

TMS320C6000 Assembly Language Tools User's Guide

Literature Number: SPRU186N
April 2004



Printed on Recycled Paper

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

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

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products		Applications	
Amplifiers	amplifier.ti.com	Audio	www.ti.com/audio
Data Converters	dataconverter.ti.com	Automotive	www.ti.com/automotive
DSP	dsp.ti.com	Broadband	www.ti.com/broadband
Interface	interface.ti.com	Digital Control	www.ti.com/digitalcontrol
Logic	logic.ti.com	Military	www.ti.com/military
Power Mgmt	power.ti.com	Optical Networking	www.ti.com/opticalnetwork
Microcontrollers	microcontroller.ti.com	Security	www.ti.com/security
		Telephony	www.ti.com/telephony
		Video & Imaging	www.ti.com/video
		Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments
Post Office Box 655303 Dallas, Texas 75265

Read This First

About This Manual

The *TMS320C6000 Assembly Language Tools User's Guide* tells you how to use these assembly language tools:

- Assembler
- Archiver
- Linker
- Cross-reference lister
- Absolute lister
- Hex conversion utility
- Name utility

Before you use this book, you should install the assembly language tools.

How to Use This Manual

This book helps you learn how to use the Texas Instruments assembly language tools designed specifically for the TMS320C6000 32-bit devices. This book consists of four parts:

- Introductory information**, consisting of Chapters 1 and 2, gives you an overview of the assembly language development tools. It also discusses common object file format (COFF), which helps you to use the TMS320C6000 tools more efficiently. Read Chapter 2, *Introduction to Common Object File Format*, before using the assembler and linker.
- Assembler description**, consisting of Chapters 3 through 5, contains detailed information about using the assembler. This portion explains how to invoke the assembler and discusses source statement format, valid constants and expressions, assembler output, and assembler directives. It also describes the macro language.

- ❑ **Additional assembly language tools**, consisting of Chapters 6 through 11, describes in detail each of the tools provided with the assembler to help you create executable object files. For example, Chapter 7 explains how to invoke the linker, how the linker operates, and how to use linker directives. Chapter 11 explains how to use the hex conversion utility.
- ❑ **Reference material**, consisting of Appendixes A through C, provides technical data about the internal format and structure of COFF object files. It discusses symbolic debugging directives that the TMS320C6000 C/C++ compiler uses. Finally, it includes hex conversion utility examples, assembler and linker error messages, and a glossary.

Notational Conventions

This document uses the following conventions:

- ❑ The TMS320C62x, 'C64x, and 'C67x core is referred to as TMS320C6000 or C6000.
- ❑ Program listings, program examples, and interactive displays are shown in a special typeface. 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:

```
1 00000000                                .data
2 00000000 0000002F x                      .byte  47
3 00000001 00000032 z                      .byte  50
4 00000000                                .text
5 00000000 010401E0                        ADD    A0, A1, A2
```

- ❑ In syntax descriptions, the instruction, command, or directive is in a **bold typeface** 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. Syntax that is entered on a command line is centered. Syntax that is used in a text file is left justified. Here is an example of command-line syntax:

```
cl6x -z [options] filename1 . . . filenamen
```

The **cl6x -z** command invokes the linker and has two parameters. The first parameter, *options*, is optional (see the next bullet for details). The second parameter, *filename*, is required and you can enter more than one.

- Square brackets ([and]) identify an optional parameter. If you use an optional parameter, you specify the information within the brackets. Unless the square brackets are in a **bold** typeface, do not enter the brackets themselves. This is an example of a command that has an optional parameter:

```
hex6x [options] filename
```

The **hex6x** command has two parameters. The second parameter, *filename*, is required. The first parameter, *options*, is optional. Since *options* is plural, you can select several options.

- In assembler syntax statements, column 1 is reserved for the first character of a label or symbol. If the label or symbol is optional, it is usually not shown. If it is a required parameter, it is shown starting against the left margin of the shaded box, as in the example below. No instruction, command, directive, or parameter other than a symbol or label can begin in column 1.

```
symbol .usect "section name", size in bytes [, alignment]
```

The *symbol* is required for the `.usect` directive and must begin in column 1. The *section name* must be enclosed in quotes and the parameter *size in bytes* must be separated from the *section name* by a comma. The *alignment* is optional and, if used, must be separated by a comma.

- 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 value1 [, ... , valuen]
```

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

- In program listings and program examples, pipe symbols (||) indicate parallel instructions, and square brackets ([]) indicate conditional instructions. This is an example of parallel and conditional instructions:

```

1                                     .global tab1, tab2
2
3 00000000 00000028!                 MVK     tab1,A0
4 00000004 00000068!                 MVKH   tab1,A0
5 00000008 008031A9                 MVK     99, A1
6 0000000c 010848C0 ||                ZERO   A2
7
8 00000010 80000212 $1:[A1] B        $1
9 00000014 01003674                 STW    A2, *A0++
10 00000018 0087E1A0                 SUB    A1,1,A1
11 0000001c 00004000                 NOP    3

```

The instruction on line five executes in parallel with instruction on line six. The instruction on line eight is conditional: the branch to \$1 only occurs if the contents of A1 are not equal to 0.

- Following are other symbols and abbreviations used throughout this document:

Symbol	Definition	Symbol	Definition
B, b	Suffix — binary integer	MSB	Most significant bit
H, h	Suffix — hexadecimal integer	0x	Prefix — hexadecimal integer
LSB	Least significant bit	Q, q	Suffix — octal integer

Related Documentation From Texas Instruments

The following books describe the TMS320C6000 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.

TMS320C6000 Optimizing Compiler User's Guide (literature number SPRU187) describes the TMS320C6000™ C compiler and the assembly optimizer. This C compiler accepts ANSI standard C source code and produces assembly language source code for the TMS320C6000 generation of devices. The assembly optimizer helps you optimize your assembly code.

Code Composer User's Guide (literature number SPRU328) explains how to use the Code Composer development environment to build and debug embedded real-time DSP applications.

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

TMS320C6000 CPU and Instruction Set Reference Guide (literature number SPRU189) describes the TMS320C6000™ CPU architecture, instruction set, pipeline, and interrupts for these digital signal processors.

TMS320C6201/C6701 Peripherals Reference Guide (literature number SPRU190) describes common peripherals available on the TMS320C6201 and TMS320C6701 digital signal processors. This book includes information on the internal data and program memories, the external memory interface (EMIF), the host port interface (HPI), multi-channel buffered serial ports (McBSPs), direct memory access (DMA), enhanced DMA (EDMA), expansion bus, clocking and phase-locked loop (PLL), and the power-down modes.

TMS320C62x/C67x Technical Brief (literature number SPRU197) gives an introduction to the TMS320C62x™ and TMS320C67x™ digital signal processors, development tools, and third-party support.

Trademarks

Windows and Windows NT are trademarks of Microsoft Corporation.

The Texas Instruments logo and Texas Instruments are registered trademarks of Texas Instruments Incorporated. Trademarks of Texas Instruments include: TI, XDS, Code Composer, Code Composer Studio, TMS320, TMS320C6000 and 320 Hotline On-line.

All other brand or product names are trademarks or registered trademarks of their respective companies or organizations.

Contents

1	Introduction to the Software Development Tools	1-1
	<i>Provides an overview of the software development tools.</i>	
1.1	Software Development Tools Overview	1-2
1.2	Tools Descriptions	1-3
1.3	Code Composer Studio and the Code Generation Tools	1-5
2	Introduction to Common Object File Format	2-1
	<i>Common object file format, or COFF, is the object file format used by the TMS320C6000 tools. This chapter discusses the basic COFF concept of sections and how they can help you use the assembler and linker more efficiently. Read this chapter before using the assembler and linker.</i>	
2.1	Sections	2-2
2.2	How the Assembler Handles Sections	2-4
2.2.1	Uninitialized Sections	2-4
2.2.2	Initialized Sections	2-6
2.2.3	Named Sections	2-6
2.2.4	Subsections	2-7
2.2.5	Section Program Counters	2-8
2.2.6	Using Sections Directives	2-8
2.3	How the Linker Handles Sections	2-11
2.3.1	Default Memory Allocation	2-12
2.3.2	Placing Sections in the Memory Map	2-13
2.4	Relocation	2-14
2.5	Run-Time Relocation	2-16
2.6	Loading a Program	2-17
2.7	Symbols in a COFF File	2-18
2.7.1	External Symbols	2-18
2.7.2	The Symbol Table	2-19

3	Assembler Description	3-1
	<i>Explains how to invoke the assembler and discusses source statement format, valid constants and expressions, and assembler output.</i>	
3.1	Assembler Overview	3-2
3.2	The Assembler's Role in the Software Development Flow	3-3
3.3	Invoking the Assembler	3-4
3.4	Naming Alternate Directories for Assembler Input	3-7
3.4.1	Using the -I Assembler Option	3-7
3.4.2	Using the C6X_A_DIR or A_DIR Environment Variable	3-8
3.5	Source Statement Format	3-10
3.5.1	Label Field	3-11
3.5.2	Mnemonic Field	3-12
3.5.3	Unit Specifier Field	3-12
3.5.4	Operand Field	3-13
3.5.5	Comment Field	3-13
3.6	Constants	3-14
3.6.1	Binary Integers	3-14
3.6.2	Octal Integers	3-14
3.6.3	Decimal Integers	3-15
3.6.4	Hexadecimal Integers	3-15
3.6.5	Character Constants	3-15
3.6.6	Assembly-Time Constants	3-16
3.7	Character Strings	3-17
3.8	Symbols	3-18
3.8.1	Labels	3-18
3.8.2	Local Labels	3-18
3.8.3	Symbolic Constants	3-21
3.8.4	Defining Symbolic Constants (-d Option)	3-21
3.8.5	Predefined Symbolic Constants	3-23
3.8.6	Substitution Symbols	3-25
3.9	Expressions	3-26
3.9.1	Operators	3-27
3.9.2	Expression Overflow and Underflow	3-27
3.9.3	Well-Defined Expressions	3-28
3.9.4	Conditional Expressions	3-28
3.9.5	Legal Expressions	3-28
3.9.6	Expression Examples	3-29
3.10	Source Listings	3-31
3.11	Debugging Assembly Source	3-34
3.12	Cross-Reference Listings	3-36

4	Assembler Directives	4-1
	<i>Describes the directives according to function and presents the directives in alphabetical order.</i>	
4.1	Directives Summary	4-2
4.2	Directives That Define Sections	4-8
4.3	Directives That Initialize Constants	4-10
4.4	Directive That Aligns the Section Program Counter	4-13
4.5	Directives That Format the Output Listings	4-14
4.6	Directives That Reference Other Files	4-16
4.7	Directives That Enable Conditional Assembly	4-17
4.8	Directives That Define Symbols at Assembly Time	4-18
4.9	Miscellaneous Directives	4-20
4.10	Directives Reference	4-21
5	Macro Language	5-1
	<i>Describes macro directives, substitution symbols used as macro parameters, and how to create macros.</i>	
5.1	Using Macros	5-2
5.2	Defining Macros	5-3
5.3	Macro Parameters/Substitution Symbols	5-5
5.3.1	Directives That Define Substitution Symbols	5-6
5.3.2	Built-In Substitution Symbol Functions	5-7
5.3.3	Recursive Substitution Symbols	5-9
5.3.4	Forced Substitution	5-9
5.3.5	Accessing Individual Characters of Subscripted Substitution Symbols	5-10
5.3.6	Substitution Symbols as Local Variables in Macros	5-12
5.4	Macro Libraries	5-13
5.5	Using Conditional Assembly in Macros	5-14
5.6	Using Labels in Macros	5-16
5.7	Producing Messages in Macros	5-17
5.8	Using Directives to Format the Output Listing	5-19
5.9	Using Recursive and Nested Macros	5-21
5.10	Macro Directives Summary	5-23
6	Archiver Description	6-1
	<i>Describes instructions for invoking the archiver, creating new archive libraries, and modifying existing libraries.</i>	
6.1	Archiver Overview	6-2
6.2	The Archiver's Role in the Software Development Flow	6-3
6.3	Invoking the Archiver	6-4
6.4	Archiver Examples	6-6

7	Linker Description	7-1
	<i>Explains how to invoke the linker, provides details about linker operation, discusses linker directives, and presents a detailed linking example.</i>	
7.1	Linker Overview	7-2
7.2	The Linker's Role in the Software Development Flow	7-3
7.3	Invoking the Linker	7-4
7.4	Linker Options	7-5
7.4.1	Relocation Capabilities (<code>-a</code> and <code>-r</code> Options)	7-6
7.4.2	Create an Absolute Listing File (<code>-abs</code> Option)	7-7
7.4.3	Allocate Memory for Use by the Loader to Pass Arguments (<code>--args</code> Option)	7-8
7.4.4	Disable Merge of Symbolic Debugging Information (<code>-b</code> Option)	7-8
7.4.5	C Language Options (<code>-c</code> and <code>-cr</code> Options)	7-9
7.4.6	Define an Entry Point (<code>-e global_symbol</code> Option)	7-9
7.4.7	Set Default Fill Value (<code>-f fill_value</code> Option)	7-10
7.4.8	Make a Symbol Global (<code>-g symbol</code> Option)	7-10
7.4.9	Make All Global Symbols Static (<code>-h</code> Option)	7-10
7.4.10	Define Heap Size (<code>-heap size</code> Option)	7-11
7.4.11	Alter the Library Search Algorithm (<code>-l</code> Option, <code>-l</code> Option, and <code>C_DIR/C6X_C_DIR</code> Environment Variables)	7-11
7.4.12	Disable Conditional Linking (<code>-j</code> Option)	7-13
7.4.13	Create a Map File (<code>-m filename</code> Option)	7-14
7.4.14	Name an Output Module (<code>-o</code> Option)	7-15
7.4.15	Strip Symbolic Information (<code>-s</code> Option)	7-16
7.4.16	Define Stack Size (<code>-stack size</code> Option)	7-16
7.4.17	Generate Far Call Trampolines (<code>--trampolines</code> Option)	7-17
7.4.18	Introduce an Unresolved Symbol (<code>-u symbol</code> Option)	7-19
7.4.19	Display a Message When an Undefined Output Section Is Created (<code>-w</code> Option)	7-20
7.4.20	Exhaustively Read and Search Libraries (<code>-x</code> and <code>-priority</code> Options)	7-20
7.4.21	Generate XML Link Information File (<code>--xml_link_info</code> Option)	7-21
7.5	Linker Command Files	7-22
7.5.1	Reserved Names in Linker Command Files	7-24
7.5.2	Constants in Linker Command Files	7-24
7.6	Object Libraries	7-25
7.7	The MEMORY Directive	7-27
7.7.1	Default Memory Model	7-27
7.7.2	MEMORY Directive Syntax	7-27

7.8	The SECTIONS Directive	7-31
7.8.1	SECTIONS Directive Syntax	7-31
7.8.2	Allocation	7-34
7.8.3	Specifying Input Sections	7-37
7.8.4	Allocation Using Multiple Memory Ranges	7-40
7.8.5	Automatic Splitting of Output Sections Among Non-Contiguous Memory Ranges	7-41
7.8.6	Allocating an Archive Member to an Output Section	7-43
7.9	Specifying a Section's Run-Time Address	7-44
7.9.1	Specifying Load and Run Addresses	7-44
7.9.2	Uninitialized Sections	7-45
7.9.3	Referring to the Load Address by Using the .label Directive	7-45
7.10	Using UNION and GROUP Statements	7-48
7.10.1	Overlaying Sections With the UNION Statement	7-48
7.10.2	Grouping Output Sections Together	7-50
7.10.3	Nesting UNIONS and GROUPS	7-50
7.10.4	Checking the Consistency of Allocators	7-51
7.11	Special Section Types (DSECT, COPY, and NOLOAD)	7-53
7.12	Default Allocation Algorithm	7-54
7.12.1	How the Allocation Algorithm Creates Output Sections	7-54
7.12.2	Reducing Memory Fragmentation	7-55
7.13	Assigning Symbols at Link Time	7-56
7.13.1	Syntax of Assignment Statements	7-56
7.13.2	Assigning the SPC to a Symbol	7-57
7.13.3	Assignment Expressions	7-57
7.13.4	Symbols Defined by the Linker	7-59
7.13.5	Assigning Exact Start, End, and Size Values of a Section to a Symbol	7-60
7.13.6	Why the Dot Operator Does Not Always Work	7-60
7.13.7	Address and Dimension Operators	7-61
7.14	Creating and Filling Holes	7-64
7.14.1	Initialized and Uninitialized Sections	7-64
7.14.2	Creating Holes	7-64
7.14.3	Filling Holes	7-66
7.14.4	Explicit Initialization of Uninitialized Sections	7-67

7.15	Linker-Generated Copy Tables	7-68
7.15.1	A Current Boot-Loaded Application Development Process	7-68
7.15.2	An Alternative Approach	7-69
7.15.3	Overlay Management Example	7-70
7.15.4	Generating Copy Tables Automatically with the Linker	7-71
7.15.5	The table() Operator	7-72
7.15.6	Boot-Time Copy Tables	7-72
7.15.7	Using the table() Operator to Manage Object Components	7-73
7.15.8	Copy Table Contents	7-74
7.15.9	General Purpose Copy Routine	7-75
7.15.10	Linker Generated Copy Table Sections and Symbols	7-76
7.15.11	Splitting Object Components and Overlay Management	7-78
7.16	Partial (Incremental) Linking	7-80
7.17	Linking C/C++ Code	7-82
7.17.1	Run-Time Initialization	7-82
7.17.2	Object Libraries and Run-Time Support	7-83
7.17.3	Setting the Size of the Stack and Heap Sections	7-83
7.17.4	Autoinitialization of Variables at Run Time	7-84
7.17.5	Initialization of Variables at Load Time	7-85
7.17.6	The -c and -cr Linker Options	7-86
7.18	Linker Example	7-87
8	Absolute Lister Description	8-1
	<i>Explains how to invoke the absolute lister to obtain a listing of the absolute addresses of an object file.</i>	
8.1	Producing an Absolute Listing	8-2
8.2	Invoking the Absolute Lister	8-3
8.3	Absolute Lister Example	8-5
9	Cross-Reference Lister Description	9-1
	<i>Explains how to invoke the cross-reference lister to obtain a listing of symbols, their definitions, and their references in the linked source files.</i>	
9.1	Producing a Cross-Reference Listing	9-2
9.2	Invoking the Cross-Reference Lister	9-3
9.3	Cross-Reference Listing Example	9-4

10	Object File Utilities Descriptions	10-1
	<i>Explains how to invoke the object file display utility , the name utility, and the strip utility.</i>	
10.1	Invoking the Object File Display Utility	10-2
10.2	XML Tag Index	10-3
10.3	Example XML Consumer	10-9
10.3.1	The Main Application	10-9
10.3.2	xml.h Declaration of the XMLEntity Object	10-12
10.3.3	xml.cpp Definition of the XMLEntity Object	10-13
10.4	Invoking the Name Utility	10-16
10.5	Invoking the Strip Utility	10-17
11	Hex Conversion Utility Description	11-1
	<i>Explains how to invoke the hex utility to convert a COFF object file into one of several standard hexadecimal formats suitable for loading into an EPROM programmer.</i>	
11.1	The Hex Conversion Utility's Role in the Software Development Flow	11-2
11.2	Invoking the Hex Conversion Utility	11-3
11.2.1	Invoking the Hex Conversion Utility From the Command Line	11-3
11.2.2	Invoking the Hex Conversion Utility With a Command File	11-5
11.3	Understanding Memory Widths	11-7
11.3.1	Target Width	11-8
11.3.2	Specifying the Memory Width	11-8
11.3.3	Partitioning Data Into Output Files	11-9
11.3.4	Specifying Word Order for Output Words	11-12
11.4	The ROMS Directive	11-13
11.4.1	When to Use the ROMS Directive	11-15
11.4.2	An Example of the ROMS Directive	11-16
11.5	The SECTIONS Directive	11-19
11.6	Excluding a Specified Section	11-21
11.7	Assigning Output Filenames	11-22
11.8	Image Mode and the -fill Option	11-24
11.8.1	Generating a Memory Image	11-24
11.8.2	Specifying a Fill Value	11-25
11.8.3	Steps to Follow in Using Image Mode	11-25
11.9	Building a Table for an On-Chip Boot Loader	11-26
11.9.1	Description of the Boot Table	11-26
11.9.2	The Boot Table Format	11-26
11.9.3	How to Build the Boot Table	11-27
11.9.4	Using the C6x Boot Loader	11-30
11.10	Controlling the ROM Device Address	11-32
11.10.1	Controlling the Starting Address	11-32

11.11	Description of the Object Formats	11-34
11.11.1	ASCII-Hex Object Format (-a Option)	11-35
11.11.2	Intel MCS-86 Object Format (-i Option)	11-36
11.11.3	Motorola Exorciser Object Format (-m Option)	11-37
11.11.4	Texas Instruments SDSMAC Object Format (-t Option)	11-38
11.11.5	Extended Tektronix Object Format (-x Option)	11-39
11.12	Hex Conversion Utility Error Messages	11-40
A	Common Object File Format	A-1
	<i>Contains supplemental technical data about the internal format and structure of COFF object files.</i>	
A.1	COFF File Structure	A-2
A.2	File Header Structure	A-4
A.3	Optional File Header Format	A-5
A.4	Section Header Structure	A-6
A.5	Structuring Relocation Information	A-9
A.6	Symbol Table Structure and Content	A-12
A.6.1	Special Symbols	A-13
A.6.2	Symbol Name Format	A-14
A.6.3	String Table Structure	A-14
A.6.4	Storage Classes	A-15
A.6.5	Symbol Values	A-15
A.6.6	Section Number	A-16
A.6.7	Auxiliary Entries	A-16
B	Symbolic Debugging Directives	B-1
	<i>Discusses symbolic debugging directives that the TMS320C6000 C/C++ compiler uses.</i>	
B.1	DWARF Debugging Format	B-2
B.2	COFF Debugging Format	B-3
B.3	Debug Directive Syntax	B-4
C	XML Link Information File Description	C-1
	<i>Discusses the xml_link_info file contents including file element types and document elements.</i>	
C.1	XML Information File Element Types	C-2
C.2	Document Elements	C-3
C.2.1	Header Elements	C-3
C.2.2	Input File List	C-4
C.2.3	Object Component List	C-5
C.2.4	Logical Group List	C-6
C.2.5	Placement Map	C-9
C.2.6	Far Call Trampoline List	C-10
C.2.7	Symbol Table	C-12
D	Glossary	D-1
	<i>Defines terms and acronyms used in this book.</i>	

Figures

1-1	TMS320C6000 Software Development Flow	1-2
2-1	Partitioning Memory Into Logical Blocks	2-3
2-2	Object Code Generated by the File in Example 2-1	2-10
2-3	Combining Input Sections to Form an Executable Object Module	2-12
3-1	The Assembler in the TMS320C6000 Software Development Flow	3-3
4-1	The .space and .bes Directives	4-10
4-2	The .field Directive	4-11
4-3	Initialization Directives	4-12
4-4	The .align Directive	4-13
4-5	Double-Precision Floating-Point Format	4-36
4-6	The .field Directive	4-43
4-7	Single-Precision Floating-Point Format	4-44
4-8	The .usect Directive	4-77
6-1	The Archiver in the TMS320C6000 Software Development Flow	6-3
7-1	The Linker in the TMS320C6000 Software Development Flow	7-3
7-2	Section Allocation Defined by Example 7-4	7-34
7-3	Run-Time Execution of Example 7-6	7-47
7-4	Memory Allocation Shown in Example 7-7 and Example 7-8	7-49
7-5	Autoinitialization at Run Time	7-84
7-6	Initialization at Load Time	7-85
8-1	Absolute Lister Development Flow	8-2
8-2	module1.lst	8-9
8-3	module2.lst	8-10
9-1	The Cross-Reference Lister in the TMS320C6000 Software Development Flow	9-2
11-1	The Hex Conversion Utility in the TMS320C6000 Software Development Flow	11-2
11-2	Hex Conversion Utility Process Flow	11-7
11-3	COFF Data and Memory Widths	11-9
11-4	Data, Memory, and ROM Widths	11-11
11-5	The infile.out File Partitioned Into Four Output Files	11-16
11-6	ASCII-Hex Object Format	11-35
11-7	Intel Hexadecimal Object Format	11-36
11-8	Motorola-S Format	11-37
11-9	TI-Tagged Object Format	11-38
11-10	Extended Tektronix Object Format	11-39
A-1	COFF File Structure	A-2
A-2	Sample COFF Object File	A-3
A-3	Section Header Pointers for the .text Section	A-8
A-4	Symbol Table Contents	A-12
A-5	String Table Entries for Sample Symbol Names	A-14

Tables

3-1	Operators Used in Expressions (Precedence)	3-27
3-2	Symbol Attributes	3-36
4-1	Assembler Directives Summary	4-2
5-1	Substitution Symbol Functions and Return Values	5-8
5-2	Creating Macros	5-23
5-3	Manipulating Substitution Symbols	5-23
5-4	Conditional Assembly	5-23
5-5	Producing Assembly-Time Messages	5-24
5-6	Formatting the Listing	5-24
7-1	Linker Options Summary	7-5
7-2	Groups of Operators Used in Expressions (Precedence)	7-58
9-1	Symbol Attributes in Cross-Reference Listing	9-5
10-1	XML Tag Index	10-3
11-1	Basic Hex Conversion Utility Options	11-4
11-2	Boot-Loader Options	11-27
11-3	Options for Specifying Hex Conversion Formats	11-34
A-1	File Header Contents	A-4
A-2	File Header Flags (Bytes 18 and 19)	A-4
A-3	Optional File Header Contents	A-5
A-4	Section Header Contents	A-6
A-5	Section Header Flags (Bytes 40 Through 43)	A-7
A-6	Relocation Entry Contents	A-9
A-7	Relocation Types (Bytes 8 and 9)	A-10
A-8	Symbol Table Entry Contents	A-13
A-9	Special Symbols in the Symbol Table	A-13
A-10	Symbol Storage Classes	A-15
A-11	Section Numbers	A-16
A-12	Section Format for Auxiliary Table Entries	A-16
B-1	Symbolic Debugging Directives	B-4

Examples

2-1	Using Sections Directives	2-9
2-2	Code That Generates Relocation Entries	2-14
2-3	Simple Assembler Listing	2-15
3-1	Local Labels of the Form \$n	3-19
3-2	Local Labels of the Form name?	3-20
3-3	Using Symbolic Constants Defined on Command Line	3-22
3-4	Assembler Listing	3-33
3-5	Viewing Assembly Variables as C Types	3-34
3-6	An Assembler Cross-Reference Listing	3-36
4-1	Sections Directives	4-9
5-1	Macro Definition, Call, and Expansion	5-4
5-2	Calling a Macro With Varying Numbers of Arguments	5-6
5-3	The .asg Directive	5-6
5-4	The .eval Directive	5-7
5-5	Using Built-In Substitution Symbol Functions	5-8
5-6	Recursive Substitution	5-9
5-7	Using the Forced Substitution Operator	5-10
5-8	Using Subscripted Substitution Symbols to Redefine an Instruction	5-11
5-9	Using Subscripted Substitution Symbols to Find Substrings	5-11
5-10	The .loop/.break/.endloop Directives	5-15
5-11	Nested Conditional Assembly Directives	5-15
5-12	Built-In Substitution Symbol Functions in a Conditional Assembly Code Block	5-15
5-13	Unique Labels in a Macro	5-16
5-14	Producing Messages in a Macro	5-18
5-15	Using Nested Macros	5-21
5-16	Using Recursive Macros	5-22
7-1	Linker Command File	7-22
7-2	Command File With Linker Directives	7-23
7-3	The MEMORY Directive	7-28
7-4	The SECTIONS Directive	7-33
7-5	The Most Common Method of Specifying Section Contents	7-38
7-6	Copying a Section From SLOW_MEM to FAST_MEM	7-46
7-7	The UNION Statement	7-48
7-8	Separate Load Addresses for UNION Sections	7-48
7-9	Allocate Sections Together	7-50
7-10	Nesting GROUP and UNION Statements	7-50

Examples

7-11	Default Allocation for TMS320C6000 Devices	7-54
7-12	Using a UNION for Memory Overlay	7-70
7-13	Produce Address for Linker Generated Copy Table	7-71
7-14	Linker Command File to Manage Object Components	7-73
7-15	TMS320C6000 cpy_tbl.h File	7-74
7-16	Run-Time-Support cpy_tbl.c File	7-76
7-17	Controlling the Placement of the Linker-Generated Copy Table Sections	7-77
7-18	Creating a Copy Table to Access a Split Object Component	7-78
7-19	Split Object Component Driver	7-79
7-20	Linker Command File, demo.cmd	7-88
7-21	Output Map File, demo.map	7-89
9-1	Cross-Reference Listing	9-4
11-1	A ROMS Directive Example	11-16
11-2	Map File Output From Example 11-1 Showing Memory Ranges	11-17
11-3	Sample Command File for Booting From a C6x EPROM	11-30
11-4	Alternative Sample Command File for Booting From a C6x EPROM	11-31
C-1	Header Element for the hi.out Output File	C-3
C-2	Input File List for the hi.out Output File	C-4
C-3	Object Component List for the fl-4 Input File	C-5
C-4	Logical Group List for the fl-4 Input File	C-8
C-5	Placement Map for the fl-4 Input File	C-10
C-6	Fall Call Trampoline List for the fl-4 Input File	C-11
C-7	Symbol Table for the fl-4 Input File	C-12

Notes

Default Sections Directive	2-4
Expression Can Not Be Larger Than Space Reserved	2-15
Labels and Comments in Not Shown Syntaxes	4-2
Directives That Initialize Constants When Used in a .struct/.endstruct Sequence	4-11
Ending a Macro	4-39
Data Size of longs	4-50
Directives That Can Appear in a .struct/.endstruct Sequence	4-68
Directives That Can Appear in a .union/.endunion Sequence	4-74
Naming Library Members	6-5
The -a and -r Options	7-6
Filling Memory Ranges	7-29
Binding is Incompatible With Alignment and Named Memory	7-36
Linker Command File Operator Equivalencies	7-61
Filling Sections	7-67
Object File Display Format	10-2
The TI-Tagged Format Is 16 Bits Wide	11-10
When the -order Option Applies	11-12
Sections Generated by the C/C++ Compiler	11-19
Using the -boot Option and the SECTIONS Directive	11-20
Defining the Ranges of Target Memory	11-24
Valid Entry Points	11-29



Introduction to the Software Development Tools

The TMS320C6000™ is supported by a set of software development tools, which includes an optimizing C/C++ compiler, an assembly optimizer, an assembler, a linker, and assorted utilities. This chapter provides an overview of these tools.

The TMS320C6000 is supported by the following assembly language development tools:

- Assembler
- Archiver
- Linker
- Absolute lister
- Cross-reference lister
- Object file display utility
- Name utility
- Strip utility
- Hex conversion utility

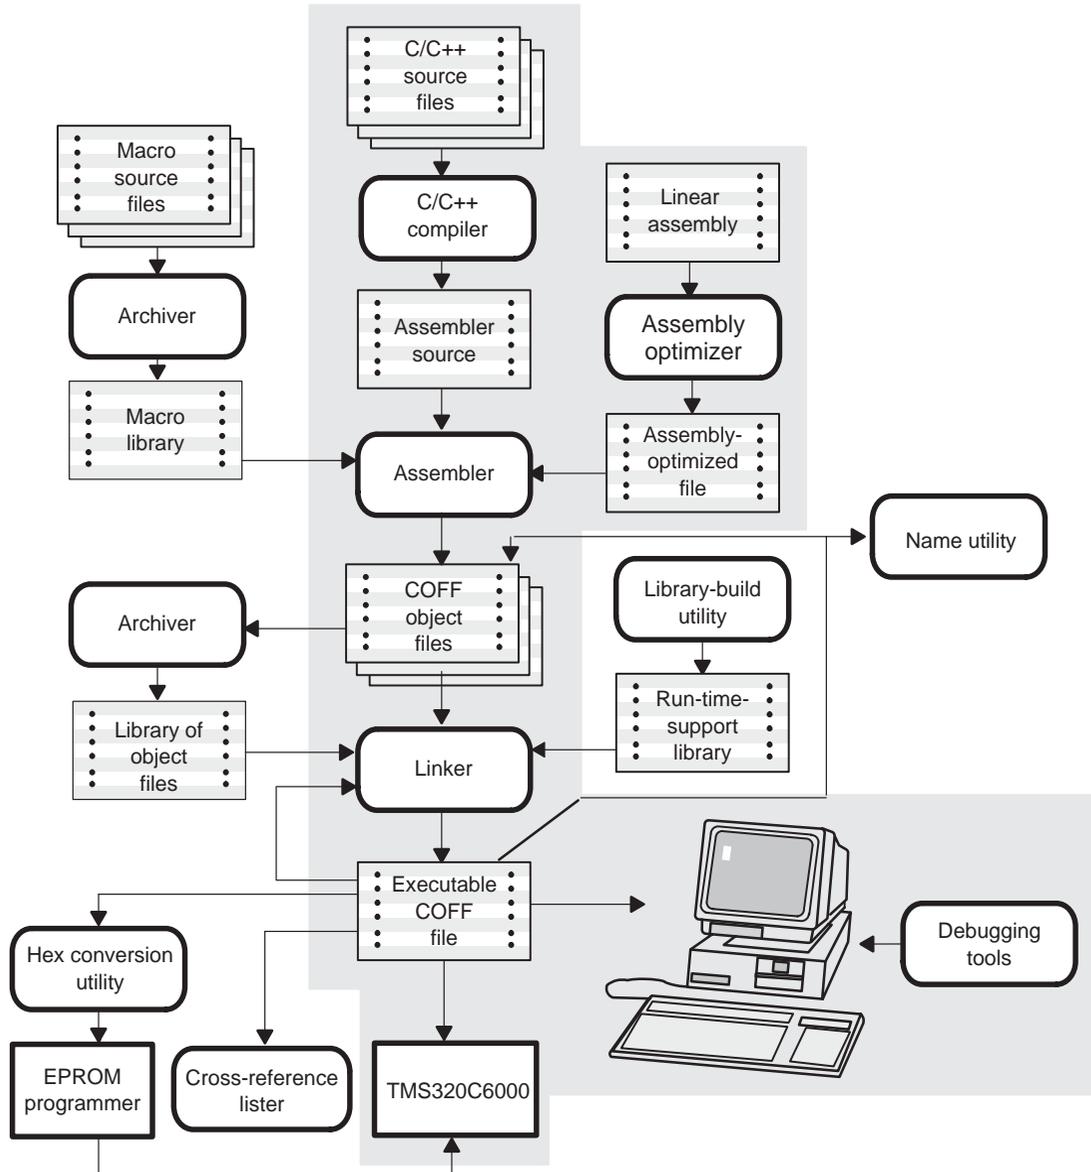
This chapter shows how these tools fit into the general software tools development flow and gives a brief description of each tool. For convenience, it also summarizes the C/C++ compiler and debugging tools. For detailed information on the compiler and debugger, and for complete descriptions of the TMS320C6000, refer to books listed in *Related Documentation From Texas Instruments* on page vii.

Topic	Page
1.1 Software Development Tools Overview	1-2
1.2 Tools Descriptions	1-3
1.3 Code Composer Studio and the Code Generation Tools	1-5

1.1 Software Development Tools Overview

Figure 1–1 shows the TMS320C6000 software development flow. The shaded portion highlights the most common development path; the other portions are optional. The other portions are peripheral functions that enhance the development process.

Figure 1–1. TMS320C6000 Software Development Flow



1.2 Tools Descriptions

The following list describes the tools that are shown in Figure 1–1:

- The **assembly optimizer** allows you to write linear assembly code without being concerned with the pipeline structure or with assigning registers. It assigns registers and uses loop optimization to turn linear assembly into highly parallel assembly that takes advantage of software pipelining.

See the *TMS320C6000 Optimizing Compiler User's Guide* for more information.

- The **C/C++ compiler** accepts C/C++ source code and produces TMS320C6000 assembly language source code. A **shell program**, an **optimizer**, and an **interlist utility** are included in the compiler package:
 - The shell program enables you to compile, assemble, and link source modules in one step.
 - The optimizer modifies code to improve the efficiency of C/C++ programs.
 - The interlist utility interlists C/C++ source statements with assembly language output to correlate code produced by the compiler with your source code.

See the *TMS320C6000 Optimizing Compiler User's Guide* for more information.

- The **assembler** translates assembly language source files into machine language COFF object files. Source files can contain instructions, assembler directives, and macro directives. You can use assembler directives to control various aspects of the assembly process, such as the source listing format, data alignment, and section content. See Chapter 3, *Assembler Description*, through Chapter 5, *Macro Language*, for more information. See the *TMS320C62x/64x/67x CPU and Instruction Set Reference Guide* for detailed information on the assembly language instruction set.
- The **linker** combines object files into a single executable COFF object module. As it creates the executable module, it performs relocation and resolves external references. The linker accepts relocatable COFF object files (created by the assembler) as input. It also accepts archiver library members and output modules created by a previous linker run. Linker directives allow you to combine object file sections, bind sections or symbols to addresses or within memory ranges, and define or redefine global symbols. See Chapter 7, *Linker Description*, for more information.

- ❑ The **archiver** allows you to collect a group of files into a single archive file, called a library. For example, you can collect several macros into a macro library. The assembler searches the library and uses the members that are called as macros by the source file. You can also use the archiver to collect a group of object files into an object library. The linker includes in the library the members that resolve external references during the link. The archiver allows you to modify a library by deleting, replacing, extracting, or adding members. See Chapter 6, *Archiver Description*, for more information.
- ❑ You can use the **library-build utility** to build your own customized run-time-support library. See the *TMS320C6000 Optimizing Compiler User's Guide* for more information.
- ❑ The **hex conversion utility** converts a COFF object file into TI-Tagged, ASCII-Hex, Intel, Motorola-S™, or Tektronix™ object format. The converted file can be downloaded to an EPROM programmer. See Chapter 11, *Hex Conversion Utility Description*, for more information.
- ❑ The **cross-reference lister** uses object files to produce a cross-reference listing showing symbols, their definition, and their references in the linked source files. See Chapter 9, *Cross-Reference Lister Description*, for more information.
- ❑ The main product of this development process is a module that can be executed in a **TMS320C6000** device. You can use one of several debugging tools to refine and correct your code. Available products include:
 - An instruction-accurate and clock-accurate software simulator
 - An XDS emulator

For information about these debugging tools, see the *Code Composer Studio User's Guide*.

1.3 Code Composer Studio and the Code Generation Tools

Code Composer Studio provides a graphical interface for using the code generation tools.

A Code Composer Studio project manages the information needed to build a target program or library. A project records:

- Filenames of source code and object libraries
- Compiler, assembler, and linker options
- Include file dependencies

When you build a project with Code Composer Studio, the appropriate code generation tools are invoked to compile, assemble, and/or link your program.

Compiler, assembler, and linker options can be specified within Code Composer Studio's Build Options dialog. Nearly all command line options are represented within this dialog. Options that are not represented can be specified by typing the option directly into the editable text box that appears at the top of the dialog.

The information in this book describes how to use the code generation tools from the command line interface. For information on using Code Composer Studio, consult the *Code Composer Studio User's Guide*. For information on setting code generation tool options within Code Composer Studio, consult the *Code Generation Tools Help*.



Introduction to Common Object File Format

The assembler and linker create object files that can be executed by a TMS320C6000™ device. The format for these object files is called common object file format (COFF).

COFF makes modular programming easier because it encourages you to think in terms of *blocks* of code and data when you write an assembly language program. These blocks are known as sections. Both the assembler and the linker provide directives that allow you to create and manipulate sections.

This chapter focuses on the concept and use of sections in assembly language programs. See Appendix A, *Common Object File Format*, for details about COFF object file structure.

Topic	Page
2.1 Sections	2-2
2.2 How the Assembler Handles Sections	2-4
2.3 How the Linker Handles Sections	2-11
2.4 Relocation	2-14
2.5 Run-Time Relocation	2-16
2.6 Loading a Program	2-17
2.7 Symbols in a COFF File	2-18

2.1 Sections

The smallest unit of an object file is called a *section*. A section is a block of code or data that occupies contiguous space in the memory map with other sections. Each section of an object file is separate and distinct. COFF object files always contain three default sections:

.text section	usually contains executable code
.data section	usually contains initialized data
.bss section	usually reserves space for uninitialized variables

In addition, the assembler and linker allow you to create, name, and link *named* sections that are used like the .data, .text, and .bss sections.

There are two basic types of sections:

Initialized sections contain data or code. The .text and .data sections are initialized; named sections created with the .sect assembler directive are also initialized.

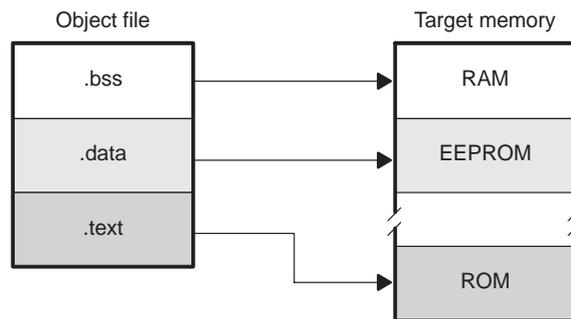
Uninitialized sections reserve space in the memory map for uninitialized data. The .bss section is uninitialized; named sections created with the .usect assembler directive are also uninitialized.

Several assembler directives allow you to associate various portions of code and data with the appropriate sections. The assembler builds these sections during the assembly process, creating an object file organized as shown in Figure 2–1.

One of the linker's functions is to relocate sections into the target system's memory map; this function is called *allocation*. Because most systems contain several types of memory, using sections can help you use target memory more efficiently. All sections are independently relocatable; you can place any section into any allocated block of target memory. For example, you can define a section that contains an initialization routine and then allocate the routine into a portion of the memory map that contains ROM.

Figure 2–1 shows the relationship between sections in an object file and a hypothetical target memory.

Figure 2–1. Partitioning Memory Into Logical Blocks



2.2 How the Assembler Handles Sections

The assembler identifies the portions of an assembly language program that belong in a given section. The assembler has five directives that support this function:

- .bss
- .usect
- .text
- .data
- .sect

The .bss and .usect directives create *uninitialized sections*; the .text, .data, and .sect directives create *initialized sections*.

You can create subsections of any section to give you tighter control of the memory map. Subsections are created using the .sect and .usect directives. Subsections are identified with the base section name and a subsection name separated by a colon. See section 2.2.4, *Subsections*, on page 2-7, for more information.

Note: Default Sections Directive

If you do not use any of the sections directives, the assembler assembles everything into the .text section.

2.2.1 Uninitialized Sections

Uninitialized sections reserve space in TMS320C6000 memory; they are usually allocated into RAM. These sections have no actual contents in the object file; they simply reserve memory. A program can use this space at run-time for creating and storing variables.

Uninitialized data areas are built by using the .bss and .usect assembler directives.

- The .bss directive reserves space in the .bss section.
- The .usect directive reserves space in a specific uninitialized named section.

Each time you invoke the .bss or .usect directive, the assembler reserves additional space in the .bss or the named section.

The syntaxes for these directives are:

```
.bss symbol, size in bytes [, alignment [, bank offset]]
symbol .usect "section name", size in bytes [, alignment [, bank offset]]
```

<i>symbol</i>	points to the first byte reserved by this invocation of the <code>.bss</code> or <code>.usect</code> directive. The <i>symbol</i> corresponds to the name of the variable that you are reserving space for. It can be referenced by any other section and can also be declared as a global symbol (with the <code>.global</code> assembler directive).
<i>size in bytes</i>	is an absolute expression. <ul style="list-style-type: none"> <input type="checkbox"/> The <code>.bss</code> directive reserves <i>size in bytes</i> bytes in the <code>.bss</code> section. You must specify a size; there is no default value. <input type="checkbox"/> The <code>.usect</code> directive reserves <i>size in bytes</i> bytes in <i>section name</i>. You must specify a size; there is no default value.
<i>alignment</i>	is an optional parameter. It specifies the minimum alignment in bytes required by the space allocated. The default value is byte aligned. The value must be power of 2.
<i>bank offset</i>	is an optional parameter. It ensures that the space allocated to the symbol occurs on a specific memory bank boundary. The <i>bank offset</i> measures the number of bytes to offset from the alignment specified before assigning the symbol to that location.
<i>section name</i>	tells the assembler which named section to reserve space in. For more information, see section 2.2.3, <i>Named Sections</i> .

The initialized section directives (`.text`, `.data`, and `.sect`) tell the assembler to stop assembling into the current section and begin assembling into the indicated section. The `.bss` and `.usect` directives, however, *do not* end the current section and begin a new one; they simply escape from the current section temporarily. The `.bss` and `.usect` directives can appear anywhere in an initialized section without affecting its contents. For an example, see section 2.2.6, *Using Sections Directives*, on page 2-8.

The assembler treats uninitialized subsections (created with the `.usect` directive) in the same manner as uninitialized sections. See section 2.2.4, *Subsections*, on page 2-7 for more information on creating subsections.

2.2.2 Initialized Sections

Initialized sections contain executable code or initialized data. The contents of these sections are stored in the object file and placed in TMS320C6000 memory when the program is loaded. Each initialized section is independently relocatable and may reference symbols that are defined in other sections. The linker automatically resolves these section-relative references.

Three directives tell the assembler to place code or data into a section. The syntaxes for these directives are:

```
.text  
.data  
.sect "section name"
```

When the assembler encounters one of these directives, it stops assembling into the current section (acting as an implied end of current section command). It then assembles subsequent code into the designated section until it encounters another `.text`, `.data`, or `.sect` directive.

Sections are built through an iterative process. For example, when the assembler first encounters a `.data` directive, the `.data` section is empty. The statements following this first `.data` directive are assembled into the `.data` section (until the assembler encounters a `.text` or `.sect` directive). If the assembler encounters subsequent `.data` directives, it adds the statements following these `.data` directives to the statements already in the `.data` section. This creates a single `.data` section that can be allocated continuously into memory.

Initialized subsections are created with the `.sect` directive. The assembler treats initialized subsections in the same manner as initialized sections. See section 2.2.4, on page 2-7 for more information on creating subsections.

2.2.3 Named Sections

Named sections are sections that *you* create. You can use them like the default `.text`, `.data`, and `.bss` sections, but they are assembled separately.

For example, repeated use of the `.text` directive builds up a single `.text` section in the object file. When linked, this `.text` section is allocated into memory as a single unit. Suppose there is a portion of executable code (perhaps an initialization routine) that you do not want allocated with `.text`. If you assemble this segment of code into a named section, it is assembled separately from `.text`, and you can allocate it into memory separately. You can also assemble initialized data that is separate from the `.data` section, and you can reserve space for uninitialized variables that is separate from the `.bss` section.

Two directives let you create named sections:

- The **.usect** directive creates uninitialized sections that are used like the `.bss` section. These sections reserve space in RAM for variables.
- The **.sect** directive creates initialized sections, like the default `.text` and `.data` sections, that can contain code or data. The `.sect` directive creates named sections with relocatable addresses.

The syntaxes for these directives are:

```
symbol .usect "section name", size in bytes [, alignment[, bank offset]]
      .sect "section name"
```

The *section name* parameter is the name of the section. Section names are significant to 200 characters. You can create up to 32 767 separate named sections. For the `.usect` and `.sect` directives, a section name can refer to a subsection; see section 2.2.4 for details.

Each time you invoke one of these directives with a new name, you create a new named section. Each time you invoke one of these directives with a name that was already used, the assembler assembles code or data (or reserves space) into the section with that name. *You cannot use the same names with different directives.* That is, you cannot create a section with the `.usect` directive and then try to use the same section with `.sect`.

2.2.4 Subsections

Subsections are smaller sections within larger sections. Like sections, subsections can be manipulated by the linker. Subsections give you tighter control of the memory map. You can create subsections by using the `.sect` or `.usect` directive. The syntaxes for a subsection name are:

```
symbol .usect "section name:subsection name", size in bytes
      [, alignment[, bank offset]]
      .sect "section name:subsection name"
```

A subsection is identified by the base section name followed by a colon and the name of the subsection. A subsection can be allocated separately or grouped with other sections using the same base name. For example, you create a subsection called `_func` within the `.text` section:

```
.sect ".text:_func"
```

Using the linker's `SECTIONS` directive, you can allocate `.text:_func` separately, or with all the `.text` sections. See section 7.8.1, *SECTIONS Directive Syntax*, on page 7-31, for an example using subsections.

You can create two types of subsections:

- ❑ Initialized subsections are created using the `.sect` directive. See section 2.2.2, *Initialized Sections*, on page 2-6.
- ❑ Uninitialized subsections are created using the `.usect` directive. See section 2.2.1, *Uninitialized Sections*, on page 2-4.

Subsections are allocated in the same manner as sections. See section 7.8, *The SECTIONS Directive*, on page 7-31, for more information.

2.2.5 Section Program Counters

The assembler maintains a separate program counter for each section. These program counters are known as *section program counters*, or *SPCs*.

An SPC represents the current address within a section of code or data. Initially, the assembler sets each SPC to 0. As the assembler fills a section with code or data, it increments the appropriate SPC. If you resume assembling into a section, the assembler remembers the appropriate SPC's previous value and continues incrementing the SPC from that value.

The assembler treats each section as if it began at address 0; the linker relocates each section according to its final location in the memory map. For more information, see section 2.4, *Relocation*, on page 2-14.

2.2.6 Using Sections Directives

Example 2-1 shows how you can build COFF sections incrementally, using the sections directives to swap back and forth between the different sections. You can use sections directives to begin assembling into a section for the first time, or to continue assembling into a section that already contains code. In the latter case, the assembler simply appends the new code to the code that is already in the section.

The format in Example 2-1 is a listing file. Example 2-1 shows how the SPCs are modified during assembly. A line in a listing file has four fields:

- Field 1** contains the source code line counter.
- Field 2** contains the section program counter.
- Field 3** contains the object code.
- Field 4** contains the original source statement.

See section 3.10, *Source Listings*, on page 3-31 for more information on interpreting the fields in a source listing.

Example 2-1. Using Sections Directives

```

1          ****
2          ** Assemble an initialized table into .data. **
3          ****
4 00000000          .data
5 00000000 00000011 coeff .word 011h,022h
6 00000004 00000022
7          ****
8          ** Reserve space in .bss for a variable. **
9          ****
9 00000000          .bss var1,4
10 00000004         .bss buffer,40
11          ****
12          ** Still in .data section **
13          ****
14 00000008 00001234 ptr .word 01234h
15          ****
16          ** Assemble code into .text section **
17          ****
18 00000000          .text
19 00000000 00800528 sum: MVK 10,A1
20 00000004 021085E0     ZERO A4
21
22 00000008 01003664 aloop: LDW *A0++,A2
23 0000000c 00004000     NOP 3
24 00000010 0087E1A0     SUB A1,1,A1
25 00000014 021041E0     ADD A2,A4,A4
26 00000018 80000112 [A1] B aloop
27 0000001c 00008000     NOP 5
28
29 00000020 0200007C-    STW A4, *+B14(var1)
30          ****
31          ** Assemble another initialized table in .data **
32          ****
33 0000000c          .data
34 0000000c 000000AA ivals .word 0aah, 0bbh, 0cch
35 00000010 000000BB
36 00000014 000000CC
37          ****
38          ** Define another section for more variables. **
39          ****
38 00000000         var2 .usect "newvars",4
39 00000004         inbuf .usect "newvars",4
40          ****
41          ** Assemble more code into the .text section. **
42          ****
43 00000024          .text
44 00000024 01003664 xmult: LDW *A0++,A2
45 00000028 00006000     NOP 4
46 0000002c 020C4480     MPYHL A2,A3,A4
47 00000030 02800028-    MVKL var2,A5
48 00000034 02800068-    MVKH var2,A5
49 00000038 02140274     STW A4,*A5
50          ****
51          ** Define a named section for interrupt vectors **
52          ****
53 00000000          .sect "vectors"
54 00000000 00000012' B sum
55 00000004 00008000     NOP 5

```

Field 1 Field 2 Field 3 Field 4

As Figure 2–2 shows, the file in Example 2–1 creates five sections:

- .text** contains 15 32-bit words of object code.
- .data** contains six words of initialized data.
- vectors** is a named section created with the `.sect` directive; it contains two words of object code.
- .bss** reserves 44 bytes in memory.
- newvars** is a named section created with the `.usect` directive; it contains eight bytes in memory.

The second column shows the object code that is assembled into these sections; the first column shows the source statements that generated the object code.

Figure 2–2. Object Code Generated by the File in Example 2–1

Line numbers	Object code	Section
19	00800528	.text
20	021085E0	
22	01003664	
23	00004000	
24	0087E1A0	
25	021041E0	
26	80000112	
27	00008000	
29	0200007C-	
44	01003664	
45	00006000	
46	020C4480	
47	02800028-	
48	02800068-	
49	02140274	
5	00000011	.data
5	00000022	
14	00001234	
34	000000AA	
34	000000BB	
34	000000CC	
54	00000000'	vectors
54	00000024'	
9	No data—	.bss
10	44 bytes reserved	
38	No data—	newvars
39	8 bytes reserved	

2.3 How the Linker Handles Sections

The linker has two main functions related to sections. First, the linker uses the sections in COFF object files as building blocks; it combines input sections (when more than one file is being linked) to create output sections in an executable COFF output module. Second, the linker chooses memory addresses for the output sections.

Two linker directives support these functions:

- ❑ The *MEMORY* directive allows you to define the memory map of a target system. You can name portions of memory and specify their starting addresses and their lengths.
- ❑ The *SECTIONS* directive tells the linker how to combine input sections into output sections and where to place these output sections in memory.

Subsections allow you to manipulate sections with greater precision. You can specify subsections with the linker's *SECTIONS* directive. If you do not specify a subsection explicitly, then the subsection is combined with the other sections with the same base section name.

It is not always necessary to use linker directives. If you do not use them, the linker uses the target processor's default allocation algorithm described in section 7.12, *Default Allocation Algorithm*. When you *do* use linker directives, you must specify them in a linker command file.

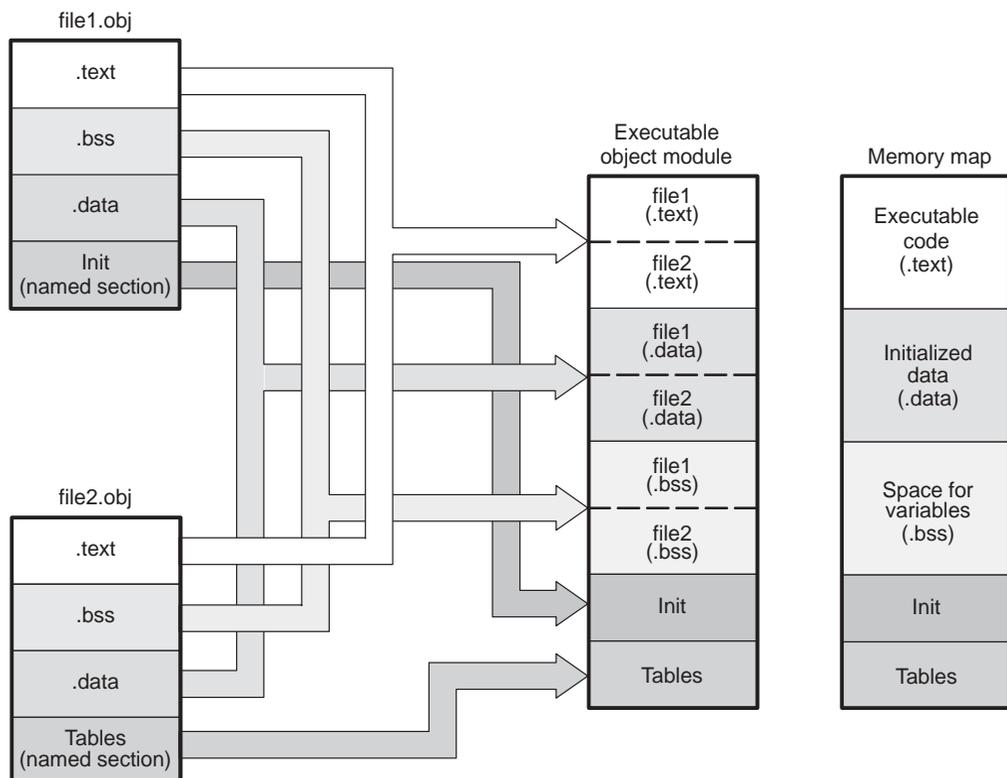
Refer to the following sections for more information about linker command files and linker directives:

Section	Page
7.5 Linker Command Files	7-22
7.7 The MEMORY Directive	7-27
7.8 The SECTIONS Directive	7-31
7.12 Default Allocation Algorithm	7-54

2.3.1 Default Memory Allocation

Figure 2–3 illustrates the process of linking two files together.

Figure 2–3. Combining Input Sections to Form an Executable Object Module



In Figure 2–3, file1.obj and file2.obj have been assembled to be used as linker input. Each contains the .text, .data, and .bss default sections; in addition, each contains a named section. The executable object module shows the combined sections. The linker combines the .text section from file1.obj and the .text section from file2.obj to form one .text section, then combines the two .data sections and the two .bss sections, and finally places the named sections at the end. The memory map shows how the sections are put into memory; by default, the linker begins at 0h and places the sections one after the other in the following order: .text, .const, .data, .bss, .cinit, and then any named sections in the order they are encountered in the input files.

The C/C++ compiler uses the .const section to store string constants, and variables or arrays that are defined as *far const*. The C/C++ compiler produces tables of data for autoinitializing global variables; these variables are stored in a named section called .cinit (see Figure 7–5 on page 7-84). For more information on the .const and .cinit sections, see the *TMS320C6000 Optimizing Compiler User's Guide*.

2.3.2 Placing Sections in the Memory Map

Figure 2–3 illustrates the linker's default method for combining sections. Sometimes you may not want to use the default setup. For example, you may not want all of the .text sections to be combined into a single .text section. Or you may want a named section placed where the .data section would normally be allocated. Most memory maps contain various types of memory (RAM, ROM, EPROM, etc.) in varying amounts; you may want to place a section in a specific type of memory.

For further explanation of section placement within the memory map, see the discussions in section 7.7, *The MEMORY Directive*, on page 7-27, and section 7.8, *The SECTIONS Directive*, on page 7-31.

2.4 Relocation

The assembler treats each section as if it began at address 0. All relocatable symbols (labels) are relative to address 0 in their sections. Of course, all sections cannot actually begin at address 0 in memory, so the linker *relocates* sections by:

- Allocating them into the memory map so that they begin at the appropriate address as defined with the linker's MEMORY directive
- Adjusting symbol values to correspond to the new section addresses
- Adjusting references to relocated symbols to reflect the adjusted symbol values

The linker uses *relocation entries* to adjust references to symbol values. The assembler creates a relocation entry each time a relocatable symbol is referenced. The linker then uses these entries to patch the references after the symbols are relocated. Example 2–2 contains a code segment for a TMS320C6000 device that generates relocation entries.

Example 2–2. Code That Generates Relocation Entries

```

1          .global X
2 00000000 00000012! Z:      B      X          ; Uses an external relocation
3 00000004 0180082A'      MVKL   Y,B3       ; Uses an internal relocation
4 00000008 0180006A'      MVKH   Y,B3       ; Uses an internal relocation
5 0000000c 00004000      NOP      3
6
7 00000010 0001E000 Y:      IDLE
8 00000014 00000212      B      Y
9 00000018 00008000      NOP      5

```

In Example 2–2, both symbols X and Y are relocatable. Y is defined in the .text section of this module; X is defined in another module. When the code is assembled, X has a value of 0 (the assembler assumes all undefined external symbols have values of 0), and Y has a value of 16 (relative to address 0 in the .text section). The assembler generates two relocation entries: one for X and one for Y. The reference to X is an external reference (indicated by the ! character in the listing). The reference to Y is to an internally defined relocatable symbol (indicated by the ' character in the listing).

After the code is linked, suppose that X is relocated to address 0x7100. Suppose also that the .text section is relocated to begin at address 0x7200; Y now has a relocated value of 0x7210. The linker uses the two relocation entries to patch the two references in the object code:

```
00000012  B      X      becomes  0ffe012
0180082A  MVKL   Y      becomes  01B9082A
0180006A  MVKH   Y      becomes  1860006A
```

Sometimes an expression contains more than one relocatable symbol, or cannot be evaluated at assembly time. In this case, the assembler encodes the entire expression in the object file. After determining the addresses of the symbols, the linker computes the value of the expression. For example:

Example 2–3. Simple Assembler Listing

```
1                                     .global  sym1, sym2
2
3  00000000 00800028%      MVKL      sym2 - sym1, A1
```

The symbols sym1 and sym2 are both externally defined. Therefore, the assembler cannot evaluate the expression sym2 – sym1, so it encodes the expression in the object file. The '%' listing character indicates a relocation expression. Suppose the linker relocates sym2 to 300h and sym1 to 200h. Then the linker computes the value of the expression to be 300h – 200h = 100h. Thus the MVKL instruction is patched to:

```
00808028      MVKL      100h,A1
```

Note: Expression Can Not Be Larger Than Space Reserved

If the value of an expression is larger, in bits, than the space reserved for it, you will receive an error message from the linker.

Each section in a COFF object file has a table of relocation entries. The table contains one relocation entry for each relocatable reference in the section. The linker usually removes relocation entries after it uses them. This prevents the output file from being relocated again (if it is relinked or when it is loaded). A file that contains no relocation entries is an *absolute* file (all its addresses are absolute addresses). If you want the linker to retain relocation entries, invoke the linker with the `-r` option (see page 7-7).

2.5 Run-Time Relocation

At times you may want to load code into one area of memory and run it in another. For example, you may have performance-critical code in an external-memory-based system. The code must be loaded into external memory, but it would run faster in internal memory.

The linker provides a simple way to handle this. Using the `SECTIONS` directive, you can optionally direct the linker to allocate a section twice: first to set its load address and again to set its run address. Use the `load` keyword for the load address and the `run` keyword for the run address.

The load address determines where a loader places the raw data for the section. Any references to the section (such as references to labels in it) refer to its run address. The application must copy the section from its load address to its run address before the first reference of the symbol is encountered at run time; this does *not* happen automatically simply because you specify a separate run address. For an example that illustrates how to move a block of code at run-time, see Example 7-6 on page 7-46.

If you provide only one allocation (either load or run) for a section, the section is allocated only once and loads and runs at the same address. If you provide both allocations, the section is actually allocated as if it were two separate sections of the same size.

Uninitialized sections (such as `.bss`) are not loaded, so the only significant address is the run address. The linker allocates uninitialized sections only once; if you specify both run and load addresses, the linker warns you and ignores the load address.

For a complete description of run-time relocation, see section 7.9, *Specifying a section's Run-Time Address*, on page 7-44.

2.6 Loading a Program

The linker produces executable COFF object modules. An executable object file has the same COFF format as object files that are used as linker input; the sections in an executable object file, however, are combined and relocated into target memory.

To run a program, the data in the executable object module must be transferred, or loaded, into target system memory. Several methods can be used for loading a program, depending on the execution environment. Three common situations are described below:

- Code Composer Studio can load an executable COFF file into a simulator or onto hardware. The CCS loader reads the executable file and copies the program into target memory.
- You can use the hex conversion utility (`hex6x`, which is shipped as part of the assembly language package) to convert the executable COFF object module into one of several object file formats. You can then use the converted file with an EPROM programmer to burn the program into an EPROM.
- A standalone simulator can be invoked by the `load6x` command and the name of the executable object file. The standalone simulator reads the executable file, copies the program into the simulator and executes it, displaying any C I/O.

2.7 Symbols in a COFF File

A COFF file contains a symbol table that stores information about symbols in the program. The linker uses this table when it performs relocation. Debugging tools can also use the symbol table to provide symbolic debugging.

2.7.1 External Symbols

External symbols are symbols that are defined in one module and referenced in another module. You can use the `.def`, `.ref`, or `.global` directive to identify symbols as external:

.def	The symbol is defined in the current module and used in another module.
.ref	The symbol is referenced in the current module, but defined in another module.
.global	The symbol may be either of the above.

The following code segment illustrates these definitions.

```
.def    x
.ref    y
.global z
.global q

q:  B      B3
   NOP    4
   MVK    1, B1
x:  MV     A0, A1
   MVKL   y, B3
   MVKH   y, B3
   B      z
   NOP    5
```

In this example, the `.def` definition of `x` says that it is an external symbol defined in this module and that other modules can reference `x`. The `.ref` definition of `y` says that it is an undefined symbol that is defined in another module. The `.global` definition of `z` says that it is defined in some module and available in this file. The `.global` definition of `q` says that it is defined in this module and that other modules can reference `q`.

The assembler places `x`, `y`, `z`, and `q` in the object file's symbol table. When the file is linked with other object files, the entries for `x` and `q` resolve references to `x` and `q` in other files. The entries for `y` and `z` cause the linker to look through the symbol tables of other files for `y`'s and `z`'s definitions.

The linker must match all references with corresponding definitions. If the linker cannot find a symbol's definition, it prints an error message about the unresolved reference. This type of error prevents the linker from creating an executable object module.

2.7.2 The Symbol Table

The assembler always generates an entry in the symbol table when it encounters an external symbol (both definitions and references defined by one of the directives in section 2.7.1). The assembler also creates special symbols that point to the beginning of each section; the linker uses these symbols to relocate references to other symbols.

The assembler does not usually create symbol table entries for any symbols other than those described above, because the linker does not use them. For example, labels are not included in the symbol table unless they are declared with the `.global` directive. For informational purposes, it is sometimes useful to have entries in the symbol table for each symbol in a program. To accomplish this, invoke the assembler with the `-as` option (see page 3-5).



Assembler Description

The TMS320C6000™ assembler translates assembly language source files into machine language object files. These files are in common object file format (COFF), which is discussed in Chapter 2, *Introduction to Common Object File Format*, and Appendix A, *Common Object File Format*. Source files can contain the following assembly language elements:

Assembler directives	described in Chapter 4
Macro directives	described in Chapter 5
Assembly language instructions	described in the <i>TMS320C6000 CPU and Instruction Set Reference Guide</i>

Topic	Page
3.1 Assembler Overview	3-2
3.2 The Assembler's Role in the Software Development Flow	3-3
3.3 Invoking the Assembler	3-4
3.4 Naming Alternate Directories for Assembler Input	3-7
3.5 Source Statement Format	3-10
3.6 Constants	3-14
3.7 Character Strings	3-17
3.8 Symbols	3-18
3.9 Expressions	3-26
3.10 Source Listings	3-31
3.11 Debugging Assembly Source	3-34
3.12 Cross-Reference Listings	3-36

3.1 Assembler Overview

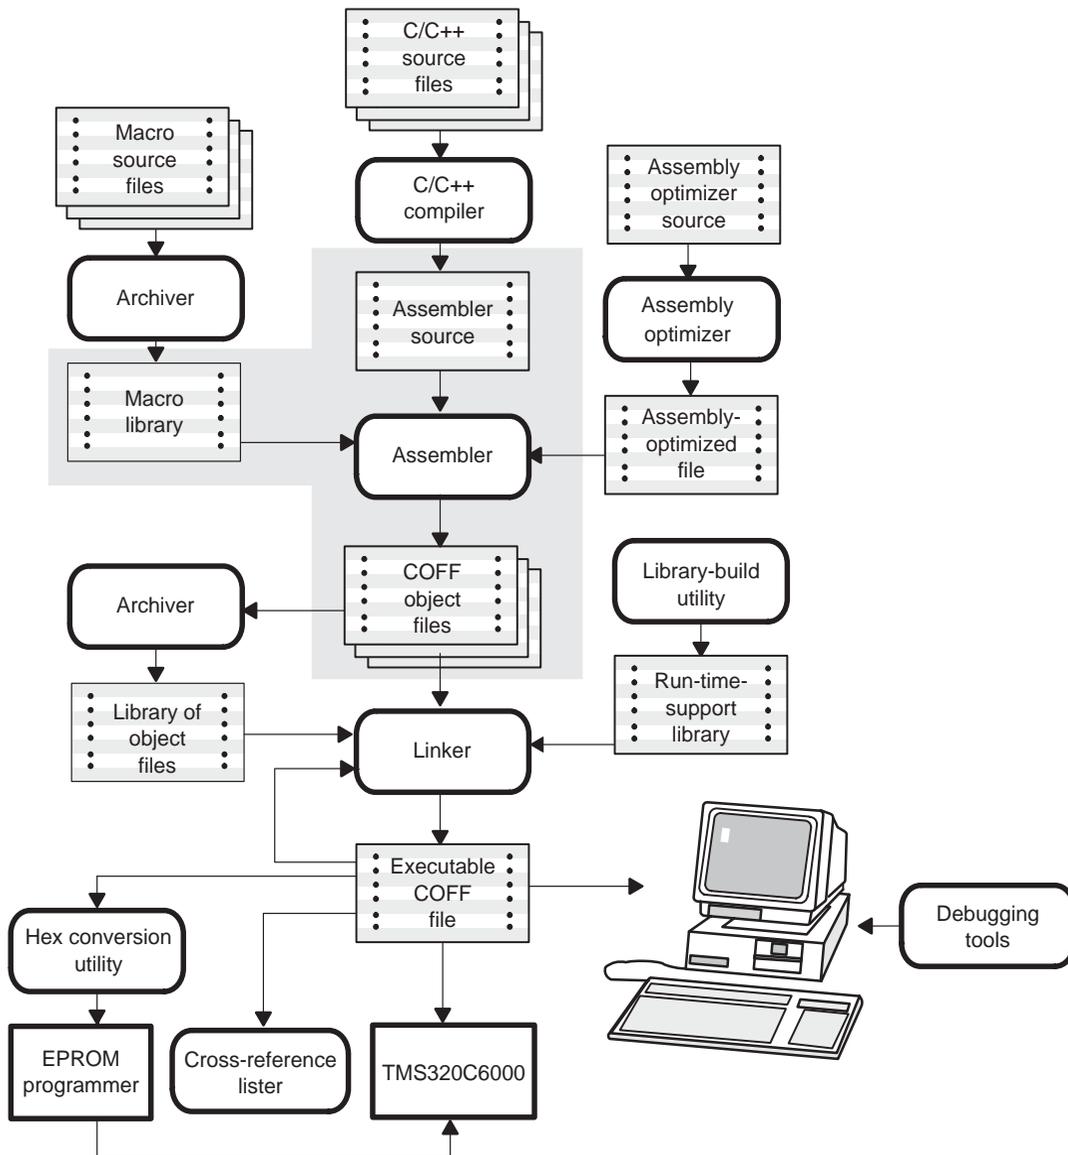
The 2-pass assembler does the following:

- Processes the source statements in a text file to produce a relocatable object file
- Produces a source listing (if requested) and provides you with control over this listing
- Allows you to segment your code into sections and maintain a section program counter (SPC) for each section of object code
- Defines and references global symbols and appends a cross-reference listing to the source listing (if requested)
- Allows conditional assembly
- Supports macros, allowing you to define macros inline or in a library

3.2 The Assembler's Role in the Software Development Flow

Figure 3-1 illustrates the assembler's role in the software development flow. The shaded portion highlights the most common assembler development path. The assembler accepts assembly language source files as input, both those you create and those created by the TMS320C6000 C/C++ compiler.

Figure 3-1. The Assembler in the TMS320C6000 Software Development Flow



3.3 Invoking the Assembler

To invoke the assembler, enter the following:

```
cl6x input file [options]
```

cl6x is the command that invokes the assembler through the compiler. The compiler considers any file with an .asm extension to be an assembly file and calls the assembler.

input file names the assembly language source file.

options identify the assembler options that you want to use. Options are not case sensitive and can appear anywhere on the command line following the command. Precede each option with a hyphen.

The valid assembler options are as follows:

-@ *-@ filename* appends the contents of a file to the command line. You can use this option to avoid limitations on command line length imposed by the host operating system. Use an asterisk or a semicolon (* or ;) at the beginning of a line in the command file to include comments. Comments that begin in any other column must begin with a semicolon.

Within the command file, filenames or option parameters containing embedded spaces or hyphens must be surrounded with quotation marks. For example: "this-file.asm"

-aa creates an absolute listing. When you use **-aa**, the assembler does not produce an object file. The **-aa** option is used in conjunction with the absolute lister.

-ac makes case insignificant in the assembly language files. For example, **-ac** makes the symbols ABC and abc equivalent. *If you do not use this option, case is significant* (default). Case significance is enforced primarily with symbol names, not with mnemonics and register names.

-ad *-adname [=value]* sets the *name* symbol. This is equivalent to inserting *name .set [value]* at the beginning of the assembly file. If *value* is omitted, the symbol is set to 1. For more information, see section 3.8.4, *Defining Symbolic Constants (-sd Option)*, on page 3-21.

-ahc *-ahcfilename* tells the assembler to copy the specified file for the assembly module. The file is inserted before source file statements. The copied file appears in the assembly listing files.

- ahi** *-ahifilename* tells the assembler to include the specified file for the assembly module. The file is included before source file statements. The included file does not appear in the assembly listing files.
- al** (lowercase L) produces a listing file with the same name as the input file with a *.lst* extension.
- apd** performs preprocessing for assembly files, but instead of writing preprocessed output, writes a list of dependency lines suitable for input to a standard make utility. The list is written to a file with the same name as the source file but with a *.ppa* extension.
- api** performs preprocessing for assembly files, but instead of writing preprocessed output, writes a list of files included with the *.include* directive. The list is written to a file with the same name as the source file but with a *.ppa* extension.
- as** puts all defined symbols in the object file's symbol table. The assembler usually puts only global symbols into the symbol table. When you use *-as*, symbols defined as labels or as assembly-time constants are also placed in the table.
- ax** produces a cross-reference table and appends it to the end of the listing file; it also adds cross-reference information to the object file for use by the cross-reference utility. If you do not request a listing file but use the *-ax* option, the assembler creates a listing file automatically, naming it with the same name as the input file with a *.lst* extension.
- au** *-uname* undefines the predefined constant *name*, which overrides any *-d* options for the specified constant.
- g** (*--symdebug:dwarf*) enables assembler source debugging in the C source debugger. Line information is output to the COFF file for every line of source in the assembly language source file. You cannot use the *-g* option on assembly code that contains *.line* directives. See section 3.11, *Debugging Assembly Source*, on page 3-34 for more information.
- I** specifies a directory where the assembler can find files named by the *.copy*, *.include*, or *.mlib* directives. The format of the *-I* option is *-I=pathname*. There is no limit to the number of directories you can specify in this manner; each pathname must be preceded by the *-I* option. For more information, see section 3.4.1, *Using the -I Assembler Option*, on page 3-7.
- me** produces object code in big-endian format.

- ml** **-mlnum** sets the processor symbols `.SMALL_MODEL`, `.LARGE_MODEL`, and `.LARGE_MODEL_OPTION`. If you are compiling with C/C++ code, the `-mlnum` information is passed to the assembler, and the model symbols are appropriately defined.
- mv** **-mv num** selects the target CPU version using the last four digits of the TMS320C6000 part number.
- q** suppresses the banner and progress information (assembler runs in quiet mode).

For more information about the `-me`, `-ml`, and `-mv` options, see the *TMS320C6000 Optimizing C Compiler User's Guide*.

3.4 Naming Alternate Directories for Assembler Input

The `.copy`, `.include`, and `.mlib` directives tell the assembler to use code from external files. The `.copy` and `.include` directives tell the assembler to read source statements from another file, and the `.mlib` directive names a library that contains macro functions. Chapter 4, *Assembler Directives*, contains examples of the `.copy`, `.include`, and `.mlib` directives. The syntax for these directives is:

```
.copy ["filename"]
.include ["filename"]
.mlib ["filename"]
```

The *filename* names a copy/include file that the assembler reads statements from or a macro library that contains macro definitions. If *filename* begins with a number the double quotes are required. The filename may be a complete pathname, a partial pathname, or a filename with no path information. The assembler searches for the file in the following locations in the order given:

- 1) The directory that contains the current source file. The current source file is the file being assembled when the `.copy`, `.include`, or `.mlib` directive is encountered.
- 2) Any directories named with the `-I` assembler option
- 3) Any directories named with the `C6X_A_DIR` or `A_DIR` environment variable

Because of this search hierarchy, you can augment the assembler's directory search algorithm by using the `-i` assembler option (described in section 3.4.1) or the `C6X_A_DIR` or `A_DIR` environment variable (described in section 3.4.2).

3.4.1 Using the `-I` Assembler Option

The `-I` assembler option names an alternate directory that contains copy/include files or macro libraries. The format of the `-I` option is as follows:

```
cl6x -I=pathname source filename [other options]
```

There is no limit to the number of `-I` options per invocation; each `-I` option names one pathname. In assembly source, you can use the `.copy`, `.include`, or `.mlib` directive without specifying path information. If the assembler does not find the file in the directory that contains the current source file, it searches the paths designated by the `-I` options.

For example, assume that a file called `source.asm` is in the current directory; `source.asm` contains the following directive statement:

```
.copy "copy.asm"
```

Assume the following paths for the `copy.asm` file:

UNIX™: `/tools/files/copy.asm`

Windows™: `c:\tools\files\copy.asm`

You could set up the search path with the commands shown below:

Operating System	Enter
UNIX	<code>cl6x -I/tools/files source.asm</code>
Windows	<code>cl6x -Ic:\tools\files source.asm</code>

The assembler first searches for `copy.asm` in the current directory because `source.asm` is in the current directory. Then the assembler searches in the directory named with the `-I` option.

3.4.2 Using the C6X_A_DIR or A_DIR Environment Variable

An environment variable is a system symbol that you define and assign a string to. The assembler uses the `C6X_A_DIR` and `A_DIR` environment variables to name alternate directories that contain copy/include files or macro libraries.

The assembler looks for the `C6X_A_DIR` environment variable first and then reads and processes it. If it does not find this variable, it reads the `A_DIR` environment variable and processes it. If both variables are set, the settings of the processor-specific variable are used. The processor-specific variable is useful when you are using Texas Instruments tools for different processors at the same time.

If the assembler does not find `C6X_A_DIR` and `A_DIR`, it then searches for `C6X_C_DIR` and `C_DIR`. See the *TMS320C6000 Optimizing Compiler User's Guide* for details on `C6X_C_DIR` and `C_DIR`.

The command syntax for assigning the environment variable is as follows:

Operating System	Enter
Windows	<code>set A_DIR= pathname₁;pathname₂; . . .</code>
UNIX (Bourne shell)	<code>A_DIR="pathname₁;pathname₂; . . ."; export A_DIR</code>

The *pathnames* are directories that contain copy/include files or macro libraries. You can separate the pathnames with a semicolon or with blanks. In assembly source, you can use the `.copy`, `.include`, or `.mlib` directive without specifying path information. If the assembler does not find the file in the directory that contains the current source file or in directories named by the `-I` option, it searches the paths named by the environment variable.

For example, assume that a file called `source.asm` contains these statements:

```
.copy "copy1.asm"
.copy "copy2.asm"
```

Assume the following paths for the files:

UNIX: /tools/files/copy1.asm
 /dsys/copy2.asm

Windows: c:\tools\files\copy1.asm
 c:\dsys\copy2.asm

You could set up the search path with the commands shown below:

Operating System	Enter
UNIX (Bourne Shell)	A_DIR="/dsys"; export A_DIR cl6x -I=/tools/files source.asm
Windows	set A_DIR=c:\dsys cl6x -Ic:\tools\files source.asm

The assembler first searches for `copy1.asm` and `copy2.asm` in the current directory because `source.asm` is in the current directory. Then the assembler searches in the directory named with the `-I` option and finds `copy1.asm`. Finally, the assembler searches the directory named with `A_DIR` and finds `copy2.asm`.

Note that the environment variable remains set until you reboot the system or reset the variable by entering one of these commands:

Operating System	Enter
Windows	set A_DIR=
UNIX	unset A_DIR

3.5 Source Statement Format

TMS320C6000 assembly language source programs consist of source statements that can contain assembler directives, assembly language instructions, macro directives, and comments. A source statement can contain five ordered fields (label, mnemonic, unit specifier, operand list, and comment). The general syntax for source statements is as follows:

```
[label[:]] [[]] [[register]] mnemonic [unit specifier] [operand list] [:comment]
```

Following are examples of source statements:

```
two      .set  2      ; Symbol Two = 2
Label:   MVK    two,A2 ; Move 2 into register A2
          .word 016h   ; Initialize a word with 016h
```

The C6000 assembler reads up to 200 characters per line. Any characters beyond 200 are truncated. Keep the operational part of your source statements (that is, everything other than comments) less than 200 characters in length for correct assembly. Your comments can extend beyond the 200-character limit, but the truncated portion is not included in the listing file.

Follow these guidelines:

- All statements must begin with a label, a blank, an asterisk, or a semicolon.
- Labels are optional; if used, they must begin in column 1.
- One or more blanks must separate each field. Tab and space characters are blanks. You must separate the operand list from the preceding field with a blank.
- Comments are optional. Comments that begin in column 1 can begin with an asterisk or a semicolon (* or ;), but comments that begin in any other column *must* begin with a semicolon.
- In a conditional instruction, the condition register must be surrounded by square brackets.
- The functional unit specifier is optional. If you do not specify the functional unit, the assembler assigns a legal functional unit based on the mnemonic field.
- A mnemonic cannot begin in column 1 or it will be interpreted as a label.

The following sections describe each of the fields.

3.5.1 Label Field

Labels are optional for all assembly language instructions and for most (but not all) assembler directives. When used, a label must begin in column 1 of a source statement. A label can contain up to 128 alphanumeric characters (A–Z, a–z, 0–9, _, and \$). Labels are case sensitive (except when the `-ac` option is used), and the first character cannot be a number. A label can be followed by a colon (:). The colon is not treated as part of the label name. If you do not use a label, the first character position must contain a blank, a semicolon, or an asterisk. You cannot use a label with an instruction that is in parallel with a previous instruction.

When you use a label, its value is the current value of the SPC. The label points to the statement it is associated with. For example, if you use the `.word` directive to initialize several words, a label points to the first word. In the following example, the label `Start` has the value `40h`.

