(39-16) = Old_Met(2*j)-LD
                                          ; AC0(15-0) = Old_met(2*j+1)+LD

    ADDSUB T3, *AR5+, AC1                ; AC1(39-16) = Old_Met(2*j)+LD
                                          ; AC1(15-0) = Old_met(2*j+1)-LD

    MAXDIFF AC0, AC1, AC2, AC3           ; Compare AC0 and AC1
    MOV AC2, *AR3+, *AR4+               ; Store the lower maxima (with AR3)
                                          ; and upper maxima (with AR4)

.ENDM

```

*Example 7–6 .Viterbi Butterflies for Channel Coding (Continued)**(b) Algebraic Instructions*

```

BFLY_DIR_ALG .MACRO
;new_metric(j)&(j+2^(K-2))
    hi(AC0) = *AR5+ + T3,          ; AC0(39-16) = Old_Met(2*j)+LD
    lo(AC0) = *AR5+ - T3          ; AC0(15-0) = Old_met(2*j+1)-LD

    hi(AC1) = *AR5+ - T3,          ; AC1(39-16) = Old_Met(2*j)-LD
    lo(AC1) = *AR5+ + T3          ; AC1(15-0) = Old_met(2*j+1)+LD

    max_diff(AC0, AC1, AC2, AC3)   ; Compare AC0 and AC1
    *AR3+ = lo(AC2), *AR4+ = hi(AC2) ; Store the lower maxima (with AR3)
                                      ; and upper maxima (with AR4)

    .ENDM

BFLY_REV_ALG .MACRO
;new_metric(j)&(j+2^(K-2))
    hi(AC0) = *AR5+ - T3,          ; AC0(39-16) = Old_Met(2*j)-LD
    lo(AC0) = *AR5+ + T3          ; AC0(15-0) = Old_met(2*j+1)+LD

    hi(AC1) = *AR5+ + T3,          ; AC1(39-16) = Old_Met(2*j)+LD
    lo(AC1) = *AR5+ - T3          ; AC1(15-0) = Old_met(2*j+1)-LD

    max_diff(AC0, AC1, AC2, AC3)   ; Compare AC0 and AC1
    *AR3+ = lo(AC2), *AR4+ = hi(AC2) ; Store the lower maxima (with AR3)
                                      ; and upper maxima (with AR4)

    .ENDM

```

Example 7–7. Viterbi Butterflies Using Instruction Parallelism**(a) Mnemonic Instructions**

```

    RPTBLOCAL end
butterfly:
    ADDSUB T3, *AR0+, AC0           ; AC0(39-16) = Old_Met(2*j)+LD
                                   ; AC0(15-0) = Old_met(2*j+1)-LD
    || MOV *AR5+, AR7

    SUBADD T3, *AR0+, AC1           ; AC1(39-16) = Old_Met(2*j)-LD
                                   ; AC1(15-0) = Old_met(2*j+1)+LD
    || MOV *AR6, T3                ; Load new local distance

    MOV AC2, *AR2+, *AR2(T0)        ; Store lower and upper maxima
                                   ; from previous MAXDIFF operation
    || MAXDIFF AC0, AC1, AC2, AC3   ; Compare AC0 and AC1

    ADDSUB T3, *AR0+, AC0           ; AC0(39-16) = Old_Met(2*j)+LD
                                   ; AC0(15-0) = Old_met(2*j+1)-LD
    || MOV *AR5+, AR6

    SUBADD T3, *AR0+, AC1           ; AC1(39-16) = Old_Met(2*j)-LD
                                   ; AC1(15-0) = Old_met(2*j+1)+LD
    || MOV *AR7, T3                ; Load new local distance

end    MOV AC2, *AR2(T0), *AR2+     ; Store lower and upper maxima
                                   ; from previous MAXDIFF operation
    || MAXDIFF AC0, AC1, AC2, AC3   ; Compare AC0 and AC1

```

*Example 7–7 .Viterbi Butterflies Using Instruction Parallelism (Continued)**(b) Algebraic Instructions*