```

.      .      .      .
.      :      :      :
.      .      .      .
      9
10 00000040 0000000A Start: .word 0Ah,3,7
      00000044 00000003
      00000048 00000007
      * Assume some code was assembled

```

A label on a line by itself is a valid statement. The label assigns the current value of the section program counter to the label; this is equivalent to the following directive statement:

```
label .equ $ ; $ provides the current value of the SPC
```

When a label appears on a line by itself, it points to the instruction on the next line (the SPC is not incremented):

```

1 00000000 Here:
2 00000000 00000003 .word 3

```

If you do not use a label, the character in column 1 must be a blank, an asterisk, or a semicolon.

3.5.2 Mnemonic Field

The mnemonic field follows the label field. The mnemonic field cannot start in column 1; if it does, it is interpreted as a label. There is one exception: the parallel bars (||) of the mnemonic field can start in column 1. The mnemonic field can begin with one of the following items:

- Pipe symbols (||) indicate instructions that are in parallel with a previous instruction. You can have up to eight instructions running in parallel. The following example demonstrates six instructions running in parallel:

```

                Inst1
||             Inst2
||             Inst3
||             Inst4
||             Inst5
||             Inst6
||             Inst7

```

} These five instructions run in parallel with the first instruction.

- Square brackets ([]) indicate conditional instructions. The machine-instruction mnemonic is executed based on the value of the register within the brackets; valid register names are A0 for C64xx only, A1, A2, B0, B1, and B2. The instruction is executed if the value of the register is nonzero. If the register name is preceded by an exclamation point (!), then the instruction is executed if the value of the register is 0. For example:

```
[A1] ZERO A2 ; If A1 is not equal to zero, A2 = 0
```

Next, the mnemonic field contains one of the following items:

- Machine-instruction mnemonic (such as ADDK, MVKH, B)
- Assembler directive (such as .data, .list, .equ)
- Macro directive (such as .macro, .var, .mexit)
- Macro call

3.5.3 Unit Specifier Field

The unit specifier field is an optional field that follows the mnemonic field for machine-instruction mnemonics. The unit specifier field begins with a period (.) followed by a functional unit specifier. In general, one instruction can be assigned to each functional unit in a single instruction cycle. There are eight functional units, two of each functional type:

.D1 and .D2	Data/addition/subtraction
.L1 and .L2	ALU/compares/long data arithmetic

.M1 and .M2	Multiply
.S1 and .S2	Shift/ALU/branch/bit field

ALU refers to an arithmetic logic unit.

There are several ways to use the unit specifier field:

- You can specify the particular functional unit (for example, .D1).
- You can specify only the functional type (for example, .M), and the assembler assigns the specific unit (for example, .M2).
- If you do not specify the functional unit, the assembler assigns the functional unit based on the mnemonic field and operand field.

For more information on functional units, including which assembly instructions require which functional type, see the *TMS320C62x, C64x, C67x CPU and Instruction Set Reference Guide*.

3.5.4 Operand Field

The operand field follows the mnemonic field and contains one or more operands. The operand field is not required for all instructions or directives. An operand consists of the following items:

- Symbols (see section 3.8 on page 3-18)
- Constants (see section 3.6 on page 3-14)
- Expressions (combination of constants and symbols; see section 3.9 on page 3-26)

You must separate operands with commas.

3.5.5 Comment Field

A comment can begin in any column and extends to the end of the source line. A comment can contain any ASCII character, including blanks. Comments are printed in the assembly source listing, but they do not affect the assembly.

A source statement that contains only a comment is valid. If it begins in column 1, it can start with a semicolon (;) or an asterisk (*). Comments that begin anywhere else on the line must begin with a semicolon. The asterisk identifies a comment only if it appears in column 1.

3.6 Constants

The assembler supports six types of constants:

- Binary integer
- Octal integer
- Decimal integer
- Hexadecimal integer
- Character
- Assembly-time

The assembler maintains each constant internally as a 32-bit quantity. Constants are not sign extended. For example, the constant 00FFh is equal to 00FF (base 16) or 255 (base 10); it *does not* equal -1 . However, when used with the `.byte` directive, -1 is equivalent to 00FFh.

3.6.1 Binary Integers

A binary integer constant is a string of up to 32 binary digits (0s and 1s) followed by the suffix B (or b). If fewer than 32 digits are specified, the assembler right justifies the value and fills the unspecified bits with zeros. These are examples of valid binary constants:

00000000B	Constant equal to 0_{10} or 0_{16}
0100000b	Constant equal to 32_{10} or 20_{16}
01b	Constant equal to 1_{10} or 1_{16}
11111000B	Constant equal to 248_{10} or $0F8_{16}$

3.6.2 Octal Integers

An octal integer constant is a string of up to 11 octal digits (0 through 7) followed by the suffix Q (or q). These are examples of valid octal constants:

10Q	Constant equal to 8_{10} or 8_{16}
010	Constant equal to 8_{10} or 8_{16} (C format)
100000Q	Constant equal to $32\ 768_{10}$ or 8000_{16}
226q	Constant equal to 150_{10} or 96_{16}

3.6.3 Decimal Integers

A decimal integer constant is a string of decimal digits ranging from $-2147\ 483\ 648$ to $4\ 294\ 967\ 295$. These are examples of valid decimal constants:

1000	Constant equal to 1000_{10} or $3E8_{16}$
-32768	Constant equal to $-32\ 768_{10}$ or 8000_{16}
25	Constant equal to 25_{10} or 19_{16}

3.6.4 Hexadecimal Integers

A hexadecimal integer constant is a string of up to eight hexadecimal digits followed by the suffix H (or h). Hexadecimal digits include the decimal values 0–9 and the letters A–F or a–f. A *hexadecimal constant must begin with a decimal value (0–9)*. If fewer than eight hexadecimal digits are specified, the assembler right justifies the bits. These are examples of valid hexadecimal constants:

78h	Constant equal to 120_{10} or 0078_{16}
0x78	Constant equal to 120_{10} or 0078_{16} (C format)
0Fh	Constant equal to 15_{10} or $000F_{16}$
37ACh	Constant equal to $14\ 252_{10}$ or $37AC_{16}$

3.6.5 Character Constants

A character constant is a single character enclosed in *single* quotes. The characters are represented internally as 8-bit ASCII characters. Two consecutive single quotes are required to represent each single quote that is part of a character constant. A character constant consisting only of two single quotes is valid and is assigned the value 0. These are examples of valid character constants:

'a'	Defines the character constant <i>a</i> and is represented internally as 61_{16}
'C'	Defines the character constant <i>C</i> and is represented internally as 43_{16}
''	Defines the character constant <i>'</i> and is represented internally as 27_{16}
''	Defines a null character and is represented internally as 00_{16}

Notice the difference between character *constants* and character *strings*. (section 3.7 discusses character strings). A character constant represents a single integer value; a string is a sequence of characters.

3.6.6 Assembly-Time Constants

If you use the `.set` directive (see page 4-62) to assign a value to a symbol, the symbol becomes a constant. To use this constant in expressions, the value that is assigned to it must be absolute. For example:

```
sym    .set 3
      MVK  sym,B1
```

You can also use the `.set` directive to assign symbolic constants for register names. In this case, the symbol becomes a synonym for the register:

```
sym    .set B1
      MVK  10,sym
```

3.7 Character Strings

A character string is a string of characters enclosed in *double* quotes. Double quotes that are part of character strings are represented by two consecutive double quotes. The maximum length of a string varies and is defined for each directive that requires a character string. Characters are represented internally as 8-bit ASCII characters.

These are examples of valid character strings:

"sample program" defines the 14-character string *sample program*.

"PLAN "C"" defines the 8-character string *PLAN "C"*.

Character strings are used for the following:

- Filenames, as in `.copy "filename"`
- Section names, as in `.sect "section name"`
- Data initialization directives, as in `.byte "charstring"`
- Operands of `.string` directives

3.8 Symbols

Symbols are used as labels, constants, and substitution symbols. A symbol name is a string of up to 200 alphanumeric characters (A–Z, a–z, 0–9, \$, and `_`). The first character in a symbol cannot be a number, and symbols cannot contain embedded blanks. The symbols you define are case sensitive; for example, the assembler recognizes ABC, Abc, and abc as three unique symbols. You can override case sensitivity with the `-ac` assembler option (see page 3-4). A symbol is valid only during the assembly in which it is defined, unless you use the `.global` directive or the `.def` directive to declare it as an external symbol (see section 2.7.1 on page 2-18).

3.8.1 Labels

Symbols used as labels become symbolic addresses that are associated with locations in the program. Labels used locally within a file must be unique. Mnemonic opcodes and assembler directive names without the `.` prefix are valid label names.

Labels can also be used as the operands of `.global`, `.ref`, `.def`, or `.bss` directives; for example:

```
.global label1

label2: MVKL    label2, B3
        MVKH    label2, B3
        B       label1
        NOP     5
```

3.8.2 Local Labels

Local labels are special labels whose scope and effect are temporary. A local label can be defined in two ways:

- `$n`, where `n` is a decimal digit in the range 0–9. For example, `$4` and `$1` are valid local labels. See Example 3–1.
- `name?`, where `name` is any legal symbol name as described above. The assembler replaces the question mark with a period followed by a unique number. When the source code is expanded, *you will not see the unique number in the listing file*. Your label appears with the question mark as it did in the source definition. You cannot declare this label as global. See Example 3–2.

Normal labels must be unique (they can be declared only once), and they can be used as constants in the operand field. Local labels, however, can be undefined and defined again. Local labels cannot be defined by directives.

A local label can be undefined or reset in one of these ways:

- By using the `.newblock` directive
- By changing sections (using a `.sect`, `.text`, or `.data` directive)
- By entering an include file (specified by the `.include` or `.copy` directive)
- By leaving an include file (specified by the `.include` or `.copy` directive)

Example 3–1. Local Labels of the Form `$n`

This is an example of code that declares and uses a local label legally:

```
$1:
    SUB    A1,1,A1
[A1] B    $1
    SUBC  A3,A0,A3
    NOP    4

    .newblock      ; undefine $1 to use it again

$1 SUB    A2,1,A2
[A2] B    $1
    MPY   A3,A3,A3
    NOP    4
```

The following code uses a local label illegally:

```
$1:
    SUB    A1,1,A1
[A1] B    $1
    SUBC  A3,A0,A3
    NOP    4
$1 SUB    A2,1,A2    ; WRONG - $1 is multiply defined
[A2] B    $1
    MPY   A3,A3,A3
    NOP    4
```

The `$1` label is not undefined before being reused by the second branch instruction. Therefore, `$1` is redefined, which is illegal.

Local labels are especially useful in macros. If a macro contains a normal label and is called more than once, the assembler issues a multiple-definition error. If you use a local label and `.newblock` within a macro, however, the local label is used and reset each time the macro is expanded.

Up to ten local labels can be in effect at one time. After you undefine a local label, you can define it and use it again. Local labels do not appear in the object code symbol table.

Because local labels are intended to be used only locally, branches to local labels are not expanded in case the branch's offset is out of range.

Example 3-2. Local Labels of the Form name?

```
*****
** First definition of local label mylab **
*****
    nop
mylab?  nop
        B mylab?
        nop 5

*****
** Include file has second definition of mylab **
*****
    .copy "a.inc"

*****
** Third definition of mylab, reset upon exit from .include **
*****
mylab?  nop
        B mylab?
        nop 5

*****
** Fourth definition of mylab in macro, macros use different **
** namespace to avoid conflicts **
*****
mymac  .macro
mylab?  nop
        B mylab?
        nop 5
        .endm

*****
** Macro invocation **
*****
    mymac

*****
** Reference to third definition of mylab. Definition is not **
** reset by macro invocation. **
*****
    B mylab?
    nop 5

*****
** Changing section, allowing fifth definition of mylab **
*****
    .sect "Sect_One"
    nop
mylab? .word 0
    nop
    nop
    B mylab?
    nop 5

*****
** The .newblock directive allows sixth definition of mylab **
*****
    .newblock
mylab? .word 0
    nop
    nop
    B mylab?
    nop 5
```

3.8.3 Symbolic Constants

Symbols can be set to constant values. By using constants, you can equate meaningful names with constant values. The `.set` and `.struct/.tag/.endstruct` directives enable you to set constants to symbolic names. Symbolic constants *cannot* be redefined. The following example shows how these directives can be used:

```
K      .set  1024                ; constant definitions
maxbuf .set  2*K

item   .struct                  ; item structure definition
value  .int                     ; value offset = 0
delta  .int                     ; delta offset = 4
i_len  .endstruct               ; item size    = 8

array  .tag  item
        .bss array, i_len*K ; declare an array of K "items"
        .text
        LDW  *+B14(array.delta + 2*i_len),A1
                ; access array [2].delta
```

The assembler also has several predefined symbolic constants; these are discussed in section 3.8.5.

3.8.4 Defining Symbolic Constants (`-ad` Option)

The `-ad` option equates a constant value with a symbol. The symbol can then be used in place of a value in assembly source. The format of the `-ad` option is as follows:

```
cl6x -adname=[value]
```

The *name* is the name of the symbol you want to define. The *value* is the value you want to assign to the symbol. If the *value* is omitted, the symbol is set to 1.

Once you have defined the name with the `-d` option, the symbol can be used in place of a constant value, a well-defined expression, or an otherwise undefined symbol used with assembly directives and instructions. For example, on the command line you enter:

```
cl6x -adSYM1=1 -adSYM2=2 -adSYM3=3 -adSYM4=4 value.asm
```

Since you have assigned values to SYM1, SYM2, SYM3, and SYM4, you can use them in source code. Example 3-3 shows how the `value.asm` file uses these symbols without defining them explicitly.

Example 3–3. Using Symbolic Constants Defined on Command Line

```

If_4:  .if      SYM4 = SYM2 * SYM2
        .byte   SYM4           ; Equal values
        .else
        .byte   SYM2 * SYM2   ; Unequal values
        .endif

IF_5:  .if      SYM1 <= 10
        .byte   10           ; Less than / equal
        .else
        .byte   SYM1         ; Greater than
        .endif

IF_6:  .if      SYM3 * SYM2 != SYM4 + SYM2
        .byte   SYM3 * SYM2  ; Unequal value
        .else
        .byte   SYM4 + SYM4   ; Equal values
        .endif

IF_7:  .if      SYM1 = SYM2           .byte   SYM1
        .elseif SYM2 + SYM3 = 5
        .byte   SYM2 + SYM3
        .endif

```

Within assembler source, you can test the symbol defined with the `-d` option with the following directives:

Type of Test	Directive Usage
Existence	<code>.if \$isdefed("name")</code>
Nonexistence	<code>.if \$isdefed("name") = 0</code>
Equal to value	<code>.if name = value</code>
Not equal to value	<code>.if name != value</code>

The argument to the `$isdefed` built-in function must be enclosed in quotes. The quotes cause the argument to be interpreted literally rather than as a substitution symbol.

3.8.5 Predefined Symbolic Constants

The assembler has several predefined symbols, including the following types:

- \$**, the dollar-sign character, represents the current value of the section program counter (SPC). \$ is a relocatable symbol.
- Register symbols**, including A0–A15 and B0–B15 for C6200 and C6700; and A16–31 and B16–31 for C6400.
- CPU control registers**, including the following:

Register	Description
AMR	Addressing mode register
CSR	Control status register
FADCR (C6700 only)	Floating-point adder configuration register
FAUCR (C6700 only)	Floating-point auxiliary configuration register
FMCR (C6700 only)	Floating-point multiplier configuration register
GFPGFR (C6400 only)	Galois field polynomial generator function register
ICR	Interrupt clear register
IER	Interrupt enable register
IFR	Interrupt flag register
NRP	Nonmaskable interrupt return pointer
IRP	Interrupt return pointer
ISR	Interrupt set register
ISTP	Interrupt service table pointer
PCE1	Program counter

Control registers can be entered as all upper-case or all lower-case characters; for example, CSR can also be entered as csr.

- Processor symbols**, including the following items:

Symbol name	Description
.TMS320C6000	Always set to 1
.TMS320C6200	Set to 1 for '6200, otherwise 0
.TMS320C6400	Set to 1 for '6400, otherwise 0
.TMS320C6700	Set to 1 for '6700, otherwise 0
.LITTLE_ENDIAN	Set to 1 if little-endian mode is selected (the <code>-me</code> assembler option is not used); otherwise 0.
.BIG_ENDIAN	Set to 1 if big-endian mode is selected (the <code>-me</code> assembler option is used); otherwise 0.

- Memory Model Symbols**

Symbol name	Description
.SMALL_MODEL	Set to 1 if a small memory model is used (does not use the <code>-ml<num></code> option). Otherwise 0
.LARGE_MODEL	Set to 1 if a large memory model is used (does not use the <code>-ml<num></code> option). Otherwise 0
.LARGE_MODEL_OPTION	Always defined. Set to the value used with the <code>-ml</code> option. The <code>-ml</code> option can be used when invoking the shell (the C/C++ compiler) or the assembler. See the <i>TMS320C600 Optimizing Compiler User's Guide</i> for more information on the <code>-ml</code> option.

- Assembler Version Symbols**

Symbol name	Description
.ASSEMBLER_VERSION	Always defined. Set to a number that consists of a major version number and a 2-digit minor version number. The number does not contain a decimal. For example, for version 5.00 of the assembler, <code>.ASSEMBLER_VERSION</code> is set to 500.

3.8.6 Substitution Symbols

Symbols can be assigned a string value (variable). This enables you to alias character strings by equating them to symbolic names. Symbols that represent character strings are called substitution symbols. When the assembler encounters a substitution symbol, its string value is substituted for the symbol name. Unlike symbolic constants, substitution symbols can be redefined.

A string can be assigned to a substitution symbol anywhere within a program; for example:

```
.global _table
.asg    "B14", PAGEPTR
.asg    "+B15(4)", LOCAL1
.asg    "+B15(8)", LOCAL2

LDW    *+PAGEPTR(_table),A0
NOP    4
STW    A0,LOCAL1
```

When you are using macros, substitution symbols are important because macro parameters are actually substitution symbols that are assigned a macro argument. The following code shows how substitution symbols are used in macros:

```
MAC .macro  src1, src2, dst ; Multiply/Accumulate macro
    MPY    src1, src2, src2
    NOP
    ADD    src2, dst, dst
    .endm

* MAC macro invocation
    MAC    A0,A1,A2
```

For more information about macros, see Chapter 5, *Macro Language*.

3.9 Expressions

An expression is a constant, a symbol, or a series of constants and symbols separated by arithmetic operators. The 32-bit ranges of valid expression values are $-2147\ 483\ 648$ to $2147\ 483\ 647$ for signed values, and 0 to $4\ 294\ 967\ 295$ for unsigned values. Three main factors influence the order of expression evaluation:

Parentheses

Expressions enclosed in parentheses are always evaluated first.

$$8 / (4 / 2) = 4, \text{ but } 8 / 4 / 2 = 1$$

You *cannot* substitute braces ({ }) or brackets ([]) for parentheses.

Precedence groups

Operators, listed in Table 3–1, are divided into nine precedence groups. When parentheses do not determine the order of expression evaluation, the highest precedence operation is evaluated first.

$$8 + 4 / 2 = 10 \text{ (} 4 / 2 \text{ is evaluated first)}$$

Left-to-right evaluation

When parentheses and precedence groups do not determine the order of expression evaluation, the expressions are evaluated from left to right, except for Group 1, which is evaluated from right to left.

$$8 / 4 * 2 = 4, \text{ but } 8 / (4 * 2) = 1$$

3.9.1 Operators

Table 3–1 lists the operators that can be used in expressions, according to precedence group.

Table 3–1. Operators Used in Expressions (Precedence)

Group	Operator	Description
1	+	Unary plus
	–	Unary minus
	~	1s complement
	!	Logical NOT
2	*	Multiplication
	/	Division
	%	Modulo
3	+	Addition
	–	Subtraction
4	<<	Shift left
	>>	Shift right
5	<	Less than
	<=	Less than or equal to
	>	Greater than
	>=	Greater than or equal to
6	=[=]	Equal to
	!=	Not equal to
7	&	Bitwise AND
8	^	Bitwise exclusive OR (XOR)
9		Bitwise OR

Note: Group 1 operators are evaluated right to left. All other operators are evaluated left to right.

3.9.2 Expression Overflow and Underflow

The assembler checks for overflow and underflow conditions when arithmetic operations are performed at assembly time. It issues a warning (the message *Value Truncated*) whenever an overflow or underflow occurs. The assembler *does not* check for overflow or underflow in multiplication.

3.9.3 Well-Defined Expressions

Some assembler directives require well-defined expressions as operands. Well-defined expressions contain only symbols or assembly-time constants that are defined before they are encountered in the expression. The evaluation of a well-defined expression must be absolute.

This is an example of a well-defined expression:

```
1000h+X
```

where X was previously defined as an absolute symbol.

3.9.4 Conditional Expressions

The assembler supports relational operators that can be used in any expression; they are especially useful for conditional assembly. Relational operators include the following:

=	Equal to	!=	Not equal to
<	Less than	<=	Less than or equal to
>	Greater than	>=	Greater than or equal to

Conditional expressions evaluate to 1 if true and 0 if false and may be used only on operands of equivalent types; for example, absolute value compared to absolute value, but not absolute value compared to relocatable value.

3.9.5 Legal Expressions

With the exception of the following expression contexts, there is no restriction on combinations of operations, constants, internally defined symbols, and externally defined symbols.

When an expression contains more than one relocatable symbol or cannot be evaluated at assembly time, the assembler encodes a relocation expression in the object file that is later evaluated by the linker. If the final value of the expression is larger in bits than the space reserved for it, you receive an error message from the linker. For more information on relocation expressions, see section 2.4 on page 2-14.

- When using the register relative addressing mode, the expression in brackets or parenthesis must be a well-defined expression, as described in section 3.9.3. For example:

```
*+A4 [15]
```

- Expressions used to describe the offset in register relative addressing mode for the registers B14 and B15, or expressions used as the operand to the branch instruction, are subject to the same limitations. For these two cases, all legal expressions can be reduced to one of two forms:

relocatable symbol \pm *absolute symbol* B (**extern_1 - 10**)

or

a well-defined expression ***+B14/B15 [14]**

3.9.6 Expression Examples

Following are examples of expressions that use relocatable and absolute symbols. These examples use four symbols that are defined in the same section:

```
.global extern_1 ; Defined in an external module
intern_1: .word "'D' ; Relocatable, defined in
                ; current module
intern_2                ; Relocatable, defined in
                ; current module
intern_3                ; Relocatable, defined in
                ; current module
```

Example 1

In these contexts, there are no limitations on how expressions can be formed.

```
.word extern_1 * intern_2 - 13 ; Legal
MVKL (intern_1 - extern_1),A1 ; Legal
```

Example 2

The first statement in the following example is valid; the statements that follow it are invalid.

```
B (extern_1 - 10) ; Legal
B (10-extern_1) ; Can't negate reloc. symbol
LDW *+B14 (-(intern_1)), A1 ; Can't negate reloc. symbol
LDW *+B14 (extern_1/10), A1 ; / not an additive operator
B (intern_1 + extern_1) ; Multiple relocatables
```

❑ **Example 3**

The first statement below is legal; although `intern_1` and `intern_2` are relocatable, their difference is absolute because they are in the same section. Subtracting one relocatable symbol from another reduces the expression to *relocatable symbol + absolute value*. The second statement is illegal because the sum of two relocatable symbols is not an absolute value.

```
B (intern_1 - intern_2 + extern_3) ; Legal
```

```
B (intern_1 + intern_2 + extern_3) ; Illegal
```

❑ **Example 4**

A relocatable symbol's placement in the expression is important to expression evaluation. Although the statement below is similar to the first statement in the previous example, it is illegal because of left-to-right operator precedence; the assembler attempts to add `intern_1` to `extern_3`.

```
B (intern_1 + extern_3 - intern_2) ; Illegal
```

3.10 Source Listings

A source listing shows source statements and the object code they produce. To obtain a listing file, invoke the assembler with the `-al` (lowercase L) option (see page 3-5).

Two banner lines, a blank line, and a title line are at the top of each source listing page. Any title supplied by the `.title` directive is printed on the title line. A page number is printed to the right of the title. If you do not use the `.title` directive, the name of the source file is printed. The assembler inserts a blank line below the title line.

Each line in the source file produces at least one line in the listing file. This line shows a source statement number, an SPC value, the object code assembled, and the source statement. Example 3-4 shows these in an actual listing file.

Field 1: Source Statement Number

Line number

The source statement number is a decimal number. The assembler numbers source lines as it encounters them in the source file; some statements increment the line counter but are not listed. (For example, `.title` statements and statements following a `.nolist` are not listed.) The difference between two consecutive source line numbers indicates the number of intervening statements in the source file that are not listed.

Include file letter

A letter preceding the line number indicates the line is assembled from the include file designated by the letter.

Nesting level number

A number preceding the line number indicates the nesting level of macro expansions or loop blocks.

Field 2: Section Program Counter

This field contains the SPC value, which is hexadecimal. All sections (`.text`, `.data`, `.bss`, and named sections) maintain separate SPCs. Some directives do not affect the SPC and leave this field blank.

Field 3: Object Code

This field contains the hexadecimal representation of the object code. All machine instructions and directives use this field to list object code. This field also indicates the relocation type associated with an operand for this line of source code. If more than one operand is relocatable, this column indicates the relocation type for the first operand. The characters that can appear in this column and their associated relocation types are listed below:

- ! undefined external reference
- ' .text relocatable
- + .sect relocatable
- " .data relocatable
- .bss, .usect relocatable
- % relocation expression

Field 4: Source Statement Field

This field contains the characters of the source statement as they were scanned by the assembler. The assembler accepts a maximum line length of 200 characters. Spacing in this field is determined by the spacing in the source statement.

Example 3-4 shows an assembler listing with each of the four fields identified.

Example 3-4. Assembler Listing

Include file letter	Nesting level number	Line number	
		1	*****
		2	** Global variables
		3	*****
		4	00000000 .bss var1, 4
		5	00000004 .bss var2, 4
		6	
		7	*****
		8	** Include multiply macro
		9	*****
		10	.copy mpy32.inc
A	1	mpy32	.macro A,B
A	2		
A	3		MPYLH.M1 A,B,A ; tmp1 = A.lo * B.hi
A	4		MPYHL.M2 A,B,B ; tmp2 = A.hi * B.lo
A	5		
A	6		MPYU.M2 A,B,B ; tmp3 = A.lo * B.lo
A	7		
A	8		ADD.L1 A,B,A ; A = tmp1 + tmp2
A	9		
A	10		SHL.S1 A,16,A ; A <= 16
A	11		
A	12		ADD.L1 B,A,A ; A = A + tmp3
A	13		.endm
		11	
		12	*****
		13	** _func multiplies 2 global ints
		14	*****
		15	00000000 .text
		16	00000000 _func
		17	00000000 0200006C- LDW *+B14(var1),A4
		18	00000004 0000016E- LDW *+B14(var2),B0
		19	00000008 00006000 NOP 4
		20	0000000c mpy32 A4,B0
1		1	
1	0000000c	02009881	MPYLH.M1 A4,B0,A4 ; tmp1 = A.lo * B.hi
1	00000010	00101882	MPYHL.M2 A4,B0,B0 ; tmp2 = A.hi * B.lo
1			
1	00000014	00101F82	MPYU.M2 A4,B0,B0 ; tmp3 = A.lo * B.lo
1			
1	00000018	02009078	ADD.L1 A4,B0,A4 ; A = tmp1 + tmp2
1			
1	0000001c	02120CA0	SHL.S1 A4,16,A4 ; A <= 16
1			
1	00000020	02009078	ADD.L1 B0,A4,A4 ; A = A + tmp3
	21	00000024 000C6362	B B3
	22	00000028 00008000	NOP 5
	23		* end _func

Field 1 Field 2 Field 3 Field 4

3.11 Debugging Assembly Source

When you invoke `cl6x` with `--symdebug:dwarf` (or `-g`) when compiling an assembly file, the assembler provides symbolic debugging information that allows you to step through your assembly code in a debugger rather than using the Disassembly window in Code Composer Studio. This enables you to view source comments and other source-code annotations while debugging.

The `.asmfunc` and `.endasmfunc` directives enable you to use C characteristics in assembly code that makes the process of debugging an assembly file more closely resemble debugging a C/C++ source file.

The `.asmfunc` and `.endasmfunc` directives (see page 4-24) allow you to name certain areas of your code, and make these areas appear in the debugger as C functions. Contiguous sections of assembly code that are not enclosed by the `.asmfunc` and `.endasmfunc` directives are automatically placed in assembler-defined functions named with this syntax:

`$filename:starting source line:ending source line$`

If you want to view your variables as a user-defined type in C code, the types must be declared and the variables must be defined in a C file. This C file can then be referenced in assembly code using the `.ref` directive (see page 4-44).

Example 3–5 shows the `cvar.c` C program that defines a variable, `svar`, as the structure type `X`. The `svar` variable is then referenced in the `addfive.asm` assembly program and 5 is added to `svar`'s second data member.

Example 3–5. Viewing Assembly Variables as C Types

(a) C Program `cvar.c`

```
typedef struct
{
    int m1;
    int m2;
} X;

X svar = { 1, 2 };
```

(b) Assembly Program addfive.asm

```

;-----
; Tell the assembler we're referencing variable "_svar", which is defined in
; another file (cvars.c).
;-----
        .ref _svar

;-----
; addfive() - Add five to the second data member of _svar
;-----
        .text
        .global addfive
addfive: .asmfunc
        LDW    .D2T2    *+B14(_svar+4),B4 ; load svar.m2 into B4
        RET    .S2      B3                ; return from function
        NOP    3        3                ; delay slots 1-3
        ADD    .D2      5,B4,B4          ; add 5 to B4 (delay slot 4)
        STW    .D2T2    B4,*+B14(_svar+4) ; store B4 back into svar.m2
                                                ; (delay slot 5)
        .endasmfunc

```

Compile both source files with the `--symdebug:dwarf` option (`-g`) and link them as follows:

```
cl6x -symdebug:dwarf cvars.c addfive.asm -z -l=lnk.cmd -l=rts6000.lib -o=addfive.out
```

When you load this program into a symbolic debugger, `addfive` appears as a C function. You can monitor the values in `svar` while stepping through `main` just as you would any regular C variable.

3.12 Cross-Reference Listings

A cross-reference listing shows symbols and their definitions. To obtain a cross-reference listing, invoke the assembler with the `-ax` option (see page 3-5) or use the `.option` directive with the `X` operand (see page 4-59). The assembler appends the cross-reference to the end of the source listing. Example 3-6 shows the four fields contained in the cross-reference listing.

Example 3-6. An Assembler Cross-Reference Listing

LABEL	VALUE	DEFN	REF
.BIG_ENDIAN	00000000	0	
.LITTLE_ENDIAN	00000001	0	
.TMS320C6200	00000001	0	
.TMS320C6700	00000000	0	
.TMS320C6X	00000001	0	
_func	00000000'	18	
var1	00000000-	4	17
var2	00000004-	5	18

Label column contains each symbol that was defined or referenced during the assembly.

Value column contains an 8-digit hexadecimal number (which is the value assigned to the symbol) or a name that describes the symbol's attributes. A value may also be preceded by a character that describes the symbol's attributes. Table 3-2 lists these characters and names.

Definition (DEFN) column contains the statement number that defines the symbol. This column is blank for undefined symbols.

Reference (REF) column lists the line numbers of statements that reference the symbol. A blank in this column indicates that the symbol was never used.

Table 3-2. Symbol Attributes

Character or Name	Meaning
REF	External reference (global symbol)
UNDF	Undefined
'	Symbol defined in a <code>.text</code> section
"	Symbol defined in a <code>.data</code> section
+	Symbol defined in a <code>.sect</code> section
-	Symbol defined in a <code>.bss</code> or <code>.usect</code> section

Assembler Directives

Assembler directives supply data to the program and control the assembly process. Assembler directives enable you to do the following:

- Assemble code and data into specified sections
- Reserve space in memory for uninitialized variables
- Control the appearance of listings
- Initialize memory
- Assemble conditional blocks
- Define global variables
- Specify libraries from which the assembler can obtain macros
- Examine symbolic debugging information

This chapter is divided into two parts: the first part (sections 4.1 through 4.9) describes the directives according to function, and the second part (section 4.10) is an alphabetical reference.

Topic	Page
4.1 Directives Summary	4-2
4.2 Directives That Define Sections	4-8
4.3 Directives That Initialize Constants	4-10
4.4 Directive That Aligns the Section Program Counter	4-13
4.5 Directives That Format the Output Listings	4-14
4.6 Directives That Reference Other Files	4-16
4.7 Directives That Enable Conditional Assembly	4-17
4.8 Directives That Define Symbols at Assembly Time	4-18
4.9 Miscellaneous Directives	4-20
4.10 Directives Reference	4-21

4.1 Directives Summary

Table 4–1 summarizes the assembler directives.

Besides the assembler directives documented here, the TMS320C6000™ software tools support the following directives:

- The assembler uses several directives for macros. Macro directives are discussed in Chapter 5, *Macro Language*; they are not discussed in this chapter.
- The assembly optimizer uses several directives that supply data and control the optimization process. Assembly optimizer directives are discussed in the *TMS320C6000 Optimizing Compiler User's Guide*; they are not discussed in this book.
- The C compiler uses directives for symbolic debugging. Unlike other directives, symbolic debugging directives are not used in most assembly language programs. Appendix B, *Symbolic Debugging Directives*, discusses these directives; they are not discussed in this chapter.

Note: Labels and Comments in Not Shown Syntaxes

Any source statement that contains a directive can also contain a label and a comment. Labels begin in the first column (they are the only elements, except comments, that can appear in the first column), and comments must be preceded by a semicolon or an asterisk if the comment is only element in the line. To improve readability, labels and comments are not shown as part of the directive syntax.

Table 4–1. Assembler Directives Summary

(a) Directives that define sections

Mnemonic and Syntax	Description	Page
.bss <i>symbol, size in bytes</i> [, <i>alignment</i> [, <i>bank offset</i>]]	Reserves <i>size</i> bytes in the .bss (uninitialized data) section	4-26
.clink [" <i>section name</i> "]	Enables conditional linking for the current or specified section.	4-28
.data	Assembles into the .data (initialized data) section	4-34
.sect " <i>section name</i> "	Assembles into a named (initialized) section	4-61
.text	Assembles into the .text (executable code) section	4-70
<i>symbol</i> .usect " <i>section name</i> ", <i>size in bytes</i> [, <i>alignment</i> [, <i>bank offset</i>]]	Reserves <i>size</i> bytes in a named (uninitialized) section	4-75

Table 4–1. Assembler Directives Summary (Continued)

(b) Directives that initialize constants (data and memory)

Mnemonic and Syntax	Description	Page
.byte <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more successive bytes in the current section	4-27
.char <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more successive bytes in the current section	4-27
.double <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 64-bit, IEEE double-precision, floating-point constants	4-35
.field <i>value</i> [, <i>size</i>]	Initializes a field of <i>size</i> bits (1–32) with <i>value</i>	4-41
.float <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit, IEEE single-precision, floating-point constants	4-43
.half <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 16-bit integers (halfword)	4-46
.uhalf <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 16-bit integers (halfword)	4-46
.int <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit integers	4-49
.uint <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit integers	4-49
.long <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit integers	4-49
.short <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 16-bit integers (halfword)	4-46
.ushort <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 16-bit integers (halfword)	4-46
.string { <i>expr</i> ₁ "string ₁ "} [, ... , { <i>expr</i> _{<i>n</i>} "string _{<i>n</i>} "}]	Initializes one or more text strings	4-66
.word <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit integers	4-49
.uword <i>value</i> ₁ [, ... , <i>value</i> _{<i>n</i>}]	Initializes one or more 32-bit integers	4-49

(c) Directives that perform alignment and reserve space

Mnemonic and Syntax	Description	Page
.align [<i>size in bytes</i>]	Aligns the SPC on a boundary specified by <i>size in bytes</i> , which must be a power of 2; defaults to byte boundary	4-21
.bes <i>size</i>	Reserves <i>size</i> bytes in the current section; a label points to the end of the reserved space	4-63
.space <i>size</i>	Reserves <i>size</i> bytes in the current section; a label points to the beginning of the reserved space	4-63

Table 4–1. Assembler Directives Summary (Continued)

(d) Directives that format the output listing

Mnemonic and Syntax	Description	Page
.drlist	Enables listing of all directive lines (default)	4-36
.drnolist	Suppresses listing of certain directive lines	4-36
.fclist	Allows false conditional code block listing (default)	4-40
.fcnolist	Suppresses false conditional code block listing	4-40
.length [<i>page length</i>]	Sets the page length of the source listing	4-51
.list	Restarts the source listing	4-52
.mlist	Allows macro listings and loop blocks (default)	4-56
.mnolist	Suppresses macro listings and loop blocks	4-56
.nolist	Stops the source listing	4-52
.option <i>option</i> ₁ [, <i>option</i> ₂ , . . .]	Selects output listing options; available options are A, B, D, H, L, M, N, O, R, T, W, and X	4-59
.page	Ejects a page in the source listing	4-61
.sslist	Allows expanded substitution symbol listing	4-65
.ssnolist	Suppresses expanded substitution symbol listing (default)	4-65
.tab <i>size</i>	Sets tab to <i>size</i> characters	4-69
.title " <i>string</i> "	Prints a title in the listing page heading	4-71
.width [<i>page width</i>]	Sets the page width of the source listing	4-51

Table 4–1. Assembler Directives Summary (Continued)

(e) Directives that reference other files

Mnemonic and Syntax	Description	Page
.copy [""]filename[""]	Includes source statements from another file	4-29
.def symbol ₁ [, ... , symbol _n]	Identifies one or more symbols that are defined in the current module and that can be used in other modules	4-44
.global symbol ₁ [, ... , symbol _n]	Identifies one or more global (external) symbols	4-44
.include [""]filename[""]	Includes source statements from another file	4-29
.mlib [""]filename[""]	Defines macro library	4-55
.ref symbol ₁ [, ... , symbol _n]	Identifies one or more symbols used in the current module that are defined in another module	4-44

(f) Directives that enable conditional assembly

Mnemonic and Syntax	Description	Page
.break [<i>well-defined expression</i>]	Ends <code>.loop</code> assembly if <i>well-defined expression</i> is true. When using the <code>.loop</code> construct, the <code>.break</code> construct is optional.	4-53
.else	Assembles code block if the <code>.if</code> <i>well-defined expression</i> is false. When using the <code>.if</code> construct, the <code>.else</code> construct is optional.	4-47
.elseif <i>well-defined expression</i>	Assembles code block if the <code>.if</code> <i>well-defined expression</i> is false and the <code>.elseif</code> condition is true. When using the <code>.if</code> construct, the <code>.elseif</code> construct is optional.	4-47
.endif	Ends <code>.if</code> code block	4-47
.endloop	Ends <code>.loop</code> code block	4-53
.if <i>well-defined expression</i>	Assembles code block if the <i>well-defined expression</i> is true	4-47
.loop [<i>well-defined expression</i>]	Begins repeatable assembly of a code block; the loop count is determined by the <i>well-defined expression</i> .	4-53

Table 4–1. Assembler Directives Summary (Continued)

(g) Structure and Union Definition Directives

Mnemonic and Syntax	Description	Page
.cunion	Acts like <code>.union</code> , but adds padding and alignment like that which is done to C structures	4-31
.cstruct	Acts like <code>.struct</code> , but adds padding and alignment like that which is done to C structures	4-31
.endunion	Ends a union definition	4-31
.endstruct	Ends a structure definition	4-31, 4-67
.struct	Begins structure definition	4-67
.tag	Assigns structure attributes to a label	4-67
.union	Begins a union definition	4-31

(h) Symbol Defining Directives

Mnemonic and Syntax	Description	Page
.label <i>symbol</i>	Defines a load-time relocatable label in a section	4-50
<i>symbol</i> .equ <i>value</i>	Equates <i>value</i> with <i>symbol</i>	4-62
<i>symbol</i> .set <i>value</i>	Equates <i>value</i> with <i>symbol</i>	4-62

(i) Substitution Symbol Directives

Mnemonic and Syntax	Description	Page
.asg <i>[""]character string[""], substitution symbol</i>	Assigns a character string to <i>substitution symbol</i>	4-22
.eval <i>well-defined expression, substitution symbol</i>	Performs arithmetic on numeric <i>substitution symbol</i>	4-22
.var	adds a local substitution symbol to a macro's parameter list	4-36

Table 4–1. Assembler Directives Summary (Continued)

(j) Miscellaneous directives

Mnemonic and Syntax	Description	Page
.emsg <i>string</i>	Sends user-defined error messages to the output device; produces no .obj file	4-37
.end	Ends program	4-39
.mmsg <i>string</i>	Sends user-defined messages to the output device	4-37
.newblock	Undefines local labels	4-58
.wmsg <i>string</i>	Sends user-defined warning messages to the output device	4-37

4.2 Directives That Define Sections

These directives associate portions of an assembly language program with the appropriate sections:

- The **.bss** directive reserves space in the .bss section for uninitialized variables.
- The **.data** directive identifies portions of code in the .data section. The .data section usually contains initialized data.
- The **.sect** directive defines an initialized named section and associates subsequent code or data with that section. A section defined with .sect can contain code or data.
- The **.text** directive identifies portions of code in the .text section. The .text section usually contains executable code.
- The **.usect** directive reserves space in an uninitialized named section. The .usect directive is similar to the .bss directive, but it allows you to reserve space separately from the .bss section.

Chapter 2, *Introduction to Common Object File Format*, discusses COFF sections in detail.

Example 4–1 shows how you can use sections directives to associate code and data with the proper sections. This is an output listing; column 1 shows line numbers, and column 2 shows the SPC values. (Each section has its own program counter, or SPC.) When code is first placed in a section, its SPC equals 0. When you resume assembling into a section after other code is assembled, the section's SPC resumes counting as if there had been no intervening code.

The directives in Example 4–1 perform the following tasks:

.text	initializes words with the values 1, 2, 3, 4, 5, 6, 7, and 8.
.data	initializes words with the values 9, 10, 11, 12, 13, 14, 15, and 16.
var_defs	initializes words with the values 17 and 18.
.bss	reserves 19 bytes.
xy	reserves 20 bytes.

The .bss and .usect directives do not end the current section or begin new sections; they reserve the specified amount of space, and then the assembler resumes assembling code or data into the current section.

Example 4-1. Sections Directives

```

1          *****
2          *      Start assembling into the .text section      *
3          *****
4 00000000          .text
5 00000000 00000001          .word   1,2
   00000004 00000002
6 00000008 00000003          .word   3,4
   0000000c 00000004
7
8          *****
9          *      Start assembling into the .data section      *
10         *****
11 00000000          .data
12 00000000 00000009          .word   9, 10
   00000004 0000000A
13 00000008 0000000B          .word   11, 12
   0000000c 0000000C
14
15         *****
16         *      Start assembling into a named,                *
17         *      initialized section, var_defs                  *
18         *****
19 00000000          .sect   "var_defs"
20 00000000 00000011          .word   17, 18
   00000004 00000012
21
22         *****
23         *      Resume assembling into the .data section      *
24         *****
25 00000010          .data
26 00000010 0000000D          .word   13, 14
   00000014 0000000E
27 00000000          .bss    sym, 19   ; Reserve space in .bss
28 00000018 0000000F          .word   15, 16   ; Still in .data
   0000001c 00000010
29
30         *****
31         *      Resume assembling into the .text section      *
32         *****
33 00000010          .text
34 00000010 00000005          .word   5, 6
   00000014 00000006
35 00000000          usym    .usect  "xy", 20   ; Reserve space in xy
36 00000018 00000007          .word   7, 8     ; Still in .text
   0000001c 00000008

```

4.3 Directives That Initialize Constants

Several directives assemble values for the current section:

- The **.bes** and **.space** directives reserve a specified number of bytes in the current section. The assembler fills these reserved bytes with 0s.
 - When you use a label with **.space**, it points to the *first* byte that contains reserved bits.
 - When you use a label with **.bes**, it points to the *last* byte that contains reserved bits.

Figure 4–1 shows how the **.space** and **.bes** directives work for the following assembled code:

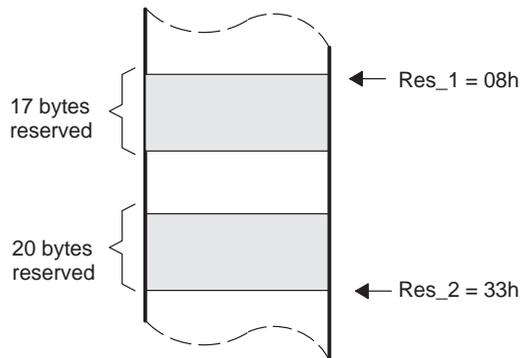
```

1
2 00000000 00000100          .word    100h, 200h
   00000004 00000200
3 00000008          Res_1:  .space    17
4 0000001c 0000000F          .word    15
5 00000033          Res_2:  .bes     20
6 00000034 000000BA          .byte   0BAh

```

Res_1 points to the first byte in the space reserved by **.space**. Res_2 points to the last byte in the space reserved by **.bes**.

Figure 4–1. The **.space** and **.bes** Directives



- The **.byte** and **.char** directives place one or more 8-bit values into consecutive bytes of the current section. These directives are similar to **.long** and **.word**, except that the width of each value is restricted to eight bits.
- The **.double** directive calculates the double-precision (64-bit) IEEE floating-point representation of one or more floating-point values and stores them in two consecutive words in the current section. The **.double** directive automatically aligns to the double-word boundary.

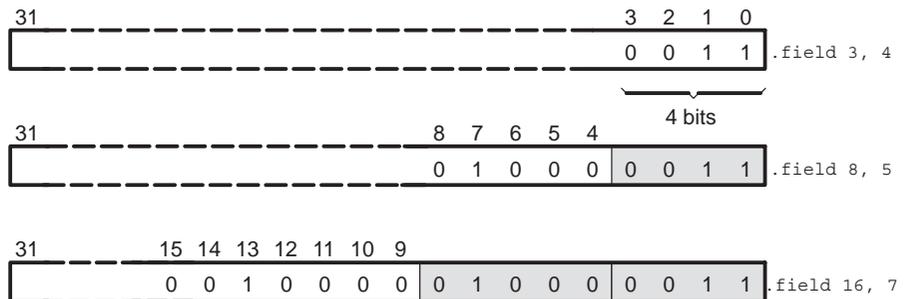
- ❑ The **.field** directive places a single value into a specified number of bits in the current word. With **.field**, you can pack multiple fields into a single word; the assembler does not increment the SPC until a word is filled.

Figure 4–2 shows how fields are packed into a word. Using the following assembled code, notice that the SPC does not change (the fields are packed into the same word):

```

1 00000000 00000003      .field 3,4
2 00000000 00000083     .field 8,5
3 00000000 00002083     .field 16,7
    
```

Figure 4–2. The **.field** Directive



- ❑ The **.float** directive calculates the single-precision (32-bit) IEEE floating-point representation of a single floating-point value and stores it in a word in the current section that is aligned to a word boundary.
- ❑ The **.half**, **.uhalf**, **.short**, and **.ushort** directives place one or more 16-bit values into consecutive 16-bit fields (halfwords) in the current section. The **.half** and **.short** directives automatically align to a short (2-byte) boundary.
- ❑ The **.int**, **.uint**, **.long**, **.word**, **.uword** directives place one or more 32-bit values into consecutive 32-bit fields (words) in the current section. The **.int**, **.long**, and **.word** directives automatically align to a word boundary.
- ❑ The **.string** directive places 8-bit characters from one or more character strings into the current section. This directive is similar to **.byte**, placing an 8-bit character in each consecutive byte of the current section.

Note: Directives That Initialize Constants When Used in a **.struct/.endstruct Sequence**

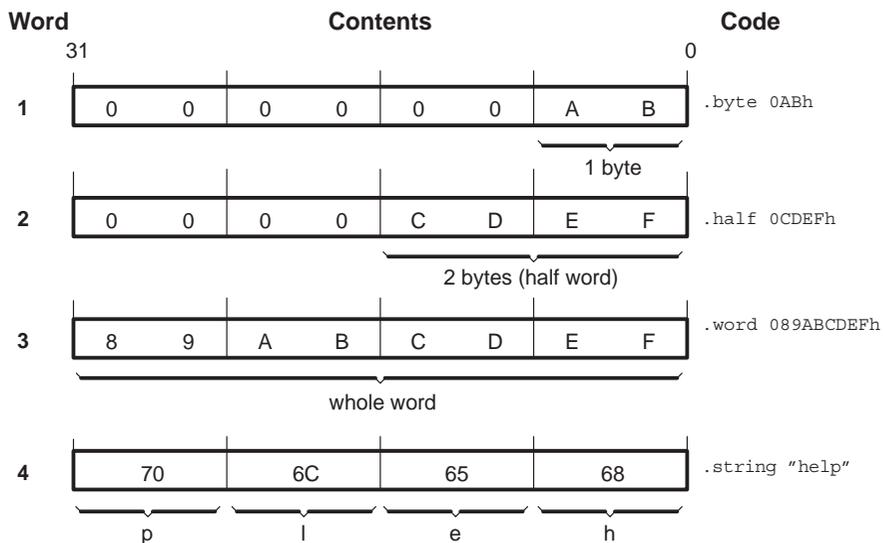
The **.byte**, **.char**, **.int**, **.long**, **.word**, **.double**, **.half**, **.short**, **.string**, **.float**, and **.field** directives *do not* initialize memory when they are part of a **.struct/.endstruct** sequence; rather, they define a member's size. For more information about the **.struct/.endstruct** directives, see page 4-67.

Figure 4–3 compares the .byte, .half, .word, and .string directives. Using the following assembled code:

```

1 00000000 000000AB      .byte  0ABh
2                               .align 4
3 00000004 0000CDEF      .half  0CDEFh
4 00000008 89ABCDEF      .word  089ABCDEFh
5 0000000c 00000068      .string "help"
   0000000d 00000065
   0000000e 0000006C
   0000000f 00000070
    
```

Figure 4–3. Initialization Directives



4.4 Directive That Aligns the Section Program Counter

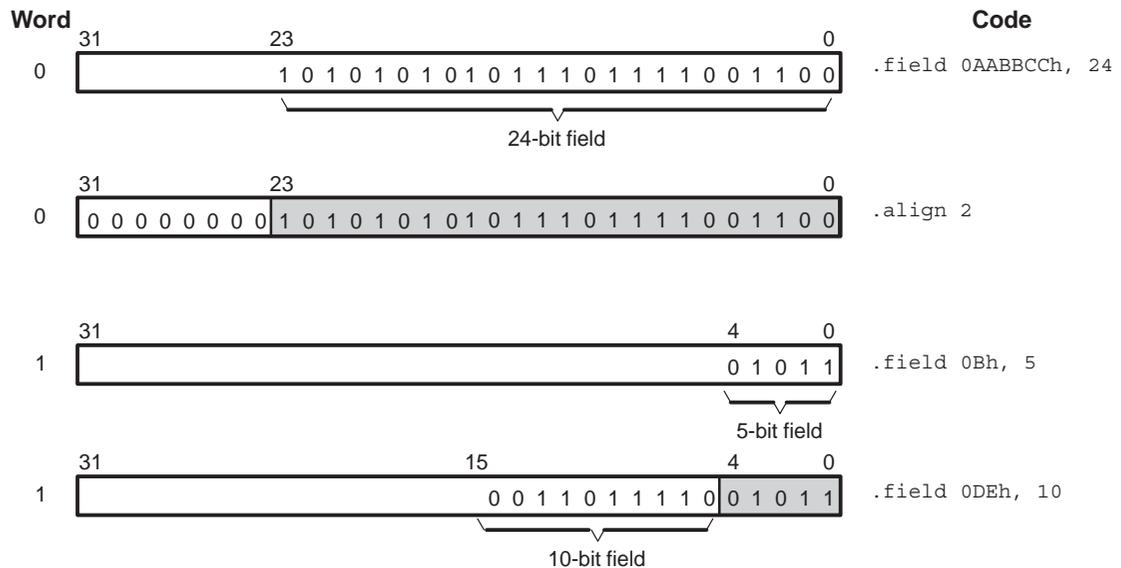
The `.align` directive aligns the SPC at the next byte boundary. This directive is useful with the `.field` directive when you do not want to pack two adjacent fields in the same byte. Figure 4–4 demonstrates the `.align` directive. Using the following assembled code:

```

1
2 00000000 00AABBCC          .field 0AABBCCch,24
3                               .align 2
4 00000000 0BAABBCC          .field 0Bh,5
5 00000004 000000DE          .field 0DEh,10

```

Figure 4–4. The `.align` Directive



4.5 Directives That Format the Output Listings

These directives format the listing file:

- The **.drlist** directive causes printing of the directive lines to the listing; the **.drnolist** directive turns it off for certain directives. You can use the **.drnolist** directive to suppress the printing of the following directives:

.asg	.eval	.length	.mnlolist	.var
.break	.fclist	.mllist	.sslist	.width
.emsg	.fcnololist	.mmsg	.ssnololist	.wmsg

You can use the **.drlist** directive to turn the listing on again.

- The source code listing includes false conditional blocks that do not generate code. The **.fclist** and **.fcnololist** directives turn this listing on and off. You can use the **.fclist** directive to list false conditional blocks exactly as they appear in the source code. You can use the **.fcnololist** directive to list only the conditional blocks that are actually assembled.
- The **.length** directive controls the page length of the listing file. You can use this directive to adjust listings for various output devices.
- The **.list** and **.nolist** directives turn the output listing on and off. You can use the **.nolist** directive to prevent the assembler from printing selected source statements in the listing file. Use the **.list** directive to turn the listing on again.
- The source code listing includes macro expansions and loop blocks. The **.mllist** and **.mnlolist** directives turn this listing on and off. You can use the **.mllist** directive to print all macro expansions and loop blocks to the listing, and the **.mnlolist** directive to suppress this listing.
- The **.option** directive controls certain features in the listing file. This directive has the following operands:
 - A** turns on listing of all directives and data, and subsequent expansions, macros, and blocks.
 - B** limits the listing of **.byte** and **.char** directives to one line.
 - D** turns off the listing of certain directives (same effect as **.drnolist**).
 - H** limits the listing of **.half** and **.short** directives to one line.
 - L** limits the listing of **.long** directives to one line.
 - M** turns off macro expansions in the listing.
 - N** turns off listing (performs **.nolist**).
 - O** turns on listing (performs **.list**).
 - R** resets the **B**, **H**, **L**, **M**, **T**, and **W** directives (turns off the limits of **B**, **H**, **L**, **M**, **T**, and **W**).

- T** limits the listing of `.string` directives to one line.
 - W** limits the listing of `.word` and `.int` directives to one line.
 - X** produces a cross-reference listing of symbols. You can also obtain a cross-reference listing by invoking the assembler with the `-x` option (see page 3-5).
- The **.page** directive causes a page eject in the output listing.
 - The source code listing includes substitution symbol expansions. The **.sslist** and **.ssnolist** directives turn this listing on and off. You can use the `.sslist` directive to print all substitution symbol expansions to the listing, and the `.ssnolist` directive to suppress this listing. These directives are useful for debugging the expansion of substitution symbols.
 - The **.tab** directive defines tab size.
 - The **.title** directive supplies a title that the assembler prints at the top of each page.
 - The **.width** directive controls the page width of the listing file. You can use this directive to adjust listings for various output devices.

4.6 Directives That Reference Other Files

These directives supply information for or about other files that can be used in the assembly of the current file:

- The **.copy** and **.include** directives tell the assembler to begin reading source statements from another file. When the assembler finishes reading the source statements in the copy/include file, it resumes reading source statements from the current file. The statements read from a copied file are printed in the listing file; the statements read from an included file are *not* printed in the listing file.
- The **.def** directive identifies a symbol that is defined in the current module and that can be used in another module. The assembler includes the symbol in the symbol table.
- The **.global** directive declares a symbol external so that it is available to other modules at link time. (For more information about global symbols, see section 2.7.1, *External Symbols*, on page 2-18). The **.global** directive does double duty, acting as a **.def** for defined symbols and as a **.ref** for undefined symbols. The linker resolves an undefined global symbol reference only if the symbol is used in the program. The **.global** directive declares a 16-bit symbol.
- The **.mlib** directive supplies the assembler with the name of an archive library that contains macro definitions. When the assembler encounters a macro that is not defined in the current module, it searches for it in the macro library specified with **.mlib**.
- The **.ref** directive identifies a symbol that is used in the current module but is defined in another module. The assembler marks the symbol as an undefined external symbol and enters it in the object symbol table so the linker can resolve its definition. The **.ref** directive forces the linker to resolve a symbol reference.

4.7 Directives That Enable Conditional Assembly

Conditional assembly directives enable you to instruct the assembler to assemble certain sections of code according to a true or false evaluation of an expression. Two sets of directives allow you to assemble conditional blocks of code:

- The **.if/.elseif/.else/.endif** directives tell the assembler to conditionally assemble a block of code according to the evaluation of an expression.

.if [*well-defined expression*] marks the beginning of a conditional block and assembles code if the *.if well-defined expression* is true.

.elseif [*well-defined expression*] marks a block of code to be assembled if the *.if well-defined expression* is false and the *.elseif* condition is true.

.else marks a block of code to be assembled if the *.if well-defined expression* is false and any *.elseif* conditions are false.

.endif marks the end of a conditional block and terminates the block.

- The **.loop/.break/.endloop** directives tell the assembler to repeatedly assemble a block of code according to the evaluation of an expression.

.loop [*well-defined expression*] marks the beginning of a repeatable block of code. The optional *expression* evaluates to the loop count.

.break [*well-defined expression*] tells the assembler to assemble repeatedly when the *.break well-defined expression* is false and to go to the code immediately after *.endloop* when the *expression* is true or omitted.

.endloop marks the end of a repeatable block.

The assembler supports several relational operators that are useful for conditional expressions. For more information about relational operators, see section 3.9.4, *Conditional Expressions*, on page 3-28.

4.8 Directives That Define Symbols at Assembly Time

Assembly-time symbol directives equate meaningful symbol names to constant values or strings.

- ❑ The **.asg** directive assigns a character string to a substitution symbol. The value is stored in the substitution symbol table. When the assembler encounters a substitution symbol, it replaces the symbol with its character string value. Substitution symbols can be redefined.

```
.asg "10, 20, 30, 40", coefficients  
  
.byte coefficients
```

- ❑ The **.cstruct/.endstruct** directives set up C structure definitions. The **.cunion/.endunion** directives set up C-like union definitions. The **.tag** directive assigns the C structure or union characteristics to a label.

The **.cstruct/.endstruct** directives allow you to organize your information into structures so that similar elements can be grouped together. The **.cunion/.endunion** directives allow you to organize your information into unions so that similar elements can be grouped together. Element offset calculation is left up to the assembler. These directives do not allocate memory. They simply create a symbolic template that can be used repeatedly.

The **.tag** directive assigns a label to a structure. This simplifies the symbolic representation and also provides the ability to define structures that contain other structures. The **.tag** directive does not allocate memory, and the structure tag (**stag**) must be defined before it is used.

- ❑ The **.eval** directive evaluates a well-defined expression, translates the results into a character string, and assigns the character string to a substitution symbol. This directive is most useful for manipulating counters:

```
.asg      1 , x  
.loop  
.byte    x*10h  
.break   x = 4  
.eval    x+1, x  
.endloop
```

- ❑ The **.label** directive defines a special symbol that refers to the load-time address within the current section. This is useful when a section loads at one address but runs at a different address. For example, you may want to load a block of performance-critical code into slower off-chip memory to save space and move the code to high-speed on-chip memory to run. See page 4-50 for an example using a load-time address label.

- The **.set** and **.equ** directives set a constant value to a symbol. The symbol is stored in the symbol table and cannot be redefined; for example:

```
bval .set 1000h
     .long bval, bval*2, bval+12
     MVK   bval, A2
```

The **.set** and **.equ** directives produce no object code. The two directives are identical and can be used interchangeably.

- The **.struct/endstruct** directives set up C-like structure definitions. The **.union/endunion** directives set up C-like union definitions. The **.tag** directive assigns the C-like structure or union characteristics to a label.

The **.struct/endstruct** directives allow you to organize your information into structures so that similar elements can be grouped together. Similarly, the **.union/endunion** directives allow you to organize your information into unions. Element offset calculation is left up to the assembler. These directives do not allocate memory. They simply create a symbolic template that can be used repeatedly.

The **.tag** directive assigns a label to a structure or union. This simplifies the symbolic representation and also provides the ability to define structures that contain other structures. The **.tag** directive does not allocate memory, and the structure tag (**stag**) must be defined before it is used.

```
COORDT .struct                               ; structure tag definition
X      .byte
Y      .byte
T_LEN .endstruct

COORD .tag COORDT                            ; declare COORD (coordinate)
     .bss COORD, T_LEN                       ; actual memory allocation

LDB   *+B14(COORD.Y), A2 ; move member Y of structure
     ; COORD into register A2.
```

4.9 Miscellaneous Directives

These directives enable miscellaneous functions or features:

- The **.asmfunc** and **.endasmfunc** directives mark function boundaries. These directives are used with the compiler `-gw` option to generate debug information for separate functions.
- The **.clink** directive sets the `STYP_CLINK` flag in the type field for the named section. The `.clink` directive can be applied to initialized or uninitialized sections. The `STYP_CLINK` flag enables conditional linking by telling the linker to leave the section out of the final COFF output of the linker if there are no references found to any symbol in the section.
- The **.end** directive terminates assembly. If you use the `.end` directive, it should be the last source statement of a program. This directive has the same effect as an end-of-file character.
- The **.newblock** directive resets local labels. Local labels are symbols of the form `$n`, where `n` is a decimal digit, or of the form `NAME?`, where you specify `NAME`. They are defined when they appear in the label field. Local labels are temporary labels that can be used as operands for jump instructions. The `.newblock` directive limits the scope of local labels by resetting them after they are used. For more information, see section 3.8.2, *Local Labels*, on page 3-18.

These three directives enable you to define your own error and warning messages:

- The **.emsg** directive sends error messages to the standard output device. The `.emsg` directive generates errors in the same manner as the assembler, incrementing the error count and preventing the assembler from producing an object file.
- The **.mmsg** directive sends assembly-time messages to the standard output device. The `.mmsg` directive functions in the same manner as the `.emsg` and `.wmsg` directives but does not set the error count or the warning count. It does not affect the creation of the object file.
- The **.wmsg** directive sends warning messages to the standard output device. The `.wmsg` directive functions in the same manner as the `.emsg` directive but increments the warning count rather than the error count. It does not affect the creation of the object file.

For more information about using the error and warning directives in macros, see section 5.7, *Producing Messages in Macros*, on page 5-17.

4.10 Directives Reference

The remainder of this chapter is a reference. Generally, the directives are organized alphabetically, one directive per page. Related directives (such as `.if/.else/.endif`), however, are presented together on one page.

.align

Align SPC on the Next Boundary

Syntax

`.align` [*size in bytes*]

Description

The `.align` directive aligns the section program counter (SPC) on the next boundary, depending on the *size in bytes* parameter. The *size* can be any power of 2, although only certain values are useful for alignment. An operand of 1 aligns the SPC on the next byte boundary, and this is the default if no *size in bytes* is given. The assembler assembles words containing null values (0) up to the next *size in bytes* boundary:

Operand of	1	aligns SPC to byte boundary
	2	aligns SPC to halfword boundary
	4	aligns SPC to word boundary
	8	aligns SPC to doubleword boundary
	128	aligns SPC to page boundary

Using the `.align` directive has two effects:

- The assembler aligns the SPC on an *x*-byte boundary *within* the current section.
- The assembler sets a flag that forces the linker to align the section so that individual alignments remain intact when a section is loaded into memory.

Example

This example shows several types of alignment, including `.align 2`, `.align 8`, and a default `.align`.

```
1 00000000 00000004      .byte      4
2                          .align      2
3 00000002 00000045      .string    "Errorcnt"
  00000003 00000072
  00000004 00000072
  00000005 0000006F
  00000006 00000072
  00000007 00000063
  00000008 0000006E
  00000009 00000074
4
5 00000008 0003746E      .align
6 00000008 002B746E      .field    3,3
7                          .field    5,4
8                          .align      2
9 0000000c 00000003      .field    3,3
10                          .align      8
11                          .field    5,4
12 00000010 00000005      .align
13 00000011 00000004      .byte      4
```

.asg/.eval

Assign a Substitution Symbol

Syntax

```
.asg ["character string"], substitution symbol
.eval well-defined expression, substitution symbol
```

Description

The `.asg` directive assigns character strings to substitution symbols. Substitution symbols are stored in the substitution symbol table. The `.asg` directive can be used in many of the same ways as the `.set` directive, but while `.set` assigns a constant value (which cannot be redefined) to a symbol, `.asg` assigns a character string (which can be redefined) to a substitution symbol.

- The assembler assigns the *character string* to the substitution symbol. The quotation marks are optional. If there are no quotation marks, the assembler reads characters up to the first comma and removes leading and trailing blanks. In either case, a character string is read and assigned to the substitution symbol.
- The *substitution symbol* must be a valid symbol name. The substitution symbol is up to 128 characters long and must begin with a letter. Remaining characters of the symbol can be a combination of alphanumeric characters, the underscore (`_`), and the dollar sign (`$`).

The `.eval` directive performs arithmetic on substitution symbols, which are stored in the substitution symbol table. This directive evaluates the *well-defined expression* and assigns the string value of the result to the substitution symbol. The `.eval` directive is especially useful as a counter in `.loop/.endloop` blocks.

- ❑ The *well-defined expression* is an alphanumeric expression in which all symbols have been previously defined in the current source module, so that the result is an absolute.
- ❑ The *substitution symbol* must be a valid symbol name. The substitution symbol is up to 128 characters long and must begin with a letter. Remaining characters of the symbol can be a combination of alphanumeric characters, the underscore (_), and the dollar sign (\$).

Example

This example shows how .asg and .eval can be used.

```
1          .sslist ; show expanded substitution symbols
2
3          .asg    *+B14(100), GLOB100
4          .asg    *+B15(4),  ARG0
5
6 00000000 003B22E4      LDW    GLOB100,A0
#          LDW    *+B14(100),A0
7 00000004 00BC22E4      LDW    ARG0,A1
#          LDW    *+B15(4),A1
8 00000008 00006000      NOP    4
9 0000000c 010401E0      ADD    A0,A1,A2
10
11          .asg    0,x
12          .loop  5
13          .word  100*x
14          .eval  x+1,x
15          .endloop
1          00000010 00000000      .word  100*x
#          .word  100*0
1          .eval  x+1,x
#          .eval  0+1,x
1          00000014 00000064      .word  100*x
#          .word  100*1
1          .eval  x+1,x
#          .eval  1+1,x
1          00000018 000000C8      .word  100*x
#          .word  100*2
1          .eval  x+1,x
#          .eval  2+1,x
1          0000001c 0000012C      .word  100*x
#          .word  100*3
1          .eval  x+1,x
#          .eval  3+1,x
1          00000020 00000190      .word  100*x
#          .word  100*4
1          .eval  x+1,x
#          .eval  4+1,x
```

.asmfunc/ .endasmfunc

Mark Function Boundaries

Syntax

symbol **.asmfunc**
.endasmfunc

Description

The **.asmfunc** and **.endasmfunc** directives mark function boundaries. These directives are used with the compiler `-g` option (`--symdebug:dwarf`) to allow sections assembly code to be debugged in the same manner as C/C++ functions.

You should not use the same directives generated by the compiler (see Appendix B) to accomplish assembly debugging; those directives should be used only by the compiler to generate symbolic debugging information for C/C++ source files.

The **.asmfunc** and **.endasmfunc** directives cannot be used when invoking the compiler with the backwards-compatibility `--symdebug:coff` option. This option instructs the compiler to use the obsolete COFF symbolic debugging format, which does not support these directives.

The *symbol* is a label that must appear in the label field.

Consecutive ranges of assembly code that are not enclosed within a pair of **.asmfunc** and **.endasmfunc** directives are given a default name in the following format:

\$filename:beginning source line:ending source line\$

Example

In this example the assembly source generates debug information for the user_func section.

```
1 00000000                                .sect   ".text"
2                                           .global userfunc
3                                           .global _printf
4
5                                           userfunc: .asmfunc
6 00000000 00000010!                       CALL    .S1      _printf
7 00000004 01BC94F6                       STW     .D2T2   B3,*B15--(16)
8 00000008 01800E2A'                       MVKHL  .S2      RL0,B3
9 0000000c 01800028+                       MVKHL  .S1      SL1+0,A3
10 00000010 01800068+                      MVKHL  .S1      SL1+0,A3
11
12 00000014 01BC22F5                       STW     .D2T1   A3,*+B15(4)
13 00000018 0180006A' ||                      MVKHL  .S2      RL0,B3
14
15 0000001c 01BC92E6  RL0:                      LDW     .D2T2   *++B15(16),B3
16 00000020 020008C0                       ZERO    .D1      A4
17 00000024 00004000                       NOP
18 00000028 000C0362                       RET     .S2      B3
19 0000002c 00008000                       NOP
20                                           .endasmfunc
21
22 00000000                                .sect   ".const"
23 00000000 00000048  SL1:                      .string "Hello World!",10,0
    00000001 00000065
    00000002 0000006C
    00000003 0000006C
    00000004 0000006F
    00000005 00000020
    00000006 00000057
    00000007 0000006F
    00000008 00000072
    00000009 0000006C
    0000000a 00000064
    0000000b 00000021
    0000000c 0000000A
    0000000d 00000000
```

.bss

Reserve Space in the .bss Section

Syntax

.bss *symbol*, *size in bytes* [, *alignment* [, *bank offset*]]

Description

The **.bss** directive reserves space for variables in the .bss section. This directive is usually used to allocate space in RAM.

- ❑ The *symbol* is a required parameter. It defines a label that points to the first location reserved by the directive. The symbol name must correspond to the variable that you are reserving space for.
- ❑ The *size in bytes* is a required parameter; it must be an absolute expression. The assembler allocates size bytes in the .bss section.
- ❑ The *alignment* is an optional parameter that ensures that the space allocated to the symbol occurs on the specified boundary. This boundary indicates the size of the slot in bytes and must be set to a power of 2. If the SPC is aligned to the specified boundary, it is not incremented.
- ❑ The *bank offset* is an optional parameter that ensures that the space allocated to the symbol occurs on a specific memory bank boundary. The bank offset value measures the number of bytes to offset from the alignment specified before assigning the symbol to that location.

For more information about COFF sections, see Chapter 2, *Introduction to Common Object File Format*.

Example

In this example, the .bss directive is used to allocate space for a variable, array. The symbol array points to 100 bytes of uninitialized space (at .bss SPC = 0). Symbols declared with the .bss directive can be referenced in the same manner as other symbols and can also be declared global.

```

1          *****
2          ** Start assembling into .text section. **
3          *****
4 00000000          .text
5 00000000 008001A0          MV      A0,A1
6
7          *****
8          ** Allocate 100 bytes in .bss.          **
9          *****
10 00000000          .bss      array,100
11
12          *****
13          ** Still in .text          **
14          *****
15 00000004 010401A0          MV      A1,A2
16
17          *****
18          ** Declare external .bss symbol          **
19          *****
20          .global array

```

.byte/.char

Initialize Byte

Syntax

```

.byte value1 [, ... , valuen]
.char value1 [, ... , valuen]

```

Description

The **.byte** and **.char** directives place one or more values into consecutive bytes of the current section. A *value* can be one of the following:

- An expression that the assembler evaluates and treats as an 8-bit signed number
- A character string enclosed in double quotes. Each character in a string represents a separate value, and values are stored in consecutive bytes. The entire string *must* be enclosed in quotes.

The first byte occupies the eight least significant bits of a full 32-bit word. The second byte occupies bits eight through 15 while the third byte occupies bits 16 through 23. The assembler truncates values greater than eight bits. You can use up to 100 value parameters, but the total line length cannot exceed 200 characters.

If you use a label, it points to the location of the first byte that is initialized.

When you use **.byte** or **.char** in a **.struct/.endstruct** sequence, **.byte** and **.char** define a member's size; they do not initialize memory. For more information about **.struct/.endstruct**, see page 4-67.

Example

In this example, 8-bit values (10, -1, abc, and a) are placed into consecutive bytes in memory with **.byte** and **.char**. The label **strx** has the value 0h, which is the location of the first initialized byte. The label **stry** has the value 6h, which is the first byte initialized by the **.char** directive.

```

1 00000000 0000000A strx .byte 10,-1,"abc",'a'
   00000001 000000FF
   00000002 00000061
   00000003 00000062
   00000004 00000063
   00000005 00000061
2 00000006 00000008 stry .char 8,-3,"def",'b'
   00000007 000000FD
   00000008 00000064
   00000009 00000065
   0000000a 00000066
   0000000b 00000062

```

Syntax**.clink** ["*section name*"]**Description**

The **.clink** directive sets up conditional linking for a section by setting the STYP_CLINK flag in the type field for *section name*. The **.clink** directive can be applied to initialized or uninitialized sections.

The *section name* identifies the section. If **.clink** is used without a section name, it applies to the current initialized section. If **.clink** is applied to an uninitialized section, the section name is required. The section name is significant to 200 characters and must be enclosed in double quotes. A section name can contain a subsection name in the form *section name:subsection name*.

The **.clink** directive tells the linker to leave the section out of the final COFF output of the linker if there are no references found in a linked section to any symbol defined in the specified section. The **-a** linker option produces the final COFF output in the form of an absolute, executable output module.

A section in which the entry point of a C program is defined cannot be marked as a conditionally linked section.

Example

In this example, the Vars and Counts sections are set for conditional linking.

```
1 00000000          .sect "Vars"
2                  .clink
3                  ; Vars section is conditionally linked
4
5 00000000 0000001A X:      .word 01Ah
6 00000004 0000001A Y:      .word 01Ah
7 00000008 0000001A Z:      .word 01Ah
8 00000000          .sect "Counts"
9                  .clink
10                 ; Counts section is conditionally linked
11
12 00000000 0000001A XCount: .word 01Ah
13 00000004 0000001A YCount: .word 01Ah
14 00000008 0000001A ZCount: .word 01Ah
15 00000000          .text
16                 ; By default, .text is unconditionally linked
17
18 00000000 00B802C4      LDH    *B14,A1
19 00000004 00000028+    MVKL   X,A0
20 00000008 00000068+    MVKH   X,A0
21                 ; These references to symbol X cause the Vars
22                 ; section to be linked into the COFF output
23 0000000c 00040AF8      CMPLT  A0,A1,A0
```

.copy/.include*Copy Source File*

Syntax

```
.copy [""]filename[""]  
.include [""]filename[""]
```

Description

The **.copy** and **.include** directives tell the assembler to read source statements from a different file. The statements that are assembled from a copy file are printed in the assembly listing. The statements that are assembled from an included file are *not* printed in the assembly listing, regardless of the number of **.list/.nolist** directives assembled.

When a **.copy** or **.include** directive is assembled, the assembler:

- 1) Stops assembling statements in the current source file
- 2) Assembles the statements in the copied/included file
- 3) Resumes assembling statements in the main source file, starting with the statement that follows the **.copy** or **.include** directive

The *filename* is a required parameter that names a source file. It can be enclosed in double quotes and must follow operating system conventions. If *filename* starts with a number the double quotes are required.

You can specify a full pathname (for example, /320tools/file1.asm). If you do not specify a full pathname, the assembler searches for the file in:

- 1) The directory that contains the current source file
- 2) Any directories named with the **-i** assembler option
- 3) Any directories specified by the **C6X_A_DIR** or **A_DIR** environment variable

For more information about the **-i** option, **C6x_A_DIR**, and **A_DIR**, see section 3.4, *Naming Alternate Directories for Assembler Input*, on page 3-7.

The **.copy** and **.include** directives can be nested within a file being copied or included. The assembler limits nesting to 32 levels; the host operating system may set additional restrictions. The assembler precedes the line numbers of copied files with a letter code to identify the level of copying. An A indicates the first copied file, B indicates a second copied file, etc.

Example 1

In this example, the `.copy` directive is used to read and assemble source statements from other files; then, the assembler resumes assembling into the current file.

The original file, `copy.asm`, contains a `.copy` statement copying the file `byte.asm`. When `copy.asm` assembles, the assembler copies `byte.asm` into its place in the listing (note listing below). The copy file `byte.asm` contains a `.copy` statement for a second file, `word.asm`.

When it encounters the `.copy` statement for `word.asm`, the assembler switches to `word.asm` to continue copying and assembling. Then the assembler returns to its place in `byte.asm` to continue copying and assembling. After completing assembly of `byte.asm`, the assembler returns to `copy.asm` to assemble its remaining statement.

copy.asm (source file)	byte.asm (first copy file)	word.asm (second copy file)
<pre>.space 29 .copy "byte.asm" **Back in original file .string "done"</pre>	<pre>** In byte.asm .byte 32,1+ 'A' .copy "word.asm" ** Back in byte.asm .byte 67h + 3q</pre>	<pre>** In word.asm .word 0ABCDh, 56q</pre>

Listing file:

```

1 00000000          .space 29
2                    .copy "byte.asm"
A 1                    ** In byte.asm
A 2 0000001d 00000020    .byte 32,1+ 'A'
   0000001e 00000042
A 3                    .copy "word.asm"
B 1                    ** In word.asm
B 2 00000020 0000ABCD    .word 0ABCDh, 56q
   00000024 0000002E
A 4                    ** Back in byte.asm
A 5 00000028 0000006A    .byte 67h + 3q
   3
4                    ** Back in original file
5 00000029 00000064    .string "done"
   0000002a 0000006F
   0000002b 0000006E
   0000002c 00000065
```

Example 2

In this example, the `.include` directive is used to read and assemble source statements from other files; then, the assembler resumes assembling into the current file. The mechanism is similar to the `.copy` directive, except that statements are not printed in the listing file.

copy.asm (source file)	byte2.asm (first include file)	word2.asm (second include file)
<pre>.space 29 .include "byte2.asm" **Back in original file .string "done"</pre>	<pre>** In byte2.asm .byte 32,1+ 'A' .include "word2.asm" ** Back in byte.asm .byte 67h + 3q</pre>	<pre>** In word2.asm .word 0ABCDh, 56q</pre>

Listing file:

```
1 00000000                .space 29
2                        .include "byte2.asm"
3
4                        ** Back in original file
5 00000029 00000064      .string "done"
0000002a 0000006F
0000002b 0000006E
0000002c 00000065
```

**.cstruct/.cunion
.endstruct/
.endunion/.tag****Declare C Structure Type**

Syntax

```
[stag]    .cstruct|.cunion  [expr]
[mem0]   element          [expr0]
[mem1]   element          [expr1]
.         .                .
.         .                .
.         .                .
[memn]   .tag stag        [exprn]
.         .                .
.         .                .
.         .                .
[memN]   element          [exprN]
[size]    .endstruct|.endunion
label    .tag              stag
```

Description

The `.cstruct` and `.cunion` directives have been added to support ease of sharing of common data structures between assembly and C code. The `.cstruct` and `.cunion` directives can be used exactly like the existing `.struct` and `.union` directives except that they are guaranteed to perform data layout matching the layout used by the C compiler for C struct and union data types.

In particular, the `.cstruct` and `.cunion` directives force the same alignment and padding as used by the C compiler when such types are nested within compound data structures.

The `.endstruct` directive terminates the structure definition. The `.endunion` directive terminates the union definition.

The `.tag` directive gives structure characteristics to a *label*, simplifying the symbolic representation and providing the ability to define structures that contain other structures. The `.tag` directive does not allocate memory. The structure tag (*stag*) of a `.tag` directive must have been previously defined.

Following are descriptions of the parameters used with the `.struct`, `.endstruct`, and `.tag` directives:

- The *element* is one of the following descriptors: `.string`, `.byte`, `.char`, `.int`, `.half`, `.short`, `.word`, `.long`, `.double`, `.float`, `.tag`, or `.field`. All of these except `.tag` are typical directives that initialize memory. Following a `.struct` directive, these directives describe the structure element's size. They do not allocate memory. A `.tag` directive is a special case because `stag` must be used (as in the definition of `stag`).
- The *expr* is an optional expression indicating the beginning offset of the structure. The default starting point for a structure is 0.
- The *expr_{n/N}* is an optional expression for the number of elements described. This value defaults to 1. A `.string` element is considered to be one byte in size, and a `.field` element is one bit.
- The *mem_{n/N}* is an optional label for a member of the structure. This label is absolute and equates to the present offset from the beginning of the structure. A label for a structure member cannot be declared global.
- The *size* is an optional label for the total size of the structure.
- The *stag* is the structure's tag. Its value is associated with the beginning of the structure. If no `stag` is present, the assembler puts the structure members in the global symbol table with the value of their absolute offset from the top of the structure. A `.stag` is optional for `.struct`, but is required for `.tag`.

Example

This example illustrates a structure in C that will be accessed in assembly code.

```
; typedef struct STRUCT1
; {
;     int i0;      /* offset 0 */
;     short s0;   /* offset 4 */
; } struct1;      /* size 8, alignment 4 */
;
; typedef struct STRUCT2
; {
;     struct1 st1; /* offset 0 */
;     short s1;   /* offset 8 */
; } struct2;      /* size 12, alignment 4 */
;
; The structure will get the following offsets once
; the C compiler lays out the structure elements according
; to the C standard rules:
;
; offsetof(struct1, i0) = 0
; offsetof(struct1, s0) = 4
; sizeof(struct1)      = 8
;
; offsetof(struct2, s1) = 0
; offsetof(struct2, i1) = 8
; sizeof(struct2)      = 12
;
;
; Attempts to replicate this structure in assembly using the
; .struct/.union directive will not create the correct offsets
; because the assembler tries to use the most compact arrangement:

struct1    .struct
i0         .int           ; bytes 0-3
s0         .short        ; bytes 4-5
struct1len .endstruct    ; size 6, alignment 4

struct2    .struct
st1        .tag struct1  ; bytes 0-5
s1         .short        ; bytes 6-7
endstruct2 .endstruct    ; size 8, alignment 4

        .sect "data1"
        .word struct1.i0      ; 0
        .word struct1.s0      ; 4
        .word struct1len     ; 6

        .sect "data2"
        .word struct2.st1     ; 0
        .word struct2.s1     ; 6
        .word endstruct2     ; 8

;
```

.data

```
; The .cstruct/.union directives will calculate
; the offsets in the same manner as the C compiler. The
; resulting assembly structure can be used to access the
; elements of the C structure. Notice the different in
; the offsets from those structures defined via .struct
; above, and compare them to the offsets for the C code.
```

```
cstruct1    .cstruct
i0          .int           ; bytes 0-3
s0          .short        ; bytes 4-5
cstruct1len .endstruct    ; size 8, alignment 4

cstruct2    .cstruct
st1         .tag cstruct1 ; bytes 0-7
s1          .short        ; bytes 8-9
cendstruct2 .endstruct    ; size 12, alignment 4

.sect "data3"
.word cstruct1.i0, struct1.i0 ; 0
.word cstruct1.s0, struct1.s0 ; 4
.word cstruct1len, struct1len ; 8

.sect "data4"
.word cstruct2.st1, struct2.st1 ; 0
.word cstruct2.s1, struct2.s1 ; 8
.word cendstruct2, endstruct2 ; 12
```

.data

Assemble Into .data Section

Syntax

.data

Description

The **.data** directive tells the assembler to begin assembling source code into the **.data** section; **.data** becomes the current section. The **.data** section is normally used to contain tables of data or preinitialized variables.

For more information about COFF sections, see Chapter 2, *Introduction to Common Object File Format*.

Example

In this example, code is assembled into the .data and .text sections.

```
1 *****
2 **      Reserve space in .data      **
3 *****
4 00000000      .data
5 00000000      .space 0CCh
6
7 *****
8 **      Assemble into .text      **
9 *****
10 00000000     .text
11 00000000 00800358     ABS      A0,A1
12
13 *****
14 **      Assemble into .data      **
15 *****
16 000000cc     table: .data
17 000000cc FFFFFFFF     .word   -1
18 000000d0 000000FF     .byte   0FFh
19
20 *****
21 **      Assemble into .text      **
22 *****
23 00000004     .text
24 00000004 008001A0     MV      A0,A1
25
26 *****
27 **      Resume assembling into the .data section **
28 *****
29 000000d1     .data
30 000000d4 00000000     coeff  .word   00h,0ah,0bh
    000000d8 0000000A
    000000dc 0000000B
```

.double*Initialize Double-Precision Floating-Point Value*

Syntax

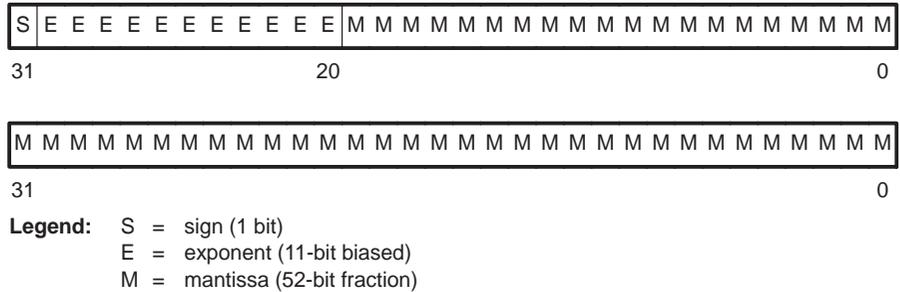
.double *value*₁ [, ... , *value*_{*n*}]

Description

The **.double** directive places the IEEE double-precision floating-point representation of one or more floating-point values into the current section. Each *value* must be a floating-point constant or a symbol that has been equated to a floating-point constant. Each constant is converted to a floating-point value in IEEE double-precision 64-bit format. Double-precision floating point constants are aligned to a double word boundary.

The 64-bit value is stored in the format shown in Figure 4–5.

Figure 4–5. Double-Precision Floating-Point Format



When you use `.double` in a `.struct/.endstruct` sequence, `.double` defines a member's size; it does not initialize memory. For more information about `.struct/.endstruct`, see page 4-67.

Example

This example shows the `.double` directive.

```
1 00000000 2C280291      .double -2.0e25
   00000004 C5308B2A
2 00000008 00000000      .double 6
   0000000c 40180000
3 00000010 00000000      .double 456
   00000014 407C8000
```

.drlist/.drnolist

Control Listing of Directives

Syntax

.drlist
.drnolist

Description

Two directives enable you to control the printing of assembler directives to the listing file:

The **.drlist** directive enables the printing of all directives to the listing file.

The **.drnolist** directive suppresses the printing of the following directives to the listing file. The `.drnolist` directive has no affect within macros.

- | | | |
|---|---|---|
| <input type="checkbox"/> <code>.asg</code> | <input type="checkbox"/> <code>.fcnolist</code> | <input type="checkbox"/> <code>.ssnolist</code> |
| <input type="checkbox"/> <code>.break</code> | <input type="checkbox"/> <code>.mlist</code> | <input type="checkbox"/> <code>.var</code> |
| <input type="checkbox"/> <code>.emsg</code> | <input type="checkbox"/> <code>.mmsg</code> | <input type="checkbox"/> <code>.wmsg</code> |
| <input type="checkbox"/> <code>.eval</code> | <input type="checkbox"/> <code>.mnolist</code> | |
| <input type="checkbox"/> <code>.fclist</code> | <input type="checkbox"/> <code>.sslist</code> | |

By default, the assembler acts as if the `.drlist` directive had been specified.

Example

This example shows how .drnolist inhibits the listing of the specified directives.

Source file:

```
.length 65
.width 85
.asg 0, x
.loop 2
.eval x+1, x
.endloop

.drnolist
.length 55
.width 95
.asg 1, x
.loop 3
.eval x+1, x
.endloop
```

Listing file:

```
3 .asg 0, x
4 .loop 2
5 .eval x+1, x
6 .endloop
1 .eval 0+1, x
1 .eval 1+1, x
7
8 .drnolist
12 .loop 3
13 .eval x+1, x
14 .endloop
```

**.emsg/.mmsg/
.wmsg**

Define Messages

Syntax

```
.emsg string
.mmsg string
.wmsg string
```

Description

These directives allow you to define your own error and warning messages. When you use these directives, the assembler tracks the number of errors and warnings it encounters and prints these numbers on the last line of the listing file.

The **.emsg** directive sends an error message to the standard output device in the same manner as the assembler. It increments the error count and prevents the assembler from producing an object file.

The **.mmsg** directive sends an assembly-time message to the standard output device in the same manner as the .emsg and .wmsg directives. It does not, however, set the error or warning counts, and it does not prevent the assembler from producing an object file.

The **.wmsg** directive sends a warning message to the standard output device in the same manner as the **.emsg** directive. It increments the warning count rather than the error count, however, and it does not prevent the assembler from producing an object file.

Example

In this example, the message **ERROR -- MISSING PARAMETER** is sent to the standard output device.

Source file:

```
.global PARAM
MSG_EX .macro parm1
      .if $symlen(parm1) = 0
      .emsg "ERROR -- MISSING PARAMETER"
      .else
      MVK parm1, A1
      .endif
      .endm
MSG_EX PARAM
MSG_EX
```

Listing file:

```
1 .global PARAM
2 MSG_EX .macro parm1
3 .if $symlen(parm1) = 0
4 .emsg "ERROR -- MISSING PARAMETER"
5 .else
6 MVK parm1, A1
7 .endif
8 .endm
9
10 00000000 MSG_EX PARAM
1 .if $symlen(parm1) = 0
1 .emsg "ERROR -- MISSING PARAMETER"
1 .else
1 00000000 00800028! MVK PARAM, A1
1 .endif
11
12 00000004 MSG_EX
1 .if $symlen(parm1) = 0
1 .emsg "ERROR -- MISSING PARAMETER"
1 ***** USER ERROR ***** - : ERROR -- MISSING PARAMETER
1 .else
1 MVK parm1, A1
1 .endif
1 Error, No Warnings
```

In addition, the following messages are sent to standard output by the assembler:

```
*** ERROR! line 12: ***** USER ERROR ***** - : ERROR -- MISSING PARAMETER
      .emsg "ERROR -- MISSING PARAMETER"
1 Assembly Error, No Assembly Warnings
Errors in source - Assembler Aborted
```

.end*End Assembly*

Syntax**.end****Description**

The **.end** directive is optional and terminates assembly. The assembler ignores any source statements that follow a **.end** directive. If you use the **.end** directive, it must be the last source statement of a program.

This directive has the same effect as an end-of-file character. You can use **.end** when you are debugging and you want to stop assembling at a specific point in your code.

Note: Ending a Macro

Do not use the **.end** directive to terminate a macro; use the **.endm** macro directive instead.

Example

This example shows how the **.end** directive terminates assembly. If any source statements follow the **.end** directive, the assembler ignores them.

Source file:

```
start:  .text
        ZERO    A0
        ZERO    A1
        ZERO    A3
        .end
        ZERO    A4
```

Listing file:

```
1 00000000          start:  .text
2 00000000 000005E0          ZERO    A0
3 00000004 008425E0          ZERO    A1
4 00000008 018C65E0          ZERO    A3
5                               .end
```


.field*Initialize Field***Syntax****.field** *value* [, *size in bits*]**Description**

The **.field** directive initializes a multiple-bit field within a single word of memory. This directive has two operands:

- ❑ The *value* is a required parameter; it is an expression that is evaluated and placed in the field. The value must be absolute.
- ❑ The *size in bits* is an optional parameter; it specifies a number from 1 to 32, which is the number of bits in the field. If you do not specify a size, the assembler assumes the size is 32 bits. If you specify a value that cannot fit in *size in bits*, the assembler truncates the value and issues a warning message. For example, `.field 3,1` causes the assembler to truncate the value 3 to 1; the assembler also prints the message:

```
*** WARNING! line 21: W0001: Field value truncated to 1
        .field 3,1
```

Successive `.field` directives pack values into the specified number of bits starting at the current 32-bit slot. Fields are packed starting at the least significant bit (bit 0), moving toward the most significant bit (bit 31) as more fields are added. If the assembler encounters a field size that does not fit in the current 32-bit word, it fills the remaining bits of the current byte with 0s, increments the SPC to the next word boundary, and begins packing fields into the next word.

You can use the `.align` directive to force the next `.field` directive to begin packing into a new word.

If you use a label, it points to the byte that contains the specified field.

When you use `.field` in a `.struct/.endstruct` sequence, `.field` defines a member's size; it does not initialize memory. For more information about `.struct/.endstruct`, see page 4-67.

Example

This example shows how fields are packed into a word. The SPC does not change until a word is filled and the next word is begun. Figure 4–6 shows how the directives in this example affect memory.

```
1 *****
2 ** Initialize a 24-bit field. **
3 *****
4 00000000 00BCCDD .field 0BCCDDh, 24
5
6 *****
7 ** Initialize a 5-bit field **
8 *****
9 00000000 0ABCCDD .field 0Ah, 5
10
11 *****
12 ** Initialize a 4-bit field **
13 ** in a new word. **
14 *****
15 00000004 0000000C .field 0Ch, 4
16
17 *****
18 ** Initialize a 3-bit field **
19 *****
20 00000004 0000001C x: .field 01h, 3
21
22 *****
23 ** Initialize a 32-bit field **
24 ** relocatable field in the **
25 ** next word **
26 *****
27 00000008 00000004' .field x
```


A symbol can be declared global for either of two reasons:

- ❑ If the symbol is *not defined in the current module* (which includes macro, copy, and include files), the `.global` or `.ref` directive tells the assembler that the symbol is defined in an external module. This prevents the assembler from issuing an unresolved reference error. At link time, the linker looks for the symbol's definition in other modules.
- ❑ If the symbol is *defined in the current module*, the `.global` or `.def` directive declares that the symbol and its definition can be used externally by other modules. These types of references are resolved at link time.

Example

This example shows four files. The `file1.lst` and `file2.lst` refer to each other for all symbols used; `file3.lst` and `file4.lst` are similarly related.

The **file1.lst** and **file3.lst** files are equivalent. Both files define the symbol `INIT` and make it available to other modules; both files use the external symbols `X`, `Y`, and `Z`. Also, `file1.lst` uses the `.global` directive to identify these global symbols; `file3.lst` uses `.ref` and `.def` to identify the symbols.

The **file2.lst** and **file4.lst** files are equivalent. Both files define the symbols `X`, `Y`, and `Z` and make them available to other modules; both files use the external symbol `INIT`. Also, `file2.lst` uses the `.global` directive to identify these global symbols; `file4.lst` uses `.ref` and `.def` to identify the symbols.

file1.lst

```
1           ; Global symbol defined in this file
2           .global INIT
3           ; Global symbols defined in file2.lst
4           .global X, Y, Z
5 00000000      INIT:
6 00000000 00902058      ADD.L1 0x01,A4,A1
7 00000004 00000000!    .word  X
8           ;
9           ;
10          ;
11          .end
```

file2.lst

```
1           ; Global symbols defined in this file
2           .global X, Y, Z
3           ; Global symbol defined in file1.lst
4           .global INIT
5          00000001 X:    .set  1
6          00000002 Y:    .set  2
7          00000003 Z:    .set  3
8 00000000 00000000!    .word  INIT
9           ;
10          ;
11          ;
12          .end
```

file3.lst

```
1           ; Global symbol defined in this file
2           .def      INIT
3           ; Global symbols defined in file4.lst
4           .ref      X, Y, Z
5 00000000      INIT:
6 00000000 00902058      ADD.L1 0x01,A4,A1
7 00000004 00000000!      .word   X
8           ;           .
9           ;           .
10          ;           .
11          .end
```

file4.lst

```
1           ; Global symbols defined in this file
2           .def      X, Y, Z
3           ; Global symbol defined in file3.lst
4           .ref      INIT
5 00000001 X:          .set     1
6 00000002 Y:          .set     2
7 00000003 Z:          .set     3
8 00000000 00000000!      .word   INIT
9           ;           .
10          ;           .
11          ;           .
12          .end
```

.half/.uhalf/.short/.ushort

Initialize 16-Bit Integers

Syntax

```
.half  value1 [, ... , valuen]  
.uhalf value1 [, ... , valuen]  
.short value1 [, ... , valuen]  
.ushort value1 [, ... , valuen]
```

Description

The **.half**, **.uhalf**, **.short**, and **.ushort** directives place one or more values into consecutive halfwords in the current section. Each value is placed in a 2-byte slot by itself. A *value* can be either:

- An expression that the assembler evaluates and treats as a 16-bit signed or unsigned number
- A character string enclosed in double quotes. Each character in a string represents a separate value and is stored alone in the least significant eight bits of a 16-bit field, which is padded with 0s.

The assembler truncates values greater than 16 bits. You can use as many values as fit on a single line, but the total line length cannot exceed 200 characters.

If you use a label with `.half` or `.short`, it points to the location where the assembler places the first byte.

The `.half` and `.short` directives perform a halfword (16-bit) alignment before data is written to the section. This guarantees that data resides on a 16-bit boundary.

When you use `.half` or `.short` in a `.struct/.endstruct` sequence, they define a member's size; they do not initialize memory. For more information about `.struct/.endstruct`, see page 4-67.

Example

In this example, `.half` is used to place 16-bit values (10, -1, abc, and a) into consecutive halfwords in memory; `.short` is used to place 16-bit values (8, -3, def, and b) into consecutive halfwords in memory. The label `STRN` has the value 100ch, which is the location of the first initialized halfword for `.short`.

```
1 00000000                                .space 100h * 16
2 00001000 0000000A                        .half 10, -1, "abc", 'a'
   00001002 0000FFFF
   00001004 00000061
   00001006 00000062
   00001008 00000063
   0000100a 00000061
3 0000100c 00000008 STRN .short 8, -3, "def", 'b'
   0000100e 0000FFFD
   00001010 00000064
   00001012 00000065
   00001014 00000066
   00001016 00000062
```

.if/.elseif/.else/.endif

Assemble Conditional Blocks

Syntax

```
.if well-defined expression
[.elseif well-defined expression]
[.else]
.endif
```

Description

Four directives provide conditional assembly:

The `.if` directive marks the beginning of a conditional block. The *well-defined expression* is a required parameter.

- If the expression evaluates to true (nonzero), the assembler assembles the code that follows the expression (up to a `.elseif`, `.else`, or `.endif`).
- If the expression evaluates to false (0), the assembler assembles code that follows a `.elseif` (if present), `.else` (if present), or `.endif` (if no `.elseif` or `.else` is present).

The **.elseif** directive identifies a block of code to be assembled when the **.if** expression is false (0) and the **.elseif** expression is true (nonzero). When the **.elseif** expression is false, the assembler continues to the next **.elseif** (if present), **.else** (if present), or **.endif** (if no **.elseif** or **.else** is present). The **.elseif** directive is optional in the conditional block, and more than one **.elseif** can be used. If an expression is false and there is no **.elseif** statement, the assembler continues with the code that follows a **.else** (if present) or a **.endif**.

The **.else** directive identifies a block of code that the assembler assembles when the **.if** expression and all **.elseif** expressions are false (0). The **.else** directive is optional in the conditional block; if an expression is false and there is no **.else** statement, the assembler continues with the code that follows the **.endif**.

The **.endif** directive terminates a conditional block.

The **.elseif** and **.else** directives can be used in the same conditional assembly block, and the **.elseif** directive can be used more than once within a conditional assembly block.

For information about relational operators, see subsection 3.9.4, *Conditional Expressions*, on page 3-28.

Example

This example shows conditional assembly:

```
1          00000001  SYM1      .set      1
2          00000002  SYM2      .set      2
3          00000003  SYM3      .set      3
4          00000004  SYM4      .set      4
5
6
7 00000000 00000004      If_4:  .if      SYM4 = SYM2 * SYM2
8                               .byte   SYM4          ; Equal values
9                               .else
10                              .byte   SYM2 * SYM2      ; Unequal values
11                              .endif
12
13 00000001 0000000A      If_5:  .if      SYM1 <= 10
14                              .byte   10             ; Less than / equal
15                              .else
16                              .byte   SYM1           ; Greater than
17                              .endif
18
19                               If_6:  .if      SYM3 * SYM2 != SYM4 + SYM2
20                               .byte   SYM3 * SYM2      ; Unequal value
21 00000002 00000008      .byte   SYM4 + SYM4      ; Equal values
22                               .endif
23
24                               If_7:  .if      SYM1 = SYM2
25                               .byte   SYM1
26                               .elseif SYM2 + SYM3 = 5
27 00000003 00000005      .byte   SYM2 + SYM3
28                               .endif
```

.int/.long/.word

Initialize 32-Bit Integer

Syntax

```

.int value1 [, ... , valuen]
.long value1 [, ... , valuen]
.word value1 [, ... , valuen]

```

Description

The **.int**, **.uint**, **.long**, **.word** and **.uword** directives place one or more values into consecutive words in the current section. Each value is placed in a 32-bit word by itself and is aligned on a word boundary. A *value* can be either:

- An expression that the assembler evaluates and treats as a 32-bit signed or unsigned number
- A character string enclosed in double quotes. Each character in a string represents a separate value and is stored alone in the least significant eight bits of a 32-bit field, which is padded with 0s.

A value can be either an absolute or a relocatable expression. If an expression is relocatable, the assembler generates a relocation entry that refers to the appropriate symbol; the linker can then correctly patch (relocate) the reference. This allows you to initialize memory with pointers to variables or labels.

You can use as many values as fit on a single line (200 characters). If you use a label with **.int**, **.long**, or **.word**, it points to the first word that is initialized.

When you use **.int**, **.long**, or **.word** directives in a **.struct**/.**endstruct** sequence, they define a member's size; they do not initialize memory. For more information about **.struct**/.**endstruct**, see page 4-67.

Example 1

This example uses the **.int** directive to initialize words. Notice that the symbol **SYMPTR** puts the symbol's address in the object code and generates a relocatable reference (indicated by the **-** character appended to the object word).

```

1 00000000          .space 73h
2 00000000          .bss  PAGE, 128
3 00000080          .bss  SYMPTR, 3
4 00000074 003C12E4 INST: LDW.D2 *++B15[0],A0
5 00000078 0000000A   .int  10, SYMPTR, -1, 35 + 'a', INST
   0000007c 00000080-
   00000080 FFFFFFFF
   00000084 00000084
   00000088 00000074'

```

Example 2

This example initializes two 32-bit fields and defines **DAT1** to point to the first location. The contents of the resulting 32-bit fields are **FFFABCDh** and **141h**.

```

1 00000000 FFFFABCD DAT1:  .long  0FFFABCDh, 'A' +100h
   00000004 00000141

```

Example 3 This example initializes five words. The symbol WordX points to the first word.

```
1 00000000 0000C80 WordX: .word 3200,1+'AB',-'AF',0F410h,'A'
00000004 00004242
00000008 FFFFB9BF
0000000c 0000F410
00000010 00000041
```

Note: Data Size of longs

For the C6000 C/C++ compiler, a long data value is 40 bits. For the C6000 assembler, a long data value is 32 bits. Therefore, the .long directive treats values assigned to it as 32-bit values.

.label

Create a Load-Time Address Label

Syntax

.label *symbol*

Description

The **.label** directive defines a special *symbol* that refers to the load-time address rather than the run-time address within the current section. Most sections created by the assembler have relocatable addresses. The assembler assembles each section as if it started at 0, and the linker relocates it to the address at which it loads and runs.

For some applications, it is desirable to have a section load at one address and run at a *different* address. For example, you may want to load a block of performance-critical code into slower memory to save space and then move the code to high-speed memory to run it. Such a section is assigned two addresses at link time: a load address and a run address. All labels defined in the section are relocated to refer to the run-time address so that references to the section (such as branches) are correct when the code runs.

The **.label** directive creates a special label that refers to the *load-time* address. This function is useful primarily to designate where the section was loaded for purposes of the code that relocates the section.

Example

This example shows the use of a load-time address label.

```
.sect ".examp"
.label examp_load ; load address of section
start:           ; run address of section
<code>
finish:         ; run address of section end
.label examp_end ; load address of section end
```

For more information about assigning run-time and load-time addresses in the linker, see section 7.9, *Specifying a Section's Run-Time Address*, on page 7-44.

.length/.width

Set Listing Page Size

Syntax

.length [*page length*]

.width [*page width*]

Description

Two directives allow you to control the size of the output listing file.

The **.length** directive sets the page length of the output listing file. It affects the current and following pages. You can reset the page length with another **.length** directive.

- Default length: 60 lines. If you do not use the **.length** directive or if you use the **.length** directive without specifying the *page length*, the output listing length defaults to 60 lines.
- Minimum length: 1 line
- Maximum length: 32 767 lines

The **.width** directive sets the page width of the output listing file. It affects the next line assembled and the lines following. You can reset the page width with another **.width** directive.

- Default width: 132 characters. If you do not use the **.width** directive or if you use the **.width** directive without specifying a *page width*, the output listing width defaults to 132 characters.
- Minimum width: 80 characters
- Maximum width: 200 characters

The width refers to a full line in a listing file; the line counter value, SPC value, and object code are counted as part of the width of a line. Comments and other portions of a source statement that extend beyond the page width are truncated in the listing.

The assembler does not list the **.width** and **.length** directives.

Example

The following example shows how to change the page length and width.

```
*****
**          Page length = 65 lines          **
**          Page width  = 85 characters     **
*****
          .length    65
          .width     85

*****
**          Page length = 55 lines          **
**          Page width  = 100 characters   **
*****
          .length    55
          .width     100
```

Syntax

.list
.nolist

Description

Two directives enable you to control the printing of the source listing:

The **.list** directive allows the printing of the source listing.

The **.nolist** directive suppresses the source listing output until a **.list** directive is encountered. The **.nolist** directive can be used to reduce assembly time and the source listing size. It can be used in macro definitions to suppress the listing of the macro expansion.

The assembler does not print the **.list** or **.nolist** directives or the source statements that appear after a **.nolist** directive. However, it continues to increment the line counter. You can nest the **.list/.nolist** directives; each **.nolist** needs a matching **.list** to restore the listing.

By default, the source listing is printed to the listing file; the assembler acts as if the **.list** directive had been used. However, if you do not request a listing file when you invoke the assembler by including the **-al** option on the command line (see page 3-5), the assembler ignores the **.list** directive.

Example

This example shows how the **.list** and **.nolist** directives turn the output listing on and off. The **.nolist**, the table: **.data** through **.byte** lines, and the **.list** directives do not appear in the listing file. Also, the line counter is incremented even when source statements are not listed.

Source file:

```
        .data
        .space  0CCh
        .text
        ABS     A0,A1

        .nolist

table:  .data
        .word   -1
        .byte   0FFh

        .list

        .text
        MV     A0,A1
        .data
coeff   .word   00h,0ah,0bh
```

Listing file:

```
1 00000000 .data
2 00000000 .space 0CCh
3 00000000 .text
4 00000000 00800358 ABS A0,A1
5
13
14 00000004 .text
15 00000004 008001A0 MV A0,A1
16 000000d1 .data
17 000000d4 00000000 coeff .word 00h,0ah,0bh
    000000d8 0000000A
    000000dc 0000000B
```

.loop/.break/.endloop

Assemble Code Block Repeatedly

Syntax

```
.loop [well-defined expression]  
.break [well-defined expression]  
.endloop
```

Description

Three directives allow you to repeatedly assemble a block of code:

The **.loop** directive begins a repeatable block of code. The optional expression evaluates to the loop count (the number of loops to be performed). If there is no *well-defined expression*, the loop count defaults to 1024, unless the assembler first encounters a **.break** directive with an expression that is true (nonzero) or omitted.

The **.break** directive, along with its expression, is optional. This means that when you use the **.loop** construct, you do not have to use the **.break** construct. The **.break** directive terminates a repeatable block of code only if the *well-defined expression* is true (nonzero) or omitted, and the assembler breaks the loop and assembles the code after the **.endloop** directive. If the expression is false (evaluates to 0), the loop continues.

The **.endloop** directive terminates a repeatable block of code; it executes when the **.break** directive is true (nonzero) or when the number of loops performed equals the loop count given by **.loop**.

Example

This example illustrates how these directives can be used with the `.eval` directive. The code in the first six lines expands to the code immediately following those six lines.

```

1          1          .eval      0,x
2          2          COEF .loop
3          3          .word     x*100
4          4          .eval     x+1, x
5          5          .break    x = 6
6          6          .endloop
1          1          00000000 00000000 .word     0*100
1          1          .eval     0+1, x
1          1          .break    1 = 6
1          1          00000004 00000064 .word     1*100
1          1          .eval     1+1, x
1          1          .break    2 = 6
1          1          00000008 000000C8 .word     2*100
1          1          .eval     2+1, x
1          1          .break    3 = 6
1          1          0000000c 0000012C .word     3*100
1          1          .eval     3+1, x
1          1          .break    4 = 6
1          1          00000010 00000190 .word     4*100
1          1          .eval     4+1, x
1          1          .break    5 = 6
1          1          00000014 000001F4 .word     5*100
1          1          .eval     5+1, x
1          1          .break    6 = 6
```

.macro/.endm

Define Macro

Syntax

```
macname      .macro [parameter1] [, ... parametern]
               model statements or macro directives
               .endm
```

Description

The `.macro` directive is used to define macros.

You can define a macro anywhere in your program, but you must define the macro before you can use it. Macros can be defined at the beginning of a source file, in an `.include/.copy` file, or in a macro library.

- macname* names the macro. You must place the name in the source statement's label field.
- `.macro` identifies the source statement as the first line of a macro definition. You must place `.macro` in the opcode field.
- [*parameters*] are optional substitution symbols that appear as operands for the `.macro` directive.
- model statements* are instructions or assembler directives that are executed each time the macro is called.

macro directives are used to control macro expansion.

.endm marks the end of the macro definition.

Macros are explained in further detail in Chapter 5, *Macro Language*.

.mlib

Define Macro Library

Syntax

```
.mlib ["filename"]
```

Description

The **.mlib** directive provides the assembler with the *filename* of a macro library. A macro library is a collection of files that contain macro definitions. The macro definition files are bound into a single file (called a library or archive) by the archiver.

Each file in a macro library contains one macro definition that corresponds to the name of the file. The *filename* of a macro library member must be the same as the macro name, and its extension must be `.asm`. The filename must follow host operating system conventions; it can be enclosed in double quotes. You can specify a full pathname (for example, `c:\320tools\macs.lib`). If you do not specify a full pathname, the assembler searches for the file in the following locations in the order given:

- 1) The directory that contains the current source file
- 2) Any directories named with the `-i` assembler option
- 3) Any directories specified by the `C6X_A_DIR` or `A_DIR` environment variable

For more information about the `-i` option, `C6X_A_DIR`, and `A_DIR`, see section 3.4, *Naming Alternate Directories for Assembler Input*, on page 3-7.

When the assembler encounters a `.mlib` directive, it opens the library specified by the *filename* and creates a table of the library's contents. The assembler enters the names of the individual library members into the opcode table as library entries. This redefines any existing opcodes or macros that have the same name. If one of these macros is called, the assembler extracts the entry from the library and loads it into the macro table. The assembler expands the library entry in the same way it expands other macros, but it does not place the source code into the listing. Only macros that are actually called from the library are extracted, and they are extracted only once.

For more information on macros and macro libraries, see Chapter 5, *Macro Language*.

Example

This example creates a macro library that defines two macros, inc1 and dec1. The file inc1.asm contains the definition of inc1, and dec1.asm contains the definition of dec1.

inc1.asm	dec1.asm
* Macro for incrementing inc1 .macro A ADD A,1,A .endm	* Macro for decrementing dec1 .macro A SUB A,1,A .endm

Use the archiver to create a macro library:

```
ar6x -a mac inc1.asm dec1.asm
```

Now you can use the .m1ib directive to reference the macro library and define the inc1 and dec1 macros:

```
1 .m1ib "mac.lib"  
2  
3 * Macro Call  
4 00000000 inc1 A0  
1 00000000 000021A0 ADD A0,1,A0  
5  
6 * Macro Call  
7 00000004 dec1 B0  
1 00000004 0003E1A2 SUB B0,1,B0
```

.m1ist/.mno1ist

Start/Stop Macro Expansion Listing

Syntax

```
.m1ist  
.mno1ist
```

Description

Two directives enable you to control the listing of macro and repeatable block expansions in the listing file:

The **.m1ist** directive allows macro and .loop/.endloop block expansions in the listing file.

The **.mno1ist** directive suppresses macro and .loop/.endloop block expansions in the listing file.

By default, the assembler behaves as if the .m1ist directive had been specified.

For more information on macros and macro libraries, see Chapter 5, *Macro Language*. For more information about .loop and .endloop, see page 4-53.

Example

This example defines a macro named STR_3. The first time the macro is called, the macro expansion is listed (by default). The second time the macro is called, the macro expansion is not listed, because a .mnolist directive was assembled. The third time the macro is called, the macro expansion is again listed because a .mlist directive was assembled.

```
1          STR_3  .macro  P1, P2, P3
2              .string ":p1:", ":p2:", ":p3:"
3              .endm
4
5 00000000          STR_3 "as", "I", "am"
1 00000000 0000003A  .string ":p1:", ":p2:", ":p3:"
00000001 00000070
00000002 00000031
00000003 0000003A
00000004 0000003A
00000005 00000070
00000006 00000032
00000007 0000003A
00000008 0000003A
00000009 00000070
0000000a 00000033
0000000b 0000003A
6
7 0000000c          .mnolist
8          STR_3 "as", "I", "am"
9          .mlist
1 00000018          STR_3 "as", "I", "am"
00000018 0000003A  .string ":p1:", ":p2:", ":p3:"
00000019 00000070
0000001a 00000031
0000001b 0000003A
0000001c 0000003A
0000001d 00000070
0000001e 00000032
0000001f 0000003A
00000020 0000003A
00000021 00000070
00000022 00000033
00000023 0000003A
```

Syntax**.newblock****Description**

The **.newblock** directive undefines any local labels currently defined. Local labels, by nature, are temporary; the **.newblock** directive resets them and terminates their scope.

A local label is a label in the form $\$n$, where n is a single decimal digit, or $name?$, where $name$ is a legal symbol name. Unlike other labels, local labels are intended to be used locally, cannot be used in expressions, and do not qualify for branch expansion if used with a branch. They can be used only as operands in 8-bit jump instructions. Local labels are not included in the symbol table.

After a local label has been defined and (perhaps) used, you should use the **.newblock** directive to reset it. The **.text**, **.data**, and **.sect** directives also reset local labels. Local labels that are defined within an include file are not valid outside of the include file.

For more information on the use of local labels, see subsection 3.8.2, *Local Labels*, on page 3-18.

Example

This example shows how the local label $\$1$ is declared, reset, and then declared again.

```
1                .global table1, table2
2
3 00000000 00000028!    MVKL    table1,A0
4 00000004 00000068!    MVKH    table1,A0
5 00000008 008031A9     MVK     99, A1
6 0000000c 010848C0    ||     ZERO   A2
7
8 00000010 80000212    $1:[A1] B     $1
9 00000014 01003674    STW     A2, *A0++
10 00000018 0087E1A0    SUB     A1,1,A1
11 0000001c 00004000    NOP     3
12
13                .newblock ; undefine $1
14
15 00000020 00000028!    MVKL    table2,A0
16 00000024 00000068!    MVKH    table2,A0
17 00000028 008031A9     MVK     99, A1
18 0000002c 010829C0    ||     SUB     A2,1,A2
19
20 00000030 80000212    $1:[A1] B     $1
21 00000034 01003674    STW     A2, *A0++
22 00000038 0087E1A0    SUB     A1,1,A1
23 0000003c 00004000    NOP     3
```

.option*Select Listing Options*

Syntax**.option** *option*₁[, *option*₂, . . .]**Description**

The **.option** directive selects options for the assembler output listing. The *options* must be separated by commas; each option selects a listing feature. These are valid options:

- A** turns on listing of all directives and data, and subsequent expansions, macros, and blocks.
- B** limits the listing of `.byte` and `.char` directives to one line.
- D** turns off the listing of certain directives (same effect as `.drnolist`).
- H** limits the listing of `.half` and `.short` directives to one line.
- L** limits the listing of `.long` directives to one line.
- M** turns off macro expansions in the listing.
- N** turns off listing (performs `.nolist`).
- O** turns on listing (performs `.list`).
- R** resets the B, H, L, M, T, and W directives (turns off the limits of B, H, L, M, T, and W).
- T** limits the listing of `.string` directives to one line.
- W** limits the listing of `.word` and `.int` directives to one line.
- X** produces a cross-reference listing of symbols. You can also obtain a cross-reference listing by invoking the assembler with the `-ax` option (see page 3-5).

Options *are not* case sensitive.

Example

This example shows how to limit the listings of the `.byte`, `.char`, `.int`, `.word`, and `.string` directives to one line each.

```
1 *****
2 ** Limit the listing of .byte, .char, **
3 ** .int, .word, and .string **
4 ** directives to 1 line each. **
5 *****
6 .option B, W, T
7 00000000 000000BD .byte -'C', 0B0h, 5
8 00000003 000000BC .char -'D', 0C0h, 6
9 00000008 0000000A .int 10, 35 + 'a', "abc"
10 0000001c AABBCDD .long 0AABBCDDh, 536 + 'A'
    00000020 00000259
11 00000024 000015AA .word 5546, 78h
12 0000002c 00000052 .string "Registers"
13
14 *****
15 ** Reset the listing options. **
16 *****
17 .option R
18 00000035 000000BD .byte -'C', 0B0h, 5
    00000036 000000B0
    00000037 00000005
19 00000038 000000BC .char -'D', 0C0h, 6
    00000039 000000C0
    0000003a 00000006
20 0000003c 0000000A .int 10, 35 + 'a', "abc"
    00000040 00000084
    00000044 00000061
    00000048 00000062
    0000004c 00000063
21 00000050 AABBCDD .long 0AABBCDDh, 536 + 'A'
    00000054 00000259
22 00000058 000015AA .word 5546, 78h
    0000005c 00000078
23 00000060 00000052 .string "Registers"
    00000061 00000065
    00000062 00000067
    00000063 00000069
    00000064 00000073
    00000065 00000074
    00000066 00000065
    00000067 00000072
    00000068 00000073
```

.page

Eject Page in Listing

Syntax

.page

Description

The **.page** directive produces a page eject in the listing file. The **.page** directive is not printed in the source listing, but the assembler increments the line counter when it encounters the **.page** directive. Using the **.page** directive to divide the source listing into logical divisions improves program readability.

Example

This example shows how the **.page** directive causes the assembler to begin a new page of the source listing.

Source file:

```

                .title    "**** Page Directive Example ****"
;               .
;               .
;               .
                .page

```

Listing file:

```

TMS320C6x COFF Assembler      Version x.xx      Tue Apr 14 17:16:51 1997
Copyright (c) 1996-1997 Texas Instruments Incorporated
**** Page Directive Example ****                                PAGE      1

                2                ;                .
                3                ;                .
                4                ;                .
TMS320C6x COFF Assembler      Version x.xx      Tue Apr 14 17:16:51 1997
Copyright (c) 1996-1997 Texas Instruments Incorporated
**** Page Directive Example ****                                PAGE      2

```

No Errors, No Warnings

.sect

Assemble Into Named Section

Syntax

.sect "section name"

Description

The **.sect** directive defines a named section that can be used like the default **.text** and **.data** sections. The **.sect** directive tells the assembler to begin assembling source code into the named section.

The *section name* identifies the section. The section name is significant to 200 characters and must be enclosed in double quotes. A section name can contain a subsection name in the form *section name:subsection name*.

For more information about COFF sections, see Chapter 2, *Introduction to Common Object File Format*.

Example

This example defines one special-purpose section, vars, and assembles code into it.

```
1 *****
2 ** Begin assembling into .text section. **
3 *****
4 00000000          .text
5 00000000 000005E0      ZERO    A0
6 00000004 008425E0      ZERO    A1
7
8 *****
9 ** Begin assembling into vars section. **
10 *****
11 00000000          .sect    "vars"
12 00000000 4048F5C3 pi    .float  3.14
13 00000004 000007D0 max   .int    2000
14 00000008 00000001 min   .int    1
15
16 *****
17 ** Resume assembling into .text section. **
18 *****
19 00000008          .text
20 00000008 010000A8      MVK     1,A2
21 0000000c 018000A8      MVK     1,A3
22
23 *****
24 ** Resume assembling into vars section. **
25 *****
26 0000000c          .sect    "vars"
27 0000000c 00000019 count .short  25
```

.set/.equ

Define Assembly-Time Constant

Syntax

```
symbol .set value
symbol .equ value
```

Description

The **.set** and **.equ** directives equate a constant value to a symbol. The symbol can then be used in place of a value in assembly source. This allows you to equate meaningful names with constants and other values. The **.set** and **.equ** directives are identical and can be used interchangeably.

- The *symbol* is a label that must appear in the label field.
- The *value* must be a well-defined expression, that is, all symbols in the expression must be previously defined in the current source module.

Undefined external symbols and symbols that are defined later in the module cannot be used in the expression. If the expression is relocatable, the symbol to which it is assigned is also relocatable.

The value of the expression appears in the object field of the listing. This value is not part of the actual object code and is not written to the output file.

Symbols defined with `.set` or `.equ` can be made externally visible with the `.def` or `.global` directive (see page 4-44). In this way, you can define global absolute constants.

Example

This example shows how symbols can be assigned with `.set` and `.equ`.

```

1          *****
2          **   Equate symbol AUX_R1 to register A1   **
3          **   and use it instead of the register.   **
4          *****
5          00000001  AUX_R1  .set   A1
6 00000000 00B802D4          STH   AUX_R1, *+B14
7
8          *****
9          **   Set symbol index to an integer expr.  **
10         **   and use it as an immediate operand.  **
11         *****
12         00000035  INDEX   .equ   100/2 +3
13 00000004 01001AD0          ADDK  INDEX, A2
14
15         *****
16         **   Set symbol SYMTAB to a relocatable expr. **
17         **   and use it as a relocatable operand.  **
18         *****
19 00000008 0000000A LABEL   .word  10
20         00000009' SYMTAB  .set   LABEL + 1
21
22         *****
23         **   Set symbol NSYMS equal to the symbol  **
24         **   INDEX and use it as you would INDEX.  **
25         *****
26         00000035  NSYMS   .set   INDEX
27 0000000c 00000035          .word  NSYMS

```

.space/.bes

Reserve Space

Syntax

`.space` *size in bytes*

`.bes` *size in bytes*

Description

The `.space` and `.bes` directives reserve the number of bytes given by *size in bytes* in the current section and fill them with 0s. The section program counter is incremented to point to the word following the reserved space.

When you use a label with the `.space` directive, it points to the *first* byte reserved. When you use a label with the `.bes` directive, it points to the *last* byte reserved.

Example

This example shows how memory is reserved with the .space and .bes directives.

```
1          *****
2          **      Begin assembling into the .text section.      **
3          *****
4 00000000          .text
5          *****
6          **      Reserve 0F0 bytes (60 words in .text section).  **
7          *****
8 00000000          .space  0F0h
9 000000f0 00000100          .word  100h, 200h
   000000f4 00000200
10         *****
11         **      Begin assembling into the .data section.      **
12         *****
13 00000000          .data
14 00000000 00000049          .string "In .data"
   00000001 0000006E
   00000002 00000020
   00000003 0000002E
   00000004 00000064
   00000005 00000061
   00000006 00000074
   00000007 00000061
15         *****
16         **      Reserve 100 bytes in the .data section;      **
17         **      RES_1 points to the first word                **
18         **      that contains reserved bytes.                **
19         *****
20 00000008          RES_1: .space  100
21 0000006c 0000000F          .word  15
22 00000070 00000008"          .word  RES_1
23         *****
24         **      Reserve 20 bytes in the .data section;      **
25         **      RES_2 points to the last word                 **
26         **      that contains reserved bytes.                **
27         *****
28 00000087          RES_2: .bes   20
29 00000088 00000036          .word  36h
30 0000008c 00000087"          .word  RES_2
```

.sslist/.ssnolist

Control Listing of Substitution Symbols

Syntax

.sslist
.ssnolist

Description

Two directives allow you to control substitution symbol expansion in the listing file:

The **.sslist** directive allows substitution symbol expansion in the listing file. The expanded line appears below the actual source line.

The **.ssnolist** directive suppresses substitution symbol expansion in the listing file.

By default, all substitution symbol expansion in the listing file is suppressed; the assembler acts as if the **.ssnolist** directive had been used.

Lines with the pound (#) character denote expanded substitution symbols.

Example

This example shows code that, by default, suppresses the listing of substitution symbol expansion, and it shows the **.sslist** directive assembled, instructing the assembler to list substitution symbol code expansion.

```

1 00000000          .bss    x,4
2 00000004          .bss    y,4
3 00000008          .bss    z,4
4
5                   addm    .macro src1,src2,dst
6                   LDW     *+B14(:src1:), A0
7                   LDW     *+B14(:src2:), A1
8                   NOP     4
9                   ADD     A0,A1,A0
10                  STW     A0,*+B14(:dst:)
11                  .endm
12
13 00000000          addm    x,y,z
1 00000000 0000006C- LDW     *+B14(x), A0
1 00000004 0080016C- LDW     *+B14(y), A1
1 00000008 00006000  NOP     4
1 0000000c 000401E0  ADD     A0,A1,A0
1 00000010 0000027C- STW     A0,*+B14(z)
14
15                  .sslist
16 00000014          addm    x,y,z
1 00000014 0000006C- LDW     *+B14(:src1:), A0
#                   LDW     *+B14(x), A0
1 00000018 0080016C- LDW     *+B14(:src2:), A1
#                   LDW     *+B14(y), A1
1 0000001c 00006000  NOP     4
1 00000020 000401E0  ADD     A0,A1,A0
1 00000024 0000027C- STW     A0,*+B14(:dst:)
#                   STW     A0,*+B14(z)
17

```

Syntax

```
.string {expr1 | "string1"} [, ... , {exprn | "stringn"}]
```

Description

The **.string** directive places 8-bit characters from a character string into the current section. The *expr* or *string* can be one of the following:

- An expression that the assembler evaluates and treats as an 8-bit signed number.
- A character string enclosed in double quotes. Each character in a string represents a separate value, and values are stored in consecutive bytes. The entire string *must* be enclosed in quotes.

The assembler truncates any values that are greater than eight bits. You can have up to 100 operands, but they must fit on a single source statement line.

If you use a label with **.string**, it points to the location of the first byte that is initialized.

When you use **.string** in a **.struct/.endstruct** sequence, **.string** defines a member's size; it does not initialize memory. For more information about **.struct/.endstruct**, see page 4-67.

Example

In this example, 8-bit values are placed into consecutive bytes in the current section. The label `Str_Ptr` has the value `0h`, which is the location of the first initialized byte.

```
1 00000000 00000041 Str_Ptr:  .string  "ABCD"
   00000001 00000042
   00000002 00000043
   00000003 00000044
2 00000004 00000041          .string  41h, 42h, 43h, 44h
   00000005 00000042
   00000006 00000043
   00000007 00000044
3 00000008 00000041          .string  "Austin", "Houston"
   00000009 00000075
   0000000a 00000073
   0000000b 00000074
   0000000c 00000069
   0000000d 0000006E
   0000000e 00000048
   0000000f 0000006F
   00000010 00000075
   00000011 00000073
   00000012 00000074
   00000013 0000006F
   00000014 0000006E
4 00000015 00000030          .string  36 + 12
```

**.struct/
.endstruct/.tag**

Declare Structure Type

Syntax

```
[stag]      .struct [expr]
[mem0]    element  [expr0]
[mem1]    element  [expr1]
.           .           .
.           .           .
.           .           .
[memn]    .tag stag   [exprn]
.           .           .
.           .           .
.           .           .
[memN]    element  [exprN]
[size]     .endstruct
label     .tag      stag
```

Description

The **.struct** directive assigns symbolic offsets to the elements of a data structure definition. This allows you to group similar data elements together and let the assembler calculate the element offset. This is similar to a C structure or a Pascal record. The **.struct** directive does not allocate memory; it merely creates a symbolic template that can be used repeatedly.

The **.endstruct** directive terminates the structure definition.

The **.tag** directive gives structure characteristics to a *label*, simplifying the symbolic representation and providing the ability to define structures that contain other structures. The **.tag** directive does not allocate memory. The structure tag (*stag*) of a **.tag** directive must have been previously defined.

Following are descriptions of the parameters used with the **.struct**, **.endstruct**, and **.tag** directives:

- The *stag* is the structure's tag. Its value is associated with the beginning of the structure. If no *stag* is present, the assembler puts the structure members in the global symbol table with the value of their absolute offset from the top of the structure. A **.stag** is optional for **.struct**, but is required for **.tag**.
- The *expr* is an optional expression indicating the beginning offset of the structure. The default starting point for a structure is 0.
- The *mem_{n/N}* is an optional label for a member of the structure. This label is absolute and equates to the present offset from the beginning of the structure. A label for a structure member cannot be declared global.

- ❑ The *element* is one of the following descriptors: `.string`, `.byte`, `.char`, `.int`, `.half`, `.short`, `.word`, `.long`, `.double`, `.float`, `.tag`, or `.field`. All of these except `.tag` are typical directives that initialize memory. Following a `.struct` directive, these directives describe the structure element's size. They do not allocate memory. A `.tag` directive is a special case because `stag` must be used (as in the definition of `stag`).
- ❑ The *expr_N/N* is an optional expression for the number of elements described. This value defaults to 1. A `.string` element is considered to be one byte in size, and a `.field` element is one bit.
- ❑ The *size* is an optional label for the total size of the structure.

Note: Directives That Can Appear in a .struct/.endstruct Sequence

The only directives that can appear in a `.struct/.endstruct` sequence are element descriptors, conditional assembly directives, and the `.align` directive, which aligns the member offsets on word boundaries. Empty structures are illegal.

These examples show various uses of the `.struct`, `.tag`, and `.endstruct` directives.

Example 1

```
1          real_rec .struct          ; stag
2      00000000 nom .int              ; member1 = 0
3      00000004 den .int              ; member2 = 1
4      00000008 real_len .endstruct  ; real_len = 2
5
6 00000000 0080016C-          LDW  *+B14(real+real_rec.den), A1
7                                          ; access structure
8
9 00000000          .bss real, real_len ; allocate mem rec
10
```

Example 2

```
11         cplx_rec .struct          ; stag
12     00000000 reali .tag real_rec   ; member1 = 0
13     00000008 imagi .tag real_rec   ; member2 = 2
14     00000010 cplx_len .endstruct  ; cplx_len = 4
15
16         complex .tag cplx_rec      ; assign structure
17                                          ; attribute
18 00000008          .bss complex, cplx_len ; allocate mem rec
19
20 00000004 0100046C-          LDW  *+B14(complex.imagi.nom), A2
21                                          ; access structure
22 00000008 0100036C-          LDW  *+B14(complex.reali.den), A2
23                                          ; access structure
24 0000000c 018C4A78          CMPEQ A2, A3, A3
```

Example 3

```

1                               .struct                ; no stag puts
2                               ; mems into global
3                               ; symbol table
4
5          00000000 X          .byte                ; create 3 dim
6          00000001 Y          .byte                ; templates
7          00000002 Z          .byte
8          00000003          .endstruct

```

Example 4

```

1          bit_rec .struct                ; stag
2          00000000 stream .string 64
3          00000040 bit7   .field 7        ; bit7 = 64
4          00000040 bit1   .field 9        ; bit9 = 64
5          00000042 bit5   .field 10       ; bit5 = 64
6          00000044 x_int  .byte           ; x_int = 68
7          00000045 bit_len .endstruct     ; length = 72
8
9          bits     .tag bit_rec
10 00000000        .bss bits, bit_len
11
12 00000000 0100106C- LDW  *+B14(bits.bit7), A2
13                                     ; load field
14 00000004 0109E7A0 AND  0Fh, A2, A2      ; mask off garbage

```

.tab

Define Tab Size

Syntax

.tab size

Description

The **.tab** directive defines the tab size. Tabs encountered in the source input are translated to *size* character spaces in the listing. The default tab size is eight spaces.

Example

In this example, each of the lines of code following a **.tab** statement consists of a single tab character followed by an NOP instruction.

Source file:

```

; default tab size
NOP
NOP
NOP

    .tab 4
NOP
NOP
NOP

    .tab 16
NOP
NOP
NOP

```

Listing file:

```
1                                     ; default tab size
2 00000000 00000000                    NOP
3 00000004 00000000                    NOP
4 00000008 00000000                    NOP
5                                     .tab 4
7 0000000c 00000000                    NOP
8 00000010 00000000                    NOP
9 00000014 00000000                    NOP
10                                    .tab 16
12 00000018 00000000                    NOP
13 0000001c 00000000                    NOP
14 00000020 00000000                    NOP
```

.text

Assemble Into .text Section

Syntax

.text

Description

The **.text** directive tells the assembler to begin assembling into the **.text** section, which usually contains executable code. The section program counter is set to 0 if nothing has yet been assembled into the **.text** section. If code has already been assembled into the **.text** section, the section program counter is restored to its previous value in the section.

The **.text** section is the default section. Therefore, at the beginning of an assembly, the assembler assembles code into the **.text** section unless you use a **.data** or **.sect** directive to specify a different section.

For more information about COFF sections, see Chapter 2, *Introduction to Common Object File Format*.

Example

This example assembles code into the `.text` and `.data` sections.

```
1          *****
2          ** Begin assembling into .data section. **
3          *****
4 00000000          .data
5 00000000 00000005          .byte  5,6
6 00000001 00000006
7
8          *****
9          ** Begin assembling into .text section. **
10         *****
11 00000000          .text
12 00000000 00000001          .byte  1
13 00000001 00000002          .byte  2,3
14 00000002 00000003
15
16         *****
17         ** Resume assembling into .data section.**
18         *****
19 00000002          .data
20 00000002 00000007          .byte  7,8
21 00000003 00000008
22
23         *****
24         ** Resume assembling into .text section.**
25         *****
26 00000003          .text
27 00000003 00000004          .byte  4
```

.title

Define Page Title

Syntax

.title "string"

Description

The **.title** directive supplies a title that is printed in the heading on each listing page. The source statement itself is not printed, but the line counter is incremented.

The *string* is a quote-enclosed title of up to 64 characters. If you supply more than 64 characters, the assembler truncates the string and issues a warning:

```
*** WARNING! line x: W0001: String is too long - will be truncated
```

The assembler prints the title on the page that follows the directive and on subsequent pages until another `.title` directive is processed. If you want a title on the first page, the first source statement must contain a `.title` directive.

.title

Example

In this example, one title is printed on the first page and a different title is printed on succeeding pages.

Source file:

```
                .title  "**** Fast Fourier Transforms ****"
;               .
;               .
;               .
                .title  "**** Floating-Point Routines ****"
                .page
```

Listing file:

```
TMS320C6x COFF Assembler   Version x.xx      Tue Apr 14 17:18:21 1997
Copyright (c) 1996-1997 Texas Instruments Incorporated
**** Fast Fourier Transforms ****                                PAGE    1

                2                ;                .
                3                ;                .
                4                ;                .
TMS320C6x COFF Assembler   Version x.xx      Tue Apr 14 17:18:21 1997
Copyright (c) 1996-1997 Texas Instruments Incorporated
**** Floating-Point Routines ****                                PAGE    2
```

No Errors, No Warnings

**.union/
.endunion/.tag**

Declare Union Type

Syntax

```

[utag]      .union [expr]
[mem0]    element  [expr0]
[mem1]    element  [expr1]
            .
            .
            .
[memn]    .tag utag   [exprn]
            .
            .
            .
[memN]    element  [exprN]
[size]     .endunion
label     .tag      utag

```

Description

The **.union** directive assigns symbolic offsets to the elements of alternate data structure definitions to be allocated in the same memory space. This enables you to define several alternate structures and then let the assembler calculate the element offset. This is similar to a C union. The **.union** directive does not allocate any memory; it merely creates a symbolic template that can be used repeatedly.

A **.struct** definition can contain a **.union** definition, and **.structs** and **.unions** can be nested.

The **.endunion** directive terminates the union definition.

The **.tag** directive gives structure or union characteristics to a *label*, simplifying the symbolic representation and providing the ability to define structures or unions that contain other structures or unions. The **.tag** directive does not allocate memory. The structure or union tag of a **.tag** directive must have been previously defined.

- The *utag* is the union's tag. Its value is associated with the beginning of the union. If no *utag* is present, the assembler puts the union members in the global symbol table with the value of their absolute offset from the top of the union. In this case, each member must have a unique name.
- The *expr* is an optional expression indicating the beginning offset of the union. Unions default to start at 0. This parameter can only be used with a top-level union. It cannot be used when defining a nested union.

- ❑ The $mem_{n/N}$ is an optional label for a member of the union. This label is absolute and equates to the present offset from the beginning of the union. A label for a union member cannot be declared global.
- ❑ The *element* is one of the following descriptors: `.byte`, `.char`, `.double`, `field`, `.float`, `.half`, `.int`, `.long`, `.short`, `.string`, `.ubyte`, `.uchar`, `.uhalf`, `.uint`, `.ulong`, `.ushort`, `.uword`, and `.word`. An element can also be a complete declaration of a nested structure or union, or a structure or union declared by its tag. Following a `.union` directive, these directives describe the element's size. They do not allocate memory.
- ❑ The $expr_{n/N}$ is an optional expression for the number of elements described. This value defaults to 1. A `.string` element is considered to be one byte in size, and a `.field` element is one bit.
- ❑ The *size* is an optional label for the total size of the union.

Note: Directives That Can Appear in a .union/.endunion Sequence

The only directives that can appear in a `.union/.endunion` sequence are element descriptors, structure and union tags, and conditional assembly directives. Empty structures are illegal.

These examples show unions with and without tags.

Example 1

```
1                                .global employid
2                                xample      .union                ; utag
3      0000 ival                 .word                ; member1 = int
4      0000 fval                 .float               ; member2 = float
5      0000 sval                 .string              ; member3 = string
6      0002 real_len           .endunion              ; real_len = 2
7
8 000000                        .bss employid, real_len ;allocate memory
9
10                               employid   .tag xample      ; name an instance
11 000000 0000-                 ADD employid.fval, A    ; access union element
```

Example 2

```
1
2                                .union                ; utag
3      0000 x                    .long                ; member1 = long
4      0000 y                    .float               ; member2 = float
5      0000 z                    .word                ; member3 = word
6      0002 size_u             .endunion              ; real_len = 2
7
```

.usect*Reserve Uninitialized Space*

Syntax

symbol **.usect** "section name", size in bytes [, alignment[, bank offset]]

Description

The **.usect** directive reserves space for variables in an uninitialized, named section. This directive is similar to the **.bss** directive; both simply reserve space for data and that space has no contents. However, **.usect** defines additional sections that can be placed anywhere in memory, independently of the **.bss** section.

- The *symbol* points to the first location reserved by this invocation of the **.usect** directive. The symbol corresponds to the name of the variable for which you are reserving space.
- The *section name* is significant to 200 characters and must be enclosed in double quotes. This parameter names the uninitialized section. A section name can contain a subsection name in the form *section name:subsection name*.
- The *size in bytes* is an expression that defines the number of bytes that are reserved in *section name*.
- The *alignment* is an optional parameter that ensures that the space allocated to the symbol occurs on the specified boundary. This boundary indicates the size of the slot in bytes and can be set to any power of 2.
- The *bank offset* is an optional parameter that ensures that the space allocated to the symbol occurs on a specific memory bank boundary. The bank offset value measures the number of bytes to offset from the alignment specified before assigning the symbol to that location.

Initialized sections directives (**.text**, **.data**, and **.sect**) end the current section and tell the assembler to begin assembling into another section. A **.usect** or **.bss** directive encountered in the current section is simply assembled, and assembly continues in the current section.

Variables that can be located contiguously in memory can be defined in the same specified section; to do so, repeat the **.usect** directive with the same section name and the subsequent symbol (variable name).

For more information about COFF sections, see Chapter 2, *Introduction to Common Object File Format*.

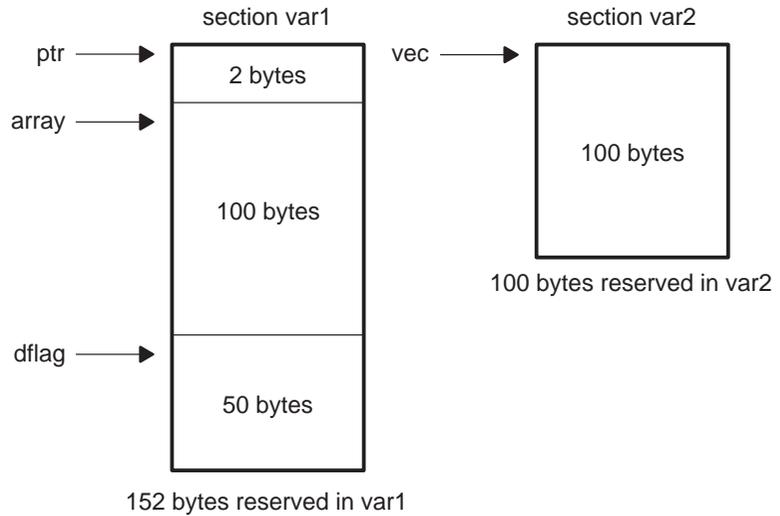
Example

This example uses the .usect directive to define two uninitialized, named sections, var1 and var2. The symbol ptr points to the first byte reserved in the var1 section. The symbol array points to the first byte in a block of 100 bytes reserved in var1, and dflag points to the first byte in a block of 50 bytes in var1. The symbol vec points to the first byte reserved in the var2 section.

Figure 4–8 shows how this example reserves space in two uninitialized sections, var1 and var2.