```

    localrepeat {
butterfly:
    hi(AC0) = *AR0+ + T3,          ; AC0(39-16) = Old_Met(2*j)+LD
    lo(AC0) = *AR0+ - T3          ; AC0(15-0) = Old_met(2*j+1)-LD
    || AR7 = *AR5+

    hi(AC1) = *AR0+ - T3,          ; AC1(39-16) = Old_Met(2*j)-LD
    lo(AC1) = *AR0+ + T3          ; AC1(15-0) = Old_met(2*j+1)+LD
    || T3 = *AR6                  ; Load new local distance

    *AR2+ = lo(AC2), *AR2(T0) = hi(AC2) ; Store lower and upper maxima
                                       ; from previous max_diff operation
    || max_diff( AC0, AC1, AC2, AC3) ; Compare AC0 and AC1

    hi(AC0) = *AR0+ + T3,          ; AC0(39-16) = Old_Met(2*j)+LD
    lo(AC0) = *AR0+ - T3          ; AC0(15-0) = Old_met(2*j+1)-LD
    || AR6 = *AR5+

    hi(AC1) = *AR0+ - T3,          ; AC1(39-16) = Old_Met(2*j)-LD
    lo(AC1) = *AR0+ + T3          ; AC1(15-0) = Old_met(2*j+1)+LD
    || T3 = *AR7                  ; Load new local distance

    *AR2(T0) = lo(AC2), *AR2+ = hi(AC2) ; Store lower and upper maxima
                                       ; from previous max_diff operation
    || max_diff( AC0, AC1, AC2, AC3) ; Compare AC0 and AC1
    }

```

TI C55x DSPLIB



The TI C55x DSPLIB is an optimized DSP function library for C programmers on TMS320C55x™ (C55x™) DSP devices. It includes over 50 C-callable assembly-optimized general-purpose signal processing routines. These routines are typically used in computationally intensive real-time applications where optimal execution speed is critical. By using these routines you can achieve execution speeds considerably faster than equivalent code written in standard ANSI C language. In addition, by providing ready-to-use DSP functions, TI DSPLIB can shorten significantly your DSP application development time.

The TI DSPLIB includes commonly used DSP routines. Source code is provided to allow you to modify the functions to match your specific needs.

Topic	Page
8.1 Features and Benefits	8-2
8.2 DSPLIB Data Types	8-2
8.3 DSPLIB Arguments	8-2
8.4 Calling a DSPLIB Function from C	8-3
8.5 Calling a DSPLIB Function from Assembly Language Source Code	8-4
8.6 Where to Find Sample Code	8-4
8.7 DSPLIB Functions	8-5

8.1 Features and Benefits

- ☐ Hand-coded assembly optimized routines
- ☐ C-callable routines fully compatible with the C55x DSP compiler
- ☐ Fractional Q15-format operands supported
- ☐ Complete set of examples on usage provided
- ☐ Benchmarks (cycles and code size) provided
- ☐ Tested against Matlab™ scripts

8.2 DSPLIB Data Types

DSPLIB functions generally operate on Q15-fractional data type elements:

- ☐ Q.15 (DATA) : A Q.15 operand is represented by a *short* data type (16 bit) that is predefined as *DATA*, in the *dsplib.h* header file.

Certain DSPLIB functions use the following data type elements:

- ☐ Q.31 (LDATA) : A Q.31 operand is represented by a *long* data type (32 bit) that is predefined as *LDATA*, in the *dsplib.h* header file.
- ☐ Q.3.12 : Contains 3 integer bits and 12 fractional bits.

8.3 DSPLIB Arguments

DSPLIB functions typically operate over vector operands for greater efficiency. Though these routines can be used to process short arrays or scalars (unless a minimum size requirement is noted), the execution times will be longer in those cases.

- ☐ **Vector stride is always equal 1:** vector operands are composed of vector elements held in consecutive memory locations (vector stride equal to 1).
- ☐ **Complex elements** are assumed to be stored in a Real-Imaginary (Re-Im) format.
- ☐ **In-place computation is allowed (unless specifically noted):** Source operand can be equal to destination operand to conserve memory.

8.4 Calling a DSPLIB Function from C

In addition to installing the DSPLIB software, to include a DSPLIB function in your code you have to:

- ☐ Include the *dsplib.h* include file
- ☐ Link your code with the DSPLIB object code library, *55xdsp.lib*.
- ☐ Use a correct linker command file describing the memory configuration available in your C55x DSP board.

For example, the following code contains a call to the *recip16* and *q15tofl* routines in DSPLIB:

```
#include "dsplib.h"

DATA x[3] = { 12398 , 23167, 564};

DATA r[NX];
DATA rexp[NX];
float rf1[NX];
float rf2[NX];

void main()
{
    short i;

    for (i=0;i<NX;i++)
    {
        r[i] =0;
        rexp[i] = 0;
    }

    recip16(x, r, rexp, NX);
    q15tofl(r, rf1, NX);

    for (i=0; i<NX; i++)
    {
        rf2[i] = (float)rexp[i] * rf1[i];
    }

    return;
}
```

In this example, the *q15tofl* DSPLIB function is used to convert Q15 fractional values to floating-point fractional values. However, in many applications, your data is always maintained in Q15 format so that the conversion between floating point and Q15 is not required.

8.5 Calling a DSPLIB Function from Assembly Language Source Code

The DSPLIB functions were written to be used from C. Calling the functions from assembly language source code is possible as long as the calling-function conforms with the C55x DSP C compiler calling conventions. Refer to the *TMS320C55x Optimizing C Compiler User's Guide*, if a more in-depth explanation is required.

Realize that the DSPLIB is not an optimal solution for assembly-only programmers. Even though DSPLIB functions can be invoked from an assembly program, the resulting execution times and code size may not be optimal due to unnecessary C-calling overhead.

8.6 Where to Find Sample Code

You can find examples on how to use every single function in DSPLIB, in the *examples* subdirectory. This subdirectory contains one subdirectory for each function. For example, the *examples/araw* subdirectory contains the following files:

- ❑ *araw.t.c*: main driver for testing the DSPLIB *acorr* (*raw*) function.
- ❑ *test.h*: contains input data(*a*) and expected output data(*yraw*) for the *acorr* (*raw*) function as. This *test.h* file is generated by using Matlab scripts.
- ❑ *test.c*: contains function used to compare the output of *araw* function with the expected output data.
- ❑ *ftest.c*: contains function used to compare two arrays of float data types.
- ❑ *ltest.c*: contains function used to compare two arrays of long data types.
- ❑ *ld3.cmd*: an example of a linker command you can use for this function .

8.7 DSPLIB Functions

8.7.1 Description of Arguments

Table 8–1 describes the arguments for each function.

Table 8–1. DSPLIB Function Argument Descriptions

Argument	Description
<i>x,y</i>	argument reflecting input data vector
<i>r</i>	argument reflecting output data vector
<i>nx,ny,nr</i>	arguments reflecting the size of vectors <i>x</i> , <i>y</i> , and <i>r</i> respectively. In functions where $nx = nr = nr$, only <i>nx</i> has been used.
<i>h</i>	Argument reflecting filter coefficient vector (filter routines only)
<i>nh</i>	Argument reflecting the size of vector <i>h</i>
<i>DATA</i>	data type definition equating a short, a 16-bit value representing a Q15 number. Usage of <i>DATA</i> instead of short is recommended to increase future portability across devices.
<i>LDATA</i>	data type definition equating a long, a 32-bit value representing a Q31 number. Usage of <i>LDATA</i> instead of long is recommended to increase future portability across devices.
<i>ushort</i>	Unsigned short (16 bit). You can use this data type directly, because it has been defined in <i>dsplib.h</i>

8.7.2 List of DSPLIB Functions

Table 8–2 lists the DSPLIB functions by these 8 functional categories:

- ☐ Fast-Fourier Transforms (FFT)
- ☐ Filtering and convolution
- ☐ Adaptive filtering
- ☐ Correlation
- ☐ Math
- ☐ Trigonometric
- ☐ Miscellaneous
- ☐ Matrix

Table 8–2. DSPLIB Functions