```
1          *****
2          ** Assemble into .text section **
3          *****
4 00000000          .text
5 00000000 008001A0          MV      A0,A1
6
7          *****
8          ** Reserve 2 bytes in var1. **
9          *****
10 00000000 ptr      .usect  "var1",2
11 00000004 0100004C-      LDH     *+B14(ptr),A2 ; still in .text
12
13          *****
14          ** Reserve 100 bytes in var1 **
15          *****
16 00000002 array   .usect  "var1",100
17 00000008 01800128-      MVK     array,A3 ; still in .text
18 0000000c 01800068-      MVKH    array,A3
19
20          *****
21          ** Reserve 50 bytes in var1 **
22          *****
23 00000066 dflag  .usect  "var1",50
24 00000010 02003328-      MVK     dflag,A4
25 00000014 02000068-      MVKH    dflag,A4
26
27          *****
28          ** Reserve 100 bytes in var1 **
29          *****
30 00000000 vec    .usect  "var2",100
31 00000018 0000002A-      MVK     vec,B0 ; still in .text
32 0000001c 0000006A-      MVKH    vec,B0
```

Figure 4–8. The .usect Directive



.var

Use Substitution Symbols as Local Variables

Syntax

`.var sym1 [,sym2, ... , symn]`

Description

The .var directive allows you to use substitution symbols as local variables within a macro. With this directive, you can define up to 32 local macro substitution symbols (including parameters) per macro.

Description

The .var directive creates temporary substitution symbols with the initial value of the null string. These symbols are not passed in as parameters, and they are lost after expansion.

Description

For more information on macros, see Chapter 5, *Macro Language*.



Macro Language

The TMS320C6000™ assembler supports a macro language that enables you to create your own instructions. This is especially useful when a program executes a particular task several times. The macro language lets you:

- Define your own macros and redefine existing macros
- Simplify long or complicated assembly code
- Access macro libraries created with the archiver
- Define conditional and repeatable blocks within a macro
- Manipulate strings within a macro
- Control expansion listing

Topic	Page
5.1 Using Macros	5-2
5.2 Defining Macros	5-3
5.3 Macro Parameters/Substitution Symbols	5-5
5.4 Macro Libraries	5-13
5.5 Using Conditional Assembly in Macros	5-14
5.6 Using Labels in Macros	5-16
5.7 Producing Messages in Macros	5-17
5.8 Using Directives to Format the Output Listing	5-19
5.9 Using Recursive and Nested Macros	5-21
5.10 Macro Directives Summary	5-23

5.1 Using Macros

Programs often contain routines that are executed several times. Instead of repeating the source statements for a routine, you can define the routine as a macro, then call the macro in the places where you would normally repeat the routine. This simplifies and shortens your source program.

If you want to call a macro several times but with different data each time, you can assign parameters within a macro. This enables you to pass different information to the macro each time you call it. The macro language supports a special symbol called a *substitution symbol*, which is used for macro parameters. See section 5.3, *Macro Parameters/Substitution Symbols*, page 5-5, for more information.

Using a macro is a 3-step process.

Step 1: Define the macro. You must define macros before you can use them in your program. There are two methods for defining macros:

- Macros can be defined at the beginning of a *source file* or in an *copy/include file*. See section 5.2, *Defining Macros*, for more information.
- Macros can also be defined in a *macro library*. A macro library is a collection of files in archive format created by the archiver. Each member of the archive file (macro library) may contain one macro definition corresponding to the member name. You can access a macro library by using the `.mlib` directive. For more information, see section 5.4, *Macro Libraries*, page 5-13.

Step 2: Call the macro. After you have defined a macro, call it by using the macro name as a mnemonic in the source program. This is referred to as a *macro call*.

Step 3: Expand the macro. The assembler expands your macros when the source program calls them. During expansion, the assembler passes arguments by variable to the macro parameters, replaces the macro call statement with the macro definition, then assembles the source code. By default, the macro expansions are printed in the listing file. You can turn off expansion listing by using the `.mno` directive. For more information, see section 5.8, *Using Directives to Format the Output Listing*, page 5-19.

When the assembler encounters a macro definition, it places the macro name in the opcode table. This redefines any previously defined macro, library entry, directive, or instruction mnemonic that has the same name as the macro. This allows you to expand the functions of directives and instructions, as well as to add new instructions.

5.2 Defining Macros

You can define a macro anywhere in your program, but you must define the macro before you can use it. Macros can be defined at the beginning of a source file or in a `.copy/.include` file (see page 4-29); they can also be defined in a macro library. For more information, see section 5.4, *Macro Libraries*, page 5-13.

Macro definitions can be nested, and they can call other macros, but all elements of the macro must be defined in the same file. Nested macros are discussed in section 5.9, *Using Recursive and Nested Macros*, page 5-21.

A macro definition is a series of source statements in the following format:

```

macname  .macro [parameter1] [, ... , parametern]
           model statements or macro directives
           [.mexit]
           .endm

```

<i>macname</i>	names the macro. You must place the name in the source statement's label field. Only the first 128 characters of a macro name are significant. The assembler places the macro name in the internal opcode table, replacing any instruction or previous macro definition with the same name.
.macro	is the directive that identifies the source statement as the first line of a macro definition. You must place <code>.macro</code> in the opcode field.
<i>parameter</i> ₁ , <i>parameter</i> _n	are optional substitution symbols that appear as operands for the <code>.macro</code> directive. Parameters are discussed in section 5.3, <i>Macro Parameters/Substitution Symbols</i> , page 5-5.
<i>model statements</i>	are instructions or assembler directives that are executed each time the macro is called.
<i>macro directives</i>	are used to control macro expansion.
.mexit	is a directive that functions as a <code>goto .endm</code> . The <code>.mexit</code> directive is useful when error testing confirms that macro expansion fails and completing the rest of the macro is unnecessary.
.endm	is the directive that terminates the macro definition.

Example 5–1 shows the definition, call, and expansion of a macro.

Example 5–1. Macro Definition, Call, and Expansion

Macro definition: The following code defines a macro, `sadd4`, with four parameters:

```

1          sadd4  .macro  r1,r2,r3,r4
2          !
3          !  sadd4  r1, r2 ,r3, r4
4          !  r1 = r1 + r2 + r3 + r4 (saturated)
5          !
6          SADD   r1,r2,r1
7          SADD   r1,r3,r1
8          SADD   r1,r4,r1
9          .endm

```

Macro call: The following code calls the `sadd4` macro with four arguments:

```

10
11 00000000          sadd4  A0,A1,A2,A3

```

Macro expansion: The following code shows the substitution of the macro definition for the macro call. The assembler substitutes `A0`, `A1`, `A2`, and `A3` for the `r1`, `r2`, `r3`, and `r4` parameters of `sadd4`.

```

1          00000000 00040278          SADD   A0,A1,A0
1          00000004 00080278          SADD   A0,A2,A0
1          00000008 000C0278          SADD   A0,A3,A0

```

If you want to include comments with your macro definition but *do not* want those comments to appear in the macro expansion, use an exclamation point to precede your comments. If you *do* want your comments to appear in the macro expansion, use an asterisk or semicolon. See section 5.7, *Producing Messages in Macros*, page 5-17, for more information about macro comments.

5.3 Macro Parameters/Substitution Symbols

If you want to call a macro several times with different data each time, you can assign parameters within the macro. The macro language supports a special symbol, called a *substitution symbol*, which is used for macro parameters.

Macro parameters are substitution symbols that represent a character string. These symbols can also be used outside of macros to equate a character string to a symbol name (see section 3.8.6, *Substitution Symbols*, page 3-25).

Valid substitution symbols can be up to 128 characters long and *must begin with a letter*. The remainder of the symbol can be a combination of alphanumeric characters, underscores, and dollar signs.

Substitution symbols used as macro parameters are local to the macro they are defined in. You can define up to 32 local substitution symbols (including substitution symbols defined with the `.var` directive) per macro. For more information about the `.var` directive, see section 5.3.6, *Substitution Symbols as Local Variables in Macros*, page 5-12.

During macro expansion, the assembler passes arguments by variable to the macro parameters. The character-string equivalent of each argument is assigned to the corresponding parameter. Parameters without corresponding arguments are set to the null string. If the number of arguments exceeds the number of parameters, the last parameter is assigned the character-string equivalent of all remaining arguments.

If you pass a list of arguments to one parameter or if you pass a comma or semicolon to a parameter, you must surround these terms with quotation marks.

At assembly time, the assembler replaces the macro parameter/substitution symbol with its corresponding character string, then translates the source code into object code.

Example 5–2 shows the expansion of a macro with varying numbers of arguments.

Example 5–2. Calling a Macro With Varying Numbers of Arguments

```

Macro definition:

Parms .macro a,b,c
;     a = :a:
;     b = :b:
;     c = :c:
      .endm

Calling the macro:

      Parms 100,label
;         a = 100
;         b = label
;         c = " "

      Parms 100, , x
;         a = 100
;         b = " "
;         c = x

      Parms "100,200,300",x,y
;         a = 100,200,300
;         b = x
;         c = y

      Parms ""string"",x,y
;         a = "string"
;         b = x
;         c = y

```

5.3.1 Directives That Define Substitution Symbols

You can manipulate substitution symbols with the **.asg** and **.eval** directives.

- The **.asg** directive assigns a character string to a substitution symbol.

The syntax of the **.asg** directive is:

```
.asg [""]character string[""], substitution symbol
```

The quotation marks are optional. If there are no quotation marks, the assembler reads characters up to the first comma and removes leading and trailing blanks. In either case, a character string is read and assigned to the substitution symbol.

Example 5–3 shows character strings being assigned to substitution symbols.

Example 5–3. The **.asg** Directive

```

.asg "A4", RETVAL           ; return value
.asg "B14", PAGEPTR        ; global page pointer
.asg ""Version 1.0"", version
.asg "p1, p2, p3", list

```

- The **.eval** directive performs arithmetic on numeric substitution symbols.

The syntax of the **.eval** directive is:

```
.eval well-defined expression, substitution symbol
```

The **.eval** directive evaluates the expression and assigns the string value of the result to the substitution symbol. If the expression is not well defined, the assembler generates an error and assigns the null string to the symbol.

Example 5–4 shows arithmetic being performed on substitution symbols.

Example 5–4. The **.eval** Directive

```
.asg 1,counter
.loop 100
.word counter
.eval counter + 1,counter
.endloop
```

In Example 5–4, the **.asg** directive could be replaced with the **.eval** directive (**.eval 1, counter**) without changing the output. In simple cases like this, you can use **.eval** and **.asg** interchangeably. However, you must use **.eval** if you want to calculate a *value* from an expression. While **.asg** only assigns a character string to a substitution symbol, **.eval** evaluates an expression and then assigns the character string equivalent to a substitution symbol.

For more information about the **.asg** and **eval** assembler directives, see page 4-22.

5.3.2 Built-In Substitution Symbol Functions

The following built-in substitution symbol functions enable you to make decisions on the basis of the string value of substitution symbols. These functions always return a value, and they can be used in expressions. Built-in substitution symbol functions are especially useful in conditional assembly expressions. Parameters of these functions are substitution symbols or character-string constants.

In the function definitions shown in Table 5–1, *a* and *b* are parameters that represent substitution symbols or character-string constants. The term *string* refers to the string value of the parameter. The symbol *ch* represents a character constant.

Table 5–1. Substitution Symbol Functions and Return Values

Function	Return Value
\$symlen (<i>a</i>)	Length of string <i>a</i>
\$symcmp (<i>a,b</i>)	< 0 if <i>a</i> < <i>b</i> ; 0 if <i>a</i> = <i>b</i> ; > 0 if <i>a</i> > <i>b</i>
\$firstch (<i>a,ch</i>)	Index of the first occurrence of character constant <i>ch</i> in string <i>a</i>
\$lastch (<i>a,ch</i>)	Index of the last occurrence of character constant <i>ch</i> in string <i>a</i>
\$isdefed (<i>a</i>)	1 if string <i>a</i> is defined in the symbol table 0 if string <i>a</i> is not defined in the symbol table
\$ismember (<i>a,b</i>)	Top member of list <i>b</i> is assigned to string <i>a</i> 0 if <i>b</i> is a null string
\$iscons (<i>a</i>)	1 if string <i>a</i> is a binary constant 2 if string <i>a</i> is an octal constant 3 if string <i>a</i> is a hexadecimal constant 4 if string <i>a</i> is a character constant 5 if string <i>a</i> is a decimal constant
\$isname (<i>a</i>)	1 if string <i>a</i> is a valid symbol name 0 if string <i>a</i> is not a valid symbol name
\$isreg (<i>a</i>) [†]	1 if string <i>a</i> is a valid predefined register name 0 if string <i>a</i> is not a valid predefined register name

[†] For more information about predefined register names, see section 3.8.5, *Predefined Symbolic Constants*, on page 3-23.

Example 5–5 shows built-in substitution symbol functions.

Example 5–5. Using Built-In Substitution Symbol Functions

```

pushx  .macro list
!
! Push more than one item
! $ismember removes the first item in the list

    .var      item
    .loop
    .break    ($ismember(item, list) = 0)
    STW      item, *B15--[1]
    .endloop
    .endm

pushx   A0, A1, A2, A3

```

5.3.3 Recursive Substitution Symbols

When the assembler encounters a substitution symbol, it attempts to substitute the corresponding character string. If that string is also a substitution symbol, the assembler performs substitution again. The assembler continues doing this until it encounters a token that is not a substitution symbol or until it encounters a substitution symbol that it has already encountered during this evaluation.

In Example 5–6, the *x* is substituted for *z*; *z* is substituted for *y*; and *y* is substituted for *x*. The assembler recognizes this as infinite recursion and ceases substitution.

Example 5–6. Recursive Substitution

```
.asg  "x",z  ; declare z and assign z = "x"
.asg  "z",y  ; declare y and assign y = "z"
.asg  "y",x  ; declare x and assign x = "y"
MVKL  x, A1
MVKH  x, A1

*   MVKL  x,A1 ; recursive expansion
*   MVKH  x,A1 ; recursive expansion
```

5.3.4 Forced Substitution

In some cases, substitution symbols are not recognizable to the assembler. The forced substitution operator, which is a set of colons surrounding the symbol, enables you to force the substitution of a symbol's character string. Simply enclose a symbol with colons to force the substitution. Do not include any spaces between the colons and the symbol.

The syntax for the forced substitution operator is:

```
:symbol:
```

The assembler expands substitution symbols surrounded by colons before expanding other substitution symbols.

You can use the forced substitution operator only inside macros, and you cannot nest a forced substitution operator within another forced substitution operator.

Example 5–7 shows how the forced substitution operator is used.

Example 5–7. Using the Forced Substitution Operator

```
force      .macro    x
           .loop     8
PORT:x:    .set      x*4
           .eval     x+1, x
           .endloop
           .endm

           .global   portbase
           force     0
```

This generates the following source code:

```
PORT0     .set      0
PORT1     .set      4
.
.
.
PORT7     .set      28
```

5.3.5 Accessing Individual Characters of Subscripted Substitution Symbols

In a macro, you can access the individual characters (substrings) of a substitution symbol with subscripted substitution symbols. You must use the forced substitution operator for clarity.

You can access substrings in two ways:

`:symbol (well-defined expression)`:

This method of subscripting evaluates to a character string with one character.

`:symbol (well-defined expression1, well-defined expression2)`:

In this method, expression₁ represents the substring's starting position, and expression₂ represents the substring's length. You can specify exactly where to begin subscripting and the exact length of the resulting character string. *The index of substring characters begins with 1, not 0.*

Example 5–8 and Example 5–9 show built-in substitution symbol functions used with subscripted substitution symbols.

Example 5–8. Using Subscripted Substitution Symbols to Redefine an Instruction

```

storex .macro      x
      .var        tmp
      .asg        :x(1):, tmp
      .if         $symcmp(tmp,"A") == 0
      STW         x,*A15--(4)
      .elseif    $symcmp(tmp,"B") == 0
      STW         x,*A15--(4)
      .elseif    $iscons(x)
      MVK         x,A0
      STW         A0,*A15--(4)
      .else
      .emsg       "Bad Macro Parameter"
      .endif
      .endm

      storex     10h
      storex     A15

```

In Example 5–8, subscripted substitution symbols redefine the STW instruction so that it handles immediate.

Example 5–9. Using Subscripted Substitution Symbols to Find Substrings

```

substr .macro      start, strg1, strg2, pos
      .var        len1, len2, i, tmp
      .if         $symlen(start) = 0
      .eval       1, start
      .endif
      .eval       0, pos
      .eval       start, i
      .eval       $symlen(strg1), len1
      .eval       $symlen(strg2), len2
      .loop
      .break     i = (len2 - len1 + 1)
      .asg       ":strg2(i, len1):", tmp
      .if         $symcmp(strg1, tmp) = 0
      .eval       i, pos
      .break
      .else
      .eval       i + 1, i
      .endif
      .endloop
      .endm

      .asg       0, pos
      .asg       "ar1 ar2 ar3 ar4", regs
      substr     1, "ar2", regs, pos
      .word     pos

```

In Example 5–9, the subscripted substitution symbol is used to find a substring strg1 beginning at position start in the string strg2. The position of the substring strg1 is assigned to the substitution symbol pos.

5.3.6 Substitution Symbols as Local Variables in Macros

If you want to use substitution symbols as local variables within a macro, you can use the **.var** directive to define up to 32 local macro substitution symbols (including parameters) per macro. The **.var** directive creates temporary substitution symbols with the initial value of the null string. These symbols are not passed in as parameters, and they are lost after expansion.

```
.var sym1 [,sym2, ... ,symn]
```

The **.var** directive is used in Example 5–8 and Example 5–9, page 5-11.

5.4 Macro Libraries

One way to define macros is by creating a macro library. A macro library is a collection of files that contain macro definitions. You must use the archiver to collect these files, or members, into a single file (called an archive). Each member of a macro library contains one macro definition. The files in a macro library must be unassembled source files. The macro name and the member name must be the same, and the macro filename's extension must be .asm. For example:

Macro Name	Filename in Macro Library
simple	simple.asm
add3	add3.asm

You can access the macro library by using the .mlib assembler directive (described on page 4-55). The syntax is:

```
.mlib filename
```

When the assembler encounters the .mlib directive, it opens the library named by filename and creates a table of the library's contents. The assembler enters the names of the individual members within the library into the opcode tables as library entries; this redefines any existing opcodes or macros that have the same name. If one of these macros is called, the assembler extracts the entry from the library and loads it into the macro table.

The assembler expands the library entry in the same way it expands other macros. (See section 5.1, *Using Macros*, on page 5-2, for how the assembler expands macros.) You can control the listing of library entry expansions with the .mlist directive. For more information about the .mlist directive, see section 5.8, *Using Directives to Format the Output Listing*, page 5-19 and the .mlist description on page 4-56. Only macros that are actually called from the library are extracted, and they are extracted only once.

You can use the archiver to create a macro library by including the desired files in an archive. A macro library is no different from any other archive, except that the assembler expects the macro library to contain macro definitions. The assembler expects *only* macro definitions in a macro library; putting object code or miscellaneous source files into the library may produce undesirable results. For information about creating a macro library archive, see Chapter 6, *Archiver Description*.

5.5 Using Conditional Assembly in Macros

The conditional assembly directives are **.if/elseif/else/endif** and **.loop/.break/endloop**. They can be nested within each other up to 32 levels deep. The format of a conditional block is:

```
.if well-defined expression
[elseif well-defined expression]
[else]
.endif
```

The **.elseif** and **.else** directives are optional in conditional assembly. The **.elseif** directive can be used more than once within a conditional assembly code block. When **.elseif** and **.else** are omitted and when the **.if** expression is false (0), the assembler continues to the code following the **.endif** directive. For more information on the **.if/elseif/else/endif** directives, see page 4-47.

The **.loop/.break/endloop** directives enable you to assemble a code block repeatedly. The format of a repeatable block is:

```
.loop [well-defined expression]
[.break [well-defined expression]]
.endloop
```

The **.loop** directive's optional *well-defined expression* evaluates to the loop count (the number of loops to be performed). If the expression is omitted, the loop count defaults to 1024 unless the assembler encounters a **.break** directive with an expression that is true (nonzero). For more information on the **.loop/.break/ .endloop** directives, see page 4-53.

The **.break** directive and its expression are optional in repetitive assembly. If the expression evaluates to false, the loop continues. The assembler breaks the loop when the **.break** expression evaluates to true or when the **.break** expression is omitted. When the loop is broken, the assembler continues with the code after the **.endloop** directive.

Example 5–10, Example 5–11, and Example 5–12 show the **.loop/.break/ .endloop** directives, properly nested conditional assembly directives, and built-in substitution symbol functions used in a conditional assembly code block.

Example 5–10. The .loop/.break/.endloop Directives

```

.asg 1,x
.loop

.break (x == 10) ; if x == 10, quit loop/break with
                ; expression

.eval  x+1,x
.endloop

```

Example 5–11. Nested Conditional Assembly Directives

```

.asg 1,x
.loop

.if (x == 10) ; if x == 10 quit loop
.break       ; force break
.endif

.eval x+1,x
.endloop

```

Example 5–12. Built-In Substitution Symbol Functions in a Conditional Assembly Code Block

```

MACK3 .macro src1, src2, sum, k
!
! dst = dst + k * (src1 * src2)

.if k = 0
MPY src1, src2, src2
NOP
ADD src2, sum, sum
.else
MPY src1,src2,src2
MVK k,src1
MPY src1,src2,src2
NOP
ADD src2,sum,sum
.endif

.endm

MACK3 A0,A1,A3,0
MACK3 A0,A1,A3,100

```

For more information, see section 4.7, *Directives That Enable Conditional Assembly*, on page 4-17.

5.7 Producing Messages in Macros

The macro language supports three directives that enable you to define your own assembly-time error and warning messages. These directives are especially useful when you want to create messages specific to your needs. The last line of the listing file shows the error and warning counts. These counts alert you to problems in your code and are especially useful during debugging.

- .emsg** sends error messages to the listing file. The `.emsg` directive generates errors in the same manner as the assembler, incrementing the error count and preventing the assembler from producing an object file.
- .mmsg** sends assembly-time messages to the listing file. The `.mmsg` directive functions in the same manner as the `.emsg` directive but does not set the error count or prevent the creation of an object file.
- .wmsg** sends warning messages to the listing file. The `.wmsg` directive functions in the same manner as the `.emsg` directive, but it increments the warning count and does not prevent the generation of an object file.

Macro comments are comments that appear in the definition of the macro *but do not show up in the expansion of the macro*. An exclamation point in column 1 identifies a macro comment. If you want your comments to appear in the macro expansion, precede your comment with an asterisk or semicolon.

Example 5–14 shows user messages in macros and macro comments that do not appear in the macro expansion.

Example 5-14. Producing Messages in a Macro

```
TEST  .macro  x,y
!
! This macro checks for the correct number of parameters.
! It generates an error message if x and y are not present.
!
! The first line tests for proper input.
!
    .if      ($symlen(x) + ||$symlen(y) == 0)
    .emsg    "ERROR --missing parameter in call to TEST"
    .mexit
    .else
    .
    .
    .endif
    .if
    .
    .
    .endif
    .endm
```

For more information about the .emsg, .mmsg, and .wmsg assembler directives, see page 4-37.

5.8 Using Directives to Format the Output Listing

Macros, substitution symbols, and conditional assembly directives may hide information. You may need to see this hidden information, so the macro language supports an expanded listing capability.

By default, the assembler shows macro expansions and false conditional blocks in the list output file. You may want to turn this listing off or on within your listing file. Four sets of directives enable you to control the listing of this information:

Macro and loop expansion listing

.mlist expands macros and `.loop/.endloop` blocks. The `.mlist` directive prints all code encountered in those blocks.

.mnolist suppresses the listing of macro expansions and `.loop/.endloop` blocks.

For macro and loop expansion listing, `.mlist` is the default.

False conditional block listing

.fclist causes the assembler to include in the listing file all conditional blocks that do not generate code (false conditional blocks). Conditional blocks appear in the listing exactly as they appear in the source code.

.fcnolist suppresses the listing of false conditional blocks. Only the code in conditional blocks that actually assemble appears in the listing. The `.if`, `.elseif`, `.else`, and `.endif` directives do not appear in the listing.

For false conditional block listing, `.fclist` is the default.

Substitution symbol expansion listing

.sslist expands substitution symbols in the listing. This is useful for debugging the expansion of substitution symbols. The expanded line appears below the actual source line.

.ssnolist turns off substitution symbol expansion in the listing.

For substitution symbol expansion listing, `.ssnolist` is the default.

□ **Directive listing**

.drlist causes the assembler to print to the listing file all directive lines.

.drnolist suppresses the printing of certain directives in the listing file. These directives are `.asg`, `.eval`, `.var`, `.sslist`, `.mlist`, `.fclist`, `.ssnolist`, `.mnolist`, `.fcnolist`, `.emsg`, `.wmsg`, `.mmsg`, `.length`, `.width`, and `.break`.

For directive listing, `.drlist` is the default.

5.9 Using Recursive and Nested Macros

The macro language supports recursive and nested macro calls. This means that you can call other macros in a macro definition. You can nest macros up to 32 levels deep. When you use recursive macros, you call a macro from its own definition (the macro calls itself).

When you create recursive or nested macros, you should pay close attention to the arguments that you pass to macro parameters because the assembler uses dynamic scoping for parameters. This means that the called macro uses the environment of the macro from which it was called.

Example 5–15 shows nested macros. The `y` in the `in_block` macro hides the `y` in the `out_block` macro. The `x` and `z` from the `out_block` macro, however, are accessible to the `in_block` macro.

Example 5–15. Using Nested Macros

```

in_block .macro y,a
        .          ; visible parameters are y,a and
        .          ;      x,z from the calling macro
        .endm

out_block .macro x,y,z
        .          ; visible parameters are x,y,z
        .
        in_block x,y ; macro call with x and y as
        .          ;      arguments
        .
        .endm
out_block      ; macro call

```

Example 5–16 shows recursive and fact macros. The `fact` macro produces assembly code necessary to calculate the factorial of `n`, where `n` is an immediate value. The result is placed in the `A1` register. The `fact` macro accomplishes this by calling `fact1`, which calls itself recursively.

Example 5-16. Using Recursive Macros

```
.fcnoelist
fact1 .macro n
    .if n == 1
        MVK globcnt, A1                ; Leave the answer in the A1 register.
    .else
        .eval n-1, temp                ; Compute the decrement of symbol n.
        .eval globcnt*temp, globcnt    ; Multiply to get a new result.
        fact1 temp                     ; Recursive call.
    .endif
.endm

fact .macro n
    .if ! $iscons(n)                  ; Test that input is a constant.
        .emsg "Parm not a constant"
    .elseif n < 1                      ; Type check input.
        MVK 0, A1
    .else
        .var temp
        .asg n, globcnt
        fact1 n                        ; Perform recursive procedure
    .endif
.endm
```

5.10 Macro Directives Summary

The following directives can be used with macros. The `.macro`, `.mexit`, `.endm` and `.var` directives are valid only with macros; the remaining directives are general assembly language directives.

Table 5–2. Creating Macros

Mnemonic and Syntax	Description	See Page	
		Macro Use	Directive Description
<code>.endm</code>	End macro definition	5-3	5-3
<code>macname .macro</code> [<i>parameter</i> ₁] [, ... , <i>parameter</i> _n]	Define macro by <i>macname</i>	5-3	5-3
<code>.mexit</code>	Go to <code>.endm</code>	5-3	5-3
<code>.mlib filename</code>	Identify library containing macro definitions	5-13	4-55

Table 5–3. Manipulating Substitution Symbols

Mnemonic and Syntax	Description	See Page	
		Macro Use	Directive Description
<code>.asg</code> [<i>character string</i> [""], <i>substitution symbol</i>]	Assign character string to substitution symbol	5-6	4-22
<code>.eval</code> <i>well-defined expression</i> , <i>substitution symbol</i>	Perform arithmetic on numeric substitution symbols	5-7	4-22
<code>.var</code> <i>sym</i> ₁ [, <i>sym</i> ₂ , ... , <i>sym</i> _n]	Define local macro symbols	5-12	5-12

Table 5–4. Conditional Assembly

Mnemonic and Syntax	Description	See Page	
		Macro Use	Directive Description
<code>.break</code> [<i>well-defined expression</i>]	Optional repeatable block assembly	5-14	4-53
<code>.endif</code>	End conditional assembly	5-14	4-47
<code>.endloop</code>	End repeatable block assembly	5-14	4-53
<code>.else</code>	Optional conditional assembly block	5-14	4-47
<code>.elseif</code> <i>well-defined expression</i>	Optional conditional assembly block	5-14	4-47
<code>.if</code> <i>well-defined expression</i>	Begin conditional assembly	5-14	4-47
<code>.loop</code> [<i>well-defined expression</i>]	Begin repeatable block assembly	5-14	4-53

Table 5–5. Producing Assembly-Time Messages

Mnemonic and Syntax	Description	See Page	
		Macro Use	Directive Description
.emsg	Send error message to standard output	5-17	4-37
.mmsg	Send assembly-time message to standard output	5-17	4-37
.wmsg	Send warning message to standard output	5-17	4-37

Table 5–6. Formatting the Listing

Mnemonic and Syntax	Description	See Page	
		Macro Use	Directive Description
.fclist	Allow false conditional code block listing (default)	5-19	4-40
.fcnolist	Suppress false conditional code block listing	5-19	4-40
.mlist	Allow macro listings (default)	5-19	4-56
.mnolist	Suppress macro listings	5-19	4-56
.sslist	Allow expanded substitution symbol listing	5-19	4-65
.ssnolist	Suppress expanded substitution symbol listing (default)	5-19	4-65

Archiver Description

The TMS320C6000™ archiver lets you combine several individual files into a single archive file. For example, you can collect several macros into a macro library. The assembler searches the library and uses the members that are called as macros by the source file. You can also use the archiver to collect a group of object files into an object library. The linker includes in the library the members that resolve external references during the link. The archiver allows you to modify a library by deleting, replacing, extracting, or adding members.

Topic	Page
6.1 Archiver Overview	6-2
6.2 The Archiver's Role in the Software Development Flow	6-3
6.3 Invoking the Archiver	6-4
6.4 Archiver Examples	6-6

6.1 Archiver Overview

You can build libraries from any type of files. Both the assembler and the linker accept archive libraries as input; the assembler can use libraries that contain individual source files, and the linker can use libraries that contain individual object files.

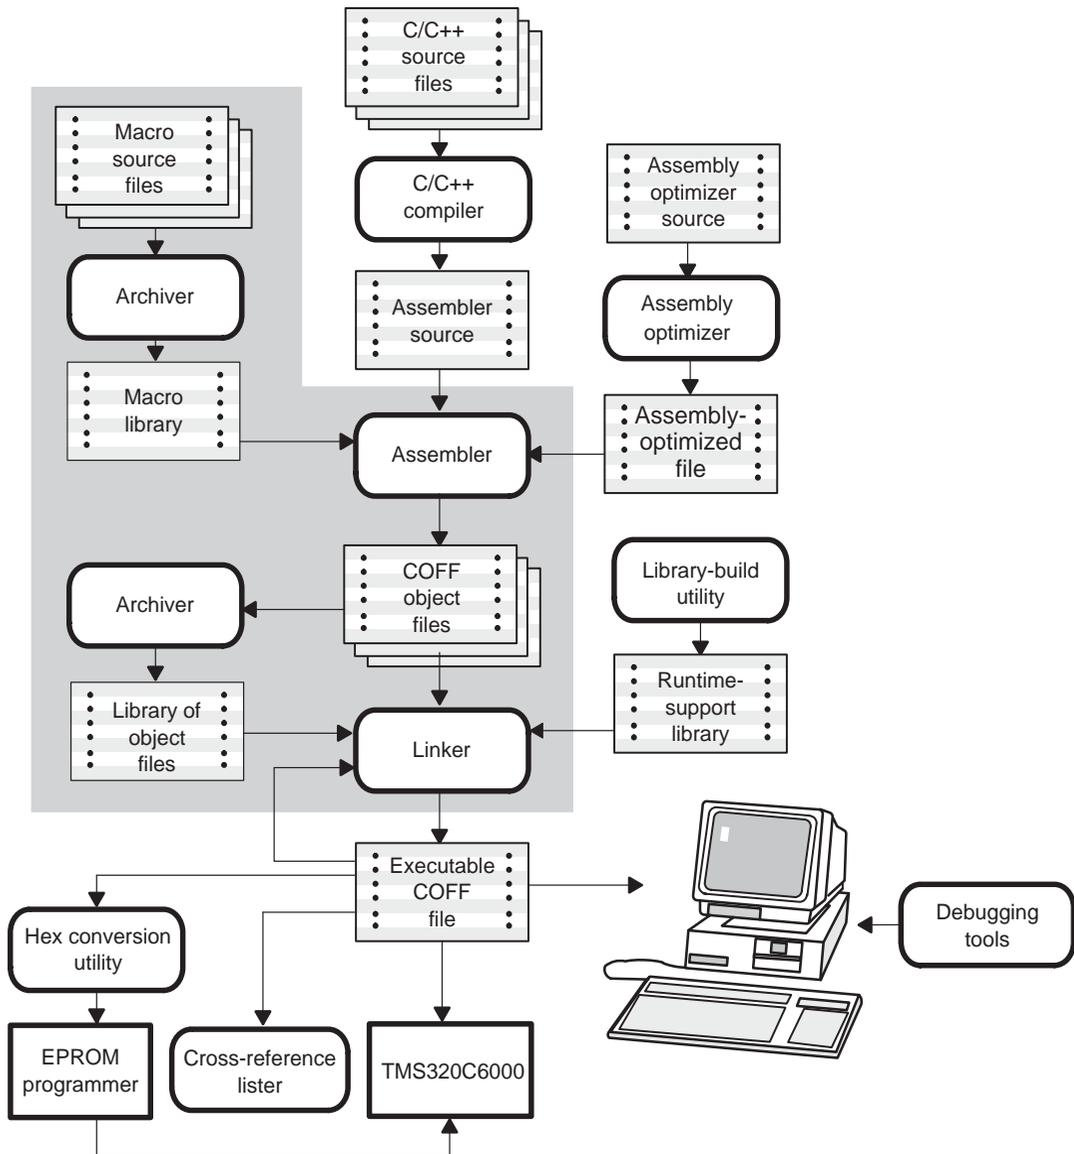
One of the most useful applications of the archiver is building libraries of object modules. For example, you can write several arithmetic routines, assemble them, and use the archiver to collect the object files into a single, logical group. You can then specify the object library as linker input. The linker searches the library and includes members that resolve external references.

You can also use the archiver to build macro libraries. You can create several source files, each of which contains a single macro, and use the archiver to collect these macros into a single, functional group. You can use the `.mlib` directive during assembly to specify that macro library to be searched for the macros that you call. Chapter 5, *Macro Language*, discusses macros and macro libraries in detail, while this chapter explains how to use the archiver to build libraries.

6.2 The Archiver's Role in the Software Development Flow

Figure 6-1 shows the archiver's role in the software development process. The shaded portion highlights the most common archiver development path. Both the assembler and the linker accept libraries as input.

Figure 6-1. The Archiver in the TMS320C6000 Software Development Flow



6.3 Invoking the Archiver

To invoke the archiver, enter:

```
ar6x [-]command [options] libname [filename1 ... filenamen]
```

ar6x is the command that invokes the archiver.

[-]command tells the archiver how to manipulate the existing library members and any specified *filenames*. A command can be preceded by an optional hyphen. You must use one of the following commands when you invoke the archiver, but you can use only one command per invocation. The archiver commands are as follows:

- @** uses the contents of the specified file instead of command line entries. You can use this command to avoid limitations on command line length imposed by the host operating system. Use a ; at the beginning of a line in the command file to include comments. (See page 6-7 for an example using an archiver command file.)
- a** adds the specified files to the library. This command does not replace an existing member that has the same name as an added file; it simply *appends* new members to the end of the archive.
- d** deletes the specified members from the library.
- r** replaces the specified members in the library. If you do not specify filenames, the archiver replaces the library members with files of the same name in the current directory. If the specified file is not found in the library, the archiver adds it instead of replacing it.
- t** prints a table of contents of the library. If you specify filenames, only those files are listed. If you do not specify any filenames, the archiver lists all the members in the specified library.
- x** extracts the specified files. If you do not specify member names, the archiver extracts all library members. When the archiver extracts a member, it simply copies the member into the current directory; it *does not* remove it from the library.

<i>options</i>	<p>In addition to one of the <i>commands</i>, you can specify options. To use options, combine them with a command; for example, to use the <i>a</i> command and the <i>s</i> option, enter <i>-as</i> or <i>as</i>. The hyphen is optional for archiver options only. These are the archiver options:</p> <ul style="list-style-type: none">-q (quiet) suppresses the banner and status messages.-s prints a list of the global symbols that are defined in the library. (This option is valid only with the <i>a</i>, <i>r</i>, and <i>d</i> commands.)-u replaces library members only if the replacement has a more recent modification date. You must use the <i>r</i> command with the <i>-u</i> option to specify which members to replace.-v (verbose) provides a file-by-file description of the creation of a new library from an old library and its members.
<i>libname</i>	<p>names the archive library to be built or modified. If you do not specify an extension for <i>libname</i>, the archiver uses the default extension <i>.lib</i>.</p>
<i>filenames</i>	<p>names individual files to be manipulated. These files can be existing library members or new files to be added to the library. When you enter a filename, you must enter a complete filename including extension, if applicable.</p>

Note: Naming Library Members

It is possible (but not desirable) for a library to contain several members with the same name. If you attempt to delete, replace, or extract a member whose name is the same as another library member, the archiver deletes, replaces, or extracts the first library member with that name.

6.4 Archiver Examples

The following are examples of typical archiver operations:

- If you want to create a library called `function.lib` that contains the files `sine.obj`, `cos.obj`, and `flt.obj`, enter:

```
ar6x -a function sine.obj cos.obj flt.obj
```

The archiver responds as follows:

```
==> new archive 'function.lib'
==> building archive 'function.lib'
```

- You can print a table of contents of `function.lib` with the `-t` command, enter:

```
ar6x -t function
```

The archiver responds as follows:

FILE NAME	SIZE	DATE
sine.obj	300	Wed Apr 16 10:00:24 1997
cos.obj	300	Wed Apr 16 10:00:30 1997
flt.obj	300	Wed Apr 16 09:59:56 1997

- If you want to add new members to the library, enter:

```
ar6x -as function atan.obj
```

The archiver responds as follows:

```
==> symbol defined: '_sin'
==> symbol defined: '$sin'
==> symbol defined: '_cos'
==> symbol defined: '$cos'
==> symbol defined: '_tan'
==> symbol defined: '$tan'
==> symbol defined: '_atan'
==> symbol defined: '$atan'
==> building archive 'function.lib'
```

Because this example does not specify an extension for the libname, the archiver adds the files to the library called `function.lib`. If `function.lib` does not exist, the archiver creates it. (The `-s` option tells the archiver to list the global symbols that are defined in the library.)

- If you want to modify a library member, you can extract it, edit it, and replace it. In this example, assume there is a library named `macros.lib` that contains the members `push.asm`, `pop.asm`, and `swap.asm`.

```
ar6x -x macros push.asm
```

The archiver makes a copy of `push.asm` and places it in the current directory; it does not remove `push.asm` from the library. Now you can edit the extracted file. To replace the copy of `push.asm` in the library with the edited copy, enter:

```
ar6x -r macros push.asm
```

- If you want to use a command file, specify the command filename after the `-@` command. For example:

```
ar6x @modules.cmd
```

The archiver responds as follows:

```
==> building archive 'modules.lib'
```

This is the `modules.cmd` command file:

```
; Command file to replace members of the
;   modules library with updated files
; Use r command and u option:
ru
; Specify library name:
modules.lib
; List filenames to be replaced if updated:
align.asm
bss.asm
data.asm
text.asm
sect.asm
clink.asm
copy.asm
double.asm
drnolist.asm
emsg.asm
end.asm
```

The `r` command specifies that the filenames given in the command file replace files of the same name in the `modules.lib` library. The `-u` option specifies that these files are replaced only when the current file has a more recent revision date than the file that is in the library.



Linker Description

The TMS320C6000™ linker creates executable modules by combining COFF object files. This chapter describes the linker options, directives, and statements used to create executable modules. Object libraries, command files, and other key concepts are discussed as well.

The concept of COFF sections is basic to linker operation; Chapter 2, *Introduction to Common Object File Format*, discusses the COFF format in detail.

Topic	Page
7.1 Linker Overview	7-2
7.2 The Linker's Role in the Software Development Flow	7-3
7.3 Invoking the Linker	7-4
7.4 Linker Options	7-5
7.5 Linker Command Files	7-22
7.6 Object Libraries	7-25
7.7 The MEMORY Directive	7-27
7.8 The SECTIONS Directive	7-31
7.9 Specifying a Section's Run-Time Address	7-44
7.10 Using UNION and GROUP Statements	7-48
7.11 Special Section Types (DSECT, COPY, and NOLOAD)	7-53
7.12 Default Allocation Algorithm	7-54
7.13 Assigning Symbols at Link Time	7-56
7.14 Creating and Filling Holes	7-64
7.15 Linker-Generated Copy Tables	7-68
7.16 Partial (Incremental) Linking	7-80
7.17 Linking C/C++ Code	7-82
7.18 Linker Example	7-87

7.1 Linker Overview

The TMS320C6000 linker allows you to configure system memory by allocating output sections efficiently into the memory map. As the linker combines object files, it performs the following tasks:

- Allocates sections into the target system's configured memory
- Relocates symbols and sections to assign them to final addresses
- Resolves undefined external references between input files

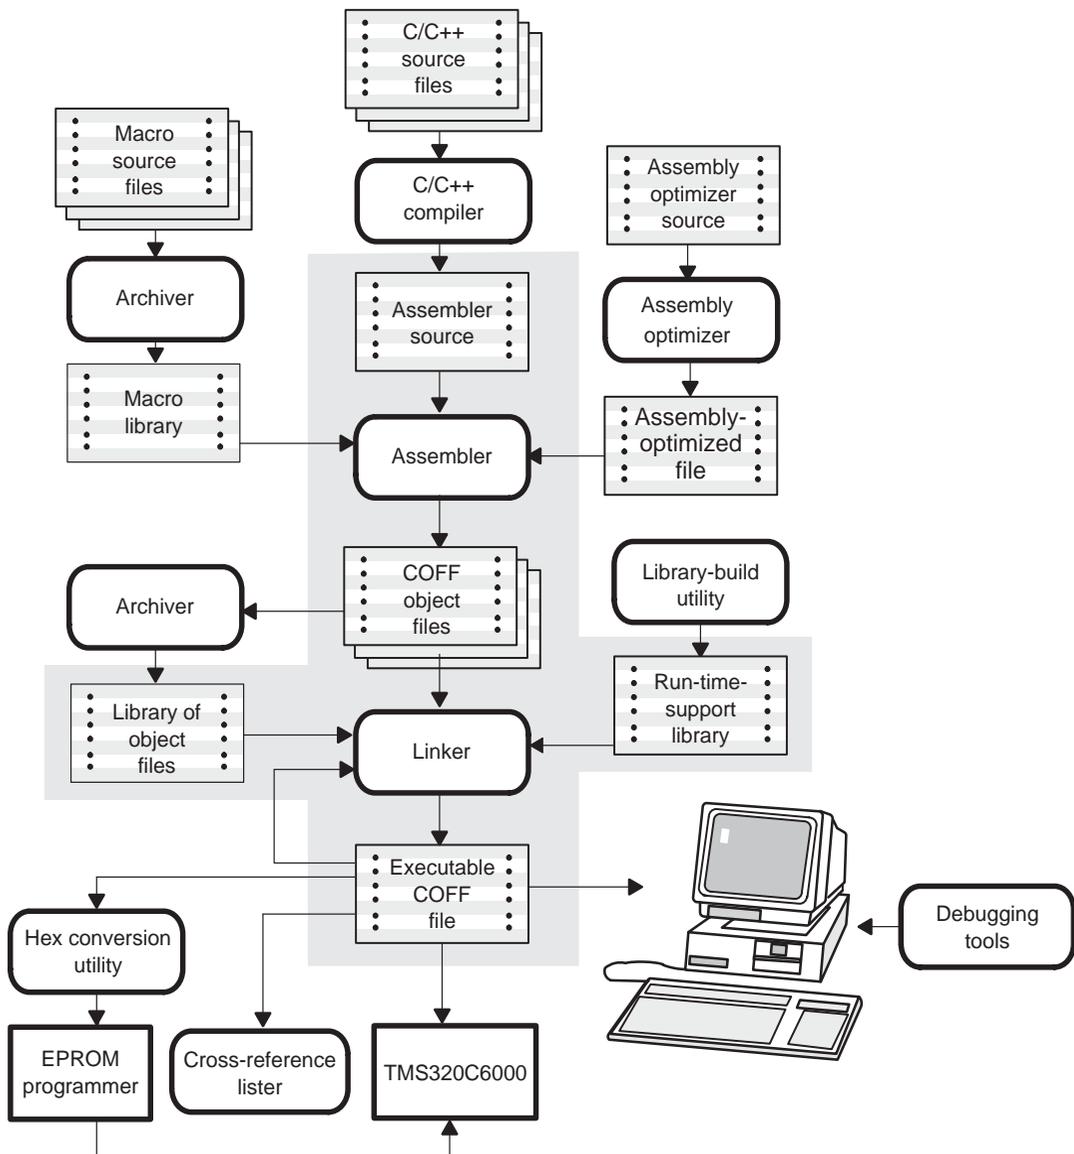
The linker command language controls memory configuration, output section definition, and address binding. The language supports expression assignment and evaluation. You configure system memory by defining and creating a memory model that you design. Two powerful directives, MEMORY and SECTIONS, allow you to:

- Allocate sections into specific areas of memory
- Combine object file sections
- Define or redefine global symbols at link time

7.2 The Linker's Role in the Software Development Flow

Figure 7-1 illustrates the linker's role in the software development process. The linker accepts several types of files as input, including object files, command files, libraries, and partially linked files. The linker creates an executable COFF object module that can be downloaded to one of several development tools or executed by a TMS320C6000 device.

Figure 7-1. The Linker in the TMS320C6000 Software Development Flow



7.3 Invoking the Linker

The general syntax for invoking the linker is:

```
cl6x -z [options] filename1 ... filenamen
```

- cl6x -z** is the command that invokes the linker.
- options* can appear anywhere on the command line or in a linker command file. (Options are discussed in section 7.4, *Linker Options*, on page 7-5.)
- filename*₁,
*filename*_n can be object files, linker command files, or archive libraries. The default extension for all input files is *.obj*; any other extension must be explicitly specified. The linker can determine whether the input file is an object or ASCII file that contains linker commands. The default output filename is *a.out*, unless you use the *-o* option to name the output file.

There are two methods for invoking the linker:

- Specify options and filenames on the command line. This example links two files, *file1.obj* and *file2.obj*, and creates an output module named *link.out*.

```
cl6x -z file1.obj file2.obj -o link.out
```

- Put filenames and options in a linker command file. Filenames that are specified inside a linker command file must begin with a letter. For example, assume the file *linker.cmd* contains the following lines:

```
-o link.out
file1.obj
file2.obj
```

Now you can invoke the linker from the command line; specify the command filename as an input file:

```
cl6x -z linker.cmd
```

When you use a command file, you can also specify other options and files on the command line. For example, you could enter:

```
cl6x -z -m link.map linker.cmd file3.obj
```

The linker reads and processes a command file as soon as it encounters the filename on the command line, so it links the files in this order: *file1.obj*, *file2.obj*, and *file3.obj*. This example creates an output file called *link.out* and a map file called *link.map*.

For information on invoking the linker for C/C++ files, see section 7.17, *Linking C/C++ Code*, on page 7-82.

7.4 Linker Options

Linker options control linking operations. They can be placed on the command line or in a command file. Linker options must be preceded by a hyphen (-). Options can be separated from arguments (if they have them) by an optional space. Table 7-1 summarizes the linker options.

Table 7-1. Linker Options Summary

Option	Description	Page
-a	Produces an absolute, executable module. This is the default; if neither -a nor -r is specified, the linker acts as if -a were specified.	7-6
-abs	Produces an absolute listing file	7-7
-ar	Produces a relocatable, executable object module	7-6
--args	Allocates memory to be used by the loader to pass arguments	7-8
-b	Disables merge of symbolic debugging information	7-8
-c	Autoinitializes variables at run time	7-9
-cr	Initializes variables at load time	7-9
-e <i>global_symbol</i>	Defines a global symbol that specifies the primary entry point for the output module	7-9
-f <i>fill_value</i>	Sets default fill values for holes within output sections; <i>fill_value</i> is a 32-bit constant	7-10
-g <i>symbol</i>	Makes <i>symbol</i> global (overrides -h)	7-10
-h	Makes all global symbols static	7-10
-heap <i>size</i>	Sets heap size (for the dynamic memory allocation in C) to <i>size</i> words and defines a global symbol that specifies the heap size. Default = 1K words	7-11
-help or -?	Produces help listing	--
-i <i>pathname</i>	Alters library-search algorithms to look in a directory named with <i>pathname</i> before looking in the default location. This option must appear before the -l option.	7-12
-j	Disables conditional linking	7-13
-l <i>filename</i>	Names an archive library or linker command <i>filename</i> as linker input	7-11
-m <i>filename</i>	Produces a map or listing of the input and output sections, including holes, and places the listing in <i>filename</i>	7-14
-o <i>filename</i>	Names the executable output module. The default filename is a.out.	7-15

Table 7-1. Linker Options Summary (Continued)

Option	Description	Page
-priority	Satisfies unresolved references by the first library that contains a definition for that symbol	7-20
-r	Produces a nonexecutable, relocatable output module	7-6
-s	Strips symbol table information and line number entries from the output module	7-16
-stack size	Sets C system stack size to <i>size</i> words and defines a global symbol that specifies the stack size. Default = 1K words	7-16
--trampolines	Generates far call trampolines	7-17
-u symbol	Places an unresolved external <i>symbol</i> into the output module's symbol table	7-19
-w	Displays a message when an undefined output section is created	7-20
-x	Forces rereading of libraries, which resolves back references	7-20
--xml_link_info file	Generates a well-formed XML <i>file</i> containing detailed information about the result of a link	7-21

7.4.1 Relocation Capabilities (-a and -r Options)

The linker performs relocation, which is the process of adjusting all references to a symbol when the symbol's address changes. The linker supports two options (-a and -r) that allow you to produce an absolute or a relocatable output module.

□ Producing an absolute output module (-a option)

When you use the -a option without the -r option, the linker produces an *absolute, executable* output module. Absolute files contain *no* relocation information. Executable files contain the following:

- Special symbols defined by the linker (section 7.13.4, on page 7-59, describes these symbols)
- An optional header that describes information such as the program entry point
- *No* unresolved references

The following example links file1.obj and file2.obj and creates an absolute output module called a.out:

```
cl6x -z -a file1.obj file2.obj
```

Note: The -a and -r Options

If you do not use the -a or the -r option, the linker acts as if you specified -a.

❑ Producing a relocatable output module (`-r` option)

When you use the `-r` option without the `-a` option, the linker retains relocation entries in the output module. If the output module is relocated (at load time) or relinked (by another linker execution), use `-r` to retain the relocation entries.

The linker produces a file that is not executable when you use the `-r` option without `-a`. A file that is not executable does not contain special linker symbols or an optional header. The file can contain unresolved references, but these references do not prevent creation of an output module.

This example links `file1.obj` and `file2.obj` and creates a relocatable output module called `a.out`:

```
cl6x -z -r file1.obj file2.obj
```

The output file `a.out` can be relinked with other object files or relocated at load time. (Linking a file that will be relinked with other files is called partial linking. For more information, see section 7.16, *Partial (Incremental) Linking*, on page 7-80.)

❑ Producing an executable relocatable output module (`-ar` option combination)

If you invoke the linker with both the `-a` and `-r` options, the linker produces an *executable, relocatable* object module. The output file contains the special linker symbols, an optional header, and all resolved symbol references; however, the relocation information is retained.

This example links `file1.obj` and `file2.obj` and creates an executable, relocatable output module called `xr.out`:

```
cl6x -z -ar file1.obj file2.obj -o xr.out
```

When the linker encounters a file that contains no relocation or symbol table information, it issues a warning message (but continues executing). Relinking an absolute file can be successful only if each input file contains no information that needs to be relocated (that is, each file has no unresolved references and is bound to the same virtual address that it was bound to when the linker created it).

7.4.2 Create an Absolute Listing File (`-abs` Option)

The `-abs` option produces an output file for each file that was linked. These files are named with the input filenames and an extension of `.abs`. Header files, however, do not generate a corresponding `.abs` file.

7.4.3 Allocate Memory for Use by the Loader to Pass Arguments (`--args` Option)

The `--args` option instructs the linker to allocate memory to be used by the loader to pass arguments from the command line of the loader to the program. The syntax of the `--args` option is:

```
--args=size
```

The `size` is a number representing the number of bytes to be allocated in target memory for command-line arguments.

By default, the linker creates the `__c_args__` symbol and sets it to `-1`. When you specify `--args=size`, the following occur:

- The linker creates an uninitialized section named `.args` of `size` bytes.
- The `__c_args__` symbol contains the address of the `.args` section.

The loader and the target boot code use the `.args` section and the `__c_args__` symbol to determine whether and how to pass arguments from the host to the target program. See the *TMS320C6000 Optimizing Compiler User's Guide* for information about the loader.

7.4.4 Disable Merge of Symbolic Debugging Information (`-b` Option)

By default, the linker eliminates duplicate entries of symbolic debugging information. Such duplicate information is commonly generated when a C program is compiled for debugging. For example:

```
-[ header.h ]-
typedef struct
{
    <define some structure members>
} XYZ;

-[ f1.c ]-
#include "header.h"
...

-[ f2.c ]-
#include "header.h"
...
```

When these files are compiled for debugging, both `f1.obj` and `f2.obj` have symbolic debugging entries to describe type `XYZ`. For the final output file, only one set of these entries is necessary. The linker eliminates the duplicate entries automatically.

Use the `-b` option if you want the linker to keep such duplicate entries. Using the `-b` option has the effect of the linker running faster and using less machine memory.

7.4.5 C Language Options (`-c` and `-cr` Options)

The `-c` and `-cr` options cause the linker to use linking conventions that are required by the C compiler.

- The `-c` option tells the linker to autoinitialize variables at run time.
- The `-cr` option tells the linker to initialize variables at load time.

For more information, see section 7.17, *Linking C Code*, on page 7-82, section 7.17.4, *Autoinitialization of Variables at Run Time*, on page 7-84, and section 7.17.5, *Initialization of Variables at Load Time*, on page 7-85.

7.4.6 Define an Entry Point (`-e global_symbol` Option)

The memory address at which a program begins executing is called the *entry point*. When a loader loads a program into target memory, the program counter (PC) must be initialized to the entry point; the PC then points to the beginning of the program.

The linker can assign one of four values to the entry point. These values are listed below in the order in which the linker tries to use them. If you use one of the first three values, it must be an external symbol in the symbol table.

- The value specified by the `-e` option. The syntax is:

```
-e global_symbol
```

where *global_symbol* defines the entry point and must be as an external symbol of the input files.

- The value of symbol `_c_int00` (if present). The `_c_int00` symbol *must* be the entry point if you are linking code produced by the C compiler.
- The value of symbol `_main` (if present)
- 0 (default value)

This example links `file1.obj` and `file2.obj`. The symbol `begin` is the entry point; `begin` must be defined as external in `file1` or `file2`.

```
cl6x -z -e begin file1.obj file2.obj
```

7.4.7 Set Default Fill Value (`-f fill_value` Option)

The `-f` option fills the holes formed within output sections. The syntax for the `-f` option is:

```
-f fill_value
```

The argument `fill_value` is a 32-bit constant (up to eight hexadecimal digits). If you do not use `-f`, the linker uses 0 as the default fill value.

This example fills holes with the hexadecimal value ABCDABCD:

```
c16x -z -f 0xABCDABCD file1.obj file2.obj
```

7.4.8 Make a Symbol Global (`-g symbol` Option)

The `-h` option makes all global symbols static. If you have a symbol that you want to remain global and you use the `-h` option, you can use the `-g` option to declare that symbol to be global. The `-g` option overrides the effect of the `-h` option for the symbol that you specify. The syntax for the `-g` option is:

```
-g global_symbol
```

7.4.9 Make All Global Symbols Static (`-h` Option)

The `-h` option makes all global symbols static. Static symbols are not visible to externally linked modules. By making global symbols static, global symbols are essentially hidden. This allows external symbols with the same name (in different files) to be treated as unique.

The `-h` option effectively nullifies all `.global` assembler directives. All symbols become local to the module in which they are defined, so no external references are possible. For example, assume `file1.obj` and `file2.obj` both define global symbols called `EXT`. By using the `-h` option, you can link these files without conflict. The symbol `EXT` defined in `file1.obj` is treated separately from the symbol `EXT` defined in `file2.obj`.

```
c16x -z -h file1.obj file2.obj
```

7.4.10 Define Heap Size (`-heap size` Option)

The C/C++ compiler uses an uninitialized section called `.systemem` for the C run-time memory pool used by `malloc()`. You can set the size of this memory pool at link time by using the `-heap` option. The syntax for the `-heap` option is:

```
-heap size
```

The *size* must be a constant. This example defines a 4K byte heap:

```
c16x -z -heap 0x1000 /* defines a 4k heap (.systemem section)*/
```

The linker creates the `.systemem` section only if there is a `.systemem` section in an input file.

The linker also creates a global symbol `__SYSTEMEM_SIZE` and assigns it a value equal to the size of the heap. The default size is 1K bytes.

For more information, see section 7.17, *Linking C/C++ Code*, on page 7-82.

7.4.11 Alter the Library Search Algorithm (`-l` Option, `-I` Option, and `C_DIR/C6X_C_DIR` Environment Variables)

Usually, when you want to specify a file as linker input, you simply enter the filename; the linker looks for the file in the current directory. For example, suppose the current directory contains the library `object.lib`. Assume that this library defines symbols that are referenced in the file `file1.obj`. This is how you link the files:

```
c16x -z file1.obj object.lib
```

If you want to use a file that is not in the current directory, use the `-l` (lowercase L) linker option. The syntax for this option is:

```
-l [pathname] filename
```

The *filename* is the name of an archive, an object file, or linker command file; the space between `-l` and the filename is optional.

The `-l` option is not required when one or more members of an object library are specified for input to an output section. For more information, see section 7.8.6, *Allocating an Archive Member to an Output Section* on page 7-43.

You can augment the linker's directory search algorithm by using the `-I` linker option or the `C_DIR` or `C6X_C_DIR` environment variables. The linker searches for object libraries and command files specified by the `-I` option in the following order:

- 1) It searches directories named with the `-I` linker option. The `-I` option must appear before the `-l` option on the command line or in a command file.
- 2) It searches directories named with `C_DIR` and `C6X_C_DIR`.
- 3) If `C_DIR` and `C6X_C_DIR` are not set, it searches directories named with the assembler's `A_DIR` or `C6X_A_DIR` environment variable.
- 4) It searches the current directory.

7.4.11.1 Name an Alternate Library Directory (`-I` pathname Option)

The `-I` option names an alternate directory that contains input files. The syntax for this option is:

```
-I pathname
```

The *pathname* names a directory that contains input files; the space between `-I` and the *pathname* is optional.

When the linker is searching for input files named with the `-l` option, it searches through directories named with `-I` first. Each `-I` option specifies only one directory, but you can several `-I` options per invocation. When you use the `-I` option to name an alternate directory, it must precede any `-l` option used on the command line or in a command file.

For example, assume that there are two archive libraries called `r.lib` and `lib2.lib`. The table below shows the directories that `r.lib` and `lib2.lib` reside in, how to set environment variable, and how to use both libraries during a link. Select the row for your operating system:

Operating System	Pathname	Enter
UNIX (Bourne shell)	/ld and /ld2	c16x -z f1.obj f2.obj -I/ld -I/ld2 -lr.lib -llib2.lib
Windows	\ld and \ld2	c16x -z f1.obj f2.obj -I\ld -I\ld2 -lr.lib -llib2.lib

7.4.11.2 Name an Alternate Library Directory (C_DIR and C6X_C_DIR Environment Variables)

An environment variable is a system symbol that you define and assign a string to. The linker uses environment variables named C6X_C_DIR and C_DIR to name alternate directories that contain object libraries. The command syntaxes for assigning the environment variable are:

Operating System	Enter
UNIX (Bourne shell)	<code>C_DIR="pathname₁;pathname₂;..."; export C_DIR</code>
Windows	<code>set C_DIR= pathname₁;pathname₂;...</code>

The *pathnames* are directories that contain input files. Use the `-l` (lowercase L) linker option on the command line or in a command file to tell the linker which library or linker command file to search for.

In the example below, assume that two archive libraries called `r.lib` and `lib2.lib` reside in `ld` and `ld2` directories. The table below shows the directories that `r.lib` and `lib2.lib` reside in, how to set the environment variable, and how to use both libraries during a link. Select the row for your operating system:

Operating System	Pathname	Invocation Command
UNIX (Bourne shell)	<code>/ld and /ld2</code>	<code>C_DIR="/ld ;/ld2"; export C_DIR cl6x -z f1.obj f2.obj -l r.lib -l lib2.lib</code>
Windows	<code>\ld and \ld2</code>	<code>set C_DIR=\ld;\ld2 cl6x -z f1.obj f2.obj -l r.lib -l lib2.lib</code>

The environment variable remains set until you reboot the system or reset the variable by entering:

Operating System	Enter
UNIX (Bourne shell)	<code>unset C_DIR</code>
Windows	<code>set C_DIR=</code>

The assembler uses an environment variable named C6X_A_DIR or A_DIR to name alternate directories that contain copy/include files or macro libraries. If C6X_C_DIR or C_DIR is not set, the linker searches for object libraries in the directories named with C6X_A_DIR or A_DIR. For more information about object libraries, see section 7.6 on page 7-25.

7.4.12 Disable Conditional Linking (`-j` Option)

The `-j` option disables removal of unreferenced sections. Only sections marked as candidates for removal with the `.clink` assembler directive are affected by conditional linking. See page 4-28 for details on setting up conditional linking using the `.clink` directive.

7.4.13 Create a Map File (`-m filename` Option)

The `-m` option creates a linker map listing and puts it in *filename*. The syntax for the `-m` option is:

```
-m filename
```

The linker map describes:

- Memory configuration
- Input and output section allocation
- The addresses of external symbols after they have been relocated

The map file contains the name of the output module and the entry point; it can also contain up to three tables:

- A table showing the new memory configuration if any nondefault memory is specified (memory configuration). The table has the following columns; this information is generated on the basis of the information in the `MEMORY` directive in the linker command file:

- **Name.** This is the name of the memory range specified with the `MEMORY` directive.
- **Origin.** This specifies the starting address of a memory range.
- **Length.** This specifies the length of a memory range.
- **Attributes.** This specifies one to four attributes associated with the named range:
 - R** specifies that the memory can be read.
 - W** specifies that the memory can be written to.
 - X** specifies that the memory can contain executable code.
 - I** specifies that the memory can be initialized.
- **Fill.** This specifies a fill character for the memory range.

For more information about the `MEMORY` directive, see section 7.7, *The MEMORY Directive*, on page 7-27.

- A table showing the linked addresses of each output section and the input sections that make up the output sections (section allocation map). This table has the following columns; this information is generated on the basis of the information in the `SECTIONS` directive in the linker command file:
 - **Output section.** This is the name of the output section specified with the `SECTIONS` directive.

- **Origin.** The first origin listed for each output section is the starting address of that output section. The indented origin value is the starting address of that portion of the output section.
- **Length.** The first length listed for each output section is the length of that output section. The indented length value is the length of that portion of the output section.
- **Attributes/input sections.** This lists the input file or value associated with an output section.

For more information about the SECTIONS directive, see section 7.8, *The SECTIONS Directive*, on page 7-31.

- A table showing each external symbol and its address sorted by symbol name.
- A table showing each external symbol and its address sorted by symbol address.

This following example links file1.obj and file2.obj and creates a map file called map.out:

```
cl6x -z file1.obj file2.obj -m map.out
```

Example 7-21 on page 7-89 shows an example of a map file.

7.4.14 Name an Output Module (-o Option)

The linker creates an output module when no errors are encountered. If you do not specify a filename for the output module, the linker gives it the default name a.out. If you want to write the output module to a different file, use the -o option. The syntax for the -o option is:

```
-o filename
```

The *filename* is the new output module name.

This example links file1.obj and file2.obj and creates an output module named run.out:

```
cl6x -z -o run.out file1.obj file2.obj
```

7.4.15 Strip Symbolic Information (`-s` Option)

The `-s` option creates a smaller output module by omitting symbol table information and line number entries. The `-s` option is useful for production applications when you do not want to disclose symbolic information to the consumer.

This example links `file1.obj` and `file2.obj` and creates an output module, stripped of line numbers and symbol table information, named `nosym.out`:

```
cl6x -z -o nosym.out -s file1.obj file2.obj
```

Using the `-s` option limits later use of a symbolic debugger.

7.4.16 Define Stack Size (`-stack size` Option)

The TMS320C6000 C/C++ compiler uses an uninitialized section, `.stack`, to allocate space for the run-time stack. You can set the size of this section in bytes at link time with the `-stack` option. The syntax for the `-stack` option is:

```
-stack size
```

The *size* must be a constant and is in bytes. This example defines a 4K byte stack:

```
cl6x -z -stack 0x1000 /* defines a 4K stack (.stack section) */
```

If you specified a different stack size in an input section, the input section stack size is ignored. Any symbols defined in the input section remain valid; only the stack size is different.

When the linker defines the `.stack` section, it also defines a global symbol, `__STACK_SIZE`, and assigns it a value equal to the size of the section. The default software stack size is 1K bytes.

7.4.17 Generate Far Call Trampolines (`--trampolines` Option)

The TMS320C6000 has PC-relative call and PC-relative branch instructions whose range is smaller than the entire address space. When these instructions are used, the destination address must be near enough to the instruction that the difference between the call and the destination fits in the available encoding bits. If the called function is too far away from the calling function, the linker generates an error.

The alternative to a PC-relative call is an absolute call, which is often implemented as an indirect call: load the called address into a register, and call that register. This is often undesirable because it takes more instructions (speed- and size-wise) and requires an extra register to contain the address.

By default, the compiler generates near calls. The `--trampolines` option causes the linker to generate a trampoline code section for each call that is linked out-of-range of its called destination. The trampoline code section contains a sequence of instructions that performs a transparent long branch to the original called address. Each calling instruction that is out-of-range from the called function is redirected to the trampoline.

For example, in a section of C code the bar function calls the foo function. The compiler generates this code for the function:

```
bar:
    ...
    call    foo        ; call the function "foo"
    ...
```

If the foo function is placed out-of-range from the call to foo that is inside of bar, then with `--trampolines` the linker changes the original call to foo into a call to `foo_trampoline` as shown:

```
bar:
    ...
    call    foo_trampoline ; call a trampoline for foo
    ...
```

The above code generates a trampoline code section called `foo_trampoline`, which contains code that executes a long branch to the original called function, foo. For example:

```
foo_trampoline:
    branch_long    foo
```

Trampolines can be shared among calls to the same called function. The only requirement is that all calls to the called function be linked near the called function's trampoline.

When the linker produces a map file (the `-m` option) and it has produced one or more trampolines, then the map file will contain statistics about what trampolines were generated to reach which functions. A list of calls for each trampoline is also provided in the map file.

7.4.17.1 Carrying Trampolines from Load Space to Run Space

It is sometimes useful to load code in one location in memory and run it in another. The linker provides the capability to specify separate load and run allocations for a section. The burden of actually copying the code from the load space to the run space is left to you.

A copy function must be executed before the real function can be executed in its run space. To facilitate this copy function, the assembler provides the `.label` directive, which allows you to define a load-time address. These load-time addresses can then be used to determine the start address and size of the code to be copied. However, this mechanism will *not* work if the code contains a call that requires a trampoline to reach its called function. This is because the trampoline code is generated at link time, after the load-time addresses associated with the `.label` directive have been defined. If the linker detects the definition of a `.label` symbol in an input section that contains a trampoline call, then a warning is generated.

To solve this problem, you can use the `START()`, `END()`, and `SIZE()` operators (see section 7.13.7, on page 7-61). These operators allow you to define symbols to represent the load-time start address and size inside the linker command file. These symbols can be referenced by the copy code, and their values are not resolved until link time, after the trampoline sections have been allocated.

Here is an example of how you could use the `START()` and `SIZE()` operators in association with an output section to copy the trampoline code section along with the code containing the calls that need trampolines:

```
SECTIONS
{
  .foo : load = ROM, run = RAM, start(foo_start), size(foo_size)
        { x.obj(.text) }

  .text: {} > ROM

  .far : { -lrts.lib(.text) } > FAR_MEM
}
```

A function in `x.obj` contains a run-time-support call. The run-time-support library is placed in far memory and so the call is out-of-range. A trampoline section will be added to the `.foo` output section by the linker. The copy code can refer to the symbols `foo_start` and `foo_size` as parameters for the load start address and size of the entire `.foo` output section. This allows the copy code to copy the trampoline section along with the original `x.obj` code in `.text` from its load space to its run space.

7.4.17.2 Disadvantages of Using Trampolines

An alternative method to creating a trampoline code section for a call that cannot reach its called function is to actually modify the source code for the call. In some cases this can be done without affecting the size of the code. However, in general, this approach is extremely difficult, especially when the size of the code is affected by the transformation.

While generating far call trampolines provides a more straightforward solution, trampolines have the disadvantage that they are somewhat slower than directly calling a function. They require both a call and a branch. Additionally, while inline code could be tailored to the environment of the call, trampolines are generated in a more general manner, and may be slightly less efficient than inline code.

7.4.18 Introduce an Unresolved Symbol (`-u symbol` Option)

The `-u` option introduces an unresolved symbol into the linker's symbol table. This forces the linker to search a library and include the member that defines the symbol. The linker must encounter the `-u` option *before* it links in the member that defines the symbol. The syntax for the `-u` option is:

```
-u symbol
```

For example, suppose a library named `rts6200.lib` contains a member that defines the symbol `syntab`; none of the object files being linked reference `syntab`. However, suppose you plan to relink the output module and you want to include the library member that defines `syntab` in this link. Using the `-u` option as shown below forces the linker to search `rts6200.lib` for the member that defines `syntab` and to link in the member.