Functions	Description
FFT	
void cfft (DATA *x, ushort nx, ushort scale)	Radix–2 complex forward FFT – MACRO
void ciff2 (DATA *x, ushort nx, ushort scale)	Radix–2 complex inverse FFT – MACRO
void cbrev (DATA *x, DATA *r, ushort n)	Complex bit–reverse function
Filtering and Convolution	
ushort fir(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nx, ushort nh)	FIR Direct form
ushort fir2(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nx, ushort nh)	FIR Direct form (Optimized to use DUAL–MAC)
ushort firs(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nx, ushort nh2)	Symmetric FIR Direct form (generic routine)
ushort cfir(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nx, ushort nh)	Complex FIR direct form
ushort convol(DATA *x, DATA *h, DATA *r, ushort nr, ushort nh)	Convolution
ushort convol1(DATA *x, DATA *h, DATA *r, ushort nr, ushort nh)	Convolution (Optimized to use DUAL–MAC)
ushort convol2(DATA *x, DATA *h, DATA *r, ushort nr, ushort nh)	Convolution (Optimized to use DUAL–MAC)
ushort iircas4(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nbiqu, ushort nx)	IIR cascade Direct Form 2. 4 coefficients per biquad.
ushort iircas5(DATA *x, DATA *h, DATA *r, DATA *dbuffer, ushort nbiqu, ushort nx)	IIR cascade Direct Form 2. 5 coefficients per biquad
ushort iircas51(DATA *x, DATA *h, DATA *r, DATA **dbuffer, ushort nbiqu, ushort nx)	IIR cascade Direct Form 1. 5 coefficients per biquad
ushort iirlat (DATA *x, DATA *h, DATA *r, DATA *pbuffer, int nx, int nh)	Lattice inverse IIR filter
ushort firlat (DATA *x, DATA *h, DATA *r, DATA *pbuffer, int nx, int nh)	Lattice forward FIR filter

Table 8–2. DSPLIB Functions (Continued)

Functions	Description
Adaptive filtering	
ushort dlms(DATA *x, DATA *h, DATA *r, DATA *des, DATA *dbuffer, DATA step, ushort nh, ushort nx)	LMS FIR (delayed version)
Correlation	
ushort acorr (DATA *x, DATA *r, ushort nx, ushort nr, type)	Auto-correlation (positive side only) – MACRO
Trigonometric	
ushort sine(DATA *x, DATA *r, ushort nx)	sine of a vector
ushort atan2_16(DATA *i, DATA *q, DATA *r, ushort nx)	4 – Quadrant Inverse Tangent of a vector
ushort atan16(DATA *x, DATA *r, ushort nx)	Arctan of a vector
Math	
ushort add (DATA *x, DATA *y, DATA *r, ushort nx, ushort scale)	Optimized vector addition
ushort expn(DATA *x, DATA *r, ushort nx)	Exponent of a vector
ushort logn(DATA *x, LDATA *r, ushort nx)	Natural log of a vector
ushort log_2(DATA *x, LDATA *r, ushort nx)	Log base 2 of a vector
ushort log_10(DATA *x, LDATA *r, ushort nx)	Log base 10 of a vector
short maxidx (DATA *x, ushort nx)	Index for maximum magnitude in a vector
short maxval (DATA *x, ushort nx)	Maximum magnitude in a vector
short minidx (DATA *x, ushort nx)	Index for minimum magnitude in a vector
short minval (DATA *x, ushort nx)	Minimum element in a vector
short neg (DATA *x, DATA *r, ushort nx)	16-bit vector negate
short neg32 (LDATA *x, LDATA *r, ushort nx)	32-bit vector negate
short power (DATA *x, LDATA *r, ushort nx)	sum of squares of a vector (power)
void recip16(DATA *x, DATA *r, DATA *rexp, ushort nx)	Vector reciprocal
void ldiv16(LDATA *x, DATA *y, DATA *r, DATA *rexp, ushort nx)	32 bit by 16-bit long division

Table 8–2. DSPLIB Functions (Continued)

Functions	Description
ushort sqrt_16(DATA *x, DATA *r, short nx)	Square root of a vector
short sub (DATA *x, DATA *y, DATA *r, ushort nx, ushort scale)	Vector subtraction
Matrix	
ushort mmul(DATA *x1,short row1,short col1,DATA *x2,short row2,short col2,DATA *r)	matrix multiply
ushort mtrans(DATA *x, short row, short col, DATA *r)	matrix transpose