```
cl6x -z -u syntab file1.obj file2.obj rts6200.lib
```

If you do not use `-u`, this member is not included, because there is no explicit reference to it in `file1.obj` or `file2.obj`.

7.4.19 Display a Message When an Undefined Output Section Is Created (`-w` Option)

In a linker command file, you can set up a `SECTIONS` directive that describes how input sections are combined into output sections. However, if the linker encounters one or more input sections that do not have a corresponding output section defined in the `SECTIONS` directive, the linker combines the input sections that have the same name into an output section with that name. By default, the linker does not display a message to tell you that this occurred.

You can use the `-w` option to cause the linker to display a message when it creates a new output section.

For more information about the `SECTIONS` directive, see section 7.8 on page 7-31. For more information about the default actions of the linker, see section 7.12 on page 7-54.

7.4.20 Exhaustively Read and Search Libraries (`-x` and `-priority` Options)

There are two ways to exhaustively search for unresolved symbols:

- Reread libraries if you cannot resolve a symbol reference (`-x`).
- Search libraries in the order that they are specified (`-priority`).

The linker normally reads input files, including archive libraries, only once when they are encountered on the command line or in the command file. When an archive is read, any members that resolve references to undefined symbols are included in the link. If an input file later references a symbol defined in a previously read archive library, the reference is not resolved.

With the `-x` option, you can force the linker to reread all libraries. The linker rereads libraries until no more references can be resolved. Linking using `-x` may be slower, so you should use it only as needed. For example, if `a.lib` contains a reference to a symbol defined in `b.lib`, and `b.lib` contains a reference to a symbol defined in `a.lib`, you can resolve the mutual dependencies by listing one of the libraries twice, as in:

```
cl6x -z -la.lib -lb.lib -la.lib
```

or you can force the linker to do it for you:

```
cl6x -z -x -la.lib -lb.lib
```

The `-priority` option provides an alternate search mechanism for libraries. Using `-priority` causes each unresolved reference to be satisfied by the first library that contains a definition for that symbol. For example:

```
objfile  references A
lib1     defines B
lib2     defines A, B; obj defining A references B

% cl6x -z objfile lib1 lib2
```

Under the existing model, objfile resolves its reference to A in lib2, pulling in a reference to B, which resolves to the B in lib2.

Under `-priority`, objfile resolves its reference to A in lib2, pulling in a reference to B, but now B is resolved by searching the libraries in order and resolves B to the first definition it finds, namely the one in lib1.

The `-priority` option is useful for libraries that provide overriding definitions for related sets of functions in other libraries without having to provide a complete version of the whole library.

For example, suppose you want to override versions of `malloc` and `free` defined in the `rts6200.lib` without providing a full replacement for `rts6200.lib`. Using `-priority` and linking your new library before `rts6200.lib` guarantees that all references to `malloc` and `free` resolve to the new library.

The `-priority` option is intended to support linking programs with DSP/BIOS where situations like the one illustrated above occur.

7.4.21 Generate XML Link Information File (`--xml_link_info` Option)

The linker supports the generation of an XML link information file via the `--xml_link_info file` option. This option causes the linker to generate a well-formed XML file containing detailed information about the result of a link. The information included in this file includes all of the information that is currently produced in a linker generated map file.

See Appendix C, *XML Link Information File Description*, for specifics on the contents of the generated file.

7.5 Linker Command Files

Linker command files allow you to put linking information in a file; this is useful when you invoke the linker often with the same information. Linker command files are also useful because they allow you to use the MEMORY and SECTIONS directives to customize your application. You must use these directives in a command file; you cannot use them on the command line.

Linker command files are ASCII files that contain one or more of the following:

- Input filenames, which specify object files, archive libraries, or other command files. (If a command file calls another command file as input, this statement must be the *last* statement in the calling command file. The linker does not return from called command files.)
- Linker options, which can be used in the command file in the same manner that they are used on the command line
- The MEMORY and SECTIONS linker directives. The MEMORY directive defines the target memory configuration (see section 7.7, on page 7-27). The SECTIONS directive controls how sections are built and allocated (see section 7.8 on page 7-31.)
- Assignment statements, which define and assign values to global symbols

To invoke the linker with a command file, enter the `cl6x -z` command and follow it with the name of the command file:

```
cl6x -z command_filename
```

The linker processes input files in the order that it encounters them. If the linker recognizes a file as an object file, it links the file. Otherwise, it assumes that a file is a command file and begins reading and processing commands from it. Command filenames are case sensitive, regardless of the system used.

Example 7-1 shows a sample linker command file called `link.cmd`.

Example 7-1. Linker Command File

```
a.obj          /* First input filename      */
b.obj          /* Second input filename       */
-o prog.out    /* Option to specify output file */
-m prog.map    /* Option to specify map file   */
```

The sample file in Example 7-1 contains only filenames and options. (You can place comments in a command file by delimiting them with `/*` and `*/`.) To invoke the linker with this command file, enter:

```
cl6x -z link.cmd
```

You can place other parameters on the command line when you use a command file:

```
c16x -z -r link.cmd c.obj d.obj
```

The linker processes the command file as soon as it encounters the filename, so a.obj and b.obj are linked into the output module before c.obj and d.obj.

You can specify multiple command files. If, for example, you have a file called names.lst that contains filenames and another file called dir.cmd that contains linker directives, you could enter:

```
c16x -z names.lst dir.cmd
```

One command file can call another command file; this type of nesting is limited to 16 levels. If a command file calls another command file as input, this statement must be the *last* statement in the calling command file.

Blanks and blank lines are insignificant in a command file except as delimiters. This also applies to the format of linker directives in a command file. Example 7-2 shows a sample command file that contains linker directives.

Example 7-2. Command File With Linker Directives

```
a.obj b.obj c.obj          /* Input filenames      */
-o prog.out -m prog.map    /* Options              */

MEMORY                    /* MEMORY directive    */
{
  FAST_MEM:  origin = 0x0100    length = 0x0100
  SLOW_MEM:  origin = 0x7000    length = 0x1000
}

SECTIONS                  /* SECTIONS directive  */
{
  .text:    > SLOW_MEM
  .data:    > SLOW_MEM
  .bss:     > FAST_MEM
}
```

For more information about the MEMORY directive, see section 7.7, *The MEMORY Directive*, on page 7-27. For more information about the SECTIONS directive, see section 7.8, *The SECTIONS Directive*, on page 7-31.

7.5.1 Reserved Names in Linker Command Files

The following names are reserved as keywords for linker directives. Do not use them as symbol or section names in a command file.

align	group	org
ALIGN	GROUP	origin
attr	l (lowercase L)	ORIGIN
ATTR	len	range
block	length	run
BLOCK	LENGTH	RUN
COPY	load	SECTIONS
DSECT	LOAD	spare
f	MEMORY	type
fill	NOLOAD	TYPE
FILL	o	UNION

7.5.2 Constants in Linker Command Files

You can specify constants with either of two syntax schemes: the scheme used for specifying decimal, octal, or hexadecimal constants used in the assembler (see section 3.6, *Constants*, on page 3-14) or the scheme used for integer constants in C syntax.

Examples:

Format	Decimal	Octal	Hexadecimal
Assembler format	32	40q	020h
C format	32	040	0x20

7.6 Object Libraries

An object library is a partitioned archive file that contains object files as members. Usually, a group of related modules are grouped together into a library. When you specify an object library as linker input, the linker includes any members of the library that define existing unresolved symbol references. You can use the archiver to build and maintain libraries. Chapter 6, *Archiver Description*, contains more information about the archiver.

Using object libraries can reduce link time and the size of the executable module. Normally, if an object file that contains a function is specified at link time, the file is linked whether the function is used or not; however, if that same function is placed in an archive library, the file is included only if the function is referenced.

The order in which libraries are specified is important, because the linker includes only those members that resolve symbols that are undefined at the time the library is searched. The same library can be specified as often as necessary; it is searched each time it is included. Alternatively, you can use the `-x` option to reread libraries until no more references can be resolved (see section 7.4.20, *Exhaustively Read and Search Libraries (-x and -priority Options)*, on page 7-20). A library has a table that lists all external symbols defined in the library; the linker searches through the table until it determines that it cannot use the library to resolve any more references.

The following examples link several files and libraries, using these assumptions:

- Input files `f1.obj` and `f2.obj` both reference an external function named `clrscr`.
- Input file `f1.obj` references the symbol `origin`.
- Input file `f2.obj` references the symbol `fillclr`.
- Member 0 of library `libc.lib` contains a definition of `origin`.
- Member 3 of library `liba.lib` contains a definition of `fillclr`.
- Member 1 of both libraries defines `clrscr`.

If you enter:

```
c16x -z f1.obj f2.obj liba.lib libc.lib
```

then:

- Member 1 of `liba.lib` satisfies the `f1.obj` and `f2.obj` references to `clrscr` because the library is searched and the definition of `clrscr` is found.
- Member 0 of `libc.lib` satisfies the reference to `origin`.
- Member 3 of `liba.lib` satisfies the reference to `fillclr`.

If, however, you enter:

```
c16x -z f1.obj f2.obj libc.lib liba.lib
```

then the references to clrscr are satisfied by member 1 of libc.lib.

If none of the linked files reference symbols defined in a library, you can use the `-u` option to force the linker to include a library member. (See section 7.4.18, *Introduce an Unresolved Symbol (-u symbol Option)*, on page 7-19.) The next example creates an undefined symbol `rot1` in the linker's global symbol table:

```
c16x -z -u rot1 libc.lib
```

If any member of `libc.lib` defines `rot1`, the linker includes that member.

Library members are allocated according to the `SECTIONS` directive default allocation algorithm. For more information, see section 7.8, *The SECTIONS Directive*, on page 7-31.

Section 7.4.11, *Alter the Library Search Algorithm (-I Option, -I Option, and C_DIR/C6X_C_DIR Environment Variables)* on page 7-11 describes methods for specifying directories that contain object libraries.

7.7 The MEMORY Directive

The linker determines where output sections are allocated into memory; it must have a model of target memory to accomplish this. The MEMORY directive allows you to specify a model of target memory so that you can define the types of memory your system contains and the address ranges they occupy. The linker maintains the model as it allocates output sections and uses it to determine which memory locations can be used for object code.

The memory configurations of TMS320C6000 systems differ from application to application. The MEMORY directive allows you to specify a variety of configurations. After you use MEMORY to define a memory model, you can use the SECTIONS directive to allocate output sections into defined memory.

For more information, see section 2.3, *How the Linker Handles Sections*, on page 2-11 and section 2.4, *Relocation*, on page 2-14.

7.7.1 Default Memory Model

If you do not use the MEMORY directive, the linker uses a default memory model that is based on the TMS320C6000 architecture. This model assumes that the full 32-bit address space (2^{32} locations) is present in the system and available for use. For more information about the default memory model, see section 7.12, *Default Allocation Algorithm*, on page 7-54.

7.7.2 MEMORY Directive Syntax

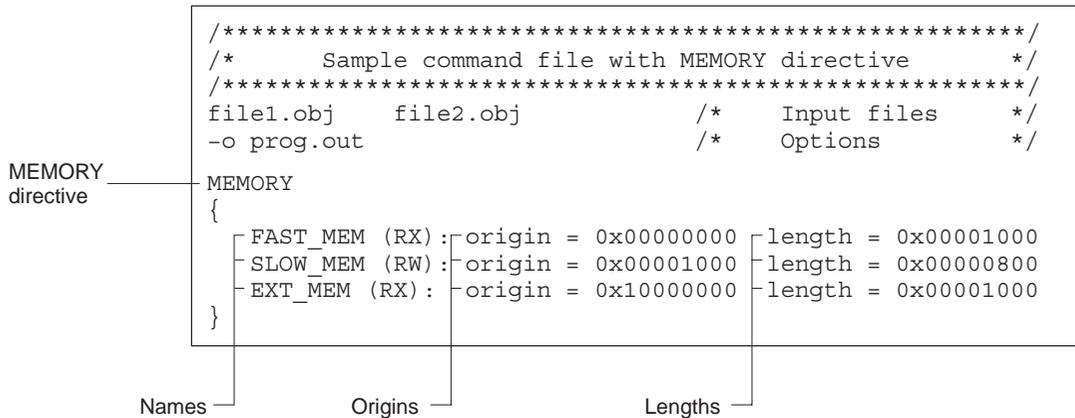
The MEMORY directive identifies ranges of memory that are physically present in the target system and can be used by a program. Each range has several characteristics:

- Name
- Starting address
- Length
- Optional set of attributes
- Optional fill specification

When you use the MEMORY directive, be sure to identify all memory ranges that are available for loading code. Memory defined by the MEMORY directive is configured; any memory that you do not explicitly account for with MEMORY is unconfigured. The linker does not place any part of a program into unconfigured memory. You can represent nonexistent memory spaces by simply not including an address range in a MEMORY directive statement.

The MEMORY directive is specified in a command file by the word MEMORY (uppercase), followed by a list of memory range specifications enclosed in braces. The MEMORY directive in Example 7–3 defines a system that has 4K bytes of fast external memory at address 0x0000 0000, 2K bytes of slow external memory at address 0x0000 1000 and 4K bytes of slow external memory at address 0x1000 0000.

Example 7–3. The MEMORY Directive



The general syntax for the MEMORY directive is:

```

MEMORY
{
  name 1 [(attr)] : origin = constant, length = constant [, fill = constant]
  .
  .
  name n [(attr)] : origin = constant, length = constant [, fill = constant]
}

```

name names a memory range. A memory name can be one to 64 characters; valid characters include A–Z, a–z, \$, ., and _. The names have no special significance to the linker; they simply identify memory ranges. Memory range names are internal to the linker and are not retained in the output file or in the symbol table. All memory ranges must have unique names and must not overlap.

<i>attr</i>	<p>specifies one to four attributes associated with the named range. Attributes are optional; when used, they must be enclosed in parentheses. Attributes restrict the allocation of output sections into certain memory ranges. If you do not use any attributes, you can allocate any output section into any range with no restrictions. Any memory for which no attributes are specified (including all memory in the default model) has all four attributes. Valid attributes are:</p> <p>R specifies that the memory can be read. W specifies that the memory can be written to. X specifies that the memory can contain executable code. I specifies that the memory can be initialized.</p>
origin	specifies the starting address of a memory range; enter as <i>origin</i> , <i>org</i> , or <i>o</i> . The value, specified in bytes, is a 32-bit constant and can be decimal, octal, or hexadecimal.
length	specifies the length of a memory range; enter as <i>length</i> , <i>len</i> , or <i>l</i> . The value, specified in bytes, is a 32-bit constant and can be decimal, octal, or hexadecimal.
fill	specifies a fill character for the memory range; enter as <i>fill</i> or <i>f</i> . Fills are optional. The value is a 32-bit integer constant and can be decimal, octal, or hexadecimal. The fill value is used to fill areas of the memory range that are not allocated to a section.

Note: Filling Memory Ranges

If you specify fill values for large memory ranges, your output file will be very large because filling a memory range (even with 0s) causes raw data to be generated for all unallocated blocks of memory in the range.

The following example specifies a memory range with the R and W attributes and a fill constant of 0FFFFFFFh:

```
MEMORY
{
    RFILE (RW) : o = 0x0020h, l = 0x1000, f = 0xFFFFFFFFh
}
```

You normally use the MEMORY directive in conjunction with the SECTIONS directive to control allocation of output sections. After you use MEMORY to specify the target system's memory model, you can use SECTIONS to allocate output sections into specific named memory ranges or into memory that has specific attributes. For example, you could allocate the .text and .data sections into the area named FAST_MEM and allocate the .bss section into the area named SLOW_MEM.

7.8 The SECTIONS Directive

The SECTIONS directive:

- Describes how input sections are combined into output sections
- Defines output sections in the executable program
- Specifies where output sections are placed in memory (in relation to each other and to the entire memory space)
- Permits renaming of output sections

For more information, see section 2.3, *How the Linker Handles Sections*, on page 2-11; section 2.4, *Relocation*, on page 2-14; and section 2.2.4, *Subsections*, on page 2-7. Subsections allow you to manipulate sections with greater precision.

If you do not specify a SECTIONS directive, the linker uses a default algorithm for combining and allocating the sections. Section 7.12, *Default Allocation Algorithm*, on page 7-54 describes this algorithm in detail.

7.8.1 SECTIONS Directive Syntax

The SECTIONS directive is specified in a command file by the word SECTIONS (uppercase), followed by a list of output section specifications enclosed in braces.

The general syntax for the SECTIONS directive is:

```
SECTIONS
{
    name : [property [, property] [, property] . . . ]
    name : [property [, property] [, property] . . . ]
    name : [property [, property] [, property] . . . ]
}
```

Each section specification, beginning with *name*, defines an output section. (An output section is a section in the output file.) A section name can be a subsection specification. After the section name is a list of properties that define the section's contents and how the section is allocated. The properties can be separated by optional commas. Possible properties for a section are as follows:

- Load allocation** defines where in memory the section is to be loaded.

Syntax: **load =** *allocation* or
allocation or
> *allocation*

- Run allocation** defines where in memory the section is to be run.

Syntax: **run =** *allocation* or
run > *allocation*

- Input sections** defines the input sections (object files) that constitute the output section.

Syntax: **{** *input_sections* **}**

- Section type** defines flags for special section types.

Syntax: **type = COPY** or
type = DSECT or
type = NOLOAD

For more information, see section 7.11, *Special Section Types (DSECT, COPY, and NOLOAD)*, on page 7-53.

- Fill value** defines the value used to fill uninitialized holes.

Syntax: **fill =** *value* or
name : [*properties*] = *value*

For more information, see section 7.14, *Creating and Filling Holes*, on page 7-64.

Example 7–4 shows a SECTIONS directive in a sample linker command file.

Example 7–4. The SECTIONS Directive

```

/*****
/* Sample command file with SECTIONS directive */
*****/
file1.obj  file2.obj      /* Input files */
-o prog.out                /* Options   */

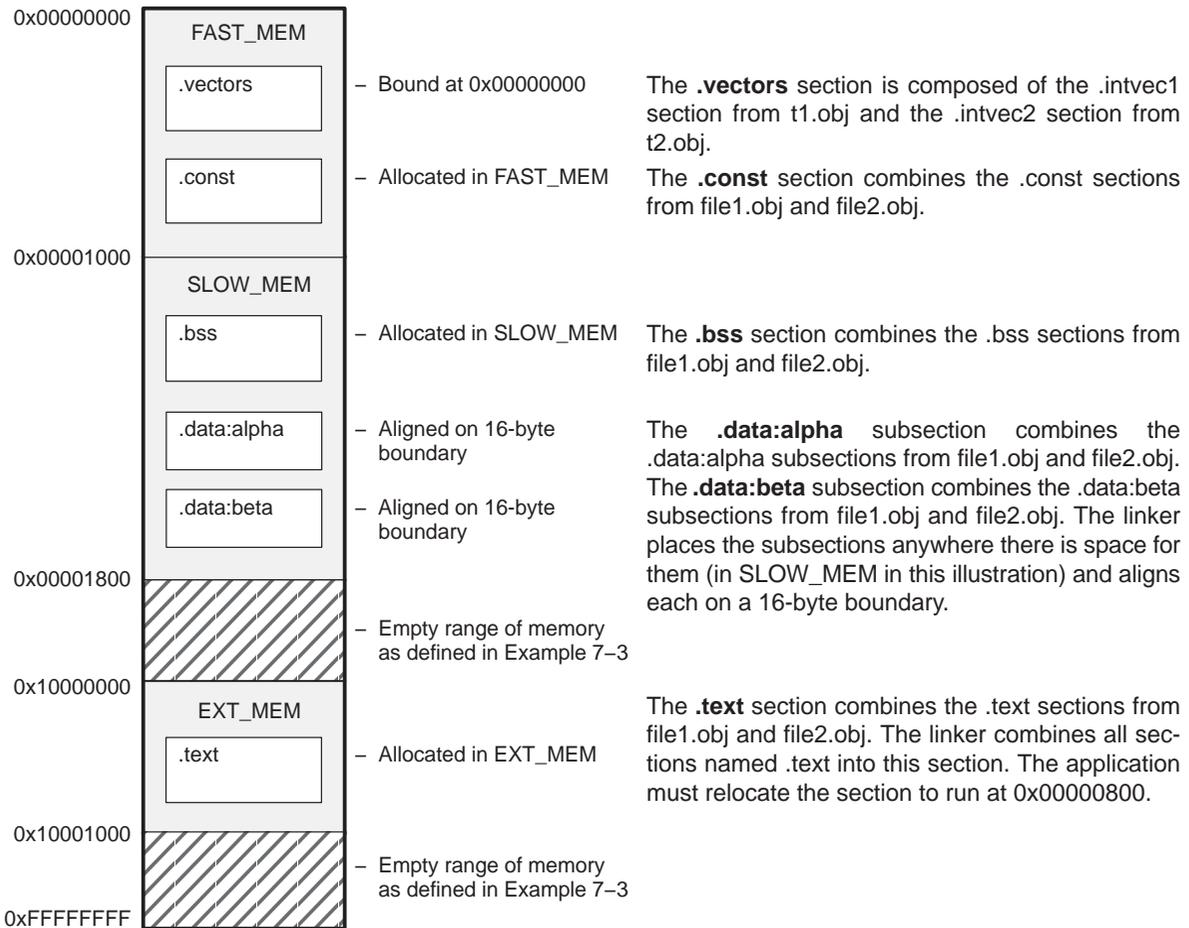
SECTIONS directive ----- SECTIONS
{
  .text:          load = EXT_MEM, run = 0x00000800
  .const:         load = FAST_MEM
  .bss:           load = SLOW_MEM
  .vectors:       load = 0x00000000
                  {
                    t1.obj(.intvec1)
                    t2.obj(.intvec2)
                    endvec = .;
                  }
  .data:alpha:    align = 16
  .data:beta:     align = 16
}

```

Section specifications -----

Figure 7–2 shows the six output sections defined by the SECTIONS directive in Example 7–4 (.vectors, .text, .const, .bss, .data:alpha, and .data:beta) and shows how these sections are allocated in memory.

Figure 7–2. Section Allocation Defined by Example 7–4



7.8.2 Allocation

The linker assigns each output section two locations in target memory: the location where the section will be loaded and the location where it will be run. Usually, these are the same, and you can think of each section as having only a single address. The process of locating the output section in the target’s memory and assigning its address(es) is called allocation. For more information about using separate load and run allocation, see section 7.9, *Specifying a Section’s Run-Time Address*, on page 7-44.

If you do not tell the linker how a section is to be allocated, it uses a default algorithm to allocate the section. Generally, the linker puts sections wherever they fit into configured memory. You can override this default allocation for a section by defining it within a SECTIONS directive and providing instructions on how to allocate it.

You control allocation by specifying one or more allocation parameters. Each parameter consists of a keyword, an optional equal sign or greater-than sign, and a value optionally enclosed in parentheses. If load and run allocation are separate, all parameters following the keyword LOAD apply to load allocation, and those following the keyword RUN apply to run allocation. The allocation parameters are:

Binding allocates a section at a specific address.

```
.text: load = 0x1000
```

Named memory allocates the section into a range defined in the MEMORY directive with the specified name (like SLOW_MEM) or attributes.

```
.text: load > SLOW_MEM
```

Alignment uses the align keyword to specify that the section must start on an address boundary.

```
.text: align = 0x100
```

Blocking uses the block keyword to specify that the section must fit between two address boundaries: if the section is too big, it starts on an address boundary.

```
.text: block(0x100)
```

For the load (usually the only) allocation, you can simply use a greater-than sign and omit the load keyword:

```
.text: > SLOW_MEM .text: {...} > SLOW_MEM
.text: > 0x4000
```

If more than one parameter is used, you can string them together as follows:

```
.text: > SLOW_MEM align 16
```

Or if you prefer, use parentheses for readability:

```
.text: load = (SLOW_MEM align(16))
```

You can also use an input section specification to identify the sections from input files that are combined to form an output section. For more information, see section 7.8.3, *Specifying Input Sections*, on page 7-37.

7.8.2.1 Binding

You can supply a specific starting address for an output section by following the section name with an address:

```
.text: 0x00001000
```

This example specifies that the `.text` section must begin at location `0x1000`. The binding address must be a 32-bit constant.

Output sections can be bound anywhere in configured memory (assuming there is enough space), but they cannot overlap. If there is not enough space to bind a section to a specified address, the linker issues an error message.

Note: Binding is Incompatible With Alignment and Named Memory

You cannot bind a section to an address if you use alignment or named memory. If you try to do this, the linker issues an error message.

7.8.2.2 Named Memory

You can allocate a section into a memory range that is defined by the `MEMORY` directive (see section 7.7, *The MEMORY Directive*, on page 7-27). This example names ranges and links sections into them:

```
MEMORY
{
    SLOW_MEM (RIX) : origin = 0x00000000, length = 0x00001000
    FAST_MEM (RWIX) : origin = 0x30000000, length = 0x00000300
}

SECTIONS
{
    .text : > SLOW_MEM
    .data : > FAST_MEM ALIGN(128)
    .bss : > FAST_MEM
}
```

In this example, the linker places `.text` into the area called `SLOW_MEM`. The `.data` and `.bss` output sections are allocated into `FAST_MEM`. You can align a section within a named memory range; the `.data` section is aligned on a 128-byte boundary within the `FAST_MEM` range.

Similarly, you can link a section into an area of memory that has particular attributes. To do this, specify a set of attributes (enclosed in parentheses) instead of a memory name. Using the same `MEMORY` directive declaration, you can specify:

```
SECTIONS
{
    .text: > (X) /* .text --> executable memory */
    .data: > (RI) /* .data --> read or init memory */
    .bss : > (RW) /* .bss --> read or write memory */
}
```

In this example, the `.text` output section can be linked into either the `SLOW_MEM` or `FAST_MEM` area because both areas have the `X` attribute. The `.data` section can also go into either `SLOW_MEM` or `FAST_MEM` because both areas have the `R` and `I` attributes. The `.bss` output section, however, must go into the `FAST_MEM` area because only `FAST_MEM` is declared with the `W` attribute.

You cannot control where in a named memory range a section is allocated, although the linker uses lower memory addresses first and avoids fragmentation when possible. In the preceding examples, assuming no conflicting assignments exist, the `.text` section starts at address 0. If a section must start on a specific address, use binding instead of named memory.

7.8.2.3 Alignment and Blocking

You can tell the linker to place an output section at an address that falls on an `n`-byte boundary, where `n` is a power of 2, by using the `align` keyword. For example:

```
.text: load = align(32)
```

allocates `.text` so that it falls on a 32-byte boundary.

Blocking is a weaker form of alignment that allocates a section anywhere *within* a block of size `n`. The specified block size must be a power of 2. For example:

```
bss: load = block(0x0080)
```

allocates `.bss` so that the entire section is contained in a single 128-byte page or begins on that boundary.

You can use alignment or blocking alone or in conjunction with a memory area, but alignment and blocking cannot be used together.

7.8.3 Specifying Input Sections

An input section specification identifies the sections from input files that are combined to form an output section. In general, the linker combines input sections by concatenating them in the order in which they are specified. However, if alignment or blocking is specified for an input section, all of the input sections within the output section are ordered as follows:

- All aligned sections, from largest to smallest
- All blocked sections, from largest to smallest
- All other sections, from largest to smallest

The size of an output section is the sum of the sizes of the input sections that it comprises.

Example 7–5 shows the most common type of section specification; note that no input sections are listed.

Example 7–5. The Most Common Method of Specifying Section Contents

```
SECTIONS
{
    .text:
    .data:
    .bss:
}
```

In Example 7–5, the linker takes all the `.text` sections from the input files and combines them into the `.text` output section. The linker concatenates the `.text` input sections in the order that it encounters them in the input files. The linker performs similar operations with the `.data` and `.bss` sections. You can use this type of specification for any output section.

You can explicitly specify the input sections that form an output section. Each input section is identified by its filename and section name:

```
SECTIONS
{
    .text :                /* Build .text output section      */
    {
        f1.obj(.text)     /* Link .text section from f1.obj   */
        f2.obj(sec1)     /* Link sec1 section from f2.obj    */
        f3.obj           /* Link ALL sections from f3.obj    */
        f4.obj(.text,sec2) /* Link .text and sec2 from f4.obj  */
    }
}
```

It is not necessary for input sections to have the same name as each other or as the output section they become part of. If a file is listed with no sections, *all* of its sections are included in the output section. If any additional input sections have the same name as an output section but are not explicitly specified by the SECTIONS directive, they are automatically linked in at the end of the output section. For example, if the linker found more `.text` sections in the preceding example and these `.text` sections *were not* specified anywhere in the SECTIONS directive, the linker would concatenate these extra sections after `f4.obj(sec2)`.

The specifications in Example 7–5 are actually a shorthand method for the following:

```
SECTIONS
{
    .text: { *(.text) }
    .data: { *(.data) }
    .bss:  { *(.bss) }
}
```

The specification `*(.text)` means *the unallocated .text sections from all the input files*. This format is useful when:

- ❑ You want the output section to contain all input sections that have a specified name, but the output section name is different from the input sections' name.
- ❑ You want the linker to allocate the input sections before it processes additional input sections or commands within the braces.

The following example illustrates the two purposes above:

```
SECTIONS
{
  .text : {
          abc.obj(xqt)
          *(.text)
        }
  .data : {
          *(.data)
          fil.obj(table)
        }
}
```

In this example, the `.text` output section contains a named section `xqt` from file `abc.obj`, which is followed by all the `.text` input sections. The `.data` section contains all the `.data` input sections, followed by a named section `table` from the file `fil.obj`. This method includes all the unallocated sections. For example, if one of the `.text` input sections was already included in another output section when the linker encountered `*(.text)`, the linker could not include that first `.text` input section in the second output section.

7.8.4 Allocation Using Multiple Memory Ranges

The linker allows you to specify an explicit list of memory ranges into which an output section can be allocated. Consider the following example:

```
MEMORY
{
    P_MEM1 : origin = 02000h, length = 01000h
    P_MEM2 : origin = 04000h, length = 01000h
    P_MEM3 : origin = 06000h, length = 01000h
    P_MEM4 : origin = 08000h, length = 01000h
}

SECTIONS
{
    .text : { } > P_MEM1 | P_MEM2 | P_MEM4
}
```

The | operator is used to specify the multiple memory ranges. The .text output section is allocated as a whole into the first memory range in which it fits. The memory ranges are accessed in the order specified. In this example, the linker first tries to allocate the section in P_MEM1. If that attempt fails, the linker tries to place the section into P_MEM2, and so on. If the output section is not successfully allocated in any of the named memory ranges, the linker issues an error message.

With this type of SECTIONS directive specification, the linker can seamlessly handle an output section that grows beyond the available space of the memory range in which it is originally allocated. Instead of modifying the linker command file, you can let the linker move the section into one of the other areas.

7.8.5 Automatic Splitting of Output Sections Among Non-Contiguous Memory Ranges

The linker can split output sections among multiple memory ranges to achieve an efficient allocation. Use the >> operator to indicate that an output section can be split, if necessary, into the specified memory ranges. For example:

```
MEMORY
{
    P_MEM1 : origin = 02000h, length = 01000h
    P_MEM2 : origin = 04000h, length = 01000h
    P_MEM3 : origin = 06000h, length = 01000h
    P_MEM4 : origin = 08000h, length = 01000h
}

SECTIONS
{
    .text: { *(.text) } >> P_MEM1 | P_MEM2 | P_MEM3 | P_MEM4
}
```

In this example, the >> operator indicates that the .text output section can be split among any of the listed memory areas. If the .text section grows beyond the available memory in P_MEM1, it is split on an input section boundary, and the remainder of the output section is allocated to P_MEM2 | P_MEM3 | P_MEM4.

The | operator is used to specify the list of multiple memory ranges.

You can also use the >> operator to indicate that an output section can be split within a single memory range. This functionality is useful when several output sections must be allocated into the same memory range, but the restrictions of one output section cause the memory range to be partitioned. Consider the following example:

```
MEMORY
{
    RAM : origin = 01000h, length = 08000h
}

SECTIONS
{
    .special: { f1.obj(.text) } = 04000h
    .text: { *(.text) } >> RAM
}
```

The .special output section is allocated near the middle of the RAM memory range. This leaves two unused areas in RAM: from 01000h to 04000h, and from the end of f1.obj(.text) to 08000h. The specification for the .text section allows the linker to split the .text section around the .special section and use the available space in RAM on either side of .special.

The >> operator can also be used to split an output section among all memory ranges that match a specified attribute combination. For example:

```
MEMORY
{
  P_MEM1 (RWX) : origin = 01000h,  length = 02000h
  P_MEM2 (RWI) : origin = 04000h,  length = 01000h
}

SECTIONS
{
  .text: { *(.text) } >> (RW)
}
```

The linker attempts to allocate all or part of the output section into any memory range whose attributes match the attributes specified in the SECTIONS directive.

This SECTIONS directive has the same effect as:

```
SECTIONS
{
  .text: { *(.text) } >> P_MEM1 | P_MEM2
}
```

Certain output sections should not be split:

- The .init section, which contains the autoinitialization table for C/C++ programs
- The .pinit section, which contains the list of global constructors for C++ programs
- An output section with separate load and run allocations. The code that copies the output section from its load-time allocation to its run-time location cannot accommodate a split in the output section.
- An output section with an input section specification that includes an expression to be evaluated. The expression may define a symbol that is used in the program to manage the output section at run time.

If you use the >> operator on any of these sections, the linker issues a warning and ignores the operator.

7.8.6 Allocating an Archive Member to an Output Section

The ability to specify an archive member of a library archive for allocation into a specific output section can be specified inside angle brackets after a library name. Any object files separated by commas or spaces from the specified archive file are legal within the angle brackets. The syntax for allocating archived library members specifically inside of a SECTIONS directive is as follows:

```
[-l] library name <member1[, member2[, ...]> [(input sections)]
```

```
SECTIONS
{
    boot >      BOOT1
    {
        -lrtsXX.lib<boot.obj> (.text)
        -lrtsXX.lib<exit.obj strcpy.obj> (.text)
    }

    .rts >      BOOT2
    {
        -lrtsXX.lib (.text)
    }

    .text >     RAM
    {
        * (.text)
    }
}
```

The above example specifies that the text sections of `boot.obj`, `exit.obj`, and `strcpy.obj` from the run-time-support library should be placed in section `.boot`. The remainder of the `.text` sections from the run-time-support library are to be placed in section `.rts`. Finally, the remainder of all other `.text` sections are to be placed in section `.text`.

The `-l` option (which normally implies a library path search be made for the named file following the option) listed before each library is optional when listing specific archive members inside `< >`. Using `< >` implies that you are referring to a library.

7.9 Specifying a Section's Run-Time Address

At times, you may want to load code into one area of memory and run it in another. For example, you may have performance-critical code in slow external memory. The code must be loaded into slow external memory, but it would run faster in fast external memory.

The linker provides a simple way to accomplish this. You can use the `SECTIONS` directive to direct the linker to allocate a section twice: once to set its load address and again to set its run address. For example:

```
.fir: load = SLOW_MEM, run = FAST_MEM
```

Use the *load* keyword for the load address and the *run* keyword for the run address.

See section 2.5, *Run-Time Relocation*, on page 2-16, for an overview on run-time relocation.

7.9.1 Specifying Load and Run Addresses

The load address determines where a loader places the raw data for the section. Any references to the section (such as labels in it) refer to its run address. The application must copy the section from its load address to its run address; this does *not* happen automatically when you specify a separate run address.

If you provide only one allocation (either load or run) for a section, the section is allocated only once and loads and runs at the same address. If you provide both allocations, the section is allocated as if it were two sections of the same size. This means that both allocations occupy space in the memory map and cannot overlay each other or other sections. (The `UNION` directive provides a way to overlay sections; see section 7.10.1, *Overlaying Sections With the UNION Statement*, on page 7-48.)

If either the load or run address has additional parameters, such as alignment or blocking, list them after the appropriate keyword. Everything related to allocation after the keyword *load* affects the load address until the keyword *run* is seen, after which, everything affects the run address. The load and run allocations are completely independent, so any qualification of one (such as alignment) has no effect on the other. You can also specify run first, then load. Use parentheses to improve readability.

The examples below specify load and run addresses:

```
.data: load = SLOW_MEM, align = 32, run = FAST_MEM
```

(align applies only to load)

```
.data: load = (SLOW_MEM align 32), run = FAST_MEM
```

(identical to previous example)

```
.data: run = FAST_MEM, align 32,
      load = align 16
```

(align 32 in FAST_MEM for run; align 16 anywhere for load)

7.9.2 Uninitialized Sections

Uninitialized sections (such as `.bss`) are not loaded, so their only significant address is the run address. The linker allocates uninitialized sections only once: if you specify both run and load addresses, the linker warns you and ignores the load address. Otherwise, if you specify only one address, the linker treats it as a run address, regardless of whether you call it load or run. This example specifies load and run addresses for an uninitialized section:

```
.bss: load = 0x1000, run = FAST_MEM
```

A warning is issued, load is ignored, and space is allocated in `FAST_MEM`. All of the following examples have the same effect. The `.bss` section is allocated in `FAST_MEM`.

```
.bss: load = FAST_MEM
.bss: run = FAST_MEM
.bss: > FAST_MEM
```

7.9.3 Referring to the Load Address by Using the `.label` Directive

Normally, any reference to a symbol in a section refers to its run-time address. However, it may be necessary at run time to refer to a load-time address. Specifically, the code that copies a section from its load address to its run address must have access to the load address. The `.label` directive defines a special symbol that refers to the section's load address. Thus, whereas normal symbols are relocated with respect to the run address, `.label` symbols are relocated with respect to the load address. For more information on the `.label` directive, see page 4-50.

Example 7-6 shows the use of the `.label` directive. Figure 7-3 illustrates the run-time execution of Example 7-6.

Example 7–6. Copying a Section From SLOW_MEM to FAST_MEM*(a) Assembly language file*

```

        .sect  ".fir"
        .align 4
        .label fir_src
fir
        ;<code here
        .label fir_end

        .text
MVKL   fir_src, A4
MVKH   fir_src, A4
MVKL   fir_end, A5
MVKH   fir_end, A5
MVKL   fir, A6
MVKH   fir, A6
SUB    A5, A4, A1

loop:
[!A1]  B      done
LDW    *A4++, B3
NOP    4
        ; branch occurs
STW    B3, *A6++
SUB    A1, 4, A1
B      loop
NOP    5
        ; branch occurs

done:
B      fir
NOP    5
        ; call occurs

```

(b) Linker command file

```

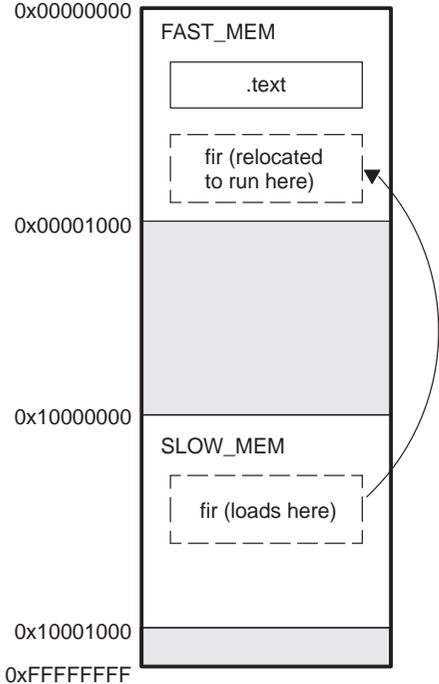
/*****
/*      PARTIAL LINKER COMMAND FILE FOR FIR EXAMPLE      */
*****/

MEMORY
{
    FAST_MEM :  origin = 0x00001000, length = 0x00001000
    SLOW_MEM  :  origin = 0x10000000, length = 0x00001000
}

SECTIONS
{
    .text: load = FAST_MEM
    .fir:  load = SLOW_MEM, run FAST_MEM
}

```

Figure 7-3. Run-Time Execution of Example 7-6



7.10 Using UNION and GROUP Statements

Two SECTIONS statements allow you to conserve memory: GROUP and UNION. Unioning sections causes the linker to allocate them to the same run address. Grouping sections causes the linker to allocate them contiguously in memory. Section names can refer to sections, subsections, or archive library members.

7.10.1 Overlaying Sections With the UNION Statement

For some applications, you may want to allocate more than one section to run at the same address. For example, you may have several routines you want in fast external memory at various stages of execution. Or you may want several data objects that are not active at the same time to share a block of memory. The UNION statement within the SECTIONS directive provides a way to allocate several sections at the same run-time address.

In Example 7–7, the .bss sections from file1.obj and file2.obj are allocated at the same address in FAST_MEM. In the memory map, the union occupies as much space as its largest component. The components of a union remain independent sections; they are simply allocated together as a unit.

Example 7–7. The UNION Statement

```

SECTIONS
{
    .text: load = SLOW_MEM
    UNION: run = FAST_MEM
    {
        .bss:part1: { file1.obj(.bss) }
        .bss:part2: { file2.obj(.bss) }
    }
    .bss:part3: run = FAST_MEM { globals.obj(.bss) }
}

```

Allocation of a section as part of a union affects only its *run* address. Under no circumstances can sections be overlaid for loading. If an initialized section is a union member (an initialized section, such as .text, has raw data), its load allocation *must* be separately specified. See Example 7–8.

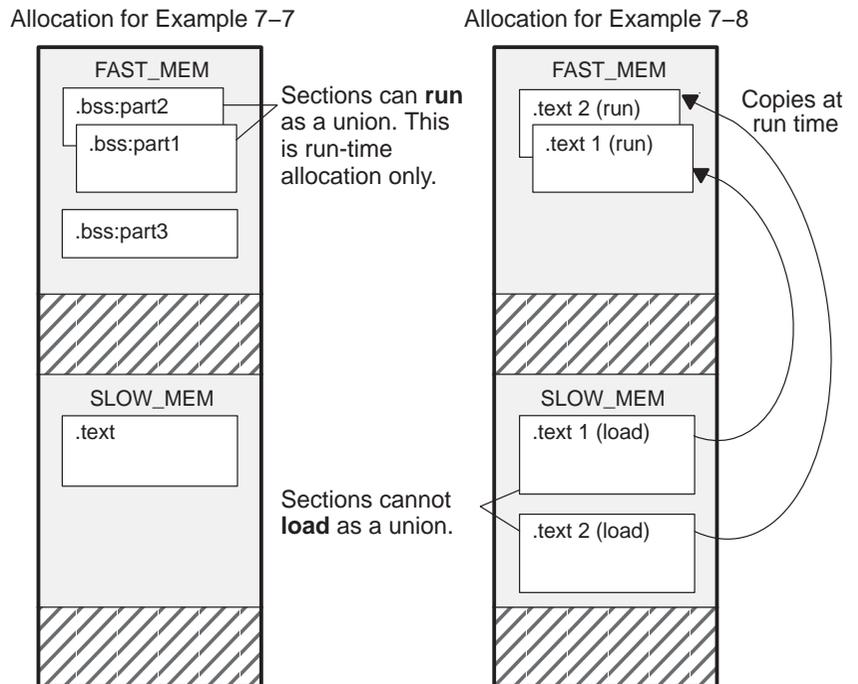
Example 7–8. Separate Load Addresses for UNION Sections

```

UNION run = FAST_MEM
{
    .text:part1: load = SLOW_MEM, { file1.obj(.text) }
    .text:part2: load = SLOW_MEM, { file2.obj(.text) }
}

```

Figure 7-4. Memory Allocation Shown in Example 7-7 and Example 7-8



Since the `.text` sections contain data, they cannot *load* as a union, although they can be *run* as a union. Therefore, each requires its own load address. If you fail to provide a load allocation for an initialized section within a UNION, the linker issues a warning and allocates load space anywhere it can in configured memory.

Uninitialized sections are not loaded and do not require load addresses.

The UNION statement applies only to allocation of run addresses, so it is meaningless to specify a load address for the union itself. For purposes of allocation, the union is treated as an uninitialized section: any one allocation specified is considered a run address, and if both run and load addresses are specified, the linker issues a warning and ignores the load address.

7.10.2 Grouping Output Sections Together

The SECTIONS directive's GROUP option forces several output sections to be allocated contiguously. For example, assume that a section named term_rec contains a termination record for a table in the .data section. You can force the linker to allocate .data and term_rec together:

Example 7–9. Allocate Sections Together

```
SECTIONS
{
    .text          /* Normal output section          */
    .bss          /* Normal output section          */
    GROUP 0x00001000 : /* Specify a group of sections    */
    {
        .data      /* First section in the group      */
        term_rec   /* Allocated immediately after .data */
    }
}
```

You can use binding, alignment, or named memory to allocate a GROUP in the same manner as a single output section. In the preceding example, the GROUP is bound to address 0x00001000. This means that .data is allocated at 0x00001000, and term_rec follows it in memory.

7.10.3 Nesting UNIONS and GROUPS

The linker allows arbitrary nesting of GROUP and UNION statements with the SECTIONS directive. By nesting GROUP and UNION statements, you can express hierarchical overlays and groupings of sections. Example 7–10 shows how two overlays can be grouped together.

Example 7–10. Nesting GROUP and UNION Statements

```
SECTIONS
{
    GROUP 1000h : run = FAST_MEM
    {
        UNION:
        {
            mysect1:load = SLOW_MEM
            mysect2: load = SLOW_MEM
        }
        UNION:
        {
            mysect3: load = SLOW_MEM
            mysect4: load = SLOW_MEM
        }
    }
}
```

For this example, the linker performs the following allocations:

- The four sections (mysect1, mysect2, mysect3, mysect4) are assigned unique, non-overlapping load addresses in the SLOW_MEM memory region. This assignment is determined by the particular load allocations given for each section.
- Sections mysect1 and mysect2 are assigned the same run address in FAST_MEM.
- Sections mysect3 and mysect4 are assigned the same run address in FAST_MEM.
- The run addresses of mysect1/mysect2 and mysect3/mysect4 are allocated contiguously, as directed by the GROUP statement (subject to alignment and blocking restrictions).

To refer to groups and unions, linker diagnostic messages use the notation:

GROUP_*n*

UNION_*n*

In this notation, *n* is a sequential number (beginning at 1) that represents the lexical ordering of the group or union in the linker control file, without regard to nesting. Groups and unions each have their own counter.

7.10.4 Checking the Consistency of Allocators

The linker checks the consistency of load and run allocations specified for unions, groups, and sections. The following rules are used:

- Run allocations are only allowed for top-level sections, groups, or unions (sections, groups, or unions that are not nested under any other groups or unions). The linker uses the run address of the top-level structure to compute the run addresses of the components within groups and unions.
- The linker does not accept a load allocation for UNIONS.
- The linker does not accept a load allocation for uninitialized sections.
- In most cases, you must provide a load allocation for an initialized section. However, the linker does not accept a load allocation for an initialized section that is located within a group that already defines a load allocator.
- As a shortcut, you can specify a load allocation for an entire group, to determine the load allocations for every initialized section or subgroup nested within the group. However, a load allocation is accepted for an entire group only if all of the following conditions are true:
 - The group is initialized (i.e., it has at least one initialized member).
 - The group is not nested inside another group that has a load allocator.
 - The group does not contain a union containing initialized sections.

If the group contains a union with initialized sections, it is necessary to specify the load allocation for each initialized section nested within the group. Consider the following example:

```
SECTIONS
{
  GROUP: load = SLOW_MEM, run = SLOW_MEM
  {
    .text1:
    UNION:
    {
      .text2:
      .text3:
    }
  }
}
```

The load allocator given for the group does not uniquely specify the load allocation for the elements within the union: `.text2` and `.text3`. In this case, the linker issues a diagnostic message to request that these load allocations be specified explicitly.

7.11 Special Section Types (DSECT, COPY, and NOLOAD)

You can assign three special types to output sections: DSECT, COPY, and NOLOAD. These types affect the way that the program is treated when it is linked and loaded. You can assign a type to a section by placing the type after the section definition. For example:

```
SECTIONS
{
  sec1: load = 0x00002000, type =    DSECT   {f1.obj}
  sec2: load = 0x00004000, type =    COPY   {f2.obj}
  sec3: load = 0x00006000, type =    NOLOAD  {f3.obj}
}
```

- The DSECT type creates a dummy section with the following characteristics:
 - It is not included in the output section memory allocation. It takes up no memory and is not included in the memory map listing.
 - It can overlay other output sections, other DSECTs, and unconfigured memory.
 - Global symbols defined in a dummy section are relocated normally. They appear in the output module's symbol table with the same value they would have if the DSECT had actually been loaded. These symbols can be referenced by other input sections.
 - Undefined external symbols found in a DSECT cause specified archive libraries to be searched.
 - The section's contents, relocation information, and line number information are not placed in the output module.

In the preceding example, none of the sections from f1.obj are allocated, but all the symbols are relocated as though the sections were linked at address 0x2000. The other sections can refer to any of the global symbols in sec1.

- A COPY section is similar to a DSECT section, except that its contents and associated information are written to the output module. The .cinit section that contains initialization tables for the TMS320C6000 C/C++ compiler has this attribute under the run-time initialization model.
- A NOLOAD section differs from a normal output section in one respect: the section's contents, relocation information, and line number information are not placed in the output module. The linker allocates space for the section, and it appears in the memory map listing.

7.12 Default Allocation Algorithm

The MEMORY and SECTIONS directives provide flexible methods for building, combining, and allocating sections. However, any memory locations or sections that you choose *not* to specify must still be handled by the linker. The linker uses default algorithms to build and allocate sections within the specifications you supply.

If you do not use the MEMORY and SECTIONS directives, the linker allocates output sections as though the definitions in Example 7–11 were specified.

Example 7–11. Default Allocation for TMS320C6000 Devices

```
MEMORY
{
    RAM      : origin = 0x00000001, length = 0xFFFFFFFF
}

SECTIONS
{
    .text   : ALIGN(32) {} > RAM
    .const  : ALIGN(8)  {} > RAM
    .data   : ALIGN(8)  {} > RAM
    .bss    : ALIGN(8)  {} > RAM
    .cinit  : ALIGN(4)  {} > RAM      ; cflag option only
    .pinit  : ALIGN(4)  {} > RAM      ; cflag option only
    .stack  : ALIGN(8)  {} > RAM      ; cflag option only
    .far    : ALIGN(8)  {} > RAM      ; cflag option only
    .systemem: ALIGN(8) {} > RAM      ; cflag option only
    .switch: ALIGN(4)  {} > RAM      ; cflag option only
    .cio    : ALIGN(4)  {} > RAM      ; cflag option only
}
```

All .text input sections are concatenated to form a .text output section in the executable output file, and all .data input sections are combined to form a .data output section.

If you use a SECTIONS directive, the linker performs *no part* of the default allocation. Allocation is performed according to the rules specified by the SECTIONS directive and the general algorithm described next in section 7.12.1.

7.12.1 How the Allocation Algorithm Creates Output Sections

An output section can be formed in one of two ways:

- Method 1** As the result of a SECTIONS directive definition
- Method 2** By combining input sections with the same name into an output section that is not defined in a SECTIONS directive

If an output section is formed as a result of a `SECTIONS` directive, this definition completely determines the section's contents. (See section 7.8, *The SECTIONS Directive*, on page 7-31 for examples of how to define an output section's content.)

If an output section is formed by combining input sections not specified by a `SECTIONS` directive, the linker combines all such input sections that have the same name into an output section with that name. For example, suppose the files `f1.obj` and `f2.obj` both contain named sections called `Vectors` and that the `SECTIONS` directive does not define an output section for them. The linker combines the two `Vectors` sections from the input files into a single output section named `Vectors`, allocates it into memory, and includes it in the output file.

By default, the linker does not display a message when it creates an output section that is not defined in the `SECTIONS` directive. You can use the `-w` linker option (see section 7.4.19, *Display a Message When an Undefined Output Section Is Created (-w Option)*, on page 7-20) to cause the linker to display a message when it creates a new output section.

After the linker determines the composition of all output sections, it must allocate them into configured memory. The `MEMORY` directive specifies which portions of memory are configured. If there is no `MEMORY` directive, the linker uses the default configuration as shown in Example 7-11. (See section 7.7, *The MEMORY Directive*, on page 7-27 for more information on configuring memory.)

7.12.2 Reducing Memory Fragmentation

The linker's allocation algorithm attempts to minimize memory fragmentation. This allows memory to be used more efficiently and increases the probability that your program will fit into memory. The algorithm comprises these steps:

- 1) Each output section for which you have supplied a specific binding address is placed in memory at that address.
- 2) Each output section that is included in a specific, named memory range or that has memory attribute restrictions is allocated. Each output section is placed into the first available space within the named area, considering alignment where necessary.
- 3) Any remaining sections are allocated in the order in which they are defined. Sections not defined in a `SECTIONS` directive are allocated in the order in which they are encountered. Each output section is placed into the first available memory space, considering alignment where necessary.

7.13 Assigning Symbols at Link Time

Linker assignment statements allow you to define external (global) symbols and assign values to them at link time. You can use this feature to initialize a variable or pointer to an allocation-dependent value.

7.13.1 Syntax of Assignment Statements

The syntax of assignment statements in the linker is similar to that of assignment statements in the C language:

```
symbol    = expression;  assigns the value of expression to symbol
symbol    += expression;  adds the value of expression to symbol
symbol    -= expression;  subtracts the value of expression from symbol
symbol    *= expression;  multiplies symbol by expression
symbol    /= expression;  divides symbol by expression
```

The symbol should be defined externally. If it is not, the linker defines a new symbol and enters it into the symbol table. The expression must follow the rules defined in section 7.13.3, *Assignment Expressions*. Assignment statements *must* terminate with a semicolon.

The linker processes assignment statements *after* it allocates all the output sections. Therefore, if an expression contains a symbol, the address used for that symbol reflects the symbol's address in the executable output file.

For example, suppose a program reads data from one of two tables identified by two external symbols, Table1 and Table2. The program uses the symbol `cur_tab` as the address of the current table. The `cur_tab` symbol must point to either Table1 or Table2. You could accomplish this in the assembly code, but you would need to reassemble the program to change tables. Instead, you can use a linker assignment statement to assign `cur_tab` at link time:

```
prog.obj          /* Input file */
cur_tab = Table1; /* Assign cur_tab to one of the tables */
```

7.13.2 Assigning the SPC to a Symbol

A special symbol, denoted by a dot (`.`), represents the current value of the section program counter (SPC) during allocation. The SPC keeps track of the current location within a section. The linker's `.` symbol is analogous to the assembler's `$` symbol. The `.` symbol can be used only in assignment statements within a `SECTIONS` directive because `.` is meaningful only during allocation and `SECTIONS` controls the allocation process. (See section 7.8, *The SECTIONS Directive*, on page 7-31.)

The `.` symbol refers to the current run address, not the current load address, of the section.

For example, suppose a program needs to know the address of the beginning of the `.data` section. By using the `.global` directive (see page 4-44), you can create an external undefined variable called `Dstart` in the program. Then, assign the value of `.` to `Dstart`:

```
SECTIONS
{
    .text:    {}
    .data:    { Dstart = .; }
    .bss :    {}
}
```

This defines `Dstart` to be the first linked address of the `.data` section. (`Dstart` is assigned *before* `.data` is allocated.) The linker relocates all references to `Dstart`.

A special type of assignment assigns a value to the `.` symbol. This adjusts the SPC within an output section and creates a hole between two input sections. Any value assigned to `.` to create a hole is relative to the beginning of the section, not to the address actually represented by the `.` symbol. Holes and assignments to `.` are described in section 7.14, *Creating and Filling Holes*, on page 7-64.

7.13.3 Assignment Expressions

These rules apply to linker expressions:

- Expressions can contain global symbols, constants, and the C language operators listed in Table 7-2.
- All numbers are treated as long (32-bit) integers.
- Constants are identified by the linker in the same way as by the assembler. That is, numbers are recognized as decimal unless they have a suffix (`H` or `h` for hexadecimal and `Q` or `q` for octal). C language prefixes are also recognized (`0` for octal and `0x` for hex). Hexadecimal constants must begin with a digit. No binary constants are allowed.

- ❑ Symbols within an expression have only the value of the symbol's *address*. No type-checking is performed.
- ❑ Linker expressions can be absolute or relocatable. If an expression contains *any* relocatable symbols (and 0 or more constants or absolute symbols), it is relocatable. Otherwise, the expression is absolute. If a symbol is assigned the value of a relocatable expression, it is relocatable; if it is assigned the value of an absolute expression, it is absolute.

The linker supports the C language operators listed in Table 7–2 in order of precedence. Operators in the same group have the same precedence. Besides the operators listed in Table 7–2, the linker also has an align operator that allows a symbol to be aligned on an n-byte boundary within an output section (n is a power of 2). For example, the expression

```
. = align(16);
```

aligns the SPC within the current section on the next 16-byte boundary. Because the align operator is a function of the current SPC, it can be used only in the same context as . —that is, within a SECTIONS directive.

Table 7–2. Groups of Operators Used in Expressions (Precedence)

Group 1 (Highest Precedence)		Group 6	
!	Logical NOT	&	Bitwise AND
~	Bitwise NOT		
-	Negation		
Group 2		Group 7	
*	Multiplication		Bitwise OR
/	Division		
%	Modulus		
Group 3		Group 8	
+	Addition	&&	Logical AND
-	Subtraction		
Group 4		Group 9	
>>	Arithmetic right shift		Logical OR
<<	Arithmetic left shift		
Group 5		Group 10 (Lowest Precedence)	
==	Equal to	=	Assignment
!=	Not equal to	+=	A += B → A = A + B
>	Greater than	-=	A -= B → A = A - B
<	Less than	*=	A *= B → A = A * B
<=	Less than or equal to	/=	A /= B → A = A / B
>=	Greater than or equal to		

7.13.4 Symbols Defined by the Linker

The linker automatically defines several symbols based on which sections are used in your assembly source. A program can use these symbols at run time to determine where a section is linked. Since these symbols are external, they appear in the linker map. Each symbol can be accessed in any assembly language module if it is declared with a `.global` directive (see page 4-44). You must have used the corresponding section in a source module for the symbol to be created. Values are assigned to these symbols as follows:

- .text** is assigned the first address of the `.text` output section. (It marks the *beginning* of executable code.)
- etext** is assigned the first address following the `.text` output section. (It marks the *end* of executable code.)
- .data** is assigned the first address of the `.data` output section. (It marks the *beginning* of initialized data tables.)
- edata** is assigned the first address following the `.data` output section. (It marks the *end* of initialized data tables.)
- .bss** is assigned the first address of the `.bss` output section. (It marks the *beginning* of uninitialized data.)
- end** is assigned the first address following the `.bss` output section. (It marks the *end* of uninitialized data.)

The following symbols are defined only for C/C++ support when the `-c` or `-cr` option is used.

- __STACK_SIZE** is assigned the size of the `.stack` section.
- __SYSTEMEM_SIZE** is assigned the size of the `.systemem` section.

7.13.5 Assigning Exact Start, End, and Size Values of a Section to a Symbol

The code generation tools currently support the ability to load program code in one area of (slow) memory and run it in another (faster) area. This is done by specifying separate load and run addresses for an output section or group in the linker command file. Then execute a sequence of instructions (the copying code in Example 7–6) that moves the program code from its load area to its run area before it is needed.

There are several responsibilities that a programmer must take on when setting up a system with this feature. One of these responsibilities is to determine the size and run-time address of the program code to be moved. The current mechanisms to do this involve use of the `.label` directives in the copying code. A simple example is illustrated Example 7–6.

This method of specifying the size and load address of the program code has limitations. While it works fine for an individual input section that is contained entirely within one source file, this method becomes more complicated if the program code is spread over several source files or if the programmer wants to copy an entire output section from load space to run space.

Another problem with this method is that it does not account for the possibility that the section being moved may have an associated far call trampoline section that needs to be moved with it.

7.13.6 Why the Dot Operator Does Not Always Work

The dot operator (`.`) is used to define symbols at link-time with a particular address inside of an output section. It is interpreted like a PC. Whatever the current offset within the current section is, that is the value associated with the dot. Consider an output section specification within a `SECTIONS` directive:

```
outsect:
{
    s1.obj(.text)
    end_of_s1 = .;
    start_of_s2 = .;
    s2.obj(.text)
    end_of_s2 = .;
}
```

This statement creates three symbols:

- `end_of_s1`—the end address of `.text` in `s1.obj`
- `start_of_s2`—the start address of `.text` in `s2.obj`
- `end_of_s2`—the end address of `.text` in `s2.obj`

Suppose there is padding between `s1.obj` and `s2.obj` that is created as a result of alignment. Then `start_of_s2` is not really the start address of the `.text` section in `s2.obj`, but it is the address before the padding needed to align the `.text` section in `s2.obj`. This is due to the linker's interpretation of the dot operator as the current PC. It is also due to the fact that the dot operator is evaluated independently of the input sections around it.

Another potential problem in the above example is that `end_of_s2` may not account for any padding that was required at the end of the output section. You cannot reliably use `end_of_s2` as the end address of the output section. One way to get around this problem is to create a dummy section immediately after the output section in question. For example:

```
GROUP
{
    outsect:
    {
        start_of_outsect = .;
        ...
    }
    dummy: { size_of_outsect = . - start_of_outsect; }
}
```

7.13.7 Address and Dimension Operators

Six new operators have been added to the linker command file syntax:

LOAD_START(<i>sym</i>) START(<i>sym</i>)	Defines <i>sym</i> with the load-time start address of related allocation unit
LOAD_END(<i>sym</i>) END(<i>sym</i>)	Defines <i>sym</i> with the load-time end address of related allocation unit
LOAD_SIZE(<i>sym</i>) SIZE(<i>sym</i>)	Defines <i>sym</i> with the load-time size of related allocation unit
RUN_START(<i>sym</i>)	Defines <i>sym</i> with the run-time start address of related allocation unit
RUN_END(<i>sym</i>)	Defines <i>sym</i> with the run-time end address of related allocation unit
RUN_SIZE(<i>sym</i>)	Defines <i>sym</i> with the run-time size of related allocation unit

Note: Linker Command File Operator Equivalencies

`LOAD_START()` and `START()` are equivalent, as are `LOAD_END()/END()` and `LOAD_SIZE()/SIZE()`.

The new address and dimension operators can be associated with several different kinds of allocation units, including input items, output sections, GROUPs, and UNIONs. The following sections provide some examples of how the operators can be used in each case.

7.13.7.1 Input Items

Consider an output section specification within a SECTIONS directive:

```
outsect:
{
    s1.obj(.text)
    end_of_s1 = .;
    start_of_s2 = .;
    s2.obj(.text)
    end_of_s2 = .;
}
```

This can be rewritten using the START and END operators as follows:

```
outsect:
{
    s1.obj(.text) { END(end_of_s1) }
    s2.obj(.text) { START(start_of_s2), END(end_of_s2) }
}
```

The values of end_of_s1 and end_of_s2 will be the same as if you had used the dot operator in the original example, but start_of_s2 would be defined after any necessary padding that needs to be added between the two .text sections. Remember that the dot operator would cause start_of_s2 to be defined before any necessary padding is inserted between the two input sections.

The syntax for using these operators in association with input sections calls for braces { } to enclose the operator list. The operators in the list are applied to the input item that occurs immediately before the list.

7.13.7.2 Output Section

The START, END, and SIZE operators can also be associated with an output section. Here is an example:

```
outsect: START(start_of_outsect), SIZE(size_of_outsect)
{
    <list of input items>
}
```

In this case, the SIZE operator defines size_of_outsect to incorporate any padding that is required in the output section to conform to any alignment requirements that are imposed.

The syntax for specifying the operators with an output section do not require braces to enclose the operator list. The operator list is simply included as part of the allocation specification for an output section.

7.13.7.3 GROUPs

Here is another use of the START and SIZE operators in the context of a GROUP specification:

```
GROUP
{
    outsect1: { ... }
    outsect2: { ... }
} load = ROM, run = RAM, START(group_start), SIZE(group_size);
```

This can be useful if the whole GROUP is to be loaded in one location and run in another. The copying code can use group_start and group_size as parameters for where to copy from and how much is to be copied. This makes the use of .label in the source code unnecessary.

7.13.7.4 UNIONS

The RUN_SIZE and LOAD_SIZE operators provide a mechanism to distinguish between the size of a UNION's load space and the size of the space where its constituents are going to be copied before they are run. Here is an example:

```
UNION: run = RAM, LOAD_START(union_load_addr),
      LOAD_SIZE(union_ld_sz), RUN_SIZE(union_run_sz)
{
    .text1: load = ROM, SIZE(text1_size) { f1.obj(.text) }
    .text2: load = ROM, SIZE(text2_size) { f2.obj(.text) }
}
```

Here union_ld_sz is going to be equal to the sum of the sizes of all output sections placed in the union. The union_run_sz value is equivalent to the largest output section in the union. Both of these symbols incorporate any padding due to blocking or alignment requirements.

7.14 Creating and Filling Holes

The linker provides you with the ability to create areas *within output sections* that have nothing linked into them. These areas are called *holes*. In special cases, uninitialized sections can also be treated as holes. This section describes how the linker handles holes and how you can fill holes (and uninitialized sections) with values.

7.14.1 Initialized and Uninitialized Sections

There are two rules to remember about the contents of output sections. An output section contains either:

- Raw data for the *entire* section
- No raw data

A section that has raw data is referred to as *initialized*. This means that the object file contains the actual memory image contents of the section. When the section is loaded, this image is loaded into memory at the section's specified starting address. The `.text` and `.data` sections *always* have raw data if anything was assembled into them. Named sections defined with the `.sect` assembler directive also have raw data.

By default, the `.bss` section (see page 4-26) and sections defined with the `.usect` directive (see page 4-75) have no raw data (they are *uninitialized*). They occupy space in the memory map but have no actual contents. Uninitialized sections typically reserve space in fast external memory for variables. In the object file, an uninitialized section has a normal section header and can have symbols defined in it; no memory image, however, is stored in the section.

7.14.2 Creating Holes

You can create a hole in an initialized output section. A hole is created when you force the linker to leave extra space between input sections within an output section. When such a hole is created, *the linker must supply raw data for the hole*.

Holes can be created only *within* output sections. Space can exist *between* output sections, but such space is not a hole. To fill the space between output sections, see section 7.7.2, *MEMORY Directive Syntax*, on page 7-27.

To create a hole in an output section, you must use a special type of linker assignment statement within an output section definition. The assignment statement modifies the SPC (denoted by `.`) by adding to it, assigning a greater value to it, or aligning it on an address boundary. The operators, expressions, and syntaxes of assignment statements are described in section 7.13, *Assigning Symbols at Link Time*, on page 7-56.

The following example uses assignment statements to create holes in output sections:

```
SECTIONS
{
    outsect:
    {
        file1.obj(.text)
        . += 0x0100      /* Create a hole with size 0x0100 */
        file2.obj(.text)
        . = align(16);  /* Create a hole to align the SPC */
        file3.obj(.text)
    }
}
```

The output section outsect is built as follows:

- 1) The .text section from file1.obj is linked in.
- 2) The linker creates a 256-byte hole.
- 3) The .text section from file2.obj is linked in after the hole.
- 4) The linker creates another hole by aligning the SPC on a 16-byte boundary.
- 5) Finally, the .text section from file3.obj is linked in.

All values assigned to the . symbol within a section refer to the *relative address within the section*. The linker handles assignments to the . symbol as if the section started at address 0 (even if you have specified a binding address). Consider the statement . = align(16) in the example. This statement effectively aligns the file3.obj .text section to start on a 16-byte boundary within outsect. If outsect is ultimately allocated to start on an address that is not aligned, the file3.obj .text section will not be aligned either.

The . symbol refers to the current run address, not the current load address, of the section.

Expressions that decrement the . symbol are illegal. For example, it is invalid to use the -= operator in an assignment to the . symbol. The most common operators used in assignments to the . symbol are += and align.

If an output section contains all input sections of a certain type (such as .text), you can use the following statements to create a hole at the beginning or end of the output section.

```
.text:    { . += 0x0100; }      /* Hole at the beginning */
.data:    {
            *(.data)
            . += 0x0100; }    /* Hole at the end          */
```

Another way to create a hole in an output section is to combine an uninitialized section with an initialized section to form a single output section. *In this case, the linker treats the uninitialized section as a hole and supplies data for it.* The following example illustrates this method:

```
SECTIONS
{
  outsect:
  {
    file1.obj(.text)
    file1.obj(.bss)          /* This becomes a hole */
  }
}
```

Because the .text section has raw data, all of outsect must also contain raw data. Therefore, the uninitialized .bss section becomes a hole.

Uninitialized sections become holes only when they are combined with initialized sections. If several uninitialized sections are linked together, the resulting output section is also uninitialized.

7.14.3 Filling Holes

When a hole exists in an initialized output section, the linker must supply raw data to fill it. The linker fills holes with a 32-bit fill value that is replicated through memory until it fills the hole. The linker determines the fill value as follows:

- 1) If the hole is formed by combining an uninitialized section with an initialized section, you can specify a fill value for the uninitialized section. Follow the section name with an = sign and a 32-bit constant. For example:

```
SECTIONS
{
  outsect:
  {
    file1.obj(.text)
    file2.obj(.bss) = 0xFF00FF00 /* Fill this hole */
  }                               /* with 0xFF00FF00 */
}
```

- 2) You can also specify a fill value for all the holes in an output section by supplying the fill value after the section definition:

```
SECTIONS
{
  outsect:fill = 0xFF00FF00
                               /* Fills holes with 0xFF00FF00 */
  {
    . += 0x0010;                /* This creates a hole */
    file1.obj(.text)
    file1.obj(.bss)             /* This creates another hole */
  }
}
```

- 3) If you do not specify an initialization value for a hole, the linker fills the hole with the value specified with the `-f` option (see section 7.4.7, *Set Default Fill Value (-f fill_value Option)*, on page 7-10). For example, suppose the command file `link.cmd` contains the following `SECTIONS` directive:

```
SECTIONS
{
    .text: { .= 0x0100; } /* Create a 100-word hole */
}
```

Now invoke the linker with the `-f` option:

```
cl6x -z -f 0xFFFFFFFF link.cmd
```

This fills the hole with `0xFFFFFFFF`.

- 4) If you do not invoke the linker with the `-f` option or otherwise specify a fill value, the linker fills holes with 0s.

Whenever a hole is created and filled in an initialized output section, the hole is identified in the link map along with the value the linker uses to fill it.

7.14.4 Explicit Initialization of Uninitialized Sections

You can force the linker to initialize an uninitialized section by specifying an explicit fill value for it in the `SECTIONS` directive. This causes the entire section to have raw data (the fill value). For example:

```
SECTIONS
{
    .bss: fill = 0x12341234 /* Fills .bss with 0x12341234 */
}
```

Note: Filling Sections

Because filling a section (even with 0s) causes raw data to be generated for the entire section in the output file, your output file will be very large if you specify fill values for large sections or holes.

7.15 Linker-Generated Copy Tables

The linker supports extensions to the linker command file syntax that enable the following:

- Make it easier for you to copy objects from load-space to run-space at boot time
- Make it easier for you to manage memory overlays at run time
- Allow you to split GROUPs and output sections that have separate load and run addresses

7.15.1 A Current Boot-Loaded Application Development Process

In some embedded applications, there is a need to copy or download code and/or data from one location to another at boot time before the application actually begins its main execution thread. For example, an application may have its code and/or data in FLASH memory and need to copy it into on-chip memory before the application begins execution.

One way you can develop an application like this is to create a copy table in assembly code that contains three elements for each block of code or data that needs to be moved from FLASH into on-chip memory at boot time:

- The load address
- The run address
- The size

The process you follow to develop such an application might look like this:

- 1) Build the application to produce a .map file that contains the load and run addresses of each section that has a separate load and run placement.
- 2) Edit the copy table (used by the boot loader) to correct the load and run addresses as well as the size of each block of code or data that needs to be moved at boot time.
- 3) Build the application again, incorporating the updated copy table.
- 4) Run the application.

This process puts a heavy burden on you to maintain the copy table (by hand, no less). Each time a piece of code or data is added or removed from the application, you must repeat the process in order to keep the contents of the copy table up to date.

7.15.2 An Alternative Approach

You can avoid some of this maintenance burden by using the `LOAD_START()`, `RUN_START()`, and `SIZE()` operators that are already part of the linker command file syntax. For example, instead of building the application to generate a `.map` file, the linker command file can be annotated:

```
SECTIONS
{
    .flashcode: { app_tasks.obj(.text) }
        load = FLASH, run = PMEM,
        LOAD_START(_flash_code_ld_start),
        RUN_START(_flash_code_rn_start),
        SIZE(_flash_code_size)

    ...
}
```

In this example, the `LOAD_START()`, `RUN_START()`, and `SIZE()` operators instruct the linker to create three symbols:

Symbol	Description
<code>_flash_code_ld_start</code>	Load address of <code>.flashcode</code> section
<code>_flash_code_rn_start</code>	Run address of <code>.flashcode</code> section
<code>_flash_code_size</code>	Size of <code>.flashcode</code> section

These symbols can then be referenced from the copy table. The actual data in the copy table will be updated automatically each time the application is linked. This approach removes step 1 of the process described in section 7.15.1.

While maintenance of the copy table is reduced markedly, you must still carry the burden of keeping the copy table contents in sync with the symbols that are defined in the linker command file. Ideally, the linker would generate the boot copy table automatically. This would avoid having to build the application twice *and* free you from having to explicitly manage the contents of the boot copy table.

For more information on the `LOAD_START()`, `RUN_START()`, and `SIZE()` operators, see section 7.13.7, *Address and Dimension Operators*, on page 7-61.

7.15.3 Overlay Management Example

Consider an application which contains a memory overlay that must be managed at run time. The memory overlay is defined using a UNION in the linker command file as illustrated in Example 7–12:

Example 7–12. Using a UNION for Memory Overlay

```
SECTIONS
{
    ...

    UNION
    {
        GROUP
        {
            .task1: { task1.obj(.text) }
            .task2: { task2.obj(.text) }

        } load = ROM, LOAD_START(_task12_load_start), SIZE(_task12_size)

        GROUP
        {
            .task3: { task3.obj(.text) }
            .task4: { task4.obj(.text) }

        } load = ROM, LOAD_START(_task34_load_start), SIZE(_task_34_size)

    } run = RAM, RUN_START(_task_run_start)

    ...
}
```

The application must manage the contents of the memory overlay at run time. That is, whenever any services from .task1 or .task2 are needed, the application must first ensure that .task1 and .task2 are resident in the memory overlay. Similarly for .task3 and .task4.

To affect a copy of .task1 and .task2 from ROM to RAM at run time, the application must first gain access to the load address of the tasks (_task12_load_start), the run address (_task_run_start), and the size (_task12_size). Then this information is used to perform the actual code copy.

7.15.4 Generating Copy Tables Automatically with the Linker

The linker supports extensions to the linker command file syntax that enable you to do the following:

- Identify any object components that may need to be copied from load space to run space at some point during the run of an application
- Instruct the linker to automatically generate a copy table that contains (at least) the load address, run address, and size of the component that needs to be copied
- Instruct the linker to generate a symbol specified by you that provides the address of a linker-generated copy table. For instance, Example 7–12 can be written as shown in Example 7–13:

Example 7–13. Produce Address for Linker Generated Copy Table

```
SECTIONS
{
    ...