Miscellaneous	
ushort fltoq15(float *x, DATA *r, ushort nx)	Float to Q15 conversion
ushort q15tofl(DATA *x, float *r, ushort nx)	Q15 to float conversion

Index

- 2s-complement addition
 - concept, 5-5
 - extended-precision, 5-10
- 2s-complement arithmetic
 - concepts, 5-5
 - extended-precision (concept), 5-8
 - extended-precision (details), 5-10, 5-17
- 2s-complement division, 5-21
- 2s-complement fractional format, 5-5
- 2s-complement integer format, 5-4
- 2s-complement multiplication
 - concept, 5-5
 - extended-precision, 5-17
- 2s-complement numbers
 - concept, 5-2
 - decimal equivalent (example), 5-3
 - negative of (example), 5-3
- 2s-complement subtraction
 - concept, 5-5
 - extended-precision, 5-10

A

- A unit, parallel optimization within, 4-29
- AC1 (access 1) stage of pipeline, 4-53
- AC2 (access 2) stage of pipeline, 4-53
- access 2 (AC2) stage of pipeline, 4-53
- access 1 (AC1) stage of pipeline, 4-53
- accesses to dual-access RAM (DARAM), 4-78
- accesses to single-access RAM (SARAM), 4-79
- AD (address) stage of pipeline, 4-53
- adaptive FIR filtering
 - concept, 7-6
 - example, 7-9
- addition
 - 2s-complement, 5-5
 - extended-precision 2s-complement, 5-10

- address (AD) stage of pipeline, 4-53
- address buses available, 1-2
- address generation units, 1-2
- addressing modes
 - ARn, 2-8
 - bit addressing by instruction, 2-8
 - DP direct, 2-8
 - k23, 2-8
- ADDSUB instruction used in Viterbi code, 7-19
- advanced parallelism rules, 4-26
- allocating data, 3-36
- application-specific instructions, 7-1
- architecture of TMS320C55x DSPs
 - features supporting parallelism, 4-21
 - overview, 1-2
- arguments to functions in TMS320C55x DSP library, 8-2
- arithmetic
 - 2s-complement, 5-5
 - extended-precision 2s-complement (concept), 5-8
 - extended-precision 2s-complement (details), 5-10, 5-17
- assembly code, 3-8
- assembly code optimization, 4-1
- asymmetric FIR filtering, 7-2
 - with FIRSN instruction, 7-4

B

- B bus (BB), use in dual-MAC operations, 4-2
- basic parallelism rules, 4-25
- BB bus, use in dual-MAC operations, 4-2
- bit field expand instruction
 - concept, 7-12
 - example, 7-13
- bit field extract instruction
 - concept, 7-14

- example, 7-15
- bit–reverse addressing, 6-1
 - in FFT algorithms, 6-4
 - in TMS320C55x DSP function library, 6-8
 - in–place versus off–place bit–reversing, 6-6
- introduction, 6-2
- syntaxes, 6-2
- typical initialization requirements, 6-5
- bit–stream (de)multiplexing, 7-11
- block–repeat instructions, identifying which type to use, 4-49
- built–in parallelism, 4-20
- buses available, 1-2
- buses supporting parallelism, 4-23
- butterfly structure (figure), 7-17
- byte extensions (parallelism rule), 4-26

C

- C bus (CB), use in dual–MAC operations, 4-2
- c code, 3-8
- C function library for TMS320C55x DSPs, 8-1
 - calling a function from assembly source, 8-4
 - calling a function from C, 8-3
 - data types used, 8-2
 - function arguments, 8-2
 - list of functions, 8-5
 - where to find sample code, 8-4
- calling a function in TMS320C55x DSP library
 - from assembly source, 8-4
 - from C, 8-3
- CARRY bit
 - affected by addition, 5-10
 - affected by subtraction, 5-13
- CB bus, use in dual–MAC operations, 4-2
- circular addressing, 3-29
- coarse granularity
 - memory–mapped register pipeline protection, 4-66
 - status register pipeline protection, 4-64
- code development flow, 1-3
- code examples. *See* examples
- CODE_SECTION, 3-40
- comments, 2-3
- compatibility with TMS320C54x DSPs, 1-2

- compiler options, 3-2

- ml, 3-4
- mn, 3-4
- mr, 3-4
- ms, 3-4
- oimize, 3-3
- on, 3-2, 3-24
- onx, 3-3
- opn, 3-3
- pm, 3-4, 3-24

- complex vector multiplication using dual–MAC hardware, 4-6

- computation blocks, 1-2

- computed single–repeat register (CSR), using with a single–repeat instruction, 4-50

- constants, 2-5

- control code, 3-27

- convolutional encoding, 7-10

- CSR, using with a single–repeat instruction, 4-50

D

- D (decode) stage of pipeline, 4-53

- D bus (DB), use in dual–MAC operations, 4-2

- D unit, parallel optimization within, 4-37

- DARAM accesses (buses and pipeline stages), 4-78

- data alignment, 3-34

- data allocation, 3-36

- data buses available, 1-2

- data types used in TMS320C55x DSP function library, 8-2

- DB bus, use in dual–MAC operations, 4-2

- decode (D) stage of pipeline, 4-53

- delayed LMS algorithm, 7-7

- directive, SECTIONS, 3-40

- directives

- def, 2-5
- init, 2-5
- MEMORY, 2-10
- sect, 2-5
- SECTIONS, 2-10
- usect, 2-5

- division, 2s–complement, 5-21

- DSP function library for TMS320C55x DSPs, 8-1

- calling a function from assembly source, 8-4
- calling a function from C, 8-3
- data types used, 8-2

- function arguments, 8-2
- list of functions, 8-5
- where to find sample code, 8-4

DSPLIB. *See* DSP function library for TMS320C55x DSPs

dual-access RAM (DARAM) accesses (buses and pipeline stages), 4-78

dual-MAC hardware/operations

- efficient use of, 4-2
- matrix multiplication, 4-18
- pointer usage in, 4-3
- taking advantage of implicit algorithm symmetry, 4-5

dynamic scaling for overflow handling, 5-32

E

encoding, convolutional, 7-10

examples

- adaptive FIR filter (delayed LMS), 7-9
- block FIR filter (dual-MAC implementation), 4-12, 4-14
- branch-on-auxiliary-register-not-zero loop construct, 4-48
- complex vector multiplication (dual-MAC implementation), 4-7
- demultiplexing bit streams (convolutional encoder), 7-15
- multiplexing bit streams (convolutional encoder), 7-13
- nested loops, 4-47
- output streams for convolutional encoder, 7-11
- parallel optimization across CPU functional units, 4-39
- parallel optimization within A unit, 4-29
- parallel optimization within CPU functional units, 4-29
- parallel optimization within D unit, 4-37
- parallel optimization within P unit, 4-35
- symmetric FIR filter, 7-5
- use of CSR for single-repeat loop, 4-51
- Viterbi butterfly, 7-19

execute (X) stage of pipeline, 4-54

extended auxiliary register usage in dual-MAC operations, 4-3

extended coefficient data pointer usage in dual-MAC operations, 4-3

extended-precision 2s-complement arithmetic

- concepts, 5-8
- details, 5-10, 5-17

F

FFT flow graph with bit-reversed input, 6-4

field expand instruction

- concept, 7-12
- example, 7-13

field extract instruction

- concept, 7-14
- example, 7-15

filtering

- adaptive FIR (concept), 7-6
- adaptive FIR (example), 7-9
- symmetric/asymmetric FIR, 7-2
- with dual-MAC hardware, 4-3, 4-5, 4-9

fine granularity

- memory-mapped register pipeline protection, 4-66
- status register pipeline protection, 4-64

FIR filtering

- adaptive (concept), 7-6
- adaptive (example), 7-9
- symmetric/asymmetric, 7-2

FIRS instruction, 7-3, 7-4

FIRSN instruction, 7-4

fixed scaling for overflow handling, 5-32

fixed-point arithmetic, 5-1

fractional representation (2s-complement), 5-5

fractions versus integers, 5-3

FUNC_EXT_CALLED, 3-8

function allocation, 3-40

function inlining, 3-10

function library for TMS320C55x DSPs, 8-1

- calling a function from assembly source, 8-4
- calling a function from C, 8-3