    UNION
    {
        GROUP
        {
            .task1: { task1.obj(.text) }
            .task2: { task2.obj(.text) }

        } load = ROM, table(_task12_copy_table)

        GROUP
        {
            .task3: { task3.obj(.text) }
            .task4: { task4.obj(.text) }

        } load = ROM, table(_task34_copy_table)

    } run = RAM

    ...
}
```

Using the SECTIONS directive from Example 7–13 in the linker command file, the linker generates two copy tables named: `_task12_copy_table` and `_task34_copy_table`. Each copy table provides the load address, run address, and size of the GROUP that is associated with the copy table. This information is accessible from application source code using the linker-generated symbols, `_task12_copy_table` and `_task34_copy_table`, which provide the addresses of the two copy tables, respectively.

Using this method, you do not have to worry about the creation or maintenance of a copy table. You can reference the address of any copy table generated by the linker in C/C++ or assembly source code, passing that value to a general purpose copy routine which will process the copy table and affect the actual copy.

7.15.5 The table() Operator

You can use the table() operator to instruct the linker to produce a copy table. A table() operator can be applied to an output section, a GROUP, or a UNION member. The copy table generated for a particular table() specification can be accessed through a symbol specified by you that is provided as an argument to the table() operator. The linker creates a symbol with this name and assigns it the address of the copy table as the value of the symbol. The copy table can then be accessed from the application using the linker-generated symbol.

Each table() specification you apply to members of a given UNION must contain a unique name. If a table() operator is applied to a GROUP, then none of that GROUP's members may be marked with a table() specification. The linker detects violations of these rules and reports them as warnings, ignoring each offending use of the table() specification. The linker does not generate a copy table for erroneous table() operator specifications.

7.15.6 Boot-Time Copy Tables

The linker supports a special copy table name, BINIT (or binit), that you can use to create a boot-time copy table. For example, the linker command file for the boot-loaded application described in section 7.15.2 can be rewritten as follows:

```
SECTIONS
{
    .flashcode: { app_tasks.obj(.text) }
    load = FLASH, run = PMEM,
    table(BINIT)
    ...
}
```

For this example, the linker creates a copy table that can be accessed through a special linker-generated symbol, `__binit__`, which contains the list of all object components that need to be copied from their load location to their run location at boot-time. If a linker command file does not contain any uses of table(BINIT), then the `__binit__` symbol is given a value of -1 to indicate that a boot-time copy table does not exist for a particular application.

You can apply the `table(BINIT)` specification to an output section, GROUP, or UNION member. If used in the context of a UNION, only one member of the UNION can be designated with `table(BINIT)`. If applied to a GROUP, then none of that GROUP's members may be marked with `table(BINIT)`. The linker detects violations of these rules and reports them as warnings, ignoring each offending use of the `table(BINIT)` specification.

7.15.7 Using the `table()` Operator to Manage Object Components

If you have several pieces of code that need to be managed together, then you can apply the same `table()` operator to several different object components. In addition, if you want to manage a particular object component in multiple ways, you can apply more than one `table()` operator to it. Consider the linker command file excerpt in Example 7–14:

Example 7–14. Linker Command File to Manage Object Components

```
SECTIONS
{
  UNION
  {
    .first: { a1.obj(.text), b1.obj(.text), c1.obj(.text) }
           load = EMEM, run = PMEM, table(BINIT), table(_first_ctbl)

    .second: { a2.obj(.text), b2.obj(.text) }
            load = EMEM, run = PMEM, table(_second_ctbl)
  }

  .extra: load = EMEM, run = PMEM, table(BINIT)

  ...
}
```

In this example, the output sections `.first` and `.extra` are copied from external memory (EMEM) into program memory (PMEM) at boot time while processing the BINIT copy table. After the application has started executing its main thread, it can then manage the contents of the overlay using the two overlay copy tables named: `_first_ctbl` and `_second_ctbl`.

7.15.8 Copy Table Contents

In order to use a copy table that is generated by the linker, you must be aware of the contents of the copy table. This information is included in a new run-time-support library header file, `cpy_tbl.h`, which contains a C source representation of the copy table data structure that is automatically generated by the linker.

Example 7–15 shows the TMS320C6000 copy table header file.

Example 7–15. TMS320C6000 `cpy_tbl.h` File

```

/*****
/* cpy_tbl.h
/*
/* Copyright (c) 2003 Texas Instruments Incorporated
/*
/* Specification of copy table data structures which can be automatically
/* generated by the linker (using the table() operator in the LCF).
/*
/*****

/*****
/* Copy Record Data Structure
/*****
typedef struct copy_record
{
    unsigned int load_addr;
    unsigned int run_addr;
    unsigned int size;
} COPY_RECORD;

/*****
/* Copy Table Data Structure
/*****
typedef struct copy_table
{
    unsigned short rec_size;
    unsigned short num_recs;
    COPY_RECORD    recs[1];
} COPY_TABLE;

/*****
/* Prototype for general purpose copy routine.
/*****
extern void copy_in(COPY_TABLE *tp);

```

For each object component that is marked for a copy, the linker creates a `COPY_RECORD` object for it. Each `COPY_RECORD` contains at least the following information for the object component:

- The load address
- The run address
- The size

The linker collects all `COPY_RECORD`s that are associated with the same copy table into a `COPY_TABLE` object. The `COPY_TABLE` object contains the size of a given `COPY_RECORD`, the number of `COPY_RECORD`s in the table, and the array of `COPY_RECORD`s in the table. For instance, in the BINIT example in section 7.15.6, the `.first` and `.extra` output sections will each have their own `COPY_RECORD` entries in the BINIT copy table. The BINIT copy table will then look like this:

```
COPY_TABLE __binit__ = { 12, 2,
                        { <load address of .first>,
                          <run address of .first>,
                          <size of .first> },
                        { <load address of .extra>,
                          <run address of .extra>,
                          <size of .extra> } };
```

7.15.9 General Purpose Copy Routine

The `cpy_tbl.h` file in Example 7–15 also contains a prototype for a general-purpose copy routine, `copy_in()`, which is provided as part of the run-time-support library. The `copy_in()` routine takes a single argument: the address of a linker-generated copy table. The routine then processes the copy table data object and performs the copy of each object component specified in the copy table.

The `copy_in()` function definition is provided in the `cpy_tbl.c` run-time-support source file shown in Example 7–16.

Example 7–16. Run-Time-Support *cpy_tbl.c* File

```

/*****
/* cpy_tbl.c
/*
/* Copyright (c) 2003 Texas Instruments Incorporated
/*
/* General purpose copy routine. Given the address of a linker-generated
/* COPY_TABLE data structure, effect the copy of all object components
/* that are designated for copy via the corresponding LCF table() operator.
/*
/*
/*****
#include <cpy_tbl.h>
#include <string.h>

/*****
/* COPY_IN()
/*****
void copy_in(COPY_TABLE *tp)
{
    unsigned short i;
    for (i = 0; i < tp->num_recs; i++)
    {
        COPY_RECORD crp = tp->recs[i];
        unsigned char *ld_addr = (unsigned char *)crp.load_addr;
        unsigned char *rn_addr = (unsigned char *)crp.run_addr;
        memcpy(rn_addr, ld_addr, crp.size);
    }
}

```

7.15.10 Linker Generated Copy Table Sections and Symbols

The linker creates and allocates a separate input section for each copy table that it generates. Each copy table symbol is defined with the address value of the input section that contains the corresponding copy table.

The linker generates a unique name for each overlay copy table input section. For example, `table(_first_ctbl)` would place the copy table for the `.first` section into an input section called `.ovly:_first_ctbl`. The linker creates a single input section, `.binit`, to contain the entire boot-time copy table.

Example 7–17 illustrates how you can control the placement of the linker-generated copy table sections using the input section names in the linker command file.

Example 7–17. Controlling the Placement of the Linker-Generated Copy Table Sections

```
SECTIONS
{
  UNION
  {
    .first: { a1.obj(.text), b1.obj(.text), c1.obj(.text) }
           load = EMEM, run = PMEM, table(BINIT), table(_first_ctbl)

    .second: { a2.obj(.text), b2.obj(.text) }
            load = EMEM, run = PMEM, table(_second_ctbl)
  }

  .extra: load = EMEM, run = PMEM, table(BINIT)

  ...

  .ovly: { } > BMEM
  .binit: { } > BMEM
}
```

For the linker command file in Example 7–17, the boot-time copy table is generated into a `.binit` input section, which is collected into the `.binit` output section, which is mapped to an address in the BMEM memory area. The `_first_ctbl` is generated into the `.ovly:_first_ctbl` input section and the `_second_ctbl` is generated into the `.ovly:_second_ctbl` input section. Since the base names of these input sections match the name of the `.ovly` output section, the input sections are collected into the `.ovly` output section, which is then mapped to an address in the BMEM memory area.

If you do not provide explicit placement instructions for the linker-generated copy table sections, they are allocated according to the linker's default placement algorithm.

The linker does not allow other types of input sections to be combined with a copy table input section in the same output section. The linker does not allow a copy table section that was created from a partial link session to be used as input to a succeeding link session.

7.15.11 Splitting Object Components and Overlay Management

In previous versions of the linker, splitting sections that have separate load and run placement instructions was not permitted. This restriction was because there was no effective mechanism for you, the developer, to gain access to the load address or run address of each one of the pieces of the split object component. Therefore, there was no effective way to write a copy routine that could move the split section from its load location to its run location.

However, the linker can access both the load address and run address of every piece of a split object component. Using the table() operator, you can tell the linker to generate this information into a copy table. The linker gives each piece of the split object component a COPY_RECORD entry in the copy table object.

For example, consider an application which has 7 tasks. Tasks 1 through 3 are overlaid with tasks 4 through 7 (using a UNION directive). The load placement of all of the tasks is split among 4 different memory areas (LMEM1, LMEM2, LMEM3, and LMEM4). The overlay is defined as part of memory area PMEM. You must move each set of tasks into the overlay at run time before any services from the set are used.

You can use table() operators in combination with splitting operators, >>, to create copy tables that have all the information needed to move either group of tasks into the memory overlay as shown in Example 7–18. Example 7–19 illustrates a possible driver for such an application.

Example 7–18. Creating a Copy Table to Access a Split Object Component

```
SECTIONS
{
  UNION
  {
    .task1to3: { *(.task1), *(.task2), *(.task3) }
               load >> LMEM1 | LMEM2 | LMEM4, table(_task13_ctbl)

    GROUP
    {
      .task4: { *(.task4) }
      .task5: { *(.task5) }
      .task6: { *(.task6) }
      .task7: { *(.task7) }

    } load >> LMEM1 | LMEM3 | LMEM4, table(_task47_ctbl)

  } run = PMEM

  ...

  .ovly: > LMEM4
}
```

Example 7–19. Split Object Component Driver

```

#include <cpy_tbl.h>

extern far COPY_TABLE task13_ctbl;
extern far COPY_TABLE task47_ctbl;

extern void task1(void);
...
extern void task7(void);

main()
{
    ...
    copy_in(&task13_ctbl);
    task1();
    task2();
    task3();
    ...

    copy_in(&task47_ctbl);
    task4();
    task5();
    task6();
    task7();
    ...
}

```

You must declare a COPY_TABLE object as *far* to allow the overlay copy table section placement to be independent from the other sections containing data objects (such as .bss).

The contents of the .task1to3 section are split in the section's load space and contiguous in its run space. The linker-generated copy table, _task13_ctbl, contains a separate COPY_RECORD for each piece of the split section .task1to3. When the address of _task13_ctbl is passed to copy_in(), each piece of .task1to3 is copied from its load location into the run location.

The contents of the GROUP containing tasks 4 through 7 are also split in load space. The linker performs the GROUP split by applying the split operator to each member of the GROUP in order. The copy table for the GROUP then contains a COPY_RECORD entry for every piece of every member of the GROUP. These pieces are copied into the memory overlay when the _task47_ctbl is processed by copy_in().

The split operator can be applied to an output section, GROUP, or the load placement of a UNION or UNION member. The linker does not permit a split operator to be applied to the run placement of either a UNION or of a UNION member. The linker detects such violations, emits a warning, and ignores the offending split operator usage.

7.16 Partial (Incremental) Linking

An output file that has been linked can be linked again with additional modules. This is known as *partial linking* or *incremental linking*. Partial linking allows you to partition large applications, link each part separately, and then link all the parts together to create the final executable program.

Follow these guidelines for producing a file that you will relink:

- The intermediate files produced by the linker *must* have relocation information. Use the `-r` option when you link the file the first time. (See section 7.4.1, *Relocation Capabilities (-a and -r Options)*, on page 7-6.)
- Intermediate files *must* have symbolic information. By default, the linker retains symbolic information in its output. Do not use the `-s` option if you plan to relink a file, because `-s` strips symbolic information from the output module. (See section 7.4.15, *Strip Symbolic Information (-s Option)*, on page 7-16.)
- Intermediate link steps should be concerned only with the formation of output sections and not with allocation. All allocation, binding, and MEMORY directives should be performed in the final link step.
- If the intermediate files have global symbols that have the same name as global symbols in other files and you want them to be treated as static (visible only within the intermediate file), you must link the files with the `-h` option (see section 7.4.9, *Make All Global Symbols Static (-h Option)*, on page 7-10).
- If you are linking C code, do not use `-c` or `-cr` until the final link step. Every time you invoke the linker with the `-c` or `-cr` option, the linker attempts to create an entry point. (See section 7.4.5, *C Language Options (-c and -cr Options)*, on page 7-9.)

The following example shows how you can use partial linking:

Step 1: Link the file `file1.com`; use the `-r` option to retain relocation information in the output file `tempout1.out`.

```
cl6x -z -r -o tempout1 file1.com
```

`file1.com` contains:

```
SECTIONS
{
    ss1:    {
            f1.obj
            f2.obj
            .
            .
            fn.obj
        }
}
```

Step 2: Link the file `file2.com`; use the `-r` option to retain relocation information in the output file `tempout2.out`.

```
cl6x -z -r -o tempout2 file2.com
```

`file2.com` contains:

```
SECTIONS
{
    ss2:    {
            g1.obj
            g2.obj
            .
            .
            gn.obj
        }
}
```

Step 3: Link `tempout1.out` and `tempout2.out`.

```
cl6x -z -m final.map -o final.out tempout1.out tempout2.out
```

7.17 Linking C/C++ Code

The C/C++ compiler produces assembly language source code that can be assembled and linked. For example, a C program consisting of modules prog1, prog2, etc., can be assembled and then linked to produce an executable file called prog.out:

```
cl6x -z -c -o prog.out prog1.obj prog2.obj ... rts6200.lib
```

The `-c` option tells the linker to use special conventions that are defined by the C/C++ environment.

The archive libraries listed below contain C/C++ run-time-support functions:

rts6200.lib	rts6400.lib	rts6700.lib
rts6200e.lib	rts6400e.lib	rts6700e.lib

C, C++, and mixed C and C++ programs can use the same run-time-support library. Run-time-support functions and variables that can be called and referenced from both C and C++ will have the same linkage.

For more information about the TMS320C6000 C/C++ language, including the run-time environment and run-time-support functions, see the *TMS320C6000 Optimizing Compiler User's Guide*.

7.17.1 Run-Time Initialization

All C/C++ programs must be linked with code to initialize and execute the program, called a *bootstrap* routine, also known as the *boot.obj* object module. The symbol `_c_int00` is defined as the program entry point and is the start of the C boot routine in *boot.obj*; referencing `_c_int00` ensures that *boot.obj* is automatically linked in from the run-time-support library. When a program begins running, it executes *boot.obj* first. The *boot.obj* symbol contains code and data for initializing the run-time environment and performs the following tasks:

- Sets up the system stack and configuration registers
- Processes the run-time *.cinit* initialization table and autoinitializes global variables (when the linker is invoked with the `-c` option)
- Disables interrupts and calls `_main`

The run-time-support object libraries contain *boot.obj*. You can:

- Use the archiver to extract *boot.obj* from the library and then link the module in directly.
- Include the appropriate run-time-support library as an input file (the linker automatically extracts *boot.obj* when you use the `-c` or `-cr` option).

7.17.2 Object Libraries and Run-Time Support

The *TMS320C6000 Optimizing C/C++ Compiler User's Guide* describes additional run-time-support functions that are included in `rts.src`. If your program uses any of these functions, you must link the appropriate run-time-support library with your object files.

You can also create your own object libraries and link them. The linker includes and links only those library members that resolve undefined references.

7.17.3 Setting the Size of the Stack and Heap Sections

The C/C++ language uses two uninitialized sections called `.system` and `.stack` for the memory pool used by the `malloc()` functions and the run-time stacks, respectively. You can set the size of these by using the `-heap` or `-stack` option and specifying the size of the section as a 4-byte constant immediately after the option. The default size for both, if the options are not used, is 1K words.

See section 7.4.10, *Define Heap Size (-heap size Option)*, on page 7-11 and section 7.4.16, *Define Stack Size (-stack size Option)*, on page 7-16 for more information on setting stack sizes.

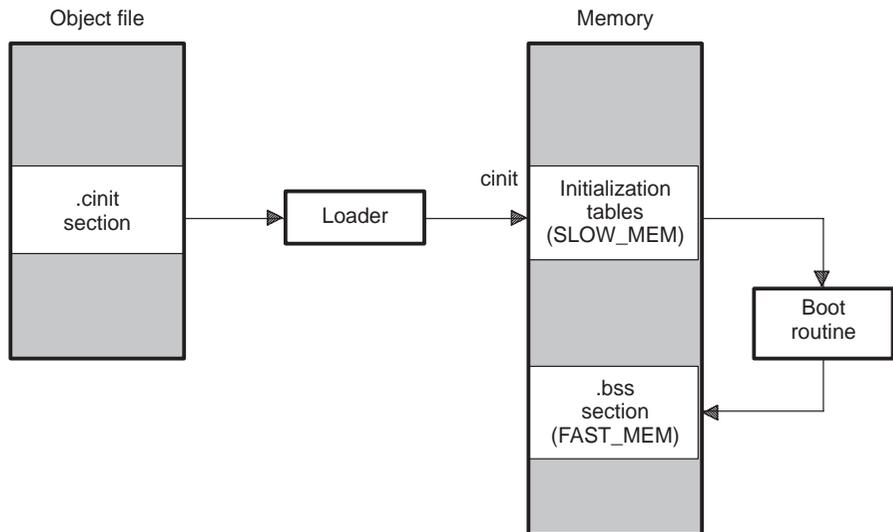
7.17.4 Autoinitialization of Variables at Run Time

Autoinitializing variables at run time is the default method of autoinitialization. To use this method, invoke the linker with the `-c` option.

Using this method, the `.cinit` section is loaded into memory along with all the other initialized sections. The linker defines a special symbol called `cinit` that points to the beginning of the initialization tables in memory. When the program begins running, the C boot routine copies data from the tables (pointed to by `.cinit`) into the specified variables in the `.bss` section. This allows initialization data to be stored in slow external memory and copied to fast external memory each time the program starts.

Figure 7–5 illustrates autoinitialization at run time. Use this method in any system where your application runs from code burned into slow external memory.

Figure 7–5. Autoinitialization at Run Time



7.17.5 Initialization of Variables at Load Time

Initialization of variables at load time enhances performance by reducing boot time and by saving the memory used by the initialization tables. To use this method, invoke the linker with the `-cr` option.

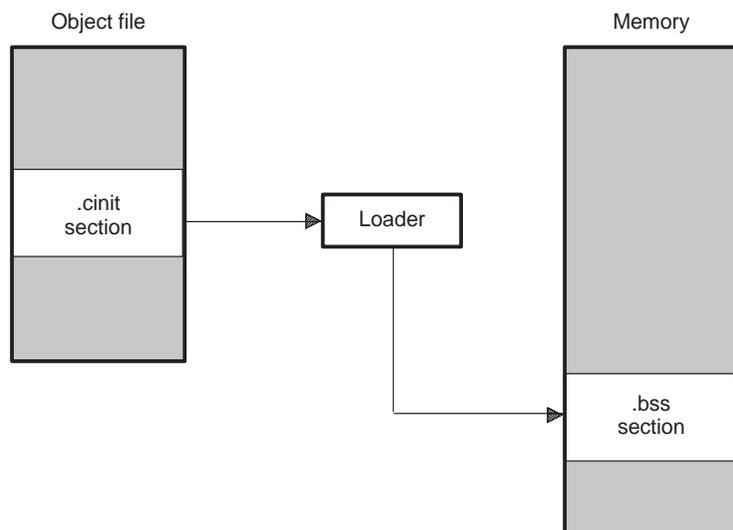
When you use the `-cr` linker option, the linker sets the `STYP_COPY` bit in the `.cinit` section's header. This tells the loader not to load the `.cinit` section into memory. (The `.cinit` section occupies no space in the memory map.) The linker also sets the `cinit` symbol to `-1` (normally, `cinit` points to the beginning of the initialization tables). This indicates to the boot routine that the initialization tables are not present in memory; accordingly, no run-time initialization is performed at boot time.

A loader must be able to perform the following tasks to use initialization at load time:

- Detect the presence of the `.cinit` section in the object file.
- Determine that `STYP_COPY` is set in the `.cinit` section header, so that it knows not to copy the `.cinit` section into memory.
- Understand the format of the initialization tables.

Figure 7–6 illustrates the initialization of variables at load time.

Figure 7–6. Initialization at Load Time



7.17.6 The `-c` and `-cr` Linker Options

The following list outlines what happens when you invoke the linker with the `-c` or `-cr` option.

- The symbol `_c_int00` is defined as the program entry point. The `_c_int00` symbol is the start of the C boot routine in `boot.obj`; referencing `_c_int00` ensures that `boot.obj` is automatically linked in from the appropriate run-time-support library.
- The `.cinit` output section is padded with a termination record to designate to the boot routine (autoinitialize at run time) or the loader (initialize at load time) when to stop reading the initialization tables.
- When you autoinitialize at run time (`-c` option), the linker defines `cinit` as the starting address of the `.cinit` section. The C boot routine uses this symbol as the starting point for autoinitialization.
- When you initialize at load time (`-cr` option):
 - The linker sets `cinit` to `-1`. This indicates that the initialization tables are not in memory, so no initialization is performed at run time.
 - The `STYP_COPY` flag (0010h) is set in the `.cinit` section header. `STYP_COPY` is the special attribute that tells the loader to perform initialization directly and not to load the `.cinit` section into memory. The linker does not allocate space in memory for the `.cinit` section.

7.18 Linker Example

This example links three object files named `demo.obj`, `ctrl.obj`, and `tables.obj` and creates a program called `demo.out`.

Assume that target memory has the following configuration:

Program Memory

Address Range	Contents
0x00000000 to 0x00001000	SLOW_MEM
0x00001000 to 0x00002000	FAST_MEM
0x08000000 to 0x08000400	EEPROM

The output sections are constructed from the following input sections:

- Executable code, contained in the `.text` sections of `demo.obj`, `ctrl.obj`, and `tables.obj`, must be linked into `FAST_MEM`.
- A set of interrupt vectors, contained in the `.intvecs` section of `tables.obj`, must be linked at address `0x00000000`.
- A table of coefficients, contained in the `.data` section of `tables.obj`, must be linked into `EEPROM`. The remainder of block `EEPROM` must be initialized to the value `0xFF00FF00`.
- A set of variables, contained in the `.bss` section of `ctrl.obj`, must be linked into `SLOW_MEM` and preinitialized to `0x00000100`.
- The `.bss` sections of `demo.obj` and `tables.obj` must be linked into `SLOW_MEM`.

Example 7–20 shows the linker command file for this example. Example 7–21 shows the map file.

Example 7–20. Linker Command File, demo.cmd

```

/*****
/****          Specify Linker Options          ****/
/*****
-e SETUP          /* Define the program entry point */
-o demo.out      /* Name the output file */
-m demo.map      /* Create an output map */
/*****
/****          Specify the Input Files          ****/
/*****
demo.obj
ctrl.obj
tables.obj
/*****
/****          Specify the Memory Configuration  ****/
/*****
MEMORY
{
    FAST_MEM   : org = 0x00000000   len = 0x00001000
    SLOW_MEM   : org = 0x00001000   len = 0x00001000
    EEPROM    : org = 0x08000000   len = 0x00000400
}
/*****
/****          Specify the Output Sections      ****/
/*****
SECTIONS
{
    .text      : {} > FAST_MEM /* Link all .text sections into ROM */
    .intvecs   : {} > 0x0      /* Link interrupt vectors at 0x0 */
    .data      :              /* Link .data sections */
    {
        tables.obj(.data)
        . = 0x400; /* Create hole at end of block */
    } = 0xFF00FF00 > EEPROM /* Fill and link into EEPROM */
    ctrl_vars: /* Create new ctrl_vars section */
    {
        ctrl.obj(.bss)
    } = 0x00000100 > SLOW_MEM /* Fill with 0x100 and link into RAM */
    .bss      : {} > SLOW_MEM /* Link remaining .bss sections into RAM */
}
/*****
/****          End of Command File          ****/
/*****

```

Invoke the linker by entering the following command:

```
c16x -z demo.cmd
```

This creates the map file shown in Example 7–21 and an output file called demo.out that can be run on a TMS320C6000.

Example 7-21. Output Map File, demo.map

```
OUTPUT FILE NAME: <demo.out>
```

```
ENTRY POINT SYMBOL: 0
```

MEMORY CONFIGURATION

name	origin	length	used	attributes	fill
FAST_MEM	00000000	000001000	00000078	RWIX	
SLOW_MEM	00001000	000001000	00000502	RWIX	
EEPROM	08000000	000000400	00000400	RWIX	

SECTION ALLOCATION MAP

output section	page	origin	length	attributes/ input sections
.text	0	00000000	00000064	
		00000000	00000030	demo.obj (.text)
		00000030	00000000	tables.obj (.text)
		00000030	00000010	--HOLE-- [fill = 00000000]
.intvecs	0	00000000	00000014	
		00000000	00000014	tables.obj (.intvecs)
.data	0	08000000	00000400	
		08000000	00000004	tables.obj (.data)
		08000004	000003fc	--HOLE-- [fill = ff00ff00]
		08000400	00000000	ctrl.obj (.data)
ctrl_vars	0	00001000	00000500	
		00001000	00000500	ctrl.obj (.bss) [fill = 00000100]
.bss	0	00001500	00000002	UNINITIALIZED
		00001500	00000002	demo.obj (.bss)
		00001502	00000000	tables.obj (.bss)

GLOBAL SYMBOLS

address	name	address	name
00001500	\$.bss	00000000	.text
00001500	.bss	00000000	_x42
08000000	.data	00000018	_SETUP
00000000	.text	00000040	_fill_tab
00000018	_SETUP	00000064	etext
00000040	_fill_tab	00001500	\$.bss
00000000	_x42	00001500	.bss
08000400	edata	00001502	end
00001502	end	08000000	gvar
00000064	etext	08000000	.data
08000000	gvar	08000400	edata

```
[11 symbols]
```



Absolute Lister Description

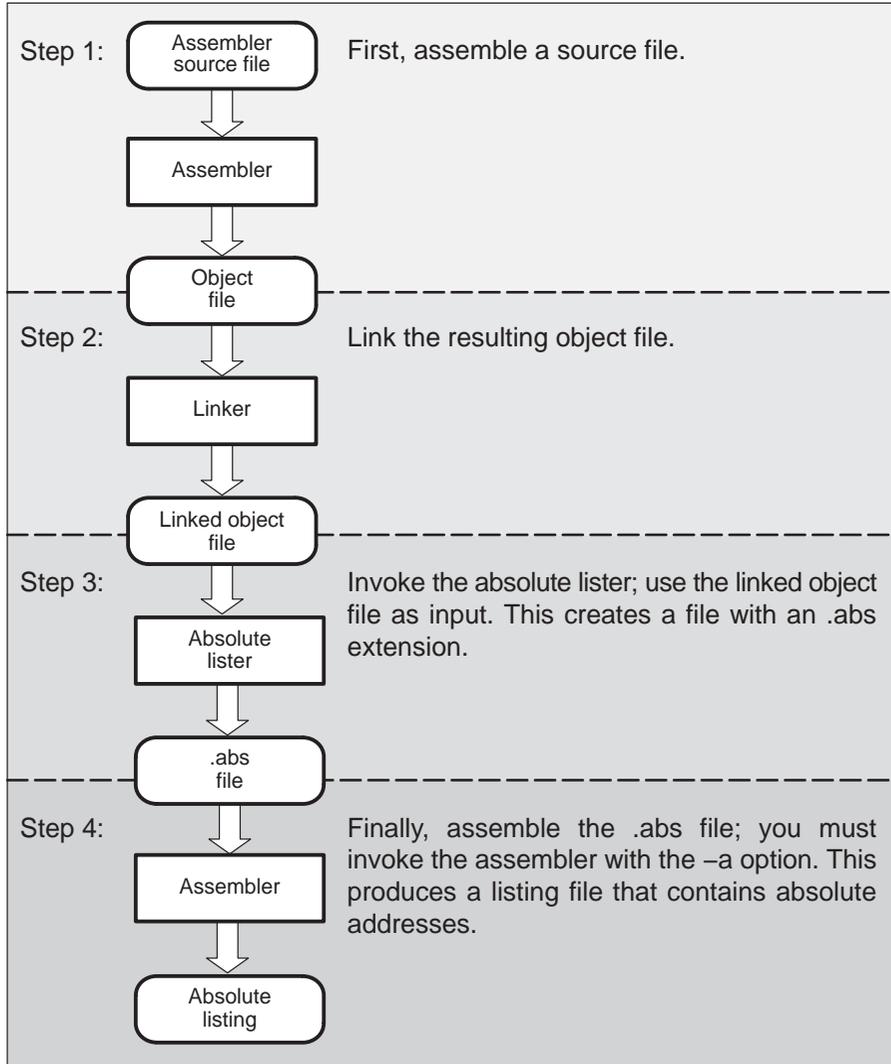
The TMS320C6000™ absolute lister is a debugging tool that accepts linked object files as input and creates .abs files as output. These .abs files can be assembled to produce a listing that shows the absolute addresses of object code. Manually, this could be a tedious process requiring many operations; however, the absolute lister utility performs these operations automatically.

Topic	Page
8.1 Producing an Absolute Listing	8-2
8.2 Invoking the Absolute Lister	8-3
8.3 Absolute Lister Example	8-5

8.1 Producing an Absolute Listing

Figure 8–1 illustrates the steps required to produce an absolute listing.

Figure 8–1. Absolute Lister Development Flow



8.2 Invoking the Absolute Lister

The syntax for invoking the absolute lister is:

```
abs6x [-options] input file
```

abs6x	is the command that invokes the absolute lister.
<i>options</i>	<p>identifies the absolute lister options that you want to use. Options are not case sensitive and can appear anywhere on the command line following the command. Precede each option with a hyphen (-). The absolute lister options are as follows:</p> <ul style="list-style-type: none"> -e enables you to change the default naming conventions for filename extensions on assembly files, C source files, and C header files. The three options are listed below. <ul style="list-style-type: none"> <input type="checkbox"/> -ea [<i>.]asmext</i> for assembly files (default is .asm) <input type="checkbox"/> -ec [<i>.]cext</i> for C source files (default is .c) <input type="checkbox"/> -eh [<i>.]hext</i> for C header files (default is .h) <p>The . in the extensions and the space between the option and the extension are optional.</p> -q (quiet) suppresses the banner and all progress information.
<i>input file</i>	names the linked object file. If you do not supply an extension, the absolute lister assumes that the input file has the default extension .out. If you do not supply an input filename when you invoke the absolute lister, the absolute lister prompts you for one.

The absolute lister produces an output file for each file that was linked. These files are named with the input filenames and an extension of .abs. Header files, however, do not generate a corresponding .abs file.

Assemble these files with the **-aa** assembler option as follows to create the absolute listing:

```
c16x -aa filename.abs
```

The **-e** options affect both the interpretation of filenames on the command line and the names of the output files. They should always precede any filename on the command line.

The `-e` options are useful when the linked object file was created from C files compiled with the debugging option (`-g` compiler option). When the debugging option is set, the resulting linked object file contains the name of the source files used to build it. In this case, the absolute lister does not generate a corresponding `.abs` file for the C header files. Also, the `.abs` file corresponding to a C source file uses the assembly file generated from the C source file rather than the C source file itself.

For example, suppose the C source file `hello.csr` is compiled with the debugging option set; the debugging option generates the assembly file `hello.s`. The `hello.csr` file includes `hello.hsr`. Assuming the executable file created is called `hello.out`, the following command generates the proper `.abs` file:

```
abs6x -ea s -ec csr -eh hsr hello.out
```

An `.abs` file is not created for `hello.hsr` (the header file), and `hello.abs` includes the assembly file `hello.s`, not the C source file `hello.csr`.

8.3 Absolute Lister Example

This example uses three source files. The files `module1.asm` and `module2.asm` both include the file `globals.def`.

`module1.asm`

```
.text
.align 4
.bss array, 100
.bss dflag, 4
.copy globals.def

MVKL offset, A0
MVKH offset, A0
LDW  *+b14(dflag), A2
nop  4
```

`module2.asm`

```
.bss offset,2
.copy globals.def

mvkl  offset,a0
mvkh  offset,a0
mvkl  array,a3
mvkh  array,a3
```

`globals.def`

```
.global dflag
.global array
.global offset
```

The following steps create absolute listings for the files `module1.asm` and `module2.asm`:

Step 1: First, assemble `module1.asm` and `module2.asm`:

```
c16x module1
c16x module2
```

This creates two object files called `module1.obj` and `module2.obj`.

Step 2: Next, link module1.obj and module2.obj using the following linker command file, called bttest.cmd:

```
-o bttest.out
-m bttest.map
module1.obj
module2.obj
MEMORY
{
    PMEM:    origin=00000000h    length=00010000h
    DMEM:    origin=80000000h    length=00010000h
}
SECTIONS
{
    .data: >DMEM
    .text: >PMEM
    .bss: >DMEM
}
```

Invoke the linker:

```
cl6x -z bttest.cmd
```

This command creates an executable object file called bttest.out; use this new file as input for the absolute lister.

Step 3: Now, invoke the absolute lister:

```
abs6x bttest.out
```

This command creates two files called module1.abs and module2.abs:

module1.abs:

```

        .nolist
array   .setsym    080000000h
dflag   .setsym    080000064h
offset  .setsym    080000068h
.data   .setsym    080000000h
__data__ .setsym    080000000h
edata   .setsym    080000000h
__edata__ .setsym    080000000h
.text   .setsym    000000000h
__text__ .setsym    000000000h
etext   .setsym    000000040h
__etext__ .setsym    000000040h
.bss    .setsym    080000000h
__bss__ .setsym    080000000h
end     .setsym    08000006ah
__end__ .setsym    08000006ah
$bss    .setsym    080000000h
        .setsect   ".text",000000020h
        .setsect   ".data",080000000h
        .setsect   ".bss",080000000h
        .list
        .text
        .copy      "module1.asm"
```

module2.abs:

```
.nolist
array      .setsym      080000000h
dflag      .setsym      080000064h
offset     .setsym      080000068h
.data      .setsym      080000000h
__data__   .setsym      080000000h
edata      .setsym      080000000h
__edata__  .setsym      080000000h
.text      .setsym      000000000h
__text__   .setsym      000000000h
etext      .setsym      000000040h
__etext__  .setsym      000000040h
.bss       .setsym      080000000h
__bss__    .setsym      080000000h
end        .setsym      08000006ah
__end__    .setsym      08000006ah
$bss       .setsym      080000000h
           .setsect     ".text",000000000h
           .setsect     ".data",080000000h
           .setsect     ".bss",080000068h
           .list
           .text
           .copy        "module2.asm"
```

These files contain the following information that the assembler needs when you invoke it in step 4:

- They contain `.setsym` directives, which equate values to global symbols. Both files contain global equates for the symbol `dflag`. The symbol `dflag` was defined in the file `globals.def`, which was included in `module1.asm` and `module2.asm`.
- They contain `.setsect` directives, which define the absolute addresses for sections.
- They contain `.copy` directives, which tell the assembler which assembly language source file to include.

The `.setsym` and `.setsect` directives are not useful in normal assembly; they are useful only for creating absolute listings.

Step 4: Finally, assemble the .abs files created by the absolute lister (remember that you must use the `-aa` option when you invoke the assembler):

```
c16x -aa module1.abs
c16x -aa module2.abs
```

This command sequence creates two listing files called `module1.lst` and `module2.lst`; no object code is produced. These listing files are similar to normal listing files; however, the addresses shown are absolute addresses.

The absolute listing files created are `module1.lst` (see Figure 8–2) and `module2.lst` (see Figure 8–3).

Figure 8–2. *module1.lst*

```
TMS320C6x COFF Assembler          Version x.xx          Mon Jan  5 11:34:00 1998
Copyright (c) 1996-1998 Texas Instruments Incorporated
module1.abs                                PAGE          1
      22 00000020          .text
      23          .copy          "module1.asm"
A      1 00000020          .text
A      2          .align 4
A      3 80000000          .bss  array, 100
A      4 80000064          .bss  dflag, 4
A      5          .copy  globals.def
B      1          .global dflag
B      2          .global array
B      3          .global offset
A      6
A      7 00000020 00003428!    MVKL          offset, A0
A      8 00000024 00400068!    MVKH          offset, A0
A      9 00000028 0100196C-    LDW          *+b14(dflag), A2
A     10 0000002c 00006000    nop          4

No Errors, No Warnings
```

Figure 8-3. module2.lst

```
TMS320C6x COFF Assembler      Version x.xx      Mon Jan  5 11:34:05 1998
Copyright (c) 1996-1998 Texas Instruments Incorporated
module2.abs                                PAGE      1

      22 00000000                .text
      23                                .copy      "module2.asm"
A     1 80000068                .bss offset,2
A     2                                .copy globals.def
B     1                                .global dflag
B     2                                .global array
B     3                                .global offset
A     3
A     4 00000000 00003428-      mvkl      offset,a0
A     5 00000004 00400068-      mvkh      offset,a0
A     6 00000008 01800028!      mvkl      array,a3
A     7 0000000c 01C00068!      mvkh      array,a3

No Errors, No Warnings
```

Cross-Reference Lister Description

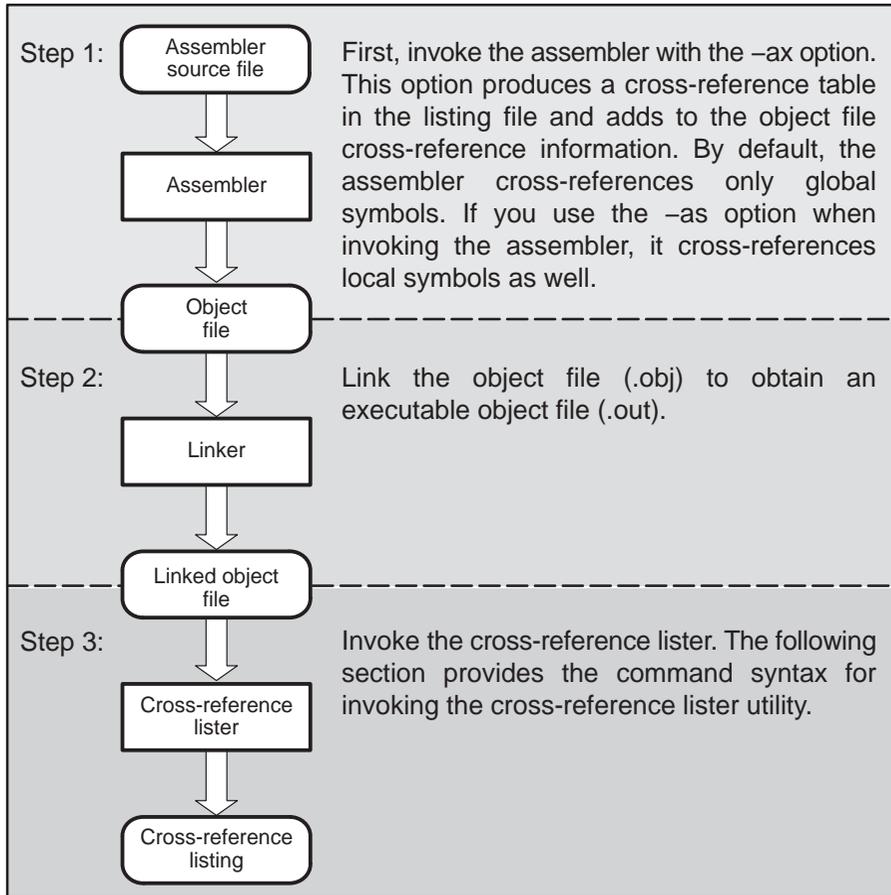
The TMS320C6000™ cross-reference lister is a debugging tool. This utility accepts linked object files as input and produces a cross-reference listing as output. This listing shows symbols, their definitions, and their references in the linked source files.

Topic	Page
9.1 Producing a Cross-Reference Listing	9-2
9.2 Invoking the Cross-Reference Lister	9-3
9.3 Cross-Reference Listing Example	9-4

9.1 Producing a Cross-Reference Listing

Figure 9–1 illustrates the steps required to produce a cross-reference listing.

Figure 9–1. The Cross-Reference Lister in the TMS320C6000 Software Development Flow



9.2 Invoking the Cross-Reference Lister

To use the cross-reference utility, the file must be assembled with the correct options and then linked into an executable file. Assemble the assembly language files with the `-ax` option. This option creates a cross-reference listing and adds cross-reference information to the object file. By default the assembler cross-references only global symbols, but if the assembler is invoked with the `-as` option, local symbols are also added. Link the object files to obtain an executable file.

To invoke the cross-reference lister, enter the following:

```
xref6x [options] [input filename [output filename]]
```

xref6x	is the command that invokes the cross-reference utility.
<i>options</i>	identifies the cross-reference lister options you want to use. Options are not case sensitive and can appear anywhere on the command line following the command. Precede each option with a hyphen (-). The cross-reference lister options are as follows: <ul style="list-style-type: none"> -l (lowercase L) specifies the number of lines per page for the output file. The format of the <code>-l</code> option is <code>-l<i>num</i></code>, where <i>num</i> is a decimal constant. For example, <code>-l30</code> sets the number of lines per page in the output file to 30. The space between the option and the decimal constant is optional. The default is 60 lines per page. -q suppresses the banner and all progress information (run quiet).
<i>input filename</i>	is a linked object file. If you omit the input filename, the utility prompts for a filename.
<i>output filename</i>	is the name of the cross-reference listing file. If you omit the output filename, the default filename is the input filename with an <code>.xrf</code> extension.

9.3 Cross-Reference Listing Example

The following is an example of cross-reference listing:

Example 9-1. Cross-Reference Listing

```

=====
Symbol: _SETUP
Filename      RTYP   AsmVal      LnkVal      DefLn      RefLn      RefLn      RefLn
-----
demo.asm      EDEF   '00000018   00000018    18         13         20
=====

Symbol: _fill_tab
Filename      RTYP   AsmVal      LnkVal      DefLn      RefLn      RefLn      RefLn
-----
ctrl.asm      EDEF   '00000000   00000040    10         5
=====

Symbol: _x42
Filename      RTYP   AsmVal      LnkVal      DefLn      RefLn      RefLn      RefLn
-----
demo.asm      EDEF   '00000000   00000000    7          4          18
=====

Symbol: gvar
Filename      RTYP   AsmVal      LnkVal      DefLn      RefLn      RefLn      RefLn
-----
tables.asm    EDEF   "00000000   08000000    11         10
=====

```

The terms defined below appear in the preceding cross-reference listing:

Symbol	Name of the symbol listed
Filename	Name of the file where the symbol appears
RTYP	The symbol's reference type in this file. The possible reference types are: <ul style="list-style-type: none"> STAT The symbol is defined in this file and is not declared as global. EDEF The symbol is defined in this file and is declared as global. EREF The symbol is not defined in this file but is referenced as global. UNDF The symbol is not defined in this file and is not declared as global.
AsmVal	This hexadecimal number is the value assigned to the symbol at assembly time. A value may also be preceded by a character that describes the symbol's attributes. Table 9–1 lists these characters and names.
LnkVal	This hexadecimal number is the value assigned to the symbol after linking.
DefLn	The statement number where the symbol is defined.
RefLn	The line number where the symbol is referenced. If the line number is followed by an asterisk (*), then that reference can modify the contents of the object. A blank in this column indicates that the symbol was never used.

Table 9–1. Symbol Attributes in Cross-Reference Listing

Character	Meaning
'	Symbol defined in a .text section
"	Symbol defined in a .data section
+	Symbol defined in a .sect section
–	Symbol defined in a .bss or .usect section



Object File Utilities Descriptions

This chapter describes how to invoke the following miscellaneous utilities:

- The **object file display utility** prints the contents of object files, executable files, and/or archive libraries in both human readable and XML formats.
- The **name utility** prints a list of names defined and referenced in a COFF object or an executable file.
- The **strip utility** removes symbol table and debugging information from object and executable files.

Topic	Page
10.1 Invoking the Object File Display Utility	10-2
10.2 XML Tag Index	10-3
10.3 Example XML Consumer	10-9
10.4 Invoking the Name Utility	10-16
10.5 Invoking the Strip Utility	10-17

10.1 Invoking the Object File Display Utility

The object file display utility, *ofd6x*, is used to print the contents of object files (.obj), executable files (.out), and/or archive libraries (.lib) in both text and XML formats.

To invoke the object file display utility, enter the following:

```
ofd6x [options] input filenames [input filenames]
```

ofd6x is the command that invokes the object file display utility.

input file- names the assembly language source file. The file name must
names contain a .asm extension.

options identify the object file display utility options that you want to use. Options are not case sensitive and can appear anywhere on the command line following the command. Precede each option with a hyphen.

- g** Appends DWARF debug information to program output.
- ofilename** Sends program output to *filename* rather than to the screen.
- x** Displays output in XML format.

If the object file display utility is invoked without any options, it displays information about the contents of the input files on the console screen.

Note: Object File Display Format

The object file display utility produces data in a text format by default. This data is not intended to be used as machine or software input.

10.2 XML Tag Index

Table 10–1 describes the XML tags that are generated by the object file display utility when invoked with the `-x` option.

Table 10–1. XML Tag Index

Tag Name	Context	Description
<addr>	<line_entry>	PC address
	<row>	PC address
	<value>	Machine address
<addr_class>	<value>	Address class
<addr_size>	<compile_unit>	Size of one machine address (octets)
	<section>	Size of one machine address (octets)
<alignment>	<section>	Alignment factor
<archive>	<ofd>	Archive file (.lib)
<attribute>	<die>	Attribute of a DWARF DIE
<aux_count>	<symbol>	Number of auxiliary entries for this symbol
<banner>	<ofd>	Tool name and version information
<block>	<section>	True if alignment is used as blocking factor
	<value>	Data block
<bss>	<section>	True if this section contains uninitialized data
<bss_size>	<optional_file_header>	Size of uninitialized data
<byte_swapped>	<file_header>	Endianness of build host is opposite of current host
<clink>	<section>	True if this section is conditionally linked
<column>	<line_entry>	Source column number
<compile_unit>	<section>	Compile unit
<const>	<value>	Constant
<copy>	<section>	True if this section is a copy section
<copyright>	<ofd>	Copyright notice
<cpu_flags>	<file_header>	CPU flags

Table 10–1. XML Tag Index (Continued)

Tag Name	Context	Description
<data>	<section>	True if this section contains initialized data
<data_size>	<optional_file_header>	Size of initialized data
<data_start>	<optional_file_header>	Beginning address of initialized data
<destination>	<register>	Destination register
<die>	<compile_unit>	DWARF debugging information entry (DIE)
<dim_bound>	<dimension>	Dimension upper-bound
<dim_num>	<dimension>	Dimension number
<dimension>	<symbol>	Array dimension
<disp>	<reloc_entry>	Extra address encoding information
<dummy>	<section>	True if this section is a dummy section
<dwarf>	<ti_coff>	DWARF information
<endian>	<file_header>	Endianness of target machine
<entry_point>	<optional_file_header>	Entry point of executable program
<exec>	<file_header>	True if this file is executable
<fde>	<section>	A DWARF frame description entry (FDE)
<field_size>	<reloc_entry>	Size of the field to relocate
<file_header>	<ti_coff>	COFF file header
<file_length>	<file_header>	Size of this file
<file_name>	<line_entry>	Name of source file
	<symbol>	Name of source file
<file_offsets>	<section>	File offsets associated with this section
<flag>	<value>	Flag
<form>	<attribute>	Attribute form
<frame_size>	<symbol>	Size of function frame
<function>	<line_numbers>	Line number entries for one function
<icode>	<section>	True if this section has I-Code associated with it

Table 10–1. XML Tag Index (Continued)

Tag Name	Context	Description
<index>	<symbol>	Index of this symbol in the symbol table
<indirect_register>	<memory>	Indirect register used for calculating destination address
<initial_location>	<fde>	Start of function referred to by the FDE
<internal>	<reloc_entry>	True if this relocation is internal
<kind>	<symbol>	Kind of symbol (defined, undefined, absolute, symbolic debug)
<length>	<symbol>	Length of section
<line>	<line_entry>	Source line number
	<symbol>	First source line associated with this symbol
<line_count>	<section>	Number of line number entries
	<symbol>	Number of line number entries
<line_entry>	<compile_unit>	Line number entry
	<line_numbers>	Line number entry
<line_numbers>	<section>	Line number entries
<line_ptr>	<file_offsets>	File offset of line number entries
	<symbol>	File offset of line number entries
<Inno_strip>	<file_header>	True if line numbers were stripped from this file
<localsym_strip>	<file_header>	True if local symbols were stripped from this file
<magic>	<optional_file_header>	Optional file header magic number (0x0108)
<math_relative>	<reloc_entry>	True if this relocation is math relative
<memory>	<row>	SOE register is saved to memory
<name>	<fde>	Name of function referred to by the FDE
	<function>	Name of the current function
	<ofd>	Name of an object or archive file
	<section>	Name of this section
	<symbol>	Name of this symbol

Table 10–1. XML Tag Index (Continued)

Tag Name	Context	Description
<next_symbol>	<symbol>	Index of next symbol after multisymbol entity
<noload>	<section>	True if this section is a no-load section
<object_file>	<ofd>	Object file (.obj, .out)
<ofd>		Object file display (OFD) document
<offset>	<memory>	Offset of destination address from indirect register
	<reloc_entry>	Offset of the field from relocatable address
<optional_file_header>	<ti_coff>	Optional file header
<padded>	<section>	True if this section has been padded (C55x only)
<page>	<section>	Memory page
<pass>	<section>	True if this section is passed through unchanged
<physical_addr>	<section>	Physical (run) address of section
<raw_data_ptr>	<file_offsets>	File offset of raw data
<raw_data_size>	<section>	Size of raw data (octets)
<ref>	<value>	Reference
<register>	<row>	SOE register is saved to register
<register_mask>	<symbol>	Mask of saved SOE registers
<regular>	<section>	True if this section is a regular section
<reloc_count>	<section>	Number of relocation entries
	<symbol>	Number of relocation entries
<reloc_entry>	<relocations>	Relocation entry
<reloc_ptr>	<file_offsets>	File offset of relocation entries
<reloc_strip>	<file_header>	True if relocation information was stripped from this file
<relocations>	<section>	Relocation entries
<return_address_register>	<fde>	Register used to pass the return address of this function
<row>	<table>	Table row

Table 10–1. XML Tag Index (Continued)

Tag Name	Context	Description
<section>	<dwarf>	DWARF section
	<symbol>	Section containing the definition of this symbol
	<ti_coff>	COFF section
<section_count>	<file_header>	Number of section headers
<size_in_addr>	<symbol>	Number of machine-address-sized units in function
<size_in_bits>	<symbol>	Size of symbol (bits)
<source>	<memory>	Source register
	<register>	Source register
<start_symbol>	<symbol>	First symbol in multi-symbol entity
<storage_class>	<symbol>	Storage class of this symbol
<storage_type>	<symbol>	Storage type of this symbol
<string>	<string_table>	String table entry
	<value>	String
<string_table>	<ti_coff>	String table
<string_table_size>	<string_table>	Size of string table
<sym_merge>	<file_header>	True if debug type-symbols were merged
<symbol>	<symbol_table>	Symbol table entry
<symbol_count>	<file_header>	Number of entries in the symbol table
<symbol_relative>	<reloc_entry>	Relocation is relative to the specified symbol
<symbol_table>	<ti_coff>	Symbol table
<table>	<fde>	FDE table
<tag>	<die>	Tag name
<tag_index>	<symbol>	Reference to user-defined type
<target_id>	<file_header>	Target ID; magic number identifying the target machine
<text>	<section>	True if this section contains code

Table 10–1. XML Tag Index (Continued)

Tag Name	Context	Description
<text_size>	<optional_file_header>	Size of executable code
<text_start>	<optional_file_header>	Beginning address of executable code
<ti_coff>	<object_file>	TI COFF file
<tool_version>	<optional_file_header>	Tool version stamp
<type>	<attribute>	Attribute type
	<reloc_entry>	Type of relocation
<type_ref>	<value>	Type reference
<value>	<attribute>	Attribute value
	<reloc_entry>	Value
	<symbol>	Value
<vector>	<section>	True if this section contains a vector table (C55x only)
<version>	<compile_unit>	DWARF version
	<file_header>	Version ID; structure version of this COFF file
<virtual_addr>	<reloc_entry>	Virtual address to be relocated
	<section>	Virtual (load) address of section
<word_size>	<reloc_entry>	Number of address-sized units containing the relocation field
<xml_version>	<dwarf>	Version of the DWARF XML language
	<ti_coff>	Version of the COFF XML language

10.3 Example XML Consumer

In this section, we present an example of a small application that uses the XML output of `ofd6x` to calculate the size of the executable code contained in an object file.

The example contains three source files: `codesize.cpp`, `xml.h`, and `xml.cpp`. When compiled into an executable named `codesize`, it can be used with `ofd6x` from the command line as follows:

```
% ofd6x -x a.out | codesize

Code Section Name: .text
Code Section Size: 44736

Code Section Name: .text2
Code Section Size: 64

Code Section Name: .text3
Code Section Size: 64

Total Code Size: 44864
```

10.3.1 The Main Application

The `codesize.cpp` file contains the main application for the object file display utility example.

```

//*****
// CODESIZE.CPP - An example application that calculates the size of the      *
// executable code in an object file using the XML output                    *
// of the OFD utility.                                                       *
//*****
#include "xml.h"
#include <iostream>

using namespace std;

static void parse_XML_prolog(istream &in);

//*****
// main() - List the names and sizes of the code sections (in octets), and   *
//          output the total code size.                                       *
//*****
int main()
{
    //-----
    // Build our tree of XML Entities from standard input (See xml.{cpp,h} for -
    // the definition of the XMLEntity object).                               -
    //-----
    parse_XML_prolog(cin);
    XMLEntity *root = new XMLEntity(cin);

```

```

//-----
// Fetch the XML Entities of the section roots. In other words, get a
// list of all the XMLEntity sub-trees named "section" that are in the
// context of "ofd->object_file->ti_coff", where "ofd" is the root of our
// XML document.
//-----
CEntityList query_result;
const char *section_query[] =
    { "ofd", "object_file", "ti_coff", "section", NULL };

query_result = root->query(section_query);

//-----
// Iterate over the section Entities, looking for code sections.
//-----
CEntityList_CIt pit;
unsigned long total_code_size = 0;

for (pit = query_result.begin(); pit != query_result.end(); ++pit)
{
    //-----
    // Query for the name, text, and raw_data_size sub-entities of each
    // section. XMLEntity::query() always returns a list, even if there
    // will only ever be a maximum of one result. If the tag is not
    // found, an empty list is returned.
    //-----
    const char *section_name_query[] = { "section", "name",          NULL };
    const char *section_text_query[] = { "section", "text",          NULL };
    const char *section_size_query[] = { "section", "raw_data_size", NULL };

    CEntityList sname_l;
    CEntityList stext_l;
    CEntityList ssize_l;

    sname_l = (*pit)->query(section_name_query);
    stext_l = (*pit)->query(section_text_query);
    ssize_l = (*pit)->query(section_size_query);
    //-----
    // If a "text" flag was found, this is a code section. Output
    // the section name and size, and add its size to our total code size
    // counter.
    //-----
    if (stext_l.size() > 0)
    {
        unsigned long size;

        size = strtoul((*ssize_l.begin())->value().c_str(), NULL, 16);

        cout << "Code Section Name: " << (*sname_l.begin())->value() << endl;
        cout << "Code Section Size: " << size << endl;
        cout << endl;

        total_code_size += size;
    }
}

```

```

}

//-----
// Output the total code size, and clean up. -
//-----
cout << "Total Code Size: " << total_code_size << endl;
delete root;

return 0;
}

//*****
// parse_XML_prolog() - Parse the XML prolog, and throw it away. *
//*****
static void parse_XML_prolog(istream &in)
{
    char c;

    while (true)
    {
        //-----
        // Look for the next tag; if it is not an XML directive, we're done. -
        //-----
        for (in.get(c); c != '<' && !in.eof(); in.get(c))
            ; // empty body

        if (in.eof()) return;
        if (in.peek() != '?') { in.unget(); return; }

        //-----
        // Otherwise, read in the directive and continue. -
        //-----
        for (in.get(c); c != '>' && !in.eof(); in.get(c))
            ; // empty body
    }
}

```

10.3.2 xml.h Declaration of the XMLEntity Object

The xml.h file contains the declaration of the XMLEntity object for the codesize.cpp application.

```
//*****
// XML.H - Declaration of the XMLEntity object. *
//*****
#ifndef XML_H
#define XML_H
#include <list>
#include <string>

//*****
// Type Declarations. *
//*****
class XMLEntity;
typedef list<XMLEntity*>          EntityList;
typedef list<const XMLEntity*>   CEntityList;
typedef CEntityList::const_iterator CEntityList_CIt;
typedef EntityList::const_iterator EntityList_CIt;
8
//*****
// CLASS XMLENTITY - A Simplified XML Entity Object. *
//*****
class XMLEntity
{
public:
    XMLEntity (istream &in, XMLEntity *parent=NULL);
    ~XMLEntity ();
    const CEntityList query (const char **context) const;
    const string      &tag   () const { return tag_m;      }
    const string      &value () const { return value_m;    }

private:
    void parse_raw_tag (const string &raw_tag);
    void sub_query     (CEntityList &result, const char **context) const;

    string      tag_m;          // Tag Name
    string      value_m;        // Value
    XMLEntity *parent_m;        // Pointer to parent in XML hierarchy
    EntityList  children_m;     // List of children in XML hierarchy
};
#endif
```

10.3.3 xml.cpp Definition of the XMLEntity Object

The xml.cpp file contains the definition of the XMLEntity object for the codesize.cpp application.

```

//*****
// XML.CPP - Definition of the XMLEntity object. *
//*****
#include "xml.h"
#include <iostream>
#include <string>
#include <list>
#include <cstdlib>

//*****
// XMLEntity::query() - Return the list of XMLEntities a list that reside *
// in the given XML context. *
//*****
const CEntityList XMLEntity::query(const char **context) const
{
    CEntityList result;

    if (!*context) return result;

    sub_query(result, context);

    return result;
}

//*****
// XMLEntity::sub_query() - Recurse through the XML tree looking for a match *
// to the current query. *
//*****
void XMLEntity::sub_query(CEntityList &result, const char **context) const
{
    if (!context[0] || tag() != context[0]) return;

    if (!context[1])
        result.push_front(this);
    else
    {
        EntityList_CIt pit;

        for (pit = children_m.begin(); pit != children_m.end(); ++pit)
            (*pit)->sub_query(result, context+1);
    }
    return;
}

//*****
// XMLEntity::parse_raw_tag() - Cut out the tag name from the complete string *
// we found between the < > brackets. This throws out any attributes. *
//*****
void XMLEntity::parse_raw_tag(const string &raw_tag)

```

```

{
    string attribute;
    int    i;

    for (i = 0; i < raw_tag.size() && raw_tag[i] != ' '; ++i)
        tag_m += raw_tag[i];
}

//*****
// XMLEntity::XMLEntity() - Recursively construct a tree of XMLEntities from *
//                          the given input stream.                          *
//*****
XMLEntity::XMLEntity(istream &in, XMLEntity *parent) :
tag_m(""), value_m(""), parent_m(parent)
{
    string raw_tag;
    char   c;
    int    i;
    //-----
    // Read in the leading '<'.
    //-----
    in.get();

    //-----
    // Store the tag name and attributes in "raw_tag", then call
    // process_raw_tag() to separate the tag name from the attributes and
    // store it in tag_m.
    //-----
    for (in.get(c); c != '>' && c != '/' && !in.eof(); in.get(c))
        raw_tag += c;

    parse_raw_tag(raw_tag);

    //-----
    // If we're reading in an end-tag, read in the closing '>' and return.
    //-----
    if (c == '/') { in.get(c); return; }

    //-----
    // Otherwise, parse our value.
    //-----
    while (true)
    {
        //-----
        // Read in the closing '>', then start reading in characters and add
        // them to value_m. Stop when we hit the beginning of a tag.
        //-----
        for (in.get(c); c != '<'; in.get(c)) value_m += c;

        //-----
        // If we're reading in a start tag, parse in the entire entity, and
        // add it to our child list (recursive constructor call).
        //-----
        if (in.peek() != '/')
            {

```

```

        //-----
        // Put back the opening '<', since XMLEntity() expects to read it. -
        //-----
in.unget();
children_m.push_front(new XMLEntity(in, this));
}
//-----
// Otherwise, read in our end tag, and exit. -
//-----
else
{
    for (in.get(c); c != '>'; in.get(c))
        ; // empty body
    break;
}
}

//-----
// Strip off leading and trailing white space from our value. -
//-----
for (i = 0; i < value_m.size(); i++)
    if (value_m[i] != ' ' && value_m[i] != '\n') break;
value_m.erase(0, i);

for (i = value_m.size()-1; i >= 0; i--)
    if (value_m[i] != ' ' && value_m[i] != '\n') break;
value_m.erase(i+1, value_m.size()-i);
}

//*****
// XMLEntity::~XMLEntity() - Delete a XMLEntity object. *
//*****
XMLEntity::~XMLEntity()
{
    EntityList_CIt pit;

    for (pit = children_m.begin(); pit != children_m.end(); ++pit)
        delete (*pit);
}

```

10.4 Invoking the Name Utility

The name utility, *nm6x*, is used to print the list of names defined and referenced in a COFF object (.obj) or an executable file (.out). The value associated with the symbol and an indication of the kind of symbol is also printed.

To invoke the name utility, enter the following:

```
nm6x [-options] [input filename]
```

nm6x	is the command that invokes the name utility.
<i>input filename</i>	is a COFF object file (.obj) or an executable file (.out).
<i>options</i>	identifies the name utility options you want to use. Options are not case sensitive and can appear anywhere on the command line following the invocation. Precede each option with a hyphen (-). The name utility options are as follows:
-a	prints all symbols.
-c	also prints C_NULL symbols.
-d	also prints debug symbols.
-f	prepends file name to each symbol.
-g	prints only global symbols.
-h	shows the current help screen.
-l	produces a detailed listing of the symbol information.
-n	sorts symbols numerically rather than alphabetically.
-o<file>	outputs to the given file.
-p	causes the name utility to not sort any symbols.
-q	(quiet mode) suppresses the banner and all progress information.
-r	sorts symbols in reverse order.
-t	also prints tag information symbols.
-u	only prints undefined symbols.

10.5 Invoking the Strip Utility

The strip utility, *strip6x*, is used to remove symbol table and debugging information from object and executable files.

To invoke the strip utility, enter the following:

```
strip6x [-p] input filename [input filename]
```

strip6x	is the command that invokes the strip utility.
<i>input filename</i>	is a COFF object file (.obj) or an executable file (.out).
<i>options</i>	identifies the strip utility options you want to use. Options are not case sensitive and can appear anywhere on the command line following the invocation. Precede each option with a hyphen (-). The name utility option is as follows:
-p	removes all information not required for execution. This option causes more information to be removed than the default behavior, but the object file is left in a state that cannot be linked. This option should be used only with executable (.out) files.

When the strip utility is invoked, the input object files are replaced with the stripped version.



Hex Conversion Utility Description

The TMS320C6000™ assembler and linker create object files that are in common object file format (COFF). COFF is a binary object file format that encourages modular programming and provides powerful and flexible methods for managing code segments and target system memory.

Most EPROM programmers do not accept COFF object files as input. The hex conversion utility converts a COFF object file into one of several standard ASCII hexadecimal formats, suitable for loading into an EPROM programmer. The utility is also useful in other applications requiring hexadecimal conversion of a COFF object file (for example, when using debuggers and loaders).

The hex conversion utility can produce these output file formats:

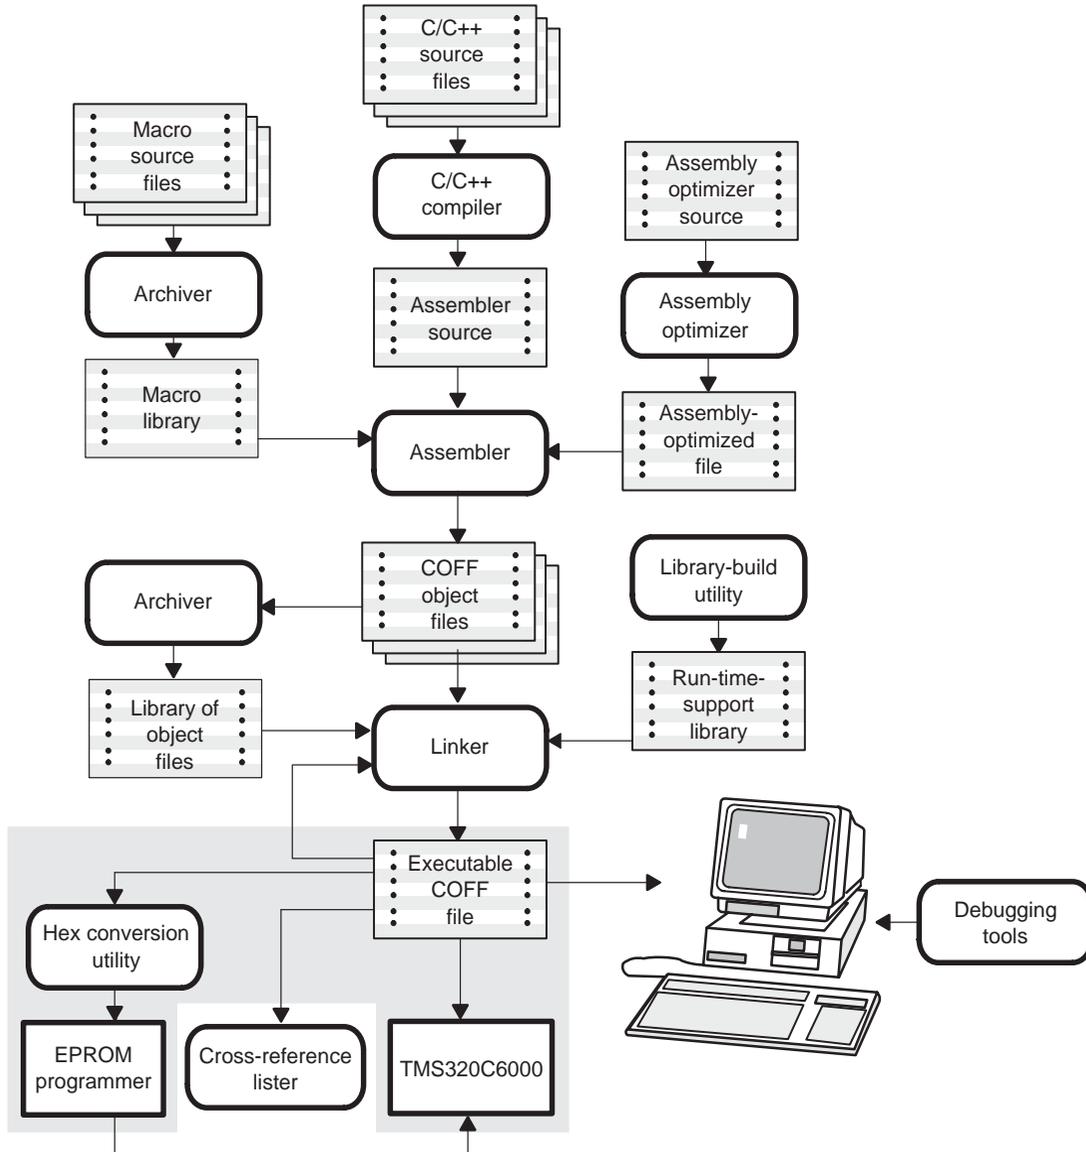
- ASCII-Hex, supporting 16-bit addresses
- Extended Tektronix (Tektronix)
- Intel MCS-86 (Intel)
- Motorola Exorciser (Motorola-S), supporting 16-bit addresses
- Texas Instruments SDSMAC (TI-Tagged), supporting 16-bit addresses

Topic	Page
11.1 The Hex Conversion Utility's Role in the Software Development Flow	11-2
11.2 Invoking the Hex Conversion Utility	11-3
11.3 Understanding Memory Widths	11-7
11.4 The ROMS Directive	11-13
11.5 The SECTIONS Directive	11-19
11.6 Excluding a Specified Section	11-21
11.7 Assigning Output Filenames	11-22
11.8 Image Mode and the -fill Option	11-24
11.9 Building a Table for an On-Chip Boot Loader	11-26
11.10 Controlling the ROM Device Address	11-32
11.11 Description of the Object Formats	11-34
11.12 Hex Conversion Utility Error Messages	11-40

11.1 The Hex Conversion Utility's Role in the Software Development Flow

Figure 11-1 highlights the role of the hex conversion utility in the software development process.

Figure 11-1. The Hex Conversion Utility in the TMS320C6000 Software Development Flow



11.2 Invoking the Hex Conversion Utility

There are two basic methods for invoking the hex conversion utility:

- ❑ **Specify the options and filenames on the command line.** The following example converts the file `firmware.out` into TI-Tagged format, producing two output files, `firm.lsb` and `firm.msb`.

```
hex6x -t firmware -o firm.lsb -o firm.msb
```

- ❑ **Specify the options and filenames in a command file.** You can create a batch file that stores command line options and filenames for invoking the hex conversion utility. The following example invokes the utility using a command file called `hexutil.cmd`:

```
hex6x hexutil.cmd
```

In addition to regular command line information, you can use the hex conversion utility `ROMS` and `SECTIONS` directives in a command file.

11.2.1 Invoking the Hex Conversion Utility From the Command Line

To invoke the hex conversion utility, enter:

```
hex6x [options] filename
```

- | | |
|-----------------|--|
| hex6x | is the command that invokes the hex conversion utility. |
| <i>options</i> | supplies additional information that controls the hex conversion process. You can use options on the command line or in a command file. Table 11–1 lists the basic options. <ul style="list-style-type: none"> ❑ All options are preceded by a hyphen and are not case sensitive. ❑ Several options have an additional parameter that must be separated from the option by at least one space. ❑ Options with multicharacter names must be spelled exactly as shown in this document; no abbreviations are allowed. ❑ Options are not affected by the order in which they are used. The exception to this rule is the <code>-q</code> (quiet) option, which must be used before any other options. |
| <i>filename</i> | names a COFF object file or a command file (for more information, see section 11.2.2, <i>Invoking the Hex Conversion Utility With a Command File</i> , on page 11-5). If you do not specify a filename, the utility prompts you for one. |

Table 11–1. Basic Hex Conversion Utility Options

General Options	Option	Description	Page
Control the overall operation of the hex conversion utility	-byte	Number output file locations by bytes rather than using target addressing	11-32
	-exclude <i>section_name</i>	Ignore specified section	11-21
	-map <i>filename</i>	Generate a map file	11-17
	-o <i>filename</i>	Specify an output filename	11-22
	-q	Run quietly (when used, it must appear <i>before</i> other options)	11-6
Image Options	Option	Description	Page
Create a continuous image of a range of target memory	-fill <i>value</i>	Fill holes with <i>value</i>	11-25
	-image	Specify image mode	11-24
	-zero	Reset the address origin to 0 in image mode	11-32
Memory Options	Option	Description	Page
Configure the memory widths for your output files	-memwidth <i>value</i>	Define the system memory word width (default 32 bits)	11-8
	-romwidth <i>value</i>	Specify the ROM device width (default depends on format used)	11-10
	-order L	Output file is in little endian format	11-12
	-order M	Output file is in big endian format	11-12
Output Options	Option	Description	Page
Specify the output format	-a	Select ASCII-Hex	11-35
	-i	Select Intel	11-36
	-m	Select Motorola-S	11-37
	-t	Select TI-Tagged	11-38
	-x	Select Tektronix (default)	11-39

Table 11–1. Basic Hex Conversion Utility Options (Continued)

Boot Options	Option	Description	Page
Control the boot loader	-boot	Convert all initialized sections into bootable form (use instead of a SECTIONS directive)	11-28
	-bootorg	Specify the source address of the boot loader table	11-28
	-bootsection <i>sectname value</i>	Specify which section contains the boot routine and where it should be placed.	11-29
	-e <i>value</i>	Specify the entry point at which to begin execution after boot loading. The <i>value</i> can be an address or a global symbol.	11-28

11.2.2 Invoking the Hex Conversion Utility With a Command File

A command file is useful if you plan to invoke the utility more than once with the same input files and options. It is also useful if you want to use the ROMS and SECTIONS hex conversion utility directives to customize the conversion process.

Command files are ASCII files that contain one or more of the following:

- Options and filenames.** These are specified in a command file in exactly the same manner as on the command line.
- ROMS directive.** The ROMS directive defines the physical memory configuration of your system as a list of address-range parameters. (For more information, see section 11.4, *The ROMS Directive*, on page 11-13.)
- SECTIONS directive.** The hex conversion utility SECTIONS directive specifies which sections from the COFF object file are selected. (For more information, see section 11.5, *The SECTIONS Directive*, on page 11-19.)
- Comments.** You can add comments to your command file by using the `/*` and `*/` delimiters. For example:

```
/* This is a comment. */
```

To invoke the utility and use the options you defined in a command file, enter:

```
hex6x command_filename
```

You can also specify other options and files on the command line. For example, you could invoke the utility by using both a command file and command line options:

```
hex6x firmware.cmd -map firmware.mxp
```

The order in which these options and filenames appear is not important. The utility reads all input from the command line and all information from the command file before starting the conversion process. However, if you are using the `-q` option, *it must appear as the first option on the command line or in a command file.*

The **-q option** suppresses the hex conversion utility's normal banner and progress information.

- Assume that a command file named `firmware.cmd` contains these lines:

```
firmware.out /* input file */
-t          /* TI-Tagged */
-o firm.lsb /* output file */
-o firm.msb /* output file */
```

You can invoke the hex conversion utility by entering:

```
hex6x firmware.cmd
```

- This example shows how to convert a file called `appl.out` into eight hex files in Intel format. Each output file is one byte wide and 4K bytes long.