- data types used, 8-2
- function arguments, 8-2
- list of functions, 8-5
- where to find sample code, 8-4

G

global symbols, 3-35

granularity

- memory-mapped register pipeline protection, 4-66

status register pipeline protection, 4-64
guard bits used for overflow handling, 5-31

H

hardware resource conflicts (parallelism rule), 4-25

I

I/O space available, 1-2
implicit algorithm symmetry, 4-5
input scaling for overflow handling, 5-32
input/output space available, 1-2
instruction buffer, 1-2
instruction length limitation (parallelism rule), 4-25
instruction pipeline
 introduced, 1-2
 minimizing delays, 4-53
 segments and stages, 4-53
instructions executed in parallel, 4-20
instructions for specific applications, 7-1
`int_norm`, 3-15
integer representation (2s-complement), 5-4
integers versus fractions, 5-3
interrupt-control logic, 1-2
intrinsics, 3-12
 `int_abss`, 3-14
 `int_inorm`, 3-15
 `int_lshrs`, 3-15
 `int_md`, 3-15
 `int_norm`, 3-15
 `int_sadd`, 3-14
 `int_shrs`, 3-15
 `int_smacr`, 3-15
 `int_smasr`, 3-15
 `int_smpy`, 3-14, 3-15
 `int_sneg`, 3-15
 `int_ssubb`, 3-14
 `int_subc`, 3-15
 `long_addc`, 3-15
 `long_labss`, 3-15
 `long_laddc`, 3-15
 `long_lsadd`, 3-14
 `long_lsmpr`, 3-14
 `long_lsneg`, 3-15
 `long_lssh`, 3-15
 `long_lsshl`, 3-15

`long_lssub`, 3-14
 `long_lsubc`, 3-15
 `long_smac`, 3-14
 `long_smas`, 3-14
 `_nassert`, 3-18, 3-21

L

labels, 2-3
least mean square (LMS) calculation for adaptive filtering, 7-6
least mean square (LMS) instruction, 7-6
linker command file, 3-39
linking, 2-10
local block-repeat instruction, when to use, 4-49
local symbols, 3-35
logical units, 3-36
long data access, 3-16
loop unrolling, 4-9
loop-control registers
 avoiding pipeline delays when accessing, 4-52
 when they are accessed in the pipeline, 4-68
loops
 implementing efficient, 4-46
 nesting of, 4-46

M

MAC, 3-28
MAC units, 1-2
map file, example, 2-12
math operations, 3-28
matrix mathematics using dual-MAC hardware, 4-18
MAXDIFF instruction used in Viterbi code, 7-19
maximum instruction length (parallelism rule), 4-25
memory, 3-34
 allocating data, 3-34
 data alignment, 3-34
 stack configuration, 3-34
 symbol declaration, 3-34
memory accesses and the pipeline, 4-77
memory available, 1-2
memory-mapped register pipeline-protection granularity, 4-66
`mmap()` qualifier (parallelism rule), 4-26
MMR pipeline-protection granularity, 4-66

- modulus operator, 3-29
- multi-algorithm applications enhanced with dual-MAC hardware, 4-4
- multi-channel applications implemented with dual-MAC hardware, 4-3
- multiplexing bit streams (convolutional encoder), 7-11
- multiplication
 - 2s-complement, 5-5
 - extended-precision 2s-complement, 5-17
- multiply-and-accumulate (MAC) units, 1-2

N

- `_nassert`, 3-18
- negative of 2s-complement number (example), 5-3
- nesting of loops, 4-46

O

- operands, 3-38
 - global, 3-38
 - local, 3-38
- operators supporting parallelism, 4-21
- optimizing assembly code, 4-1
- options, 3-2
 - `-pm`, 3-5
- overflow flags used for overflow handling, 5-31
- Overflow handling, 5-31
 - FFTs, 5-33
 - FIR filters, 5-32
 - hardware features for, 5-31
 - IIR filters, 5-32
 - techniques for, 5-31

P

- P unit, parallel optimization within, 4-35
- parallel enable bit, 4-25
- parallel execution features, 4-20
- parallel optimization, examples, 4-29, 4-39
- parallelism
 - architectural features, 4-21
 - built-in, 4-20
 - user-defined, 4-20
- parallelism rules for user-defined parallelism, 4-24

- parallelism tips, 4-28
- pipeline
 - introduced, 1-2
 - segments and stages, 4-53
- pipeline conflicts, process to resolve, 4-55
- pipeline delays
 - minimizing, 4-53
 - recommendations for preventing, 4-56
 - when accessing loop-control registers, 4-52
- pipeline-protection granularity
 - for memory-mapped registers, 4-66
 - for status registers (ST0_55-ST3_55), 4-64
- pointer usage in dual-MAC operations, 4-3
- `port()` qualifier (parallelism rule), 4-26
- pragma
 - `CODE_SECTION`, 3-40
 - `DATA_SECTION`, 3-37
 - `FUNC_EXT_CALLED`, 3-8
- procedure for efficient code, 1-3
- process for user-defined parallelism, 4-27
- process to resolve pipeline conflicts, 4-55
- processor initialization, 2-3, 2-7
- Program-level Optimization, 3-5

R

- R (read) stage of pipeline, 4-54
- read (R) stage of pipeline, 4-54
- recommendations for preventing pipeline delays, 4-56
- registers, when they are accessed in the pipeline, 4-58, 4-68
- resource conflicts (parallelism rule), 4-25
- rules for user-defined parallelism, 4-24

S

- SARAM accesses (buses and pipeline stages), 4-79
- saturation for overflow handling, 5-32
- saturation mode bits used for overflow handling, 5-31
- scaling
 - dynamic, 5-32
 - fixed, 5-32
 - input, 5-32
 - methods for FFTs, 5-33
 - methods for FIR filters, 5-32

- methods for IIR filters, 5-32
- scaling methods (overflow handling)
 - FFTs, 5-33
 - FIR filters, 5-32
 - IIR filters, 5-32
- section allocation, 2-5
 - example, 2-6
- sections, 3-36
 - .bss, 3-36
 - .cinit, 3-36
 - .const, 3-36
 - .do, 3-36
 - .iport, 3-36
 - .stack, 3-36
 - .switch, 3-36
 - .sysmem, 3-36
 - .sysstack, 3-36
 - .text, 3-36
- single-access RAM (SARAM) accesses (buses and pipeline stages), 4-79
- single-repeat instruction, when to use, 4-49
- soft dual encoding, 4-25
- ST0_55–ST3_55 pipeline-protection granularity, 4-64
- stack configuration, 3-35
- stacks available, 1-2
- standard block-repeat instruction, when to use, 4-49
- status register pipeline-protection granularity, 4-64
- SUBADD instruction used in Viterbi code, 7-19
- subtraction
 - 2s-complement, 5-5
 - extended-precision 2s-complement, 5-10
- symbol declarations, 3-35
- symmetric FIR filtering, 7-2
 - with FIRS instruction (concept), 7-3
 - with FIRS instruction (example), 7-4

T

- table to help generate optional application mapping, 4-81
- TI C55x DSPLIB, 8-1
 - calling a function in assembly source, 8-4
 - calling a function in C, 8-3
 - data types used, 8-2
 - function arguments, 8-2
 - list of functions, 8-5

- where to find sample code, 8-4
- tips
 - applying parallelism, 4-28
 - nesting loops, 4-49
 - preventing pipeline delays, 4-56
 - producing efficient code, 1-3
 - resolving pipeline conflicts, 4-55
- TMS320C54x-compatible mode, 1-2
- TMS320C55x DSP function library, 8-1
 - calling a function from assembly source, 8-4
 - calling a function from C, 8-3
 - data types used, 8-2
 - function arguments, 8-2
 - list of functions, 8-5
 - where to find sample code, 8-4
- transition registers (TRN0, TRN1) used in Viterbi algorithm, 7-18
- trip count, 3-18
 - unsigned integer types, 3-19
- tutorial, 2-1

U

- user-defined parallelism, 4-20
 - process, 4-27
 - rules, 4-24

V

- variables, 2-5
- vector multiplication using dual-MAC hardware, 4-6
- Viterbi algorithm for channel decoding, 7-16
- Viterbi butterfly
 - examples, 7-19
 - figure, 7-17

W

- W (write) stage of pipeline, 4-54
- write (W) stage of pipeline, 4-54
- writing assembly code, 2-3

X

- X (execute) stage of pipeline, 4-54
- x2, 3-2
- XARn, example, 2-9
- XARn usage in dual-MAC operations, 4-3
- XCDP usage in dual-MAC operations, 4-3