```
appl.out /* input file */
-i /* Intel format */
-map appl.mxp /* map file */

ROMS
{
  ROW1: origin=0x00000000 len=0x4000 romwidth=8
        files={ appl.u0 appl.u1 appl.u2 appl.u3 }
  ROW2: origin=0x00004000 len=0x4000 romwidth=8
        files={ appl.u4 appl.u5 appl.u6 appl.u7 }
}

SECTIONS
{
  .text, .data, .cinit, .sect1, .vectors, .const:
}
```

11.3 Understanding Memory Widths

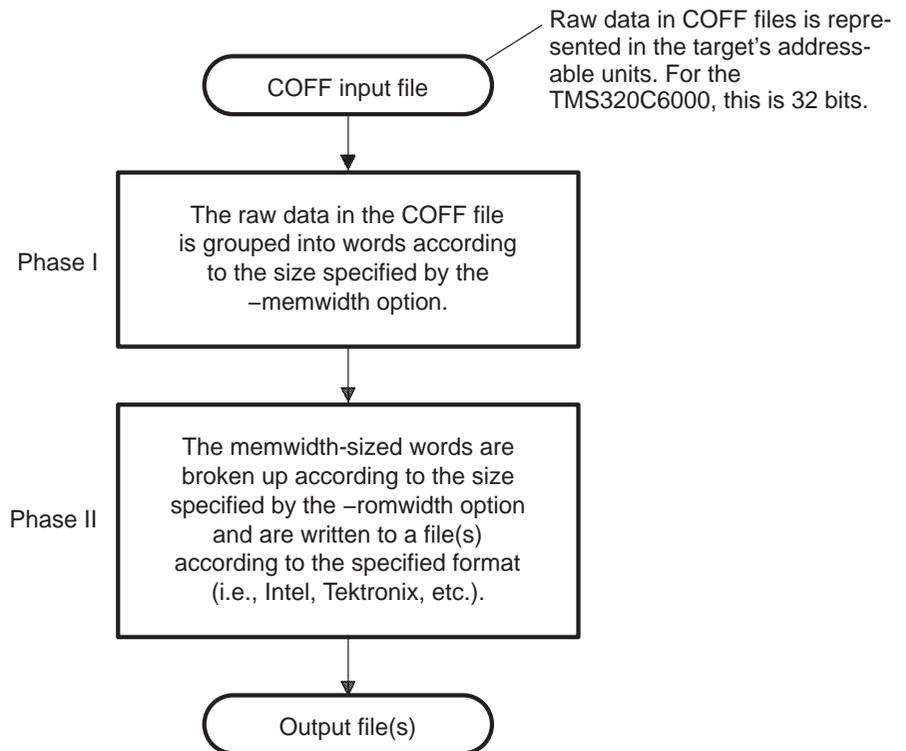
The hex conversion utility makes your memory architecture more flexible by allowing you to specify memory and ROM widths. In order to use the hex conversion utility, *you must understand how the utility treats word widths*. Three widths are important in the conversion process:

- Target width
- Memory width
- ROM width

The terms target word, memory word, and ROM word refer to a word of such a width.

Figure 11–2 illustrates the two separate and distinct phases of the hex conversion utility's process flow.

Figure 11–2. Hex Conversion Utility Process Flow



11.3.1 Target Width

Target width is the unit size (in bits) of the target processor's word. The unit size corresponds to the data bus size on the target processor. The width is fixed for each target and cannot be changed. The TMS320C6000 targets have a width of 32 bits.

11.3.2 Specifying the Memory Width

Memory width is the physical width (in bits) of the memory system. Usually, the memory system is physically the same width as the target processor width: a 32-bit processor has a 32-bit memory architecture. However, some applications require target words to be broken into multiple, consecutive, narrower memory words.

The hex conversion utility defaults memory width to the target width (in this case, 32 bits).

You can change the memory width by:

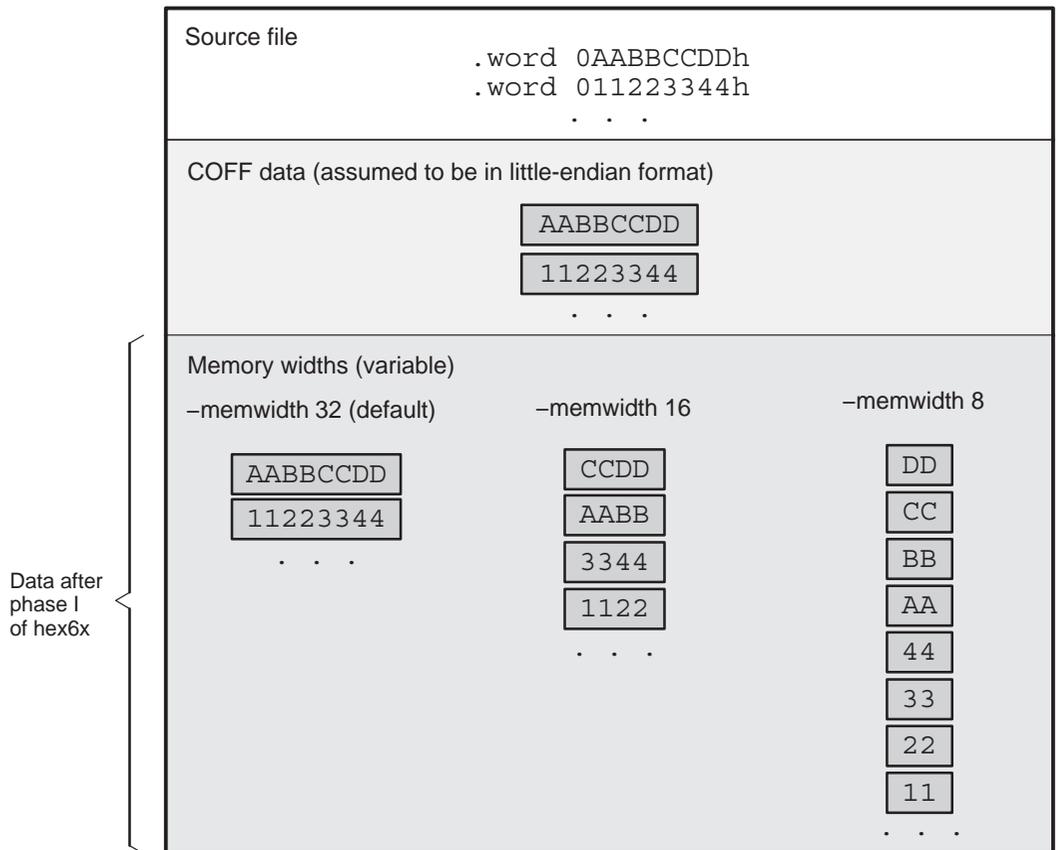
- Using the **-memwidth** option. This changes the memory width value for the entire file.
- Setting the **memwidth** parameter of the ROMS directive. This changes the memory width value for the address range specified in the ROMS directive and overrides the **-memwidth** option for that range. See section 11.4, *The ROMS Directive*, on page 11-13.

For both methods, use a value that is a power of 2 greater than or equal to 8.

You should change the memory width default value of 32 only when you need to break single target words into consecutive, narrower memory words.

Figure 11-3 demonstrates how the memory width is related to COFF data.

Figure 11–3. COFF Data and Memory Widths



11.3.3 Partitioning Data Into Output Files

ROM width specifies the physical width (in bits) of each ROM device and corresponding output file (usually one byte or eight bits). The ROM width determines how the hex conversion utility partitions the data into output files. After the COFF data is mapped to the memory words, the memory words are broken into one or more output files. The number of output files is determined by the following formulas:

- If memory width \geq ROM width:
number of files = memory width \div ROM width
- If memory width $<$ ROM width:
number of files = 1

For example, for a memory width of 32, you could specify a ROM width value of 32 and get a single output file containing 32-bit words. Or you can use a ROM width value of 16 to get two files, each containing 16 bits of each word.

The default ROM width that the hex conversion utility uses depends on the output format:

- All hex formats except TI-Tagged are configured as lists of 8-bit bytes; the default ROM width for these formats is 8 bits.
- TI-Tagged is a 16-bit format; the default ROM width for TI-Tagged is 16 bits.

Note: The TI-Tagged Format Is 16 Bits Wide

You cannot change the ROM width of the TI-Tagged format. The TI-Tagged format supports a 16-bit ROM width only.

You can change ROM width (except for TI-Tagged format) by:

- Using the **-romwidth** option. This option changes the ROM width value for the entire COFF file.
- Setting the **romwidth** parameter of the ROMS directive. This parameter changes the ROM width value for a specific ROM address range and overrides the **-romwidth** option for that range. See section 11.4, *The ROMS Directive*, on page 11-13.

For both methods, use a value that is a power of 2 greater than or equal to 8.

If you select a ROM width that is wider than the natural size of the output format (16 bits for TI-Tagged or 8 bits for all others), the utility simply writes multibyte fields into the file.

Figure 11-4 illustrates how the COFF data, memory, and ROM widths are related to one another.

Memory width and ROM width are used only for grouping the COFF data; they do not represent values. Thus, the byte ordering of the COFF data is maintained throughout the conversion process. To refer to the partitions within a memory word, the bits of the memory word are always numbered from right to left as follows:

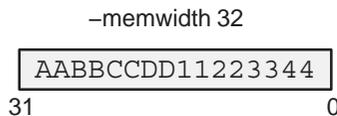
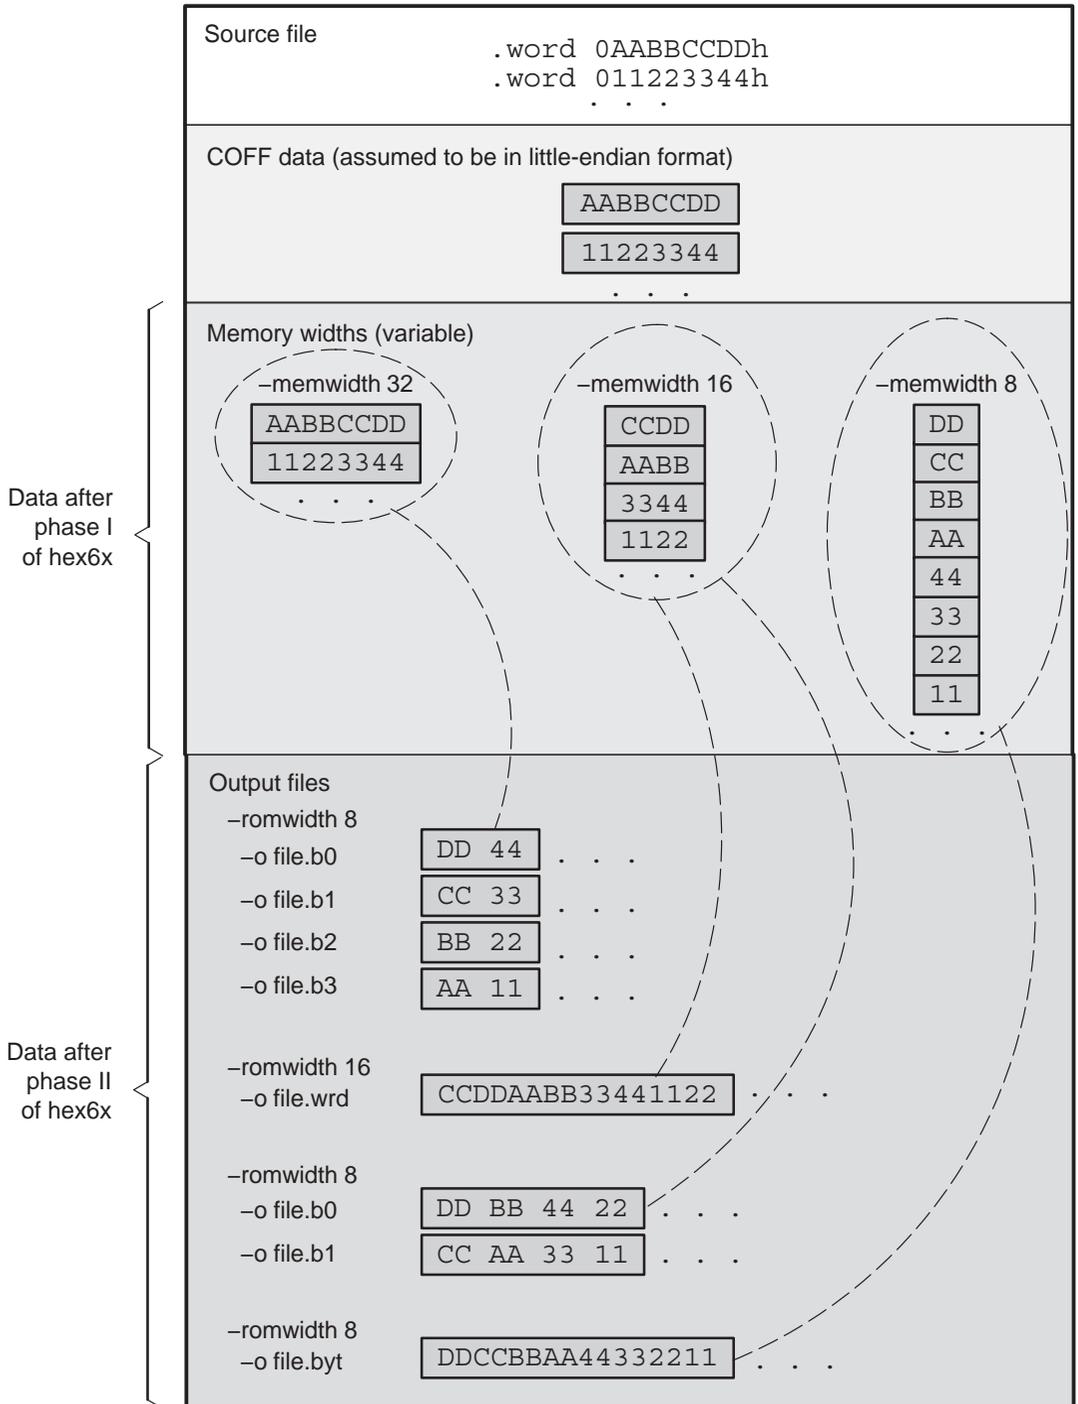


Figure 11–4. Data, Memory, and ROM Widths



11.3.4 Specifying Word Order for Output Words

There are two ways to split a wide word into consecutive memory locations in the same hex conversion utility output file:

- order M** specifies **big-endian** ordering, in which the most significant part of the wide word occupies the first of the consecutive locations
- order L** specifies **little-endian** ordering, in which the the least significant part of the wide word occupies the first of the consecutive locations

By default, the utility uses little-endian format. Unless your boot loader program expects big-endian format, avoid using `-order M`.

Note: When the `-order` Option Applies

- This option applies only when you use a memory width with a value of 32 (`-memwidth32`). Otherwise, the hex utility does not have access to the entire 32-bit word and cannot perform the byte swapping necessary to change the endianness; `-order` is ignored.
 - This option does not affect the way memory words are split into output files. Think of the files as a set: the set contains a least significant file and a most significant file, but there is no ordering over the set. When you list filenames for a set of files, you *always* list the least significant first, regardless of the `-order` option.
-

11.4 The ROMS Directive

The ROMS directive specifies the physical memory configuration of your system as a list of address-range parameters.

Each address range produces one set of files containing the hex conversion utility output data that corresponds to that address range. Each file can be used to program one single ROM device.

The ROMS directive is similar to the MEMORY directive of the TMS320C6000 linker: both define the memory map of the target address space. Each line entry in the ROMS directive defines a specific address range. The general syntax is:

```

ROMS
{
    romname: [origin=value,] [length=value,] [romwidth=value,]
             [memwidth=value,] [fill=value,]
             [files={filename1, filename2, ...}]

    romname: [origin=value,] [length=value,] [romwidth=value,]
             [memwidth=value,] [fill=value,]
             [files={filename1, filename2, ...}]

    ...
}

```

ROMS begins the directive definition.

romname identifies a memory range. The name of the memory range can be one to eight characters in length. The name has no significance to the program; it simply identifies the range. (Duplicate memory range names are allowed.)

origin specifies the starting address of a memory range. It can be entered as *origin*, *org*, or *o*. The associated value must be a decimal, octal, or hexadecimal constant. If you omit the origin value, the origin defaults to 0.

The following table summarizes the notation you can use to specify a decimal, octal, or hexadecimal constant:

Constant	Notation	Example
Hexadecimal	0x prefix or h suffix	0x77 or 077h
Octal	0 prefix	077
Decimal	No prefix or suffix	77

- length** specifies the length of a memory range as the physical length of the ROM device. It can be entered as `length`, `len`, or `l`. The value must be a decimal, octal, or hexadecimal constant. If you omit the length value, it defaults to the length of the entire address space.
- romwidth** specifies the physical ROM width of the range in bits (see section 11.3.3, *Partitioning Data Into Output Files*, on page 11-9). Any value you specify here overrides the `-romwidth` option. The value must be a decimal, octal, or hexadecimal constant that is a power of 2 greater than or equal to 8.
- memwidth** specifies the memory width of the range in bits (see section 11.3.2, *Specifying the Memory Width*, on page 11-8). Any value you specify here overrides the `-memwidth` option. The value must be a decimal, octal, or hexadecimal constant that is a power of 2 greater than or equal to 8. *When using the memwidth parameter, you must also specify the paddr parameter for each section in the SECTIONS directive.* (See section 11.5, *The SECTIONS Directive*, on page 11-19.)
- fill** specifies a fill value to use for the range. In image mode, the hex conversion utility uses this value to fill any holes between sections in a range. A hole is an area between the input sections that comprises an output section that contains no actual code or data.
- The fill value must be a decimal, octal, or hexadecimal constant with a width equal to the target width. Any value you specify here overrides the `-fill` option. When using `fill`, you must also use the `-image` command line option. See section 11.8.2, *Specifying a Fill Value*, on page 11-25.
- files** identifies the names of the output files that correspond to this range. Enclose the list of names in curly braces and order them from *least significant* to *most significant* output file, where the bits of the memory word are numbered from right to left.
- The number of file names must equal the number of output files that the range generates. To calculate the number of output files, refer to section 11.3.3, *Partitioning Data Into Output Files*, on page 11-9. The utility warns you if you list too many or too few filenames.

Unless you are using the `-image` option, all of the parameters that define a range are optional; the commas and equal signs are also optional. A range with no origin or length defines the entire address space. In image mode, an origin and length are required for all ranges.

Ranges must not overlap and must be listed in order of ascending address.

11.4.1 When to Use the ROMS Directive

If you do not use a ROMS directive, the utility defines a single default range that includes the entire address space. This is equivalent to a ROMS directive with a single range without origin or length.

Use the ROMS directive when you want to:

- Program large amounts of data into fixed-size ROMs.** When you specify memory ranges corresponding to the length of your ROMs, the utility automatically breaks the output into blocks that fit into the ROMs.
- Restrict output to certain segments.** You can also use the ROMS directive to restrict the conversion to a certain segment or segments of the target address space. The utility does not convert the data that falls outside of the ranges defined by the ROMS directive. Sections can span range boundaries; the utility splits them at the boundary into multiple ranges. If a section falls completely outside any of the ranges you define, the utility does not convert that section and issues no messages or warnings. In this way, you can exclude sections without listing them by name with the `SECTIONS` directive. However, if a section falls partially in a range and partially in unconfigured memory, the utility issues a warning and converts only the part within the range.
- Use image mode.** When you use the `-image` option, you must use a ROMS directive. Each range is filled completely so that each output file in a range contains data for the whole range. Holes before, between, or after sections are filled with the fill value from the ROMS directive, with the value specified with the `-fill` option, or with the default value of 0.

11.4.2 An Example of the ROMS Directive

The ROMS directive in Example 11–1 shows how 16K bytes of 16-bit memory could be partitioned for two 8K-byte × 8-bit EPROMs. Figure 11–5 illustrates the input and output files.

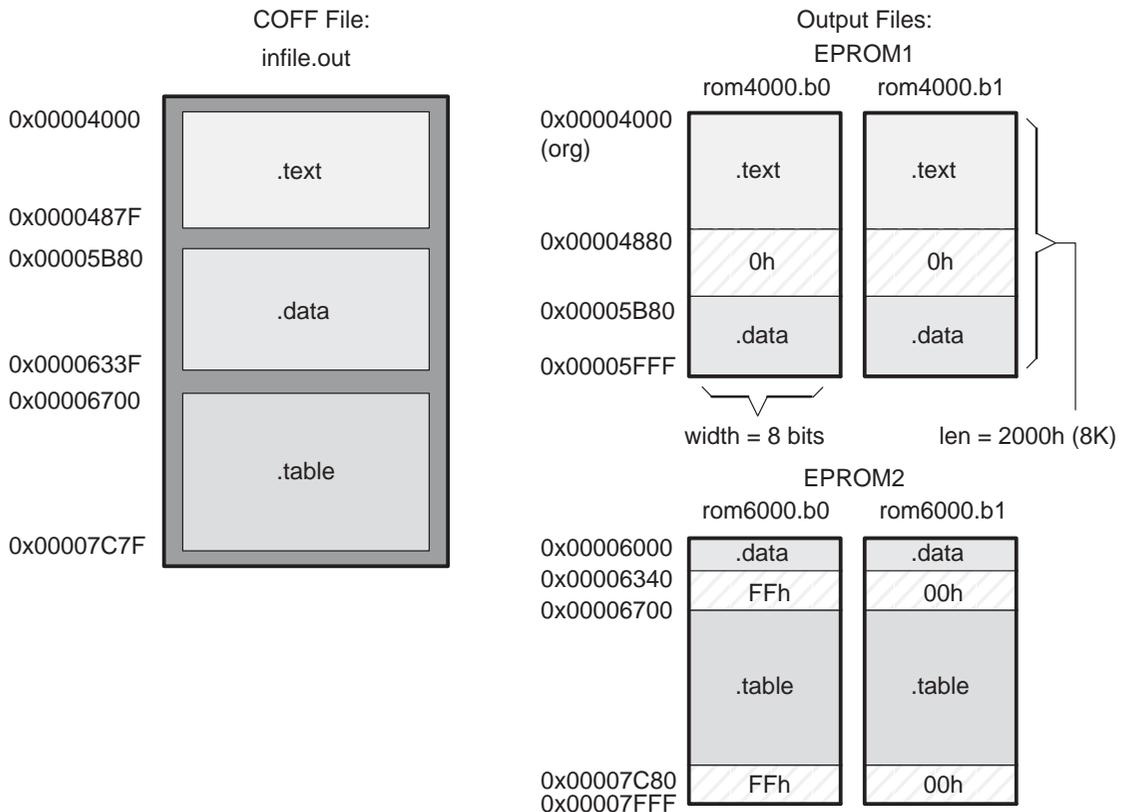
Example 11–1. A ROMS Directive Example

```
infile.out
-image
-memwidth 16

ROMS
{
  EPROM1: org = 0x00004000, len = 0x2000, romwidth = 8
        files = { rom4000.b0, rom4000.b1}

  EPROM2: org = 0x00006000, len = 0x2000, romwidth = 8,
        fill = 0xFF00FF00,
        files = { rom6000.b0, rom6000.b1}
}
```

Figure 11–5. The infile.out File Partitioned Into Four Output Files



The map file (specified with the `-map` option) is advantageous when you use the ROMS directive with multiple ranges. The map file shows each range, its parameters, names of associated output files, and a list of contents (section names and fill values) broken down by address. Example 11–2 is a segment of the map file resulting from the example in Example 11–1.

Example 11–2. Map File Output From Example 11–1 Showing Memory Ranges

```
-----
00004000..00005fff Page=0 Width=8 "EPROM1"
-----
      OUTPUT FILES:   rom4000.b0   [b0..b7]
                      rom4000.b1   [b8..b15]

      CONTENTS: 00004000..0000487f .text
                 00004880..00005b7f FILL = 00000000
                 00005b80..00005fff .data
-----
00006000..00007fff Page=0 Width=8 "EPROM2"
-----
      OUTPUT FILES:   rom6000.b0   [b0..b7]
                      rom6000.b1   [b8..b15]

      CONTENTS: 00006000..0000633f .data
                 00006340..000066ff FILL = ff00ff00
                 00006700..00007c7f .table
                 00007c80..00007fff FILL = ff00ff00
```

EPROM1 defines the address range from 0x00004000 through 0x00005FFF. The range contains the following sections:

This section ...	Has this range ...
.text	0x00004000 through 0x0000487F
.data	0x00005B80 through 0x00005FFF

The rest of the range is filled with 0h (the default fill value). The data from this range is converted into two output files:

- rom4000.b0 contains bits 0 through 7
- rom4000.b1 contains bits 8 through 15

EPROM2 defines the address range from 0x00006000 through 0x00007FFF. The range contains the following sections:

This section ...	Has this range ...
.data	0x00006000 through 0x0000633F
.table	0x00006700 through 0x00007C7F

The rest of the range is filled with 0xFF00FF00 (from the specified fill value). The data from this range is converted into two output files:

- rom6000.b0 contains bits 0 through 7
- rom6000.b1 contains bits 8 through 15

11.5 The SECTIONS Directive

You can convert specific sections of the COFF file by name with the hex conversion utility SECTIONS directive. You can also specify those sections that you want to locate in ROM at a different address than the *load* address specified in the linker command file. If you:

- Use a SECTIONS directive, the utility converts only the sections that you list in the directive and ignores all other sections in the COFF file
- Do not use a SECTIONS directive, the utility converts all initialized sections that fall within the configured memory. The TMS320C6000 compiler-generated initialized sections are `.text`, `.const`, and `.cinit`

Uninitialized sections are *never* converted, whether or not you specify them in a SECTIONS directive.

Note: Sections Generated by the C/C++ Compiler

The TMS320C6000 C/C++ compiler automatically generates these sections:

- Initialized sections:** `.text`, `.const`, `.cinit`, and `.switch`
- Uninitialized sections:** `.bss`, `.stack`, and `.system`

Use the SECTIONS directive in a command file. (For more information, see section 11.2.2, *Invoking the Hex Conversion Utility With a Command File*, on page 11-5.) The general syntax for the SECTIONS directive is:

```
SECTIONS
{
  sname[:] [paddr=value]
  sname[:] [paddr=boot]
  sname[:] [boot]
  ...
}
```

SECTIONS begins the directive definition.

sname identifies a section in the COFF input file. If you specify a section that does not exist, the utility issues a warning and ignores the name.

paddr=value specifies the physical ROM address at which this section should be located. This value overrides the section load address given by the linker. This value must be a decimal, octal, or hexadecimal constant. It can also be the word **boot** (to indicate a boot table section for use with a boot loader). *If your file contains multiple sections, and if one section uses a paddr parameter, then all sections must use a paddr parameter.*

boot configures a section for loading by a boot loader. This is equivalent to using **paddr=boot**. Boot sections have a physical address determined by the location of the boot table. The origin of the boot table is specified with the `-bootorg` option.

For more similarity with the linker's SECTIONS directive, you can use colons after the section names (in place of the equal sign on the boot keyboard). For example, the following statements are equivalent:

```
SECTIONS { .text: .data: boot }
SECTIONS { .text: .data = boot }
```

In the example below, the COFF file contains six initialized sections: `.text`, `.data`, `.const`, `.vectors`, `.coeff`, and `.tables`. Suppose you want only `.text` and `.data` to be converted. Use a SECTIONS directive to specify this:

```
SECTIONS { .text: .data: }
```

To configure both of these sections for boot loading, add the boot keyword:

```
SECTIONS { .text = boot .data = boot }
```

Note: Using the `-boot` Option and the SECTIONS Directive

When you use the SECTIONS directive with the boot table (`-boot`) option, the `-boot` option is ignored. You must explicitly specify any boot sections in the SECTIONS directive. For more information about `-boot` and other command line options associated with boot tables, see section 11.2, Invoking the Hex Conversion Utility, found on page 11-3.

11.6 Excluding a Specified Section

The `-exclude section_name` option can be used to inform the hex utility to ignore the specified section. If a `SECTIONS` directive is used, it overrides the `-exclude` option.

For example, if a `SECTIONS` directive containing the section name `mysect` is used and an `-exclude mysect` is specified, the `SECTIONS` directive takes precedence and `mysect` is not excluded.

The `-exclude` option has a limited wildcard capability. The `*` character can be placed at the beginning or end of the name specifier to indicate a suffix or prefix, respectively. For example, `-exclude sect*` disqualifies all sections that begin with the characters `sect`.

If you specify the `-exclude` option on the command line with the `*` wildcard, enter quotes around the section name and wildcard. For example, `-exclude"sect*"`. Using quotes prevents the `*` from being interpreted by the hex conversion utility. If `-exclude` is in a command file, then the quotes should not be specified.

11.7 Assigning Output Filenames

When the hex conversion utility translates your COFF object file into a data format, it partitions the data into one or more output files. When multiple files are formed by splitting memory words into ROM words, *filenames are always assigned in order from least to most significant*, where bits in the memory words are numbered from right to left. This is true, regardless of target or COFF endian ordering.

The hex conversion utility follows this sequence when assigning output filenames:

- 1) **It looks for the ROMS directive.** If a file is associated with a range in the ROMS directive and you have included a list of files (`files = { . . }`) on that range, the utility takes the filename from the list.

For example, assume that the target data is 32-bit words being converted to four files, each eight bits wide. To name the output files using the ROMS directive, you could specify:

```
ROMS
{
  RANGE1: romwidth=8, files={ xyz.b0 xyz.b1 xyz.b2 xyz.b3 }
}
```

The utility creates the output files by writing the least significant bits to `xyz.b0` and the most significant bits to `xyz.b3`.

- 2) **It looks for the `-o` options.** You can specify names for the output files by using the `-o` option. If no filenames are listed in the ROMS directive and you use `-o` options, the utility takes the filename from the list of `-o` options. The following line has the same effect as the example above using the ROMS directive:

```
-o xyz.b0 -o xyz.b1 -o xyz.b2 -o xyz.b3
```

If both the ROMS directive and `-o` options are used together, the ROMS directive overrides the `-o` options.

3) **It assigns a default filename.** If you specify no filenames or fewer names than output files, the utility assigns a default filename. A default filename consists of the base name from the COFF input file plus a 2- to 3-character extension. The extension has three parts:

a) A format character, based on the output format:

- a** for ASCII-Hex
- i** for Intel
- m** for Motorola-S
- t** for TI-Tagged
- x** for Tektronix

See section 11.11, *Description of the Object Formats*, on page 11-34 for more information.

b) The range number in the ROMS directive. Ranges are numbered starting with 0. If there is no ROMS directive, or only one range, the utility omits this character.

c) The file number in the set of files for the range, starting with 0 for the least significant file.

For example, assume `coff.out` is for a 32-bit target processor and you are creating Intel format output. With no output filenames specified, the utility produces four output files named `coff.i0`, `coff.i1`, `coff.i2`, `coff.i3`.

If you include the following ROMS directive when you invoke the hex conversion utility, you would have eight output files:

```
ROMS
{
  range1: o = 0x00001000 l = 0x1000
  range2: o = 0x00002000 l = 0x1000
}
```

These output files ...	Contain data in these locations ...
<code>coff.i00</code> , <code>coff.i01</code> , <code>coff.i01</code>	0x00001000 through 0x00001FFF
<code>coff.i02</code> , <code>coff.i03</code>	0x00002000 through 0x00002FFF

11.8 Image Mode and the `-fill` Option

This section points out the advantages of operating in image mode and describes how to produce output files with a precise, continuous image of a target memory range.

11.8.1 Generating a Memory Image

With the `-image` option, the utility generates a memory image by completely filling all of the mapped ranges specified in the ROMS directive.

A COFF file consists of blocks of memory (sections) with assigned memory locations. Typically, all sections are not adjacent: there are holes between sections in the address space for which there is no data. When such a file is converted *without* the use of image mode, the hex conversion utility bridges these holes by using the address records in the output file to skip ahead to the start of the next section. In other words, there may be discontinuities in the output file addresses. Some EPROM programmers do not support address discontinuities.

In image mode, there are no discontinuities. Each output file contains a continuous stream of data that corresponds exactly to an address range in target memory. Any holes before, between, or after sections are filled with a fill value that you supply.

An output file converted by using image mode still has address records, because many of the hexadecimal formats require an address on each line. However, in image mode, these addresses are always contiguous.

Note: Defining the Ranges of Target Memory

If you use image mode, you must also use a ROMS directive. In image mode, each output file corresponds directly to a range of target memory. You must define the ranges. If you do not supply the ranges of target memory, the utility tries to build a memory image of the entire target processor address space—potentially a huge amount of output data. To prevent this situation, the utility requires you to explicitly restrict the address space with the ROMS directive.

11.8.2 Specifying a Fill Value

The `-fill` option specifies a value for filling the holes between sections. The fill value must be specified as an integer constant following the `-fill` option. The width of the constant is assumed to be that of a word on the target processor. For example, specifying `-fill 0FFFFh` results in a fill pattern of `0000FFFFh`. The constant value is not sign extended.

The hex conversion utility uses a default fill value of 0 if you do not specify a value with the fill option. *The `-fill` option is valid only when you use `-image`; otherwise, it is ignored.*

11.8.3 Steps to Follow in Using Image Mode

- Step 1:** Define the ranges of target memory with a ROMS directive. See section 11.4, *The ROMS Directive*, on page 11-13 for details.
- Step 2:** Invoke the hex conversion utility with the `-image` option. You can optionally use the `-zero` option to reset the address origin to 0 for each output file. If you do not specify a fill value with the ROMS directive and you want a value other than the default of 0, use the `-fill` option.

11.9 Building a Table for an On-Chip Boot Loader

On the C621x, C671x, and C64x devices, a ROM boot process is supported where the EDMA copies 1K bytes from the beginning of CE1 (EMIFB CE1 on C64x) to address 0, using default ROM timings. After the transfer, the CPU begins executing from address 0. In this mode a second level boot load typically occurs, due to the limited amount of memory copied at boot.

The hex conversion utility supports the second level boot loader by automatically building the boot table.

11.9.1 Description of the Boot Table

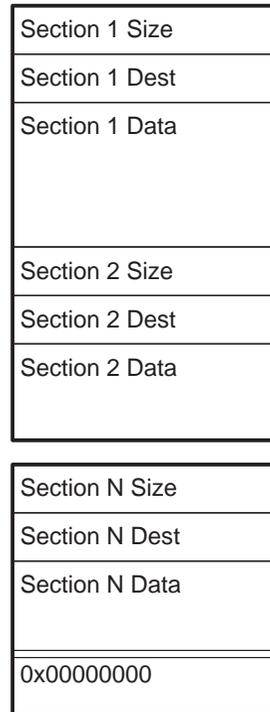
The input for a boot loader is the boot table. The boot table contains records that instruct the boot loader to copy blocks of data contained in the table to specified destination addresses. The hex conversion utility automatically builds the boot table for the boot loader. Using the utility, you specify the COFF sections you want the boot loader to initialize through the boot table, the table location, and the name of the section containing the boot loader and where it should be located. The hex conversion utility builds a complete image of the table and converts it into hexadecimal in the output files. Then, you can burn the table into ROM.

11.9.2 The Boot Table Format

The boot table format is simple. There is a header record containing a 4 byte field that indicates where the boot loader should branch after it has completed copying data. After the header, each COFF section that is to be included in the boot table will have a:

- 4 byte field containing the size of the section,
- 4 byte field containing the destination address for the copy, and
- The actual data to be copied.

Multiple sections can be entered; a termination block containing a 4 byte field of zeros follows the last COFF section.



11.9.3 How to Build the Boot Table

Table 11–2 summarizes the hex conversion utility options available for the boot loader.

Table 11–2. Boot-Loader Options

Option	Description
-boot	Convert all sections into bootable form (use instead of a SECTIONS directive)
-bootorg <i>value</i>	Specify the source address of the boot loader table.
-bootsection <i>sectname value</i>	Specify the section name <i>sectname</i> containing the boot loader routine. The <i>value</i> argument tells the hex utility where to place the boot loader routine.
-e <i>value</i>	Specify the entry point at which to begin execution after boot loading. The <i>value</i> can be an address or a global symbol.

11.9.3.1 Building the Boot Table

To build the boot table, follow these steps:

Step 1: Link the file. Each block of the boot table data corresponds to an initialized section in the COFF file. Uninitialized sections are not converted by the hex conversion utility (see Section 11.5, *The SECTIONS Directive*, on page 11-19).

You must link into your application a boot loader routine that will read the boot table and perform the copy operations. It should be linked to its eventual run-time address.

When you select a section for placement in a boot-loader table, the hex conversion utility places the section's *load address* in the destination address field for the block in the boot table. The section content is then treated as raw data for that block. *The hex conversion utility does not use the section run address.* When linking, you need not worry about the ROM address or the construction of the boot table—the hex conversion utility handles this.

Step 2: Identify the bootable sections. You can use the `-boot` option to tell the hex conversion utility to configure all sections for boot loading. Or, you can use a `SECTIONS` directive to select specific sections to be configured (see Section 11.5, *The SECTIONS Directive*, on page 11-19). Note that if you use a `SECTIONS` directive, the `-boot` option is ignored.

Step 3: Set the ROM address of the boot table. Use the `-bootorg` option to set the source address of the complete table. For example, if you are using the C6711 and booting from memory location 0x90000400, specify `-bootorg 0x90000400`. The address field for the boot table in the the hex conversion utility output file will then start at 0x90000400.

If you do not use the `-bootorg` option at all, the utility places the table at the origin of the first memory range in a `ROMS` directive. If you do not use a `ROMS` directive, the table will start at the first section load address.

Step 4: Set boot-loader-specific options. Set entry point. If `-e` is not used to set the entry point, then it will default to the entry point indicated in the COFF object file.

Step 5: Describe the boot routine. If the boot option is used, then you should use the `-bootsection` option to indicate to the hex utility which COFF section contains the boot routine. This option will prevent the boot routine from being in the boot table. The `-bootsection` option also indicates to the hex utility where the routine should be placed in ROM. For the C621x, C671x, and C64x devices, this address would typically be the beginning of CE1 (EMIFB CE1 on C64x).

This option is ignored if `-boot` is not used.

When the `SECTIONS` directive is used to explicitly identify which sections should exist in the boot table, use the `PADDR` section option to indicate where the boot routine section will exist.

Step 6: Describe your system memory configuration. See Section 11.3, *Understanding Memory Widths*, on page 11-7 and Section 11.4, *The ROMS Directive*, on page 11-13 for details.

11.9.3.2 Leaving Room for the Boot Table

The complete boot table is similar to a single section containing all of the header records and data for the boot loader. The address of this “section” is the boot table origin. As part of the normal conversion process, the hex conversion utility converts the boot table to hexadecimal format and maps it into the output files like any other section.

Be sure to leave room in your system memory for the boot table, especially when you are using the `ROMS` directive. The boot table cannot overlap other nonboot sections or unconfigured memory. Usually, this is not a problem; typically, a portion of memory in your system is reserved for the boot table. Simply configure this memory as one or more ranges in the `ROMS` directive, and use the `-bootorg` option to specify the starting address.

11.9.3.3 Setting the Entry Point for the Boot Table

After the boot routine finishes copying data, it branches to the entry point defined in the COFF object file. By using the `-e` option with the hex conversion utility, you can set the entry point to a different address.

For example, if you want your program to start running at address 0123h after loading, specify `-e 0123h` on the command line or in a command file. You can determine the `-e` address by looking at the map file that the linker generates.

Note: Valid Entry Points

The value can be a constant, or it can be a symbol that is externally defined (for example, with a `.global`) in the assembly source.

11.9.4 Using the C6x Boot Loader

This subsection explains how to use the hex conversion utility with the boot loader for C6x devices through sample hex utility command files. Example 11–3 uses the SECTIONS directive to specify exactly which COFF sections will be placed in the boot table.

Example 11–3. Sample Command File for Booting From a C6x EPROM

```
abc.out          /* input file          */
-a              /* ascii format        */
-image         /* create complete Rom image */
-zero         /* reset address origin to 0 */
-memwidth 8    /* 8-bit memory        */
-map abchex.map /* create a hex map file  */
-bootorg 0x90000400 /* external memory boot */

ROMS
{
    FLASH: org=0x90000000, len=0x20000, romwidth=8, files={abc.hex}
}

SECTIONS
{
    .boot_load:  PADDR=0x90000000
    .text:       BOOT
    .cinit:      BOOT
    .const:      BOOT
}
```

Example 11–4 does not explicitly name the boot sections with the SECTIONS directive. Instead, it uses the `-boot` option to indicate that all initialized sections should be placed in the boot table. It also uses the `-bootsection` option to distinguish the section containing the boot routine.

Example 11–4. Alternative Sample Command File for Booting From a C6x EPROM

```
abc.out          /* input file          */
-a              /* ascii format          */
-image          /* create complete Rom image */
-zero          /* reset address origin to 0 */
-memwidth 8     /* 8-bit memory          */
-map abchex.map /* create a hex map file   */
-boot          /* create boot table       */
-bootorg 0x90000400 /* external memory boot   */
-bootsection .boot_load 0x90000000 /* give boot section & addr */
ROMS
{
  FLASH: org=0x90000000, len=0x20000, romwidth=8, files={abc.hex}
}
```

11.10 Controlling the ROM Device Address

The hex conversion utility output address field corresponds to the ROM device address. The EPROM programmer burns the data into the location specified by the hex conversion utility output file address field. The hex conversion utility offers some mechanisms to control the starting address in ROM of each section. However, many EPROM programmers offer direct control of the location in ROM in which the data is burned.

11.10.1 Controlling the Starting Address

Depending on whether or not you are using the boot loader, the hex conversion utility output file controlling mechanisms are different.

Non-boot loader mode. The address field of the hex conversion utility output file is controlled by the following mechanisms listed from low to high priority:

- 1) **The linker command file.** By default, the address field of the hex conversion utility output file is the load address (as given in the linker command file).
- 2) **The `paddr` parameter of the `SECTIONS` directive.** When the `paddr` parameter is specified for a section, the hex conversion utility bypasses the section load address and places the section in the address specified by `paddr`.
- 3) **The `-zero` option.** When you use the `-zero` option, the utility resets the address origin to 0 for each output file. Since each file starts at 0 and counts upward, any address records represent offsets from the beginning of the file (the address within the ROM) rather than actual target addresses of the data.

You must use the `-zero` option in conjunction with the `-image` option to force the starting address in each output file to be zero. If you specify the `-zero` option without the `-image` option, the utility issues a warning and ignores the `-zero` option.

Boot-Loader Mode. When the boot loader is used, the hex conversion utility places the different COFF sections that are in the boot table into consecutive memory locations. Each COFF section becomes a boot table block whose destination address is equal to the linker-assigned section load address.

In a boot table, the address field of the hex conversion utility output file is not related to the section load addresses assigned by the linker. The address fields of the boot table are simply offsets to the beginning of the table. The section load addresses assigned by the linker will be encoded into the boot table along with the size of the section and the data contained within the section. These addresses will be used to store the data into memory during the boot load process.

The beginning of the boot table defaults to the linked load address of the first bootable section in the COFF input file, unless you use one of the following mechanisms, listed here from low to high priority. Higher priority mechanisms override the values set by low priority options in an overlapping range.

- 1) **The ROM origin specified in the ROMS directive.** The hex conversion utility places the boot table at the origin of the first memory range in a ROMS directive.
- 2) **The `-bootorg` option.** The hex conversion utility places the boot table at the address specified by the `-bootorg` option if you select boot loading from memory.

11.11 Description of the Object Formats

The hex conversion utility has options that identify each format and Table 11–3 specifies the format options. They are described in the following sections.

- You need to use only one of these options on the command line. If you use more than one option, the last one you list overrides the others.
- The default format is Tektronix (–x option).

Table 11–3. Options for Specifying Hex Conversion Formats

Option	Format	Address Bits	Default Width
–a	ASCII-Hex	16	8
–i	Intel	32	8
–m	Motorola-S	32	8
–t	TI-Tagged	16	16
–x	Tektronix	32	8

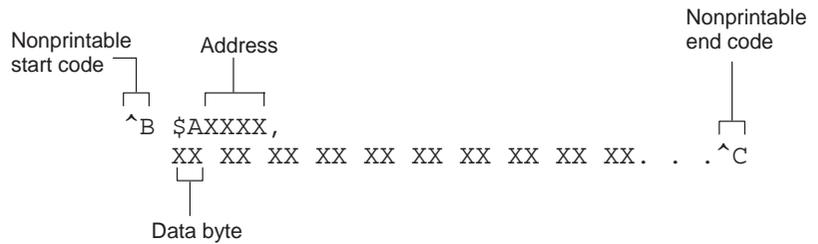
Address bits determine how many bits of the address information the format supports. Formats with 16-bit addresses support addresses up to 64K only. The utility truncates target addresses to fit in the number of available bits.

The **default width** determines the default output width of the format. You can change the default width by using the `–romwidth` option or by using the `romwidth` parameter in the `ROMS` directive. You cannot change the default width of the TI-Tagged format, which supports a 16-bit width only.

11.11.1 ASCII-Hex Object Format (`-a` Option)

The ASCII-Hex object format supports 16-bit addresses. The format consists of a byte stream with bytes separated by spaces. Figure 11–6 illustrates the ASCII-Hex format.

Figure 11–6. ASCII-Hex Object Format



The file begins with an ASCII STX character (ctrl-B, 02h) and ends with an ASCII ETX character (ctrl-C, 03h). Address records are indicated with \$AXXXX, in which XXXX is a 4-digit (16-bit) hexadecimal address. The address records are present only in the following situations:

- When discontinuities occur
- When the byte stream does not begin at address 0

You can avoid all discontinuities and any address records by using the `-image` and `-zero` options. This creates output that is simply a list of byte values.

11.11.2 Intel MCS-86 Object Format (-i Option)

The Intel object format supports 16-bit addresses and 32-bit extended addresses. Intel format consists of a 9-character (4-field) prefix—which defines the start of record, byte count, load address, and record type—the data, and a 2-character checksum suffix.

The 9-character prefix represents three record types:

Record Type	Description
00	Data record
01	End-of-file record
04	Extended linear address record

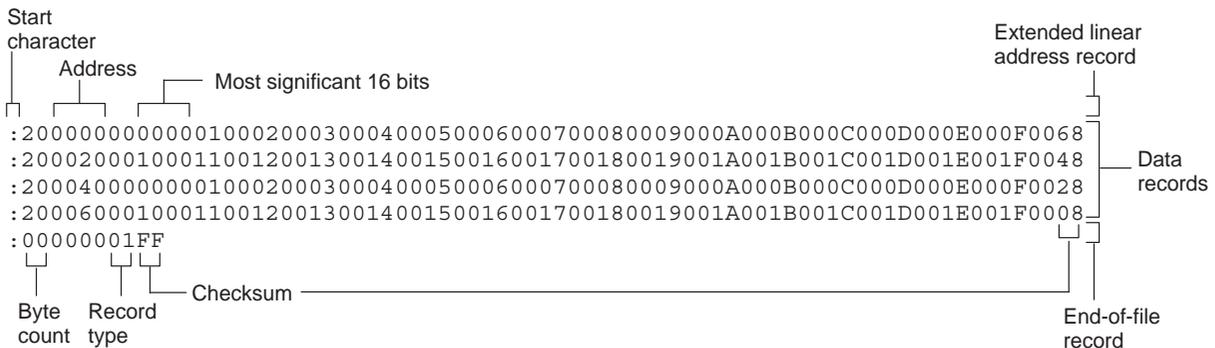
Record type 00, the data record, begins with a colon (:) and is followed by the byte count, the address of the first data byte, the record type (00), and the checksum. The address is the least significant 16 bits of a 32-bit address; this value is concatenated with the value from the most recent 04 (extended linear address) record to create a full 32-bit address. The checksum is the 2s complement (in binary form) of the preceding bytes in the record, including byte count, address, and data bytes.

Record type 01, the end-of-file record, also begins with a colon (:), followed by the byte count, the address, the record type (01), and the checksum.

Record type 04, the extended linear address record, specifies the upper 16 address bits. It begins with a colon (:), followed by the byte count, a dummy address of 0h, the record type (04), the most significant 16 bits of the address, and the checksum. The subsequent address fields in the data records contain the least significant bytes of the address.

Figure 11–7 illustrates the Intel hexadecimal object format.

Figure 11–7. Intel Hexadecimal Object Format



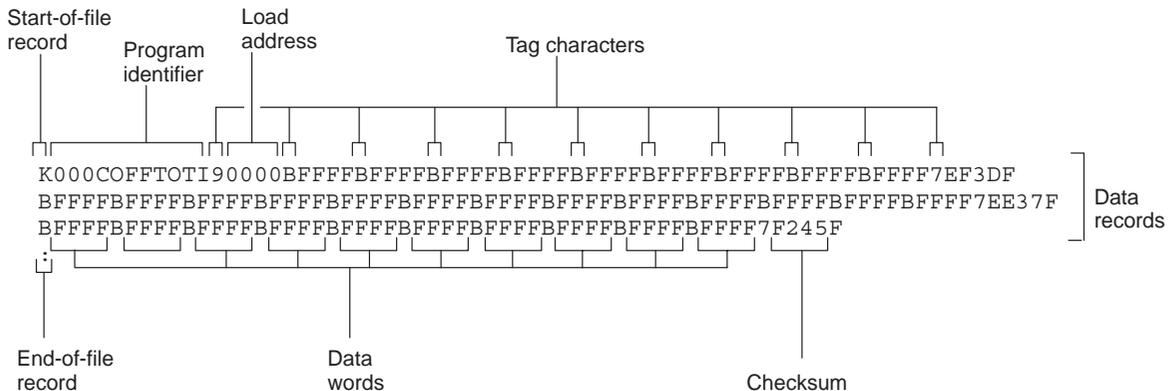
11.11.4 Texas Instruments SDSMAC Object Format (-t Option)

The Texas Instruments SDSMAC (TI-Tagged) object format supports 16-bit addresses. It consists of a start-of-file record, data records, and end-of-file record. Each of the data records consists of a series of small fields and is signified by a tag character. The significant tag characters are:

Tag Character	Description
K	Followed by the program identifier
7	Followed by a checksum
8	Followed by a dummy checksum (ignored)
9	Followed by a 16-bit load address
B	Followed by a data word (four characters)
F	Identifies the end of a data record
*	Followed by a data byte (two characters)

Figure 11-9 illustrates the tag characters and fields in TI-Tagged object format.

Figure 11-9. TI-Tagged Object Format



If any data fields appear before the first address, the first field is assigned address 0000h. Address fields may be expressed for any data byte, but none is required. The checksum field, which is preceded by the tag character 7, is the 2s complement of the sum of the 8-bit ASCII values of characters, beginning with the first tag character and ending with the checksum tag character (7 or 8). The end-of-file record is a colon (:).

11.12 Hex Conversion Utility Error Messages

section mapped to reserved memory message

Description A section is mapped into a reserved memory area as listed in the processor memory map.

Action Correct the section's allocation or boot-loader address. For valid memory locations, refer to the *TMS320C6200 CPU and Instruction Set Reference Guide*.

sections overlapping

Description Two or more COFF section load addresses overlap or a boot table address overlaps another section.

Action This problem may be caused by an incorrect translation (from the load address to the hexadecimal output file address) that is performed by the hex conversion utility when the memory width is less than the data width. See section 11.3, *Understanding Memory Widths*, on page 11-7 and section 11.10, *Controlling the ROM Device Address*, on page 11-32.

unconfigured memory error

Description The COFF file contains a section whose load address falls outside the memory range defined in the ROMS directive.

Action Correct the ROM range as defined by the ROMS directive to cover the memory range as needed, or modify the section load address. Remember that if the ROMS directive is not used, the memory range defaults to the entire processor address space. For this reason, removing the ROMS directive could also be a workaround.

Common Object File Format

The assembler and linker create object files in common object file format (COFF). COFF is an implementation of an object file format of the same name that was developed by AT&T for use on UNIX-based systems. This format is used because it encourages modular programming and provides powerful and flexible methods for managing code segments and target system memory.

Sections are a basic COFF concept. Chapter 2, *Introduction to Common Object File Format*, discusses COFF sections in detail. If you understand section operation, you can use the assembly language tools more efficiently.

This appendix contains technical details about the TMS320C6000™ COFF object file structure. Much of this information pertains to the symbolic debugging information that is produced by the C compiler. The purpose of this appendix is to provide supplementary information on the internal format of COFF object files.

Topic	Page
A.1 COFF File Structure	A-2
A.2 File Header Structure	A-4
A.3 Optional File Header Format	A-5
A.4 Section Header Structure	A-6
A.5 Structuring Relocation Information	A-9
A.6 Symbol Table Structure and Content	A-12

A.1 COFF File Structure

The elements of a COFF object file describe the file's sections and symbolic debugging information. These elements include:

- A file header
- Optional header information
- A table of section headers
- Raw data for each initialized section
- Relocation information for each initialized section
- A symbol table
- A string table

The assembler and linker produce object files with the same COFF structure; however, a program that is linked for the final time does not usually contain relocation entries. Figure A-1 illustrates the object file structure.

Figure A-1. COFF File Structure

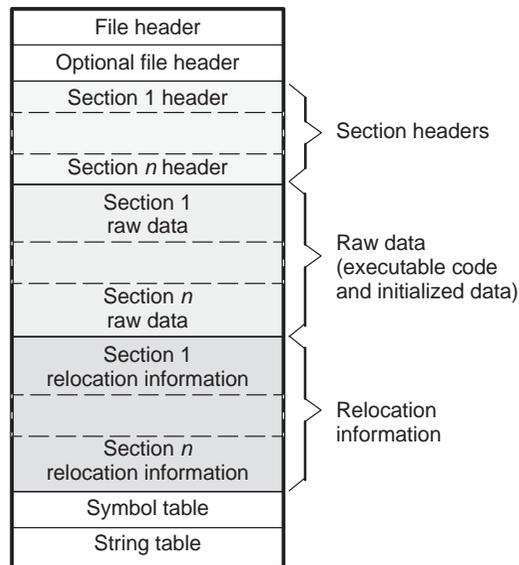
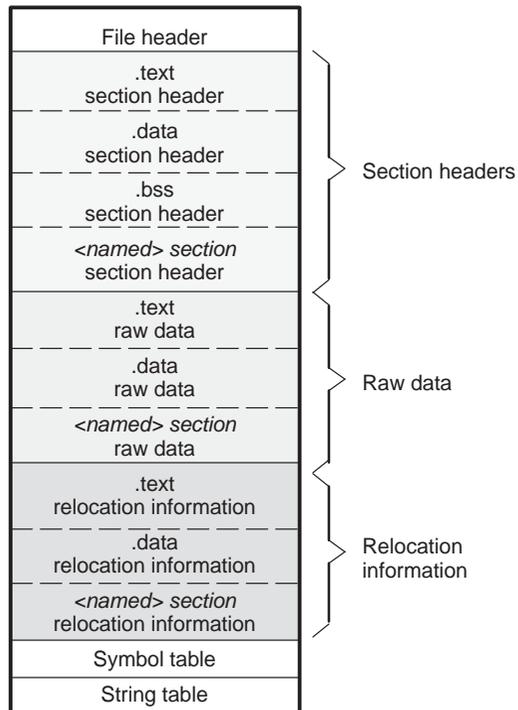


Figure A-2 shows a typical example of a COFF object file that contains the three default sections, .text, .data, and .bss, and a named section (referred to as <named>). By default, the tools place sections into the object file in the following order: .text, .data, initialized named sections, .bss, and uninitialized named sections. Although uninitialized sections have section headers, notice that they have no raw data, relocation information, or line number entries. This is because the .bss and .usect directives simply reserve space for uninitialized data; uninitialized sections contain no actual code.

Figure A-2. Sample COFF Object File



A.2 File Header Structure

The file header contains 22 bytes of information that describe the general format of an object file. Table A–1 shows the structure of the C6000 COFF file header.

Table A–1. File Header Contents

Byte Number	Type	Description
0–1	Unsigned short	Version ID; indicates version of COFF file structure
2–3	Unsigned short	Number of section headers
4–7	Integer	Time and date stamp; indicates when the file was created
8–11	Integer	File pointer; contains the symbol table's starting address
12–15	Integer	Number of entries in the symbol table
16–17	Unsigned short	Number of bytes in the optional header. This field is either 0 or 28; if it is 0, there is no optional file header.
18–19	Unsigned short	Flags (see Table A–2)
20–21	Unsigned short	Target ID; magic number (0099h) indicates the file can be executed in a C6000 system

Table A–2 lists the flags that can appear in bytes 18 and 19 of the file header. Any number and combination of these flags can be set at the same time (for example, if bytes 18 and 19 are set to 0003h, both F_RELFLG and F_EXEC are set.)

Table A–2. File Header Flags (Bytes 18 and 19)

Mnemonic	Flag	Description
F_RELFLG	0001h	Relocation information was stripped from the file.
F_EXEC	0002h	The file is relocatable (it contains no unresolved external references).
	0004h	Reserved
F_LSYMS	0008h	Local symbols were stripped from the file.
F_LITTLE	0100h	The target is a little-endian device.
F_BIG	0200h	The target is a big-endian device.

A.3 Optional File Header Format

The linker creates the optional file header and uses it to perform relocation at download time. Partially linked files do not contain optional file headers. Table A–3 illustrates the optional file header format.

Table A–3. *Optional File Header Contents*

Byte Number	Type	Description
0–1	Short	Optional file header magic number (0108h)
2–3	Short	Version stamp
4–7	Integer	Size (in bytes) of executable code
8–11	Integer	Size (in bytes) of initialized data
12–15	Integer	Size (in bytes) of uninitialized data
16–19	Integer	Entry point
20–23	Integer	Beginning address of executable code
24–27	Integer	Beginning address of initialized data

A.4 Section Header Structure

COFF object files contain a table of section headers that define where each section begins in the object file. Each section has its own section header. Table A-4 shows the structure of each section header.

Table A-4. Section Header Contents

Byte Number	Type	Description
0-7	Character	This field contains one of the following: 1) An 8-character section name padded with nulls 2) A pointer into the string table if the symbol name is longer than eight characters
8-11	Integer	Section's physical address
12-15	Integer	Section's virtual address
16-19	Integer	Section size in bytes
20-23	Integer	File pointer to raw data
24-27	Integer	File pointer to relocation entries
28-31	Integer	Reserved
32-35	Unsigned integer	Number of relocation entries
36-39	Unsigned integer	Reserved
40-43	Unsigned integer	Flags (see Table A-5)
44-45	Unsigned short	Reserved
46-47	Unsigned short	Memory page number

Table A-5 lists the flags that can appear in bytes 36 through 39 of the section header.

Table A–5. Section Header Flags (Bytes 40 Through 43)

Mnemonic	Flag	Description
STYP_REG	00000000h	Regular section (allocated, relocated, loaded)
STYP_DSECT	00000001h	Dummy section (relocated, not allocated, not loaded)
STYP_NOLOAD	00000002h	Noload section (allocated, relocated, not loaded)
STYP_COPY	00000010h	Copy section (relocated, loaded, but not allocated; relocation entries are processed normally)
STYP_TEXT	00000020h	Section contains executable code
STYP_DATA	00000040h	Section contains initialized data
STYP_BSS	00000080h	Section contains uninitialized data
STYP_BLOCK	00001000h	Alignment used as a blocking factor
STYP_PASS	00002000h	Section should pass through unchanged
STYP_CLINK	00004000h	Section requires conditional linking
STYP_VECTOR	00008000h	Section contains vector table
STYP_PADDED	00010000h	Section has been padded

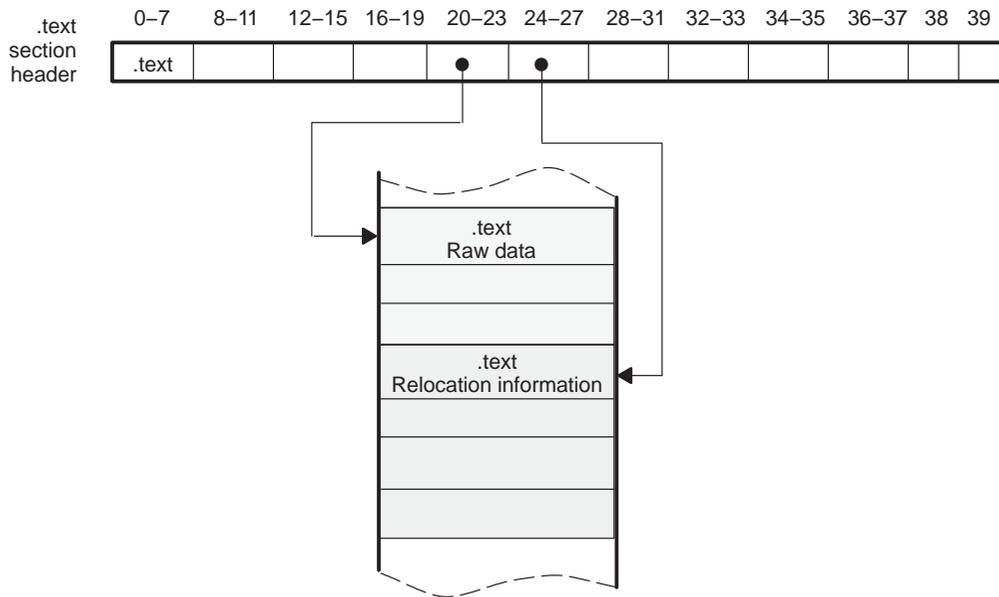
Note: The term *loaded* means that the raw data for this section appears in the object file. Only allocated sections are written to target memory.

The flags listed in Table A–5 can be combined; for example, if the flag's word is set to 060h, both STYP_DATA and STYP_TEXT are set.

Bits 8–11 of the section header flags are used for defining the alignment. The alignment is defined to be $2^{(\text{value of bits 8–11})}$. For example if bits 8–11 are 0101b (decimal integer 5), then the alignment is 32 (2^5).

Figure A–3 illustrates how the pointers in a section header point to the elements in an object file that are associated with the .text section.

Figure A-3. Section Header Pointers for the .text Section



As Figure A-2 on page A-3 shows, uninitialized sections (created with the `.bss` and `.usect` directives) vary from this format. Although uninitialized sections have section headers, they have no raw data or relocation information. They occupy no actual space in the object file. Therefore, the number of relocation entries, the number of line number entries, and the file pointers are 0 for an uninitialized section. The header of an uninitialized section simply tells the linker how much space for variables it should reserve in the memory map.

A.5 Structuring Relocation Information

A COFF object file has one relocation entry for each relocatable reference. The assembler automatically generates relocation entries. The linker reads the relocation entries as it reads each input section and performs relocation. The relocation entries determine how references within each input section are treated.

COFF file relocation information entries use the 10-byte format shown in Table A–6.

Table A–6. Relocation Entry Contents

Byte Number	Type	Description
0–3	Integer	Virtual address of the reference
4–5	Short	Symbol table index (0–65 535)
6–7	Unsigned short	Reserved
8–9	Unsigned short	Relocation type (see Table A–7)

The **virtual address** is the symbol's address in the current section *before* relocation; it specifies *where* a relocation must occur. (This is the address of the field in the object code that must be patched.)

Following is an example of code that generates a relocation entry:

```
2                               .global X
3 00000000 !00000012          b          X
```

In this example, the virtual address of the relocatable field is 0001.

The **symbol table index** is the index of the referenced symbol. In the preceding example, this field contains the index of X in the symbol table. The amount of the relocation is the difference between the symbol's current address in the section and its assembly-time address. The relocatable field must be relocated by the same amount as the referenced symbol. In the example, X has a value of 0 before relocation. Suppose X is relocated to address 2000h. This is the relocation amount (2000h – 0 = 2000h), so the relocation field at address 1 is patched by adding 2000h to it.

You can determine a symbol's relocated address if you know which section it is defined in. For example, if X is defined in .data and .data is relocated by 2000h, X is relocated by 2000h.

If the symbol table index in a relocation entry is –1 (0FFFFh), this is called an *internal relocation*. In this case, the relocation amount is simply the amount by which the current section is being relocated.

The **relocation type** specifies the size of the field to be patched and describes how the patched value is calculated. The type field depends on the addressing mode that was used to generate the relocatable reference. In the preceding example, the actual address of the referenced symbol X is placed in an 8-bit field in the object code. This is an 8-bit address, so the relocation type is R_RELBYTE. Table A-7 lists the relocation types.

Table A-7. Relocation Types (Bytes 8 and 9)

Mnemonic	Flag	Relocation Type
R_ABS	0000h	No relocation
R_RELBYTE	000Fh	8-bit direct reference to symbol's address
R_RELWORD	0010h	16-bit direct reference to symbol's address
R_RELLONG	0011h	32-bit direct reference to symbol's address
R_C60BASE	0050h	Data page pointer-based offset
R_C60DIR15	0051h	Load or store long displacement
R_C60PCR21	0052h	21-bit packet, PC relative
R_C60LO16	0054h	MVK instruction low half register
R_C60HI16	0055h	MVKH or MVKLH high half register
R_C60SECT	0056h	Section-based offset
R_C60PCR10	0053h	10-bit Packet PC Relative (BDEC, BPOS)
R_C60S16	0057h	Signed 16-bit offset for MVK
R_C60PCR7	0070h	7-bit Packet PC Relative (ADDKPC)
R_C60PCR12	0071h	12-bit Packet PC Relative (BNOP)
RE_ADD	4000h	Operator instruction +
RE_SUB	4001h	Operator instruction -
RE_NEG	4002h	Operator instruction unary -
RE_MPY	4003h	Operator instruction *
RE_DIV	4004h	Operator instruction /
RE_MOD	4005h	Operator instruction %
RE_SR	4006h	Unsigned shift right
RE_ASR	4007h	Signed shift right

Table A-7. Relocation Types (Bytes 8 and 9 (Continued))

Mnemonic	Flag	Relocation Type
RE_SL	4008h	Shift left
RE_AND	4009h	AND function
RE_OR	400Ah	OR function
RE_XOR	400Bh	Exclusive OR function
RE_NOTB	400Ch	~
RE_ULDFLD	400Dh	Unsigned relocation field load
RE_SLDFLD	400Eh	Signed relocation field load
RE_USTFLD	400Fh	Unsigned relocation field store
RE_SSTFLD	4010h	Signed relocation field store
RE_XSTFLD	4016h	Signed state is not relevant
RE_PUSH	4011h	Push symbol on the stack
RE_PUSHSV	c011h	Push symbol: SEGVALUE flag is set
RE_PUSHSK	4012h	Push signed constant on the stack
RE_PUSHUK	4013h	Push unsigned constant on the stack
RE_PUSHPC	4014h	Push current section PC on the stack
RE_DUP	4015h	Duplicate tos and push copy

A.6 Symbol Table Structure and Content

The order of symbols in the symbol table is very important; they appear in the sequence shown in Figure A-4.

Figure A-4. Symbol Table Contents

Static variables
⋮
Defined global symbols
Undefined global symbols

Static variables refer to symbols defined in C/C++ that have storage class static outside any function. If you have several modules that use symbols with the same name, making them static confines the scope of each symbol to the module that defines it (this eliminates multiple-definition conflicts).

The entry for each symbol in the symbol table contains the symbol's:

- Name (or an offset into the string table)
- Type
- Value
- Section it was defined in
- Storage class

Section names are also defined in the symbol table.

All symbol entries, regardless of class and type, have the same format in the symbol table. Each symbol table entry contains the 18 bytes of information listed in Table A-8. Each symbol may also have an 18-byte auxiliary entry; the special symbols listed in Table A-9, page A-13, always have an auxiliary entry. Some symbols may not have all the characteristics listed above; if a particular field is not set, it is set to null.

Table A–8. Symbol Table Entry Contents

Byte Number	Type	Description
0–7	Char	This field contains one of the following: <ol style="list-style-type: none"> 1) An 8-character symbol name, padded with nulls 2) A pointer into the string table if the symbol name is longer than eight characters
8–11	Integer	Symbol value; storage class dependent
12–13	Short	Section number of the symbol
14–15	Unsigned short	Reserved
16	Char	Storage class of the symbol
17	Char	Number of auxiliary entries (always 0 or 1)

A.6.1 Special Symbols

The symbol table contains some special symbols that are generated by the compiler, assembler, and linker. Each special symbol contains ordinary symbol table information as well as an auxiliary entry. Table A–9 lists these symbols.

Table A–9. Special Symbols in the Symbol Table

Symbol	Description
.text	Address of the .text section
.data	Address of the .data section
.bss	Address of the .bss section
etext	Next available address after the end of the .text output section
edata	Next available address after the end of the .data output section
end	Next available address after the end of the .bss output section

A.6.2 Symbol Name Format

The first eight bytes of a symbol table entry (bytes 0–7) indicate a symbol’s name:

- If the symbol name is eight characters or less, this field has type *character*. The name is padded with nulls (if necessary) and stored in bytes 0–7.
- If the symbol name is greater than eight characters, this field is treated as two integers. The entire symbol name is stored in the string table. Bytes 0–3 contain 0, and bytes 4–7 are an offset into the string table.

A.6.3 String Table Structure

Symbol names that are longer than eight characters are stored in the string table. The field in the symbol table entry that would normally contain the symbol’s name contains, instead, a pointer to the symbol’s name in the string table. Names are stored contiguously in the string table, delimited by a null byte. The first four bytes of the string table contain the size of the string table in bytes; thus, offsets into the string table are greater than or equal to 4.

Figure A–5 is a string table that contains two symbol names, *Adaptive-Filter* and *Fourier-Transform*. The index in the string table is 4 for Adaptive-Filter and 20 for Fourier-Transform.

Figure A–5. String Table Entries for Sample Symbol Names

38 bytes

4 bytes			
'A'	'd'	'a'	'p'
't'	'i'	'v'	'e'
'.'	'F'	'i'	'l'
't'	'e'	'r'	'\0'
'F'	'o'	'u'	'r'
'i'	'e'	'r'	'.'
'T'	'r'	'a'	'n'
's'	'f'	'o'	'r'
'm'	'\0'		

A.6.4 Storage Classes

Byte 16 of the symbol table entry indicates the storage class of the symbol. Storage classes refer to the method in which the C/C++ compiler accesses a symbol. Table A–10 lists valid storage classes.

Table A–10. Symbol Storage Classes

Mnemonic	Value	Storage Class	Mnemonic	Value	Storage Class
C_NULL	0	No storage class	C_USTATIC	14	Undefined static
C_AUTO	1	Reserved	C_ENTAG	15	Reserved
C_EXT	2	External definition	C_MOE	16	Reserved
C_STAT	3	Static	C_REGPARAM	17	Reserved
C_REG	4	Reserved	C_FIELD	18	Reserved
C_EXTREF	5	External reference	C_UEXT	19	Tentative external definition
C_LABEL	6	Label	C_STATLAB	20	Static load time label
C_ULABEL	7	Undefined label	C_EXTLAB	21	External load time label
C_MOS	8	Reserved	C_VARARG	27	Last declared parameter of a function with variable number of arguments
C_ARG	9	Reserved	C_BLOCK	100	Reserved
C_STRTAG	10	Reserved	C_FCN	101	Reserved
C_MOU	11	Reserved	C_EOS	102	Reserved
C_UNTAG	12	Reserved	C_FILE	103	Reserved
C_TPDEF	13	Reserved	C_LINE	104	Used only by utility programs

The .text, .data, and .bss symbols are restricted to the C_STAT storage class.

A.6.5 Symbol Values

Bytes 8–11 of a symbol table entry indicate a symbol's value. The C_EXT, C_STAT, and C_LABEL storage classes hold relocatable addresses.

The value of a relocatable symbol is its virtual address. When the linker relocates a section, the value of a relocatable symbol changes accordingly.

A.6.6 Section Number

Bytes 12–13 of a symbol table entry contain a number that indicates which section the symbol was defined in. Table A–11 lists these numbers and the sections they indicate.

Table A–11. Section Numbers

Mnemonic	Section Number	Description
None	–2	Reserved
N_ABS	–1	Absolute symbol
N_UNDEF	0	Undefined external symbol
None	1	.text section (typical)
None	2	.data section (typical)
None	3	.bss section (typical)
None	4–32 767	Section number of a named section, in the order in which the named sections are encountered

If there were no .text, .data, or .bss sections, the numbering of named sections would begin with 1.

If a symbol has a section number of 0, –1, or –2, it is not defined in a section. A section number of –1 indicates that the symbol has a value but is not relocatable. A section number of 0 indicates a relocatable external symbol that is not defined in the current file.

A.6.7 Auxiliary Entries

Each symbol table entry can have *one* or *no* auxiliary entry. An auxiliary symbol table entry contains the same number of bytes as a symbol table entry (18). Table A–12 illustrates the format of auxiliary table entries.

Table A–12. Section Format for Auxiliary Table Entries

Byte Number	Type	Description
0–3	Integer	Section length
4–5	Unsigned short	Number of relocation entries
6–7	Unsigned short	Number of line number entries
8–17	—	Not used (zero filled)

Symbolic Debugging Directives

The assembler supports several directives that the TMS320C6000 C/C++ compiler uses for symbolic debugging. These directives differ for the two debugging formats, DWARF and COFF.

These directives are not meant for use by assembly-language programmers. They require arguments that can be difficult to calculate manually, and their usage must conform to a predetermined agreement between the compiler, the assembler, and the debugger. This appendix documents these directives for informational purposes only.

Topic	Page
B.1 DWARF Debugging Format	B-2
B.2 COFF Debugging Format	B-3
B.3 Debug Directive Syntax	B-4

B.1 DWARF Debugging Format

A subset of the DWARF symbolic debugging directives are always listed in the assembly language file that the compiler creates for program analysis purposes. To list the complete set used for full symbolic debug, invoke the compiler with the `-g` option, as shown below:

```
cl6x -g -k input_file
```

The `-k` option instructs the compiler to retain the generated assembly file.

To disable the generation of all symbolic debug directives, invoke the compiler with the `--symdebug:none` option:

```
cl6x --symdebug:none -k input_file
```

The DWARF debugging format consists of the following directives:

- The **.dwtag** and **.dwendtag** directives define a Debug Information Entry (DIE) in the `.debug_info` section.
- The **.dwattr** directive adds an attribute to an existing DIE.
- The **.dwpsn** directive identifies the source position of a C/C++ statement.
- The **.dwcie** and **.dwentry** directives define a Common Information Entry (CIE) in the `.debug_frame` section.
- The **.dwfde** and **.dwentry** directives define a Frame Description Entry (FDE) in the `.debug_frame` section.
- The **.dwcfa** directive defines a call frame instruction for a CIE or FDE.

B.2 COFF Debugging Format

COFF symbolic debug is now obsolete. These directives are supported for backwards-compatibility only. The decision to switch to DWARF as the symbolic debug format was made to overcome many limitations of COFF symbolic debug, including the absence of C++ support.

The COFF debugging format consists of the following directives:

- The **.sym** directive defines a global variable, a local variable, or a function. Several parameters allow you to associate various debugging information with the variable or function.
- The **.stag**, **.etag**, and **.utag** directives define structures, enumerations, and unions, respectively. The **.member** directive specifies a member of a structure, enumeration, or union. The **.eos** directive ends a structure, enumeration, or union definition.
- The **.func** and **.endfunc** directives specify the beginning and ending lines of a C/C++ function.
- The **.block** and **.endblock** directives specify the bounds of C/C++ blocks.
- The **.file** directive defines a symbol in the symbol table that identifies the current source filename.
- The **.line** directive identifies the line number of a C/C++ source statement.

B.3 Debug Directive Syntax

Table B–1 is an alphabetical listing of the symbolic debugging directives. For information on the C/C++ compiler, refer to the *TMS320C6000 Optimizing Compiler User's Guide*.

Table B–1. Symbolic Debugging Directives

Label	Directive	Arguments
	.block	[beginning line number]
	.dwattr	DIE label, DIE attribute name(DIE attribute value)[,DIE attribute name(attribute value) [, ...]
	.dwcfa	call frame instruction opcode[, operand[, operand]]
CIE label	.dwcie	version, return address register
	.dwentry	
	.dwendtag	
	.dwfde	CIE label
	.dwpsn	“filename”, line number, column number
DIE label	.dwtag	DIE tag name, DIE attribute name(DIE attribute value)[,DIE attribute name(attribute value) [, ...]
	.endblock	[ending line number]
	.endfunc	[ending line number[, register mask[, frame size]]]
	.eos	
	.etag	name[, size]
	.file	“filename”
	.func	[beginning line number]
	.line	line number[, address]
	.member	name, value[, type, storage class, size, tag, dims]
	.stag	name[, size]
	.sym	name, value[, type, storage class, size, tag, dims]
	.utag	name[, size]

XML Link Information File Description

The linker supports the generation of an XML link information file via the `--xml_link_info file` option. This option causes the linker to generate a well-formed XML file containing detailed information about the result of a link. The information included in this file includes all of the information that is currently produced in a linker-generated map file.

As the linker evolves, the XML link information file may be extended to include additional information that could be useful for static analysis of linker results.

This appendix enumerates all of the elements that are generated by the linker into the XML link information file.

Topic	Page
C.1 XML Information File Element Types	C-2
C.2 Document Elements	C-3

C.1 XML Information File Element Types

These element types will be generated by the linker:

- Container elements** represent an object that contains other elements that describe the object. Container elements have an id attribute that makes them accessible from other elements.
- String elements** contain a string representation of their value.
- Constant elements** contain a 32-bit unsigned long representation of their value (with a 0x prefix).
- Reference elements** are empty elements that contain an idref attribute that specifies a link to another container element.

In section C.2, the element type is specified for each element in parentheses following the element description. For instance, the <link_time> element lists the time of the link execution (string).

C.2 Document Elements

The root element, or the document element, is **<link_info>**. All other elements contained in the XML link information file are children of the **<link_info>** element. The following sections describe the elements that an XML information file can contain.

C.2.1 Header Elements

The first elements in the XML link information file provide general information about the linker and the link session:

- The **<banner>** element lists the name of the executable and the version information (string).
- The **<copyright>** element lists the TI copyright information (string).
- The **<link_time>** element lists the time of the link execution (string).
- The **<link_timestamp>** is a timestamp representation of the link time (unsigned 32-bit int)
- The **<output_file>** element lists the name of the linked output file generated (string).
- The **<entry_point>** element specifies the program entry point, as determined by the linker (container) with two entries:
 - The **<name>** is the entry point symbol name, if any (string).
 - The **<address>** is the entry point address (constant).

Example C–1. Header Element for the hi.out Output File

```
<banner>TMS320Cxx COFF Linker      Version x.xx (Jan 6 2003)</banner>
<copyright>Copyright (c) 1996-2003 Texas Instruments Incorporated</copyright>
<link_time>Mon Jan 6 15:38:18 2003</link_time>
<output_file>hi.out</output_file>
<entry_point>
  <name>_c_int00</name>
  <address>0xaf80</address>
</entry_point>
```

C.2.2 Input File List

The next section of the XML link information file is the input file list, which is delimited with a **<input_file_list>** container element. The **<input_file_list>** can contain any number of **<input_file>** elements.

Each **<input_file>** instance specifies the input file involved in the link. Each **<input_file>** has an **id** attribute that can be referenced by other elements, such as an **<object_component>**. An **<input_file>** is a container element enclosing the following elements:

- ❑ The **<path>** element names a directory path, if applicable (string).
- ❑ The **<kind>** element specifies a file type, either archive or object (string).
- ❑ The **<file>** element specifies an archive name or filename (string).
- ❑ The **<name>** element specifies an object file name, or archive member name (string).

Example C-2. Input File List for the hi.out Output File

```
<input_file_list>
  <input_file id="fl-1">
    <kind>object</kind>
    <file>hi.obj</file>
    <name>hi.obj</name>
  </input_file>
  <input_file id="fl-2">
    <path>/tools/lib/</path>
    <kind>archive</kind>
    <file>rtsxxx.lib</file>
    <name>boot.obj</name>
  </input_file>
  <input_file id="fl-3">
    <path>/tools/lib/</path>
    <kind>archive</kind>
    <file>rtsxxx.lib</file>
    <name>exit.obj</name>
  </input_file>
  <input_file id="fl-4">
    <path>/tools/lib/</path>
    <kind>archive</kind>
    <file>rtsxxx.lib</file>
    <name>printf.obj</name>
  </input_file>
  ...
</input_file_list>
```

C.2.3 Object Component List

The next section of the XML link information file contains a specification of all of the object components that are involved in the link. An example of an object component is an input section. In general, an object component is the smallest piece of object that can be manipulated by the linker.

The **<object_component_list>** is a container element enclosing any number of **<object_component>** elements.

Each **<object_component>** specifies a single object component. Each **<object_component>** has an **id** attribute so that it can be referenced directly from other elements, such as a **<logical_group>**. An **<object_component>** is a container element enclosing the following elements:

- The **<name>** element names the object component (string).
- The **<load_address>** element specifies the load-time address of the object component (constant).
- The **<run_address>** element specifies the run-time address of the object component (constant).
- The **<size>** element specifies the size of the object component (constant).
- The **<input_file_ref>** element specifies the source file where the object component originated (reference).

Example C-3. Object Component List for the fl-4 Input File

```

<object_component id="oc-20">
  <name>.text</name>
  <load_address>0xac00</load_address>
  <run_address>0xac00</run_address>
  <size>0xc0</size>
  <input_file_ref idref="fl-4"/>
</object_component>
<object_component id="oc-21">
  <name>.data</name>
  <load_address>0x80000000</load_address>
  <run_address>0x80000000</run_address>
  <size>0x0</size>
  <input_file_ref idref="fl-4"/>
</object_component>
<object_component id="oc-22">
  <name>.bss</name>
  <load_address>0x80000000</load_address>
  <run_address>0x80000000</run_address>
  <size>0x0</size>
  <input_file_ref idref="fl-4"/>
</object_component>

```

C.2.4 Logical Group List

The **<logical_group_list>** section of the XML link information file is similar to the output section listing in a linker generated map file. However, the XML link information file contains a specification of GROUP and UNION output sections, which are not represented in a map file. There are three kinds of list items that can occur in a **<logical_group_list>**:

- The **<logical_group>** is the specification of a section or GROUP that contains a list of object components or logical group members. Each **<logical_group>** element is given an id so that it may be referenced from other elements. Each **<logical_group>** is a container element enclosing the following elements:
 - The **<name>** element names the logical group (string).
 - The **<load_address>** element specifies the load-time address of the logical group (constant).
 - The **<run_address>** element specifies the run-time address of the logical group (constant).
 - The **<size>** element specifies the size of the logical group (constant).
 - The **<contents>** element lists elements contained in this logical group (container). These elements refer to each of the member objects contained in this logical group:
 - The **<object_component_ref>** is an object component that is contained in this logical group (reference).
 - The **<logical_group_ref>** is a logical group that is contained in this logical group (reference).
- The **<overlay>** is a special kind of logical group that represents a UNION, or a set of objects that share the same memory space (container). Each **<overlay>** element is given an id so that it may be referenced from other elements (like from an **<allocated_space>** element in the placement map). Each **<overlay>** contains the following elements:
 - The **<name>** element names the overlay (string).
 - The **<run_address>** element specifies the run-time address of overlay (constant).
 - The **<size>** element specifies the size of logical group (constant).

- The **<contents>** container element lists elements contained in this overlay. These elements refer to each of the member objects contained in this logical group:
 - The **<object_component_ref>** is an object component that is contained in this overlay (reference).
 - The **<logical_group_ref>** is a logical group that is contained in this overlay (reference).
- The **<split_section>** is another special kind of logical group which represents a collection of logical groups that is split among multiple memory areas. Each **<split_section>** element is given an id so that it may be referenced from other elements. The id consists of the following elements.
 - The **<name>** element names the split section (string).
 - The **<contents>** container element lists elements contained in this split section. The **<logical_group_ref>** elements refer to each of the member objects contained in this split section, and each element referenced is a logical group that is contained in this split section (reference).

Example C-4. Logical Group List for the fl-4 Input File

```
<logical_group_list>
  ...
  <logical_group id="lg-7">
    <name>.text</name>
    <load_address>0x20</load_address>
    <run_address>0x20</run_address>
    <size>0xb240</size>
    <contents>
      <object_component_ref idref="oc-34"/>
      <object_component_ref idref="oc-108"/>
      <object_component_ref idref="oc-e2"/>
    ...
    </contents>
  </logical_group>
  ...
  <overlay id="lg-b">
    <name>UNION_1</name>
    <run_address>0xb600</run_address>
    <size>0xc0</size>
    <contents>
      <object_component_ref idref="oc-45"/>
      <logical_group_ref idref="lg-8"/>
    </contents>
  </overlay>
  ...
  <split_section id="lg-12">
    <name>.task_scn</name>
    <size>0x120</size>
    <contents>
      <logical_group_ref idref="lg-10"/>
      <logical_group_ref idref="lg-11"/>
    </contents>
  ...
</logical_group_list>
```

C.2.5 Placement Map

The **<placement_map>** element describes the memory placement details of all named memory areas in the application, including unused spaces between logical groups that have been placed in a particular memory area.

The **<memory_area>** is a description of the placement details within a named memory area (container). The description consists of these items:

- The **<name>** names the memory area (string).
- The **<page_id>** gives the id of the memory page in which this memory area is defined (constant).
- The **<origin>** specifies the beginning address of the memory area (constant).
- The **<length>** specifies the length of the memory area (constant).
- The **<used_space>** specifies the amount of allocated space in this area (constant).
- The **<unused_space>** specifies the amount of available space in this area (constant).
- The **<attributes>** lists the RWX attributes that are associated with this area, if any (string).
- The **<fill_value>** specifies the fill value that is to be placed in unused space, if the fill directive is specified with the memory area (constant).
- The **<usage_details>** lists details of each allocated or available fragment in this memory area. If the fragment is allocated to a logical group, then a **<logical_group_ref>** element is provided to facilitate access to the details of that logical group. All fragment specifications include **<start_address>** and **<size>** elements.
 - The **<allocated_space>** element provides details of an allocated fragment within this memory area (container):
 - The **<start_address>** specifies the address of the fragment (constant).
 - The **<size>** specifies the size of the fragment (constant).
 - The **<logical_group_ref>** provides a reference to the logical group that is allocated to this fragment (reference).
 - The **<available_space>** element provides details of an available fragment within this memory area (container):
 - The **<start_address>** specifies the address of the fragment (constant).
 - The **<size>** specifies the size of the fragment (constant).

Example C–5. Placement Map for the fl–4 Input File

```

<placement_map>
  <memory_area>
    <name>PMEM</name>
    <page_id>0x0</page_id>
    <origin>0x20</origin>
    <length>0x100000</length>
    <used_space>0xb240</used_space>
    <unused_space>0xf4dc0</unused_space>
    <attributes>RWXI</attributes>
    <usage_details>
      <allocated_space>
        <start_address>0x20</start_address>
        <size>0xb240</size>
        <logical_group_ref idref="lg-7" />
      </allocated_space>
      <available_space>
        <start_address>0xb260</start_address>
        <size>0xf4dc0</size>
      </available_space>
    </usage_details>
  </memory_area>
  ...
</placement_map>

```

C.2.6 Far Call Trampoline List

The **<far_call_trampoline_list>** is a list of **<far_call_trampoline>** elements. The C6000 linker supports the generation of far call trampolines to help a call site reach a destination that is out of range. A far call trampoline function is guaranteed to reach the called function (callee) as it may utilize an indirect call to the called function.

The **<far_call_trampoline_list>** enumerates all of the far call trampolines that are generated by the linker for a particular link. The **<far_call_trampoline_list>** can contain any number of **<far_call_trampoline>** elements. Each **<far_call_trampoline>** is a container enclosing the following elements:

- The **<callee_name>** element names the destination function (string).
- The **<callee_address>** is the address of the called function (constant).
- The **<trampoline_object_component_ref>** is a reference to an object component that contains the definition of the trampoline function (reference).
- The **<trampoline_address>** is the address of the trampoline function (constant).

- The **<caller_list>** enumerates all call sites that utilize this trampoline to reach the called function (container).
- The **<trampoline_call_site>** provides the details of a trampoline call site (container) and consists of these items:
 - The **<caller_address>** specifies the call site address (constant) .
 - The **<caller_object_component_ref>** is the object component where the call site resides (reference).

Example C–6. Fall Call Trampoline List for the fl–4 Input File

```

<far_call_trampoline_list>
  ...
  <far_call_trampoline>
    <callee_name>_foo</callee_name>
    <callee_address>0x08000030</callee_address>
    <trampoline_object_component_ref idref="oc-123"/>
    <trampoline_address>0x2020</trampoline_address>
    <caller_list>
      <call_site>
        <caller_address>0x1800</caller_address>
        <caller_object_component_ref idref="oc-23"/>
      </call_site>
      <call_site>
        <caller_address>0x1810</caller_address>
        <caller_object_component_ref idref="oc-23"/>
      </call_site>
    </caller_list>
  </far_call_trampoline>
  ...
</far_call_trampoline_list>

```

C.2.7 Symbol Table

The **<symbol_table>** contains a list of all of the global symbols that are included in the link. The list provides information about a symbol's name and value. In the future, the `symbol_table` list may provide type information, the object component in which the symbol is defined, storage class, etc.

The **<symbol>** is a container element that specifies the name and value of a symbol with these elements:

- The **<name>** element specifies the symbol name (string)
- The **<value>** element specifies the symbol value (constant)

Example C-7. Symbol Table for the fl-4 Input File

```
<symbol_table>
  <symbol>
    <name>_c_int00</name>
    <value>0xaf80</value>
  </symbol>
  <symbol>
    <name>_main</name>
    <value>0xb1e0</value>
  </symbol>
  <symbol>
    <name>_printf</name>
    <value>0xac00</value>
  </symbol>
  ...
</symbol_table>
```

Glossary

A

absolute address: An address that is permanently assigned to a TMS320C6000 memory location.

alignment: A process in which the linker places an output section at an address that falls on an n -byte boundary, where n is a power of 2. You can specify alignment with the SECTIONS linker directive.

allocation: A process in which the linker calculates the final memory addresses of output sections.

American Standard Code for Information Interchange (ASCII): A standard computer code for representing and exchanging alphanumeric information.

archive library: A collection of individual files that have been grouped into a single file.

archiver: A software program that allows you to collect several individual files into a single file called an archive library. The archiver also allows you to delete, extract, or replace members of the archive library, as well as to add new members.

assembler: A software program that creates a machine-language program from a source file that contains assembly language instructions, directives, and macro directives. The assembler substitutes absolute operation codes for symbolic operation codes, and absolute or relocatable addresses for symbolic addresses.

assembly-time constant: A symbol that is assigned a constant value with the .set directive.

assignment statement: A statement that assigns a value to a variable.

autoinitialization: The process of initializing global C variables (contained in the .cinit section) before beginning program execution.

auxiliary entry: The extra entry that a symbol may have in the symbol table and that contains additional information about the symbol (whether it is a filename, a section name, a function name, etc.).

B

binding: A process in which you specify a distinct address for an output section or a symbol.

big endian: An addressing protocol in which bytes are numbered from left to right within a word. More significant bytes in a word have lower numbered addresses. Endian ordering is hardware-specific and is determined at reset. See also *little endian*

block: A set of declarations and statements that are grouped together with braces.

.bss: One of the default COFF sections. You can use the `.bss` directive to reserve a specified amount of space in the memory map that can later be used for storing data. The `.bss` section is uninitialized.

byte: A sequence of eight adjacent bits operated upon as a unit.

C

C/C++ compiler: A program that translates C/C++ source statements into assembly language source statements.

command file: A file that contains options, filenames, directives, or commands for the linker or hex conversion utility.

comment: A source statement (or portion of a source statement) that is used to document or improve readability of a source file. Comments are not compiled, assembled, or linked; they have no effect on the object file.

common object file format (COFF): A binary object file format configured by a standard developed by AT&T. All COFF sections are independently relocatable in memory space; you can place any section into any allocated block of target memory.

conditional processing: A method of processing one block of source code or an alternate block of source code, according to the evaluation of a specified expression.

configured memory: Memory that the linker has specified for allocation.

constant: A numeric value that does not change and that can be used as an operand.

cross-reference listing: An output file created by the assembler and appended to the end of the listing file. The cross reference information lists the symbols that were defined, what line they were defined on, which lines referenced them, and the values as determined by the input assembly source file.

D

.data: One of the default COFF sections. The .data section is an initialized section that contains initialized data. You can use the .data directive to assemble code into the .data section.

directives: Special-purpose commands that control the actions and functions of a software tool (as opposed to assembly language instructions, which control the actions of a device).

E

emulator: A hardware development system that emulates TMS320C6200 operation.

entry point: The starting execution point in target memory.

executable module: An object file that has been linked and can be executed in a TMS320C6000 system.

expression: A constant, a symbol, or a series of constants and symbols separated by arithmetic operators.

external symbol: A symbol that is used in the current program module but is defined in a different program module.

F

field: For the TMS320C6000, a software-configurable data type whose length can be programmed to be any value in the range of 1–32 bits.

file header: A portion of a COFF object file that contains general information about the object file, such as the number of section headers, the type of system the object file can be downloaded to, the number of symbols in the symbol table, and the symbol table's starting address.

G

global symbol: A kind of symbol that is either 1) defined in the current module and accessed in another, or 2) accessed in the current module but defined in another.

GROUP: An option of the SECTIONS directive that forces specified output sections to be allocated contiguously (as a group).

H

hex conversion utility: A program that accepts COFF files and converts them into one of several standard ASCII hexadecimal formats suitable for loading into an EPROM programmer.

high-level language debugging: The ability of a compiler to retain symbolic and high-level language information (such as type and function definitions) so that a debugging tool can use this information.

hole: An area containing no actual code or data. This area is between the input sections that compose an output section.

I

incremental linking: Linking files in several passes. Incremental linking is useful for large applications, because you can partition the application, link the parts separately, and then link all of the parts together.

initialized section: A COFF section that contains executable code or initialized data. An initialized section can be built up with the .data, .text, or .sect directive.

input section: A section from an object file that will be linked into an executable module.

L

label: A symbol that begins in column 1 of a source statement and corresponds to the address of that statement.

line-number entry: An entry in a COFF output module that maps lines of assembly code back to the original C source file that created them.

linker: A software tool that combines object files to form an object module that can be allocated into TMS320C6000 system memory and executed by the device.

listing file: An output file, created by the assembler, that lists source statements, their line numbers, and their effects on the SPC.

little endian: An addressing protocol in which bytes are numbered from right to left within a word. More significant bytes in a word have higher numbered addresses. Endian ordering is hardware-specific and is determined at reset. See also *big endian*

loader: A device that loads an executable module into TMS320C6000 system memory.

M

macro: A user-defined routine that can be used as an instruction.

macro call: The process of invoking a macro.

macro definition: A block of source statements that define the name and the code that make up a macro.

macro expansion: The source statements that are substituted for the macro call and are subsequently assembled.

macro library: An archive library composed of macros. Each file in the library must contain one macro; its name must be the same as the macro name it defines, and it must have an extension of .asm.

magic number: A COFF file header entry that identifies an object file as a module that can be executed by the TMS320C6000.

map file: An output file, created by the linker, that shows the memory configuration, section composition, and section allocation, as well as symbols and the addresses at which they were defined.

member: The elements or variables of a structure, union, archive, or enumeration.

memory map: A map of target system memory space that is partitioned into functional blocks.

mnemonic: An instruction name that the assembler translates into machine code.

model statement: Instructions or assembler directives in a macro definition that are assembled each time a macro is invoked.

N

named section: An initialized section that is defined with a `.sect` directive.

O

object file: A file that has been assembled or linked and contains machine-language object code.

object library: An archive library made up of individual object files.

operands: The arguments, or parameters, of an assembly language instruction, assembler directive, or macro directive.

optional header: A portion of a COFF object file that the linker uses to perform relocation at download time.

options: Command parameters that allow you to request additional or specific functions when you invoke a software tool.

output module: A linked, executable object file that can be downloaded and executed on a target system.

output section: A final, allocated section in a linked, executable module.

P

partial linking: Linking files in several passes. Incremental linking is useful for large applications because you can partition the application, link the parts separately, and then link all of the parts together.

Q

quiet run: An option that suppresses the normal banner and the progress information.

R

raw data: Executable code or initialized data in an output section.

relocation: A process in which the linker adjusts all the references to a symbol when the symbol's address changes.

run address: The address where a section runs.

S

section: A relocatable block of code or data that will ultimately occupy contiguous space in the TMS320C6000 memory map.

section header: A portion of a COFF object file that contains information about a section in the file. Each section has its own header; the header points to the section's starting address, contains the section's size, etc.

section program counter (SPC): An element that keeps track of the current location within a section; each section has its own SPC.

sign extend: To fill the unused MSBs of a value with the value's sign bit.

simulator: A software development system that simulates TMS320C6000 operation.

source file: A file that contains C code or assembly language code that will be compiled or assembled to form an object file.

static variable: An element whose scope is confined to a function or a program. The values of static variables are not discarded when the function or program is exited; the previous value is resumed when the function or program is reentered.

storage class: Any entry in the symbol table that indicates how a symbol is accessed.

string table: A table that stores symbol names that are longer than eight characters (symbol names of eight characters or longer cannot be stored in the symbol table; instead, they are stored in the string table). The name portion of the symbol's entry points to the location of the string in the string table.

structure: A collection of one or more variables grouped together under a single name.

subsection: A relocatable block of code or data that will ultimately occupy continuous space in the TMS320C6000 memory map. Subsections are smaller sections within larger sections. Subsections give you tighter control of the memory map.

symbol: A string of alphanumeric characters that represents an address or a value.

symbolic debugging: The ability of a software tool to retain symbolic information so that it can be used by a debugging tool, such as a simulator or an emulator.

symbol table: A portion of a COFF object file that contains information about the symbols that are defined and used by the file.

T

tag: An optional *type* name that can be assigned to a structure, union, or enumeration.

target memory: Physical memory in a TMS320C6000 system into which executable object code is loaded.

.text: One of the default COFF sections. The .text section is an initialized section that contains executable code. You can use the .text directive to assemble code into the .text section.

U

unconfigured memory: Memory that is not defined as part of the memory map and cannot be loaded with code or data.

uninitialized section: A COFF section that reserves space in the memory map but that has no actual contents. These sections are built up with the .bss and .usect directives.

UNION: An option of the SECTIONS directive that causes the linker to allocate the same address to multiple sections.

union: A variable that can hold objects of different types and sizes.

unsigned value: An element that is treated as a positive number, regardless of its actual sign.

W

well-defined expression: A term or group of terms that contains only symbols or assembly-time constants that have been defined before they appear in the expression.

word: A 16-bit addressable location in target memory.

; in assembly language source 3-13
\$ symbol for SPC 3-23
@ archiver command 6-4
-@ assembler option 3-4
* in assembly language source 3-13

A

a archiver command 6-4
-a hex conversion utility option 11-4, 11-35
-a linker option 7-6
-a name utility option 10-16
A operand of .option directive 4-14, 4-59
A_DIR environment variable 3-8, 7-12, 7-13
-aa assembler option 3-4
-abs linker option 7-7
absolute lister
 creating the absolute listing file 8-2
 development flow 8-2
 example 8-5 to 8-10
 invoking 8-3
 options 8-3
absolute listing
 -aa assembler option 3-4
 -abs linker option 7-7
 producing 8-2
absolute output module 7-6
-ac assembler option 3-4
-ad assembler option 3-5, 3-21
address XML entry C-3
-ahc assembler option 3-4
-ahi assembler option 3-4
-al assembler option 3-5
 source listing format 3-31
.align directive 4-13, 4-21
alignment 4-13, 4-21, 7-37
 defined D-1
allocated_space XML element C-9
allocation 2-2, 4-26, 7-34 to 7-43
 alignment 4-21, 7-37
 allocating output sections 7-30
 binding 7-35
 blocking 7-37
 checking consistency of load and run 7-51
 default algorithm 7-54 to 7-55
 defined D-1
 GROUP 7-50
 memory, default 2-12, 7-36 to 7-37
 UNION 7-48
alternate directories 3-7 to 3-9, 7-12
 naming with -i option 3-7
 naming with A_DIR 3-8
-apd assembler option 3-5
-api assembler option 3-5
-ar linker option 7-7
ar6x command 6-4
archive libraries 4-55 to 4-56, 7-11, 7-25 to 7-26
 back referencing 7-20
 defined D-1
 exhaustively reading 7-20
 types of files 6-2
archiver 1-4, 6-1 to 6-8
 allocating member to output section 7-43
 commands
 @ 6-4
 a 6-4
 d 6-4
 r 6-4
 t 6-4
 u 6-5
 x 6-4
 defined D-1
 examples 6-6
 in the development flow 6-3
 invoking 6-4

- archiver (continued)
 - options
 - q 6-5
 - s 6-5
 - v 6-5
 - args linker option 7-8
 - arguments, passing to the loader 7-8
 - arithmetic operators 3-27
 - as assembler option 3-5
 - ASCII-Hex object format 11-1, 11-35
 - .asg directive 4-18, 4-22
 - listing control 4-14, 4-36
 - use in macros 5-6
 - .asmfunc directive 4-20, 4-24
 - AsmVal entry in cross-reference listing 9-5
 - assembler 1-3, 3-1 to 3-36
 - character strings 3-17
 - constants 3-14 to 3-16
 - cross-reference listings 3-5, 3-36
 - defined D-1
 - expressions 3-26 to 3-30
 - handling COFF sections 2-4 to 2-10
 - in the development flow 3-3
 - invoking 3-4
 - macros 5-1 to 5-24
 - marking function boundaries 4-24
 - options
 - @ 3-4
 - aa 3-4
 - ac 3-4
 - ad 3-5, 3-21
 - ahc 3-4
 - ahi 3-4
 - al 3-5, 3-31
 - apd 3-5
 - api 3-5
 - as 3-5
 - ax 3-5, 3-36
 - g 3-5
 - i 3-5, 3-7
 - me 3-5
 - ml 3-6
 - q 3-6
 - u 3-6
 - output listing 3-33, 4-14 to 4-15
 - directive listing 4-14, 4-36
 - enabling 4-52
 - false conditional block listing 4-14, 4-40
 - list options 4-14 to 4-15, 4-59
 - assembler, output listing (continued)
 - macro listing 4-14, 4-55 to 4-56
 - page eject 4-15, 4-61
 - page length 4-14, 4-51
 - page width 4-51
 - substitution symbol listing 4-15, 4-65
 - suppressing 4-14, 4-52
 - tab size 4-15, 4-69
 - title 4-15, 4-71
 - overview 3-2
 - relocation 2-14 to 2-15, 7-6 to 7-7
 - run-time relocation 2-16
 - sections directives 2-4 to 2-10
 - source listings 3-31 to 3-33
 - source statement format 3-10 to 3-13
 - symbols 3-18 to 3-25
- assembler directives 4-1 to 4-76
 - absolute lister 8-8
 - aligning the section program counter (SPC)
 - .align 4-13, 4-21
 - default directive 2-4
 - defining assembly-time symbols 4-18 to 4-19
 - .asg 4-18, 4-22
 - .cstruct 4-18, 4-31
 - .cunion 4-18, 4-31
 - .endstruct 4-18, 4-19, 4-31, 4-67
 - .endunion 4-18, 4-19, 4-31, 4-73
 - .equ 4-19, 4-62
 - .eval 4-18, 4-22
 - .label 4-18, 4-50
 - .set 4-19, 4-62
 - .struct 4-19, 4-67
 - .tag 4-18, 4-19, 4-31, 4-67, 4-73
 - .union 4-19, 4-73
 - defining sections 4-8 to 4-9
 - .bss 2-4, 4-8, 4-26
 - .data 2-4, 4-8, 4-34
 - .sect 2-4, 4-8, 4-61
 - .text 2-4, 4-8, 4-70
 - .usect 2-4, 4-8, 4-75
 - enabling conditional assembly 4-17
 - .break 4-17, 4-53
 - .else 4-17, 4-47
 - .elseif 4-17, 4-47
 - .endif 4-17, 4-47
 - .endloop 4-17, 4-53
 - .if 4-17, 4-47
 - .loop 4-17, 4-53

assembler directives (continued)

formatting the output listing 4-14 to 4-15

- .drlist* 4-14, 4-36
- .dmnolist* 4-14, 4-36
- enabling 4-14
- .fclist* 4-14, 4-40
- .fcnolist* 4-14, 4-40
- .length* 4-14, 4-51
- .list* 4-14, 4-52
- .mlist* 4-14, 4-56
- .mnolist* 4-14, 4-56
- .nolist* 4-14, 4-52
- .option* 4-14 to 4-15, 4-59
- .page* 4-15, 4-61
- page width 4-15
- .sslist* 4-15, 4-65
- .ssnolist* 4-15, 4-65
- .tab* 4-15, 4-69
- .title* 4-15, 4-71
- .width* 4-15, 4-51

initializing constants 4-10 to 4-12

- .bes* 4-10, 4-63
- .byte* 4-10, 4-27
- .char* 4-10, 4-27
- .double* 4-10, 4-35
- .field* 4-11, 4-41
- .float* 4-11, 4-43
- .half* 4-11, 4-46
- .int* 4-11, 4-49
- .long* 4-11, 4-49
- .short* 4-11, 4-46
- .space* 4-10, 4-63
- .string* 4-11, 4-66
- .uhalf* 4-11, 4-46
- .uint* 4-11, 4-49
- .ushort* 4-11, 4-46
- .uword* 4-11, 4-49
- .word* 4-11, 4-49

miscellaneous directives 4-20

- .asmfunc* 4-20, 4-24
- .clink* 4-20, 4-28
- .emsg* 4-20, 4-37
- .end* 4-20, 4-39
- .endasmfunc* 4-20, 4-24
- .mmsg* 4-20, 4-37
- .newblock* 4-20, 4-58
- .wmsg* 4-20, 4-37

assembler directives (continued)

referencing other files 4-16

- .copy* 4-16, 4-29
- .def* 4-16, 4-44
- .global* 4-16, 4-44
- .include* 4-16, 4-29
- .mlib* 4-16, 4-55
- .ref* 4-16, 4-44

summary table 4-2 to 4-7

assembly language development flow 1-2, 3-3, 6-3, 7-3

assembly-time constants 3-16, 4-62

defined D-1

assigning a value to a symbol 4-62

assigning substitution symbols with directives 4-22

assignment expressions 7-57 to 7-58

attributes 3-36, 7-29

attributes XML element C-9

autoinitialization

at load time 7-9

described 7-85

at run time 7-9

described 7-84

defined D-1

auxiliary entries A-16

defined D-1

available_space XML element C-9

-ax assembler option 3-5

cross-reference listing 3-36

B

-b linker option 7-8

B operand of *.option* directive 4-14, 4-59

banner XML element C-3

.bes directive 4-10, 4-63

big endian

defined D-2

object code 3-5

ordering 11-12

binary integer constants 3-14

binding 7-35

defined D-2

.block COFF debugging directive B-3

blocking 7-37

-boot hex conversion utility option 11-5, 11-28

boot SECTIONS specification 11-20

boot.obj module 7-82, 7-86
 boot-time copy table generated by linker 7-72 to 7-73
 -bootorg hex conversion utility option 11-5, 11-28
 -bootsection hex conversion utility option 11-5, 11-29
 .break directive 4-17, 4-53
 listing control 4-14, 4-36
 use in macros 5-14 to 5-15
 .bss directive 2-4, 4-8, 4-26
 linker definition 7-59
 .bss section 4-8, 4-26, A-3
 defined D-2
 holes 7-66 to 7-67
 initializing 7-67
 .byte directive 4-10, 4-27
 limiting listing with the .option directive 4-14, 4-59
 -byte hex conversion utility option 11-4

C

C code, linking 7-82 to 7-86
 example 7-87 to 7-89
 C hardware stack 7-83
 -c linker option 7-9, 7-59, 7-84
 C memory pool 7-11, 7-83
 -c name utility option 10-16
 C software stack 7-83
 C system stack 7-16
 C/C++ compiler 1-3
 defined D-2
 linking conventions 7-9
 special symbols A-13
 storage classes A-15
 symbol table entries A-13
 symbolic debugging A-1
 symbolic debugging directives B-1 to B-4
 C_DIR environment variable 7-12, 7-13
 _c_int00 7-9, 7-86
 callee_address XML element C-10
 callee_name XML element C-10
 caller_address XML entry C-11
 caller_list XML element C-11
 caller_object_component_ref XML entry C-11
 .char directive 4-10, 4-27
 character constants 3-15
 character strings 3-17
 cl6x -z command 7-4
 cl6x command 3-4
 .clink directive 4-20, 4-28
 COFF 2-1 to 2-20
 auxiliary entries A-16
 conversion to hexadecimal format 11-1 to 11-40
 default allocation 7-54 to 7-55
 defined D-2
 file headers A-4
 file structure A-2 to A-3
 initialized sections 2-6
 linker 7-1
 loading a program 2-17
 object file example A-3
 optional file header A-5
 relocation 2-14 to 2-15, A-9 to A-11
 relocation type A-10
 run-time relocation 2-16
 symbol table index A-9
 virtual address A-9
 section headers A-6 to A-8
 sections 2-2 to 2-3
 allocation 2-2
 assembler 2-4 to 2-10
 initialized 2-6
 linker 2-11 to 2-13
 named 2-6 to 2-7, 7-64
 special types 7-53
 uninitialized 2-4 to 2-5
 special symbols A-13
 storage classes A-15
 string table A-14
 symbol table 2-18 to 2-20, A-12 to A-16
 symbol values A-15
 symbolic debugging directives B-3
 syntax B-4
 technical details A-1 to A-16
 uninitialized sections 2-4 to 2-5
 command files
 appending to command line 3-4
 defined D-2
 hex conversion utility 11-5 to 11-6
 linker 7-4, 7-22 to 7-24
 constants in 7-24
 example 7-88
 reserved words 7-24
 comment field 3-13

- comments
 - defined D-2
 - extending past page width 4-51
 - in a linker command file 7-22
 - in assembly language source code 3-13
 - in macros 5-17
 - source statement format 3-13
 - common object file format
 - See *also* COFF
 - defined D-2
 - conditional blocks 4-47, 5-14 to 5-15
 - assembly directives 4-17, 4-47
 - in macros 5-14 to 5-15*
 - maximum nesting levels 5-14*
 - listing of false conditional blocks 4-40
 - conditional expressions 3-28
 - conditional linking
 - .clink assembler directive 4-28
 - j linker option 7-13
 - conditional processing, defined D-2
 - configured memory 7-55
 - defined D-2
 - constant elements in XML C-2
 - constants 3-14 to 3-16, 3-21 to 3-22
 - assembly-time 3-16, 4-62
 - binary integers 3-14
 - character 3-15
 - decimal integers 3-15
 - defined D-3
 - floating-point 4-35, 4-43
 - hexadecimal integers 3-15
 - in command files 7-24
 - octal integers 3-14
 - symbolic 3-23
 - \$ 3-23
 - processor symbols 3-24*
 - register symbols 3-23*
 - status registers 3-23*
 - symbols as 3-21
 - container elements in XML C-2
 - contents XML container element C-7
 - contents XML element C-6
 - split_section logical group C-7
 - .copy directive 3-7, 4-16, 4-29
 - copy files 4-29
 - ahc assembler option 3-4
 - .copy assembler directive 3-7
 - copy routine, general-purpose 7-75 to 7-76
 - COPY section 7-53
 - copy tables automatically generated by linker 7-71 to 7-72
 - contents 7-74 to 7-75
 - sections and symbols 7-76 to 7-77
 - copyright XML element C-3
 - cr linker option 7-9, 7-59, 7-85
 - creating holes 7-64 to 7-66
 - cross-reference lister 1-4, 9-1 to 9-6
 - creating the cross-reference listing 9-2
 - development flow 9-2
 - example 9-4
 - invoking 9-3
 - listings 3-5, 3-36
 - defined D-3*
 - producing with the .option directive 4-14 to 4-15, 4-59 to 4-60*
 - options
 - l 9-3
 - q 9-3
 - symbol attributes 9-5
 - xref6x command 9-3
 - .cstruct directive 4-18, 4-31
 - .cunion directive 4-18, 4-31
- ## D
- d archiver command 6-4
 - d name utility option 10-16
 - D operand of .option directive 4-14, 4-59
 - .data directive 2-4, 4-8, 4-34
 - linker definition 7-59
 - .data section 4-8, 4-34, A-3
 - defined D-3
 - decimal integer constants 3-15
 - .def directive 4-16, 4-44
 - identifying external symbols 2-18
 - default
 - allocation 7-54 to 7-55
 - fill value for holes 7-10
 - memory allocation 2-12
 - MEMORY configuration 7-54 to 7-55
 - MEMORY model 7-27
 - SECTIONS configuration 7-31, 7-54 to 7-55
 - defining macros 5-3 to 5-4
 - DefLn entry in cross-reference listing 9-5
 - development tools overview 1-2
 - directives

- assembler. See assembler directives
 - defined D-3
 - hex conversion utility. See ROMS directive; SECTIONS hex conversion utility directive
 - linker. See MEMORY directive; SECTIONS directive
 - directory search algorithm
 - assembler 3-7 to 3-9
 - linker 7-12
 - document element in XML C-3
 - far call trampoline list C-10
 - header elements C-3
 - input file list C-4
 - logical group list C-6
 - object component list C-5
 - placement map C-9
 - symbol table C-12
 - .double directive 4-10, 4-35
 - .drlist directive 4-14, 4-36
 - use in macros 5-20
 - .drnolist directive 4-14, 4-36
 - use in macros 5-20
 - DSECT section 7-53
 - dummy section 7-53
 - DWARF symbolic debugging directives
 - format B-2
 - syntax B-4
 - .dwattr DWARF debugging directive B-2
 - .dwcfa DWARF debugging directive B-2
 - .dwcie DWARF debugging directive B-2
 - .dwendentry DWARF debugging directive B-2
 - .dwendtag DWARF debugging directive B-2
 - .dwfde DWARF debugging directive B-2
 - .dwpsn DWARF debugging directive B-2
 - .dwtag DWARF debugging directive B-2
- E**
- e absolute lister option 8-3
 - e hex conversion utility option 11-5, 11-28
 - e linker option 7-9
 - edata linker symbol 7-59
 - .else directive 4-17, 4-47
 - use in macros 5-14 to 5-15
 - .elseif directive 4-17, 4-47
 - use in macros 5-14 to 5-15
 - .emsg directive 4-20, 4-37, 5-17
 - listing control 4-14, 4-36
 - end assembly 4-39
 - .end directive 4-20, 4-39
 - end linker symbol 7-59
 - END() linker operator 7-61
 - .endasmfunc directive 4-20, 4-24
 - .endblock COFF debugging directive B-3
 - .endfunc COFF debugging directive B-3
 - .endif directive 4-17, 4-47
 - use in macros 5-14 to 5-15
 - .endloop directive 4-17, 4-53
 - use in macros 5-14 to 5-15
 - .endm directive 5-3
 - .endstruct directive 4-18, 4-19, 4-31, 4-67
 - .endunion directive 4-18, 4-19, 4-31, 4-73
 - entry points
 - assigning values to 7-9
 - __c_int00 7-9, 7-86
 - default value 7-9
 - defined D-3
 - for C code 7-86
 - for the linker 7-9
 - __main 7-9
 - entry_point XML element C-3
 - environment variables
 - A_DIR 3-8, 7-12
 - C_DIR 7-11 to 7-13
 - .eos COFF debugging directive B-3
 - EPROM programmer 1-4
 - .equ directive 4-19, 4-62
 - error messages
 - generating 4-20
 - hex conversion utility 11-40
 - producing in macros 5-17
 - .etag COFF debugging directive B-3
 - etext linker symbol 7-59
 - .eval directive 4-18, 4-22
 - listing control 4-14, 4-36
 - use in macros 5-7
 - exclude hex conversion utility option 11-4
 - executable module, defined D-3
 - executable output 7-6
 - relocatable 7-7
 - expressions 3-26 to 3-30
 - absolute and relocatable 3-28 to 3-30
 - examples 3-29 to 3-30

- arithmetic operators 3-27
- conditional 3-28
- conditional operators 3-28
- defined D-3
- left-to-right evaluation 3-26
- linker 7-57 to 7-58
- overflow 3-27
- parentheses effect on evaluation 3-26
- precedence of operators 3-26
- relocatable symbols 3-28 to 3-30
- underflow 3-27
- well-defined 3-28
- external symbols 2-18, 3-28, 4-44
 - defined D-3

F

- f linker option 7-10
- f name utility option 10-16
- far calls 7-17
- far_call_trampoline XML container element C-10
- far_call_trampoline_list XML container element C-10
- .fclist directive 4-14, 4-40
 - listing control 4-14, 4-36
 - use in macros 5-19
- .fcnolist directive 4-14, 4-40
 - listing control 4-14, 4-36
 - use in macros 5-19
- .field directive 4-11, 4-41
- file
 - copy 3-4
 - include 3-4
- .file COFF debugging directive B-3
- file headers A-4
 - defined D-3
- file XML element C-4
- filenames
 - as character strings 3-17
 - copy/include files 3-7
 - extensions, changing defaults 8-3
 - macros, in macro libraries 5-13
- files ROMS specification 11-14
- fill hex conversion utility option 11-4, 11-25
- fill MEMORY specification 7-29
- fill ROMS specification 11-14
- fill value 7-66 to 7-67

- default 7-10
- setting 7-10
- fill_value XML element C-9
- filling holes 7-66 to 7-67
- .float directive 4-11, 4-43
- floating-point constants 4-35, 4-43
- .func COFF debugging directive B-3

G

- g assembler option 3-5
- g linker option 7-10
- g name utility option 10-16
- g object file display option 10-2
- .global directive 4-16, 4-44
 - identifying external symbols 2-18
- global symbols 7-10
 - defined D-4
 - making static with -h option 7-10
 - overriding -h option 7-10
- GROUP statement 7-50
 - defined D-4

H

- h linker option 7-10
- h name utility option 10-16
- H operand of .option directive 4-14, 4-59
- .half directive 4-11, 4-46
- hardware stack, C language 7-83
- heap linker option 7-11
 - .system section 7-11, 7-83
- help linker option 7-5
- hex conversion utility 1-4, 11-1 to 11-40
 - command files 11-5 to 11-6
 - invoking* 11-3, 11-5
 - ROMS directive* 11-5
 - SECTIONS directive* 11-5
 - configuring memory widths
 - defining memory word width (mem-width)* 11-4
 - specifying output width (romwidth)* 11-4

- hex conversion utility (continued)
 - controlling the boot table
 - describing the boot routine* `-bootsection` 11-5, 11-29
 - identifying bootable sections* `-boot` 11-5, 11-28
 - setting the entry point* `-e` 11-5, 11-28
 - setting the ROM address* `-bootorg` 11-5, 11-28
 - defined D-4
 - error messages 11-40
 - generating a map file 11-4
 - generating a quiet run 11-4
 - hex6x command 11-3
 - ignore specified section 11-4
 - image mode
 - defining the target memory* 11-25
 - filling holes* 11-4, 11-25
 - invoking* 11-4, 11-24
 - numbering output locations by bytes* 11-4
 - resetting address origin to 0* 11-4
 - in the development flow 11-2
 - invoking 11-3 to 11-6
 - from the command line* 11-3
 - in a command file* 11-3
 - memory width (`memwidth`) 11-8 to 11-9
 - exceptions* 11-8
 - options
 - `-a` 11-35
 - `-fill` 11-25
 - `-i` 11-36
 - `-image` 11-24
 - `-m` 11-37
 - `-map` 11-17 to 11-18
 - `-memwidth` 11-8
 - `-o` 11-22
 - order, restrictions* 11-12
 - `-q` 11-6
 - romwidth* 11-10
 - summary table* 11-4
 - `-t` 11-38
 - `-x` 11-39
 - ordering memory words 11-12
 - big-endian ordering* 11-4, 11-12
 - little-endian ordering* 11-4, 11-12
 - output filenames 11-4, 11-22
 - default filenames* 11-22
 - ROMS directive* 11-6
 - ROM width (`romwidth`) 11-9 to 11-11
 - hex conversion utility (continued)
 - ROMS directive 11-13 to 11-18
 - creating a map file of* 11-17 to 11-40
 - defining the target memory* 11-25
 - example* 11-16 to 11-18
 - parameters* 11-13 to 11-14
 - specifying output filenames* 11-6
 - SECTIONS directive 11-19 to 11-20
 - parameters* 11-19 to 11-20
 - target width 11-8
 - hex6x command 11-3
 - hexadecimal integers 3-15
 - holes 7-10, 7-64 to 7-67
 - creating 7-64 to 7-66
 - defined D-4
 - fill value 7-32, 11-14, 11-25
 - filling 7-66 to 7-67, 11-25
 - in output sections 7-64 to 7-67
 - in uninitialized sections 7-67
- ## I
- `-i` assembler option 3-5, 3-7
 - examples by operating system 3-8
 - maximum number per invocation 3-7
 - `-i` hex conversion utility option 11-4, 11-36
 - `-i` linker option 7-12
 - I MEMORY attribute 7-29
 - `.if` directive 4-17, 4-47
 - use in macros 5-14 to 5-15
 - `-image` hex conversion utility option 11-4, 11-24
 - `.include` directive 3-7, 4-16, 4-29
 - include files 3-4, 3-7, 4-29
 - incremental linking 7-80 to 7-81
 - defined D-4
 - initialized sections 2-6, 7-64
 - `.data` section 2-6, 4-34
 - defined D-4
 - `.sect` section 2-6, 4-61
 - subsections 2-6
 - `.text` section 2-6, 4-70
 - initializing
 - 16-bit integers 4-46
 - 32-bit integers 4-49
 - bytes 4-27
 - double-precision floating-point values 4-35
 - multiple-bit field 4-41
 - single-precision floating-point values 4-43
 - text 4-66

input
 linker 7-3, 7-25 to 7-26
 sections 7-37 to 7-39
 defined D-4
 input_file XML element C-4
 input_file_list XML container element C-4
 input_file_ref XML element C-5
 .int directive 4-11, 4-49
 Intel object format 11-1, 11-36
 invoking
 archiver 6-4
 assembler 3-4
 cross-reference lister 9-3
 hex conversion utility 11-3 to 11-6
 linker 7-4
 object file display utility 10-2
 name utility 10-16
 strip utility 10-17

J

-j linker option 7-13

K

keywords
 allocation parameters 7-35
 load 2-16, 7-35, 7-44 to 7-45
 run 2-16, 7-35, 7-44 to 7-45
 kind XML element C-4

L

-l cross-reference lister option 9-3
 -l linker option 7-11
 -l name utility option 10-16
 L operand of .option directive 4-14, 4-59
 label, case sensitivity 3-4
 .label directive 4-18, 4-50
 linker description 7-45
 label field 3-11
 labels 3-18
 defined D-4
 defined and referenced (cross-reference list) 3-36
 in assembly language source 3-11
 in macros 5-16

labels (continued)
 local 3-18 to 3-20, 4-58
 symbols used as 3-18
 syntax 3-10, 3-11
 using with .byte directive 4-27
 left-to-right evaluation (of expressions) 3-26
 legal expressions 3-28 to 3-30
 .length directive 4-14, 4-51
 listing control 4-14
 length MEMORY specification 7-29
 length ROMS specification 11-14
 length XML element C-9
 library search
 algorithm 7-11 to 7-13
 using alternate mechanism 7-20
 library-build utility 1-4
 .line COFF debugging directive B-3
 line-number table, line-number entries, *defined* D-4
 link_time XML element C-3
 link_timestamp XML element C-3
 linker 1-3, 7-1 to 7-90
 angle bracket operator 7-43
 >> operator 7-41
 | operator 7-40
 allocating archive member to output section 7-43
 allocation to multiple memory ranges 7-40
 assigning symbols 7-56
 assignment expressions 7-57 to 7-58
 automatic splitting of output sections 7-41
 C code 7-82 to 7-86
 checking consistency of run and load allocators 7-51
 cl6x -z command 7-4
 COFF 7-1
 command files 7-4, 7-22 to 7-24
 example 7-88
 configured memory 7-55
 defined D-4
 END() operator 7-61
 example 7-87 to 7-90
 generated copy tables. *See* linker-generated copy tables
 GROUP statement 7-48, 7-50
 handling COFF sections 2-11 to 2-13
 in the development flow 7-3
 input 7-3, 7-22 to 7-24
 invoking 7-4
 keywords 7-24, 7-44 to 7-47

linker (continued)

- linking C code 7-9, 7-82 to 7-86
- LOAD_END() operator 7-61
- LOAD_SIZE() operator 7-61
- LOAD_START() operator 7-61
- loading a program 2-17
- MEMORY directive 2-11, 7-27 to 7-30
- nesting UNIONS and GROUPs 7-50
- object libraries 7-25 to 7-26
- operators 7-58
- options
 - a 7-6
 - abs 7-7
 - ar 7-7
 - args 7-8
 - b 7-8
 - c 7-9, 7-84
 - cr 7-9, 7-85
 - e 7-9
 - f 7-10
 - g 7-10
 - h 7-10
 - heap 7-11
 - help 7-5
 - i 7-12
 - j 7-13
 - l 7-11
 - m 7-14 to 7-15
 - o 7-15
 - priority 7-20
 - r 7-7
 - s 7-16
 - stack 7-16
 - summary table 7-5
 - trampolines 7-17
 - u 7-19
 - w 7-20
 - x 7-20
 - xml_link_info 7-21
- output 7-3, 7-15, 7-87
- overview 7-2
- partial linking 7-80 to 7-81
- RUN_END() operator 7-61
- RUN_SIZE() operator 7-61
- RUN_START() operator 7-61
- section run-time address 7-44 to 7-47
- sections 2-13
 - output 7-54
 - special 7-53
- SECTIONS directive 2-11, 7-31 to 7-43

linker (continued)

- SIZE() operator 7-61
- START() operator 7-61
- symbols 2-18 to 2-20, 7-59
- table() operator 7-72, 7-73
- unconfigured memory, overlaying 7-53
- UNION statement 7-48 to 7-49, 7-70

linker directives

- MEMORY 2-11, 7-27 to 7-30
- SECTIONS 2-11, 7-31 to 7-43

linker-generated copy tables 7-68 to 7-79

- automatic 7-71 to 7-72
- boot-loaded application process 7-68
 - alternative approach 7-69
- boot-time copy table 7-72 to 7-73
- contents 7-74 to 7-75
- general-purpose copy routine 7-75 to 7-76
- overlay management 7-78 to 7-79
- overlay management example 7-70
- sections and symbols 7-76 to 7-77
- splitting object components 7-78 to 7-79
- table() operator 7-72
 - manage object components 7-73

.list directive 4-14, 4-52

lister

- absolute 8-1 to 8-10
- cross-reference 9-1 to 9-6

listing

- control 4-14 to 4-15, 4-52, 4-56, 4-59, 4-61, 4-71
- cross-reference listing 4-14, 4-59
- file 4-14 to 4-15
 - creating with the -al option 3-5
 - defined D-5
 - format 3-31 to 3-33
- page eject 4-15
- page size 4-14, 4-51

little endian

- defined D-5
- ordering 11-12

LnkVal entry in cross-reference listing 9-5

load address of a section 7-44 to 7-45

- referring to with a label 7-45 to 7-47

load linker keyword 2-16, 7-44 to 7-45

load_address XML element

- logical_group logical group C-6
- object component list C-5

LOAD_END() linker operator 7-61

LOAD_SIZE() linker operator 7-61
 LOAD_START() linker operator 7-61, 7-69
 load6x command 2-17
 loader, defined D-5
 loading a program 2-17
 local labels 3-18 to 3-20
 logical operators 3-27
 logical_group XML container element C-6
 logical_group_list XML container element C-6
 logical_group_ref XML element C-6, C-7, C-9
 .long directive 4-11, 4-49
 limiting listing with the .option directive 4-14 to
 4-15, 4-59 to 4-60
 .loop directive 4-17, 4-53
 use in macros 5-14 to 5-15

M

–m hex conversion utility option 11-4, 11-37
 –m linker option 7-14 to 7-15
 M operand of .option directive 4-14, 4-59
 .macro directive 4-54, 5-3 to 5-4
 summary table 5-23 to 5-24
 macros 5-1 to 5-24
 conditional assembly 5-14 to 5-15
 defined, macro D-5
 defining a macro 5-3 to 5-4
 description 5-2
 directives summary 5-23 to 5-24
 disabling macro expansion listing 4-14, 4-59
 formatting the output listing 5-19 to 5-20
 labels 5-16
 macro comments 5-4, 5-17
 macro libraries 5-13, 6-2
 defined D-5
 nested macros 5-21 to 5-22
 parameters 5-5 to 5-12
 producing messages 5-17 to 5-18
 recursive macros 5-21 to 5-22
 substitution symbols 5-5 to 5-12
 terms defined D-5
 using a macro 5-2
 magic number, defined D-5
 _main 7-9
 malloc() function 7-11, 7-83
 map file 7-14 to 7-15, 11-17 to 11-18
 defined D-5
 example 7-89, 11-17
 –map hex conversion utility option 11-4
 –me assembler option 3-5
 .member COFF debugging directive B-3
 memory
 allocation 7-54 to 7-55
 default 2-12
 map 2-13
 defined D-5
 model 7-27
 named 7-36 to 7-37
 pool, C language 7-11, 7-83
 unconfigured 7-27
 MEMORY linker directive
 COFF overview 2-11
 default model 7-27, 7-54 to 7-55
 syntax 7-27 to 7-30
 memory ranges, allocation to multiple 7-40
 memory widths
 memory width (memwidth) 11-8 to 11-9
 exceptions 11-8
 ordering memory words 11-12
 big-endian ordering 11-12
 little-endian ordering 11-12
 ROM width (romwidth) 11-9 to 11-11
 target width 11-8
 memory words, ordering 11-12
 big-endian 11-12
 little-endian 11-12
 memory_area XML container element C-9
 –memwidth hex conversion utility option 11-4, 11-8
 memwidth ROMS specification 11-8, 11-14
 .mexit directive 5-3
 –ml assembler option 3-6
 .mlib directive 4-16, 4-55 to 4-56, 5-13
 use in macros 3-7
 .mlist directive 4-14, 4-56
 listing control 4-14, 4-36
 use in macros 5-19
 .mmsg directive 4-20, 4-37, 5-17
 listing control 4-14, 4-36
 mnemonic, defined D-5
 mnemonic field 3-12
 syntax 3-10
 .mnlst directive 4-14, 4-56
 listing control 4-14, 4-36
 use in macros 5-19

model statement 5-3
 defined D-5
Motorola-S object format 11-1, 11-37
-mv assembler option 3-6

N

-n name utility option 10-16
N operand of .option directive 4-14, 4-59
name MEMORY specification 7-28
name utility
 invoking 10-16
 options 10-16
name XML element
 entry_point header element C-3
 input file list C-4
 logical_group logical group C-6
 object component list C-5
 overlay logical group C-6
 placement map C-9
 split_section logical group C-7
 symbol table C-12
named memory 7-36 to 7-37
named sections 2-6 to 2-7, A-3
 defined D-6
 .sect directive 2-7, 4-61
 .usect directive 2-7, 4-75
nested macros 5-21 to 5-22
.newblock directive 4-20, 4-58
nm6x command 10-16
.nolist directive 4-14, 4-52
NOLOAD section 7-53

O

-o hex conversion utility option 11-4
-o linker option 7-15
-o name utility option 10-16
-o object file display option 10-2
O operand of .option directive 4-14, 4-59
object code (source listing) 3-32
object file
 defined D-6
 library 7-25 to 7-26

object file display utility
 invoking 10-2
 options
 -g 10-2
 -o 10-2
 -x 10-2
object formats
 address bits 11-34
 ASCII-Hex 11-1, 11-35
 selecting 11-4
 Intel 11-1, 11-36
 selecting 11-4
 Motorola-S 11-1, 11-37
 selecting 11-4
 output width 11-34
 Tektronix 11-1, 11-39
 selecting 11-4
 TI-Tagged 11-1, 11-38
 selecting 11-4
object libraries 7-11 to 7-13, 7-25 to 7-26, 7-83
 defined D-6
 using the archiver to build 6-2
object_component XML element C-5
object_component_list XML container element C-5
object_component_ref XML element C-6, C-7
octal integer constants 3-14
ofd6x command 10-2
operands
 defined D-6
 field 3-13
 label 3-18
 local label 3-18 to 3-20
 source statement format 3-13
operator precedence order 3-27
.option directive 4-14 to 4-15, 4-59
optional file header A-5
 defined D-6
options
 absolute lister 8-3
 archiver 6-4
 assembler 3-4
 cross-reference lister 9-3
 defined D-6
 hex conversion utility 11-3 to 11-4
 linker 7-5 to 7-21
 name utility 10-16
 object file display utility 10-2
 strip utility 10-17

- order hex conversion utility option, restrictions 11-12
- order L hex conversion utility option 11-4
- order M hex conversion utility option 11-4
- ordering memory words 11-12
 - big-endian ordering 11-12
 - little-endian ordering 11-12
- origin MEMORY specification 7-29
- origin ROMS specification 11-13
- origin XML element C-9
- output, cross-reference lister 9-3
- output
 - absolute lister 8-3
 - archiver 6-2
 - assembler 3-1
 - executable 7-6
 - relocatable* 7-7
 - hex conversion utility 11-4, 11-16
 - linker 7-3, 7-15, 7-87
 - listing 4-14 to 4-15
 - module, defined D-6
 - module name (linker) 7-15
 - sections
 - allocating archive member to* 7-43
 - allocation* 7-34 to 7-43
 - defined* D-6
 - displaying a message* 7-20
 - methods* 7-54 to 7-55
 - splitting* 7-41
- output_file XML element C-3
- overflow (in expression) 3-27
- overlay XML logical group C-6
- overlying sections 7-48 to 7-49
 - managing linker-generated copy tables 7-78 to 7-79

P

- p name utility option 10-16
- p strip utility option 10-17
- paddr SECTIONS specification 11-20
- page
 - eject 4-61
 - length 4-51
 - title 4-71
 - width 4-51
- .page directive 4-15, 4-61

- page_id XML element C-9
- parentheses in expressions 3-26
- partial linking 7-80 to 7-81
 - defined D-6
- path XML element C-4
- perform preprocessing on assembly files
 - write list of dependency lines 3-5
 - write list of files included with .include 3-5
- placement_map XML container element C-9
- precedence groups 3-26
 - linker 7-58
- predefined names
 - d assembler option 3-5
 - undefining with –u assembler option 3-6
- priority linker option 7-20
- processor symbols 3-24

Q

- q absolute lister option 8-3
- q archiver option 6-5
- q assembler option 3-6
- q cross-reference lister option 9-3
- q hex conversion utility option 11-4, 11-6
- q name utility option 10-16
- quiet run
 - absolute lister 8-3
 - archiver 6-5
 - assembler 3-6
 - cross-reference lister 9-3
 - defined D-6
 - hex conversion utility 11-6

R

- r archiver command 6-4
- r linker option 7-7, 7-80 to 7-81
- R MEMORY attribute 7-29
- r name utility option 10-16
- R operand of .option directive 4-14, 4-59
- recursive macros 5-21 to 5-22
- .ref directive 4-16, 4-44
 - identifying external symbols 2-18
- reference elements in XML C-2
- RefLn entry in cross-reference listing 9-5
- register symbols 3-23

- relational operators, in conditional expressions 3-28
 - relocatable output module 7-7
 - executable 7-7
 - relocation 2-14 to 2-15, 7-6 to 7-7
 - at run time 2-16
 - capabilities 7-6 to 7-7
 - defined D-6
 - information A-9 to A-11
 - reserved words, linker 7-24
 - resetting local labels 4-58
 - ROM device address 11-32
 - ROM width (romwidth) 11-9 to 11-11
 - romname ROMS specification 11-13
 - ROMS directive 11-13 to 11-18
 - creating map file of 11-17 to 11-18
 - example 11-16 to 11-18
 - parameters 11-13 to 11-14
 - romwidth hex conversion utility option 11-4, 11-10
 - romwidth ROMS specification 11-10, 11-14
 - RTYP entry in cross-reference listing 9-5
 - run address of a section 7-44 to 7-45
 - run linker keyword 2-16, 7-44 to 7-45
 - run time
 - initialization 7-82
 - support 7-83
 - run-time-support library 7-82, 7-86
 - run_address XML element
 - logical_group logical group C-6
 - object component list C-5
 - overlay logical group C-6
 - RUN_END() linker operator 7-61
 - RUN_SIZE() linker operator 7-61
 - RUN_START() linker operator 7-61, 7-69
- S**
- s archiver option 6-5
 - s linker option 7-16, 7-80
 - search libraries
 - using –priority linker option 7-20
 - using alternate mechanism 7-20
 - .sect directive 2-4, 4-8, 4-61
 - .sect section 4-8, 4-61
 - section
 - allocation into memory 7-54 to 7-55
 - COFF 2-1 to 2-20
 - creating your own 2-6 to 2-7
 - default allocation 7-54 to 7-55
 - defined D-7
 - description 2-2 to 2-3
 - directives 2-8 to 2-10
 - default 2-4
 - header A-6 to A-8
 - defined D-7
 - input sections 7-32
 - initialized 2-6
 - named 2-2, 2-6 to 2-7
 - number A-16
 - overlying with UNION statement 7-48 to 7-49
 - relocation 2-14 to 2-15
 - at run time 2-16
 - special types 7-53
 - specification 7-32
 - specifying a run-time address 7-44 to 7-45
 - specifying linker input sections 7-37 to 7-39
 - uninitialized 2-4 to 2-5
 - initializing 7-67
 - specifying a run address 7-45
 - SECTIONS hex conversion utility directive 11-19 to 11-20
 - SECTIONS linker directive 7-31 to 7-43
 - alignment 7-37
 - allocating archive member to output section 7-43
 - allocation 7-34 to 7-43
 - allocation using multiple memory ranges 7-40
 - binding 7-35
 - blocking 7-37
 - COFF overview 2-11
 - default allocation 7-54 to 7-55
 - fill value 7-32
 - GROUP 7-50
 - input sections 7-32, 7-37 to 7-39
 - .label directive 7-45 to 7-47
 - load allocation 7-32
 - memory 7-36 to 7-37
 - named memory 7-36 to 7-37
 - reserved words 7-24
 - run allocation 7-32
 - section specification 7-32
 - section type 7-32

- SECTIONS linker directive (continued)
 - specifying
 - run-time address* 2-16, 7-44 to 7-47
 - two addresses* 2-16, 7-44 to 7-45
 - splitting of output sections 7-41
 - syntax 7-31 to 7-32
 - uninitialized sections 7-45
 - UNION 7-48 to 7-52
 - use with MEMORY directive 7-27
- .set directive 4-19, 4-62
- .setsect assembler directive 8-8
- .setsym assembler directive 8-8
- .short directive 4-11, 4-46
- sign-extend, defined D-7
- size XML element
 - logical_group logical group C-6
 - object component list C-5
 - overlay logical group C-6
- size XML entry C-9
- SIZE() linker operator 7-61, 7-69
- sname SECTIONS specification 11-19
- source file
 - assembler for object file display utility 10-2
 - defined D-7
 - directory 3-7 to 3-9
- source listings 3-31 to 3-33
- source statement
 - field (source listing) 3-32
 - format 3-10
 - comment field* 3-13
 - label field* 3-11
 - mnemonic field* 3-12
 - operand field* 3-13
 - unit specifier field* 3-12
 - number (source listing) 3-31 to 3-33
- .space directive 4-10, 4-63
- SPC (section program counter) 2-8
 - aligning
 - by creating a hole* 7-64
 - to byte boundaries* 4-13
 - to word boundaries* 4-21
 - assembler's effect on 2-8 to 2-10
 - assigning label 3-11
 - defined D-7
 - linker symbol 7-57, 7-64
 - predefined symbol for 3-23
 - SPC (section program counter) 2-8
 - value
 - associated with labels* 3-11
 - shown in source listings* 3-31
 - special section types 7-53
 - special symbols in the symbol table A-13
 - split_section XML logical group C-7
 - .sslist directive 4-15, 4-65
 - listing control 4-14, 4-36
 - use in macros 5-19
 - .ssnolist directive 4-15, 4-65
 - listing control 4-14, 4-36
 - use in macros 5-19
 - stack linker option 7-16
 - .stack section 7-83
 - __STACK_SIZE 7-16, 7-59
 - .stag COFF debugging directive B-3
 - stag structure tag 4-18, 4-19, 4-31, 4-67
 - START() linker operator 7-61
 - start_address XML entry C-9
 - static symbols, creating with -h option 7-10
 - static variables A-12
 - defined D-7
 - status registers 3-23
 - storage classes A-15
 - defined D-7
 - .string directive 4-11, 4-66
 - limiting listing with the .option directive 4-14, 4-59
 - string elements in XML C-2
 - string functions (substitution symbols)
 - \$firstch 5-8
 - \$iscons 5-8
 - \$isdefed 5-8
 - \$ismember 5-8
 - \$isname 5-8
 - \$ispreg 5-8
 - \$isreg 5-8
 - \$isrreg 5-8
 - \$lastch 5-8
 - \$symcmp 5-8
 - \$symlen 5-8
 - string table A-14
 - defined D-7
 - strip utility
 - invoking 10-17
 - option 10-17
 - strip6x command 10-17

- stripping
 - line number entries 7-16
 - symbol table information 7-16, 10-17
 - symbolic debugging information 10-17
- .struct directive 4-19, 4-67
- structure
 - defined D-7
 - stag 4-18, 4-19, 4-31, 4-67
- style and symbol conventions iv to vi
- subsection
 - defined D-7
 - initialized 2-6
 - overview 2-7
- substitution symbols 3-25
 - arithmetic operations on 4-18, 5-7
 - as local variables in macros 5-12
 - assigning character strings to 3-25, 4-18, 4-22
 - built-in functions 5-7 to 5-8
 - directives that define 5-6
 - expansion listing 4-15, 4-65
 - forcing substitution 5-9 to 5-10
 - in macros 5-5 to 5-12
 - maximum number per macro 5-5
 - passing commas and semicolons 5-5
 - recursive substitution 5-9
 - subscripted substitution 5-10
 - .var directive 5-12
- .sym COFF debugging directive B-3
- symbol
 - assembler-defined 2-18 to 2-20, 3-5
 - assembly language usage 3-18 to 3-25
 - assigning values to 4-62, 4-73
 - at link time 7-56 to 7-63
 - attributes 3-36
 - case 3-4
 - character strings 3-17
 - cross-reference lister 9-5
 - defined D-7
 - defined only for C support 7-59
 - definitions (cross-reference list) 3-36
 - external 2-18, 4-44
 - global 7-10
 - in COFF file 2-18 to 2-20
 - linker-defined 7-59
 - names A-14
 - number of statements that reference 3-36
 - predefined 3-23
 - reserved words 7-24
 - setting to a constant value 3-21
- symbol (continued)
 - statement number that defines 3-36
 - substitution 3-25
 - table 2-19
 - creating entries 2-19
 - defined D-8
 - index A-9
 - placing unresolved symbols in 7-19
 - special symbols used in A-13
 - stripping entries 7-16
 - structure and content A-12 to A-16
 - symbol values A-15
 - undefining assembler-defined symbols 3-6
 - unresolved 7-19
 - used as labels 3-18
 - value assigned 3-36
- symbol XML container element C-12
- symbol_table XML container element C-12
- symbolic constants 3-23
 - \$ 3-23
 - defining 3-21
 - processor symbols 3-24
 - register symbols 3-23
 - status registers 3-23
- symbolic debugging B-1 to B-4
 - assembly source 3-34
 - COFF directives B-3
 - COFF directives syntax B-4
 - defined D-8
 - directives B-1 to B-4
 - disable merge for linker (-b option) 7-8
 - DWARF directives B-2
 - DWARF directives syntax B-4
 - producing error messages in macros 5-17
 - put all symbols in symbol table (-as assembler option) 3-5
 - stripping symbol table 7-16
- syntax of assignment statements 7-56
- __SYSTEMEM_SIZE 7-11, 7-59
- system stack, C language 7-16, 7-83

T

- t archiver command 6-4
- t hex conversion utility option 11-4, 11-38
- t name utility option 10-16
- T operand of .option directive 4-15, 4-59
- .tab directive 4-15, 4-69

table() linker operator 7-72
 used to manage object components 7-73

.tag directive 4-18, 4-19, 4-31, 4-67, 4-73

target memory
 configuration 7-22
 defined D-8
 loading a program into 7-9
 model 7-27

target width 11-8

Tektronix object format 11-1, 11-39

.text directive 2-4, 4-8, 4-70
 linker definition 7-59

.text section 4-8, 4-70, A-3
 defined D-8

TI-Tagged object format 11-1, 11-38

.title directive 4-15, 4-71

trampoline_address XML element C-10

trampoline_call_site XML element C-11

trampoline_object_component_ref XML element C-10

trampolines, generating 7-17

--trampolines linker option 7-17

U

u archiver command 6-5

-u assembler option 3-6

-u linker option 7-19

-u name utility option 10-16

.uhalf directive 4-11, 4-46

.uint directive 4-11, 4-49

unconfigured memory 7-27
 defined D-8
 overlaying 7-53

underflow (in expression) 3-27

uninitialized sections 2-4 to 2-5, 7-64
 .bss section 2-5, 4-26
 defined D-8
 initialization of 7-67
 specifying a run address 7-45
 .usect section 2-5, 4-75

union
 .tag 4-73
 .cunion directive 4-31
 .endunion directive 4-31
 .tag directive 4-31

.union directive 4-19, 4-73

UNION statement 7-48 to 7-52
 defined D-8
 memory overlay example 7-70

unit specifier
 field 3-12
 source statement format 3-12

unused_space XML element C-9

usage_details XML element C-9

.usect directive 2-4, 4-8, 4-75

used_space XML element C-9

user messages defined with directives 4-37

.ushort directive 4-11, 4-46

.utag COFF debugging directive B-3

.uword directive 4-11, 4-49

V

-v archiver option 6-5

value XML element C-12

.var directive 4-77, 5-12
 listing control 4-14, 4-36

variables, local, substitution symbols used as 5-12

W

-w linker option 7-20

W MEMORY attribute 7-29

W operand of .option directive 4-15, 4-59

well-defined expressions 3-28
 defined D-8

.width directive 4-15, 4-51
 listing control 4-14

.wmsg directive 4-20, 4-37, 5-17
 listing control 4-14, 4-36

word, defined D-8

word alignment 4-21

.word directive 4-11, 4-49
 limiting listing with the .option directive 4-14 to 4-15, 4-59 to 4-60

X

x archiver command 6-4

-x hex conversion utility option 11-4, 11-39

-x linker option 7-20

- X MEMORY attribute 7-29
- x object file display option 10-2
- X operand of .option directive 4-15, 4-59
- XML element C-12
 - address entry C-3
 - allocated_space element C-9
 - attributes element C-9
 - available_space element C-9
 - banner element C-3
 - callee_address element C-10
 - callee_name element C-10
 - caller_address entry C-11
 - caller_list element C-11
 - caller_object_component_ref entry C-11
 - contents container element
 - overlay logical group* C-7
 - split_section logical group* C-7
 - contents element, logical_group logical group C-6
 - copyright element C-3
 - entry_point element C-3
 - far_call_trampoline container element C-10
 - far_call_trampoline_list container element C-10
 - file element C-4
 - fill_value element C-9
 - input_file element C-4
 - input_file_list container element C-4
 - input_file_ref element C-5
 - kind element C-4
 - length element C-9
 - link_info C-3
 - link_time element C-3
 - link_timestamp element C-3
 - load_address element
 - logical_group logical group* C-6
 - object component list* C-5
 - logical_group container element C-6
 - logical_group_list container element C-6
 - logical_group_ref element C-6
 - overlay logical group* C-7
 - logical_group_ref entry C-9
 - memory_area container element C-9
 - name element
 - entry_point header element* C-3
 - input file list* C-4
 - object component list* C-5
 - overlay logical group* C-6
 - placement map* C-9
 - split_section logical group* C-7
 - symbol table* C-12
- XML element (continued)
 - object_component element C-5
 - object_component_list container element C-5
 - object_component_ref element C-6
 - overlay logical group* C-7
 - origin element C-9
 - output_file element C-3
 - overlay logical group C-6
 - page_id element C-9
 - path element C-4
 - placement_map container element C-9
 - run_address element
 - logical_group logical group* C-6
 - object component list* C-5
 - overlay logical group* C-6
 - size element
 - logical_group logical group* C-6
 - overlay logical group* C-6
 - size entry C-9
 - size XML element, object component list C-5
 - split_section logical group C-7
 - start_address entry C-9
 - symbol_table container element C-12
 - trampoline_address element C-10
 - trampoline_call_site element C-11
 - trampoline_object_component_ref element C-10
 - unused_space element C-9
 - usage_details element C-9
 - used_space element C-9
 - value element C-12
- XML link information file
 - description C-1
 - document elements detailed C-3
 - far call trampoline list* C-10
 - header elements* C-3
 - input file list* C-4
 - logical group list* C-6
 - object component list* C-5
 - placement map* C-9
 - symbol table* C-12
 - element types C-2
 - name XML element, logical_group logical group C-6
 - xml_link_info linker option* 7-21, C-1
- xref6x command 9-3

Z

- zero hex conversion utility option 11-4