

TMS320C6000 Optimizing Compiler v7.4

User's Guide



Literature Number: SPRU187U
July 2012

Preface	11
1 Introduction to the Software Development Tools	14
1.1 Software Development Tools Overview	15
1.2 C/C++ Compiler Overview	16
1.2.1 ANSI/ISO Standard	16
1.2.2 Output Files	17
1.2.3 Compiler Interface	17
1.2.4 Utilities	17
2 Using the C/C++ Compiler	18
2.1 About the Compiler	19
2.2 Invoking the C/C++ Compiler	19
2.3 Changing the Compiler's Behavior With Options	20
2.3.1 Frequently Used Options	30
2.3.2 Miscellaneous Useful Options	31
2.3.3 Run-Time Model Options	32
2.3.4 Selecting Target CPU Version (--silicon_version Option)	33
2.3.5 Symbolic Debugging and Profiling Options	34
2.3.6 Specifying Filenames	35
2.3.7 Changing How the Compiler Interprets Filenames	35
2.3.8 Changing How the Compiler Processes C Files	36
2.3.9 Changing How the Compiler Interprets and Names Extensions	36
2.3.10 Specifying Directories	36
2.3.11 Assembler Options	37
2.3.12 Dynamic Linking	38
2.3.13 Deprecated Options	39
2.4 Controlling the Compiler Through Environment Variables	39
2.4.1 Setting Default Compiler Options (C6X_C_OPTION)	39
2.4.2 Naming an Alternate Directory (C6X_C_DIR)	40
2.5 Precompiled Header Support	41
2.5.1 Automatic Precompiled Header	41
2.5.2 Manual Precompiled Header	41
2.5.3 Additional Precompiled Header Options	41
2.6 Controlling the Preprocessor	42
2.6.1 Predefined Macro Names	42
2.6.2 The Search Path for #include Files	43
2.6.3 Generating a Preprocessed Listing File (--preproc_only Option)	44
2.6.4 Continuing Compilation After Preprocessing (--preproc_with_compile Option)	44
2.6.5 Generating a Preprocessed Listing File With Comments (--preproc_with_comment Option)	44
2.6.6 Generating a Preprocessed Listing File With Line-Control Information (--preproc_with_line Option)	44
2.6.7 Generating Preprocessed Output for a Make Utility (--preproc_dependency Option)	45
2.6.8 Generating a List of Files Included With the #include Directive (--preproc_includes Option)	45
2.6.9 Generating a List of Macros in a File (--preproc_macros Option)	45
2.7 Understanding Diagnostic Messages	45
2.7.1 Controlling Diagnostics	46

2.7.2	How You Can Use Diagnostic Suppression Options	47
2.8	Other Messages	48
2.9	Generating Cross-Reference Listing Information (--gen_acp_xref Option)	48
2.10	Generating a Raw Listing File (--gen_acp_raw Option)	48
2.11	Using Inline Function Expansion	50
2.11.1	Inlining Intrinsic Operators	50
2.11.2	Automatic Inlining	50
2.11.3	Unguarded Definition-Controlled Inlining	50
2.11.4	Guarded Inlining and the _INLINE Preprocessor Symbol	51
2.11.5	Inlining Restrictions	52
2.12	Interrupt Flexibility Options (--interrupt_threshold Option)	53
2.13	Linking C6400 Code With C6200/C6700/Older C6400 Object Code	54
2.14	Using Interlist	54
2.15	Generating and Using Performance Advice	55
2.16	Controlling Application Binary Interface	55
2.17	Enabling Entry Hook and Exit Hook Functions	57
3	Optimizing Your Code	58
3.1	Invoking Optimization	59
3.2	Optimizing Software Pipelining	60
3.2.1	Turn Off Software Pipelining (--disable_software_pipelining Option)	61
3.2.2	Software Pipelining Information	61
3.2.3	Collapsing Prologs and Epilogs for Improved Performance and Code Size	65
3.3	Redundant Loops	67
3.4	Utilizing the Loop Buffer Using SPLOOP on C6400+, C6740, and C6600	68
3.5	Reducing Code Size (--opt_for_space (or -ms) Option)	68
3.6	Performing File-Level Optimization (--opt_level=3 option)	70
3.6.1	Controlling File-Level Optimization (--std_lib_func_def Options)	70
3.6.2	Creating an Optimization Information File (--gen_opt_info Option)	70
3.7	Performing Program-Level Optimization (--program_level_compile and --opt_level=3 options)	71
3.7.1	Controlling Program-Level Optimization (--call_assumptions Option)	71
3.7.2	Optimization Considerations When Mixing C/C++ and Assembly	72
3.8	Using Feedback Directed Optimization	73
3.8.1	Feedback Directed Optimization	73
3.8.2	Profile Data Decoder	75
3.8.3	Feedback Directed Optimization API	76
3.8.4	Feedback Directed Optimization Summary	76
3.9	Using Profile Information to Get Better Program Cache Layout and Analyze Code Coverage	77
3.9.1	Background and Motivation	77
3.9.2	Code Coverage	78
3.9.3	What Performance Improvements Can You Expect to See?	79
3.9.4	Program Cache Layout Related Features and Capabilities	79
3.9.5	Program Instruction Cache Layout Development Flow	80
3.9.6	Comma-Separated Values (CSV) Files with Weighted Call Graph (WCG) Information	83
3.9.7	Linker Command File Operator - unordered()	83
3.9.8	Things To Be Aware Of	86
3.10	Indicating Whether Certain Aliasing Techniques Are Used	87
3.10.1	Use the --aliased_variables Option When Certain Aliases are Used	87
3.10.2	Use the --no_bad_aliases Option to Indicate That These Techniques Are Not Used	87
3.10.3	Using the --no_bad_aliases Option With the Assembly Optimizer	88
3.11	Prevent Reordering of Associative Floating-Point Operations	88
3.12	Use Caution With asm Statements in Optimized Code	89
3.13	Automatic Inline Expansion (--auto_inline Option)	89
3.14	Using Performance Advice to Optimize your Code	90

3.14.1	Types of Advice	91
3.14.2	Using the Interlist Feature With Optimization	96
3.14.3	Debugging and Profiling Optimized Code	99
3.14.4	Controlling Code Size Versus Speed	100
3.14.5	What Kind of Optimization Is Being Performed?	101
4	Using the Assembly Optimizer	104
4.1	Code Development Flow to Increase Performance	105
4.2	About the Assembly Optimizer	106
4.3	What You Need to Know to Write Linear Assembly	107
4.3.1	Linear Assembly Source Statement Format	108
4.3.2	Register Specification for Linear Assembly	109
4.3.3	Functional Unit Specification for Linear Assembly	111
4.3.4	Using Linear Assembly Source Comments	111
4.3.5	Assembly File Retains Your Symbolic Register Names	112
4.4	Assembly Optimizer Directives	113
4.4.1	Instructions That Are Not Allowed in Procedures	126
4.5	Avoiding Memory Bank Conflicts With the Assembly Optimizer	127
4.5.1	Preventing Memory Bank Conflicts	128
4.5.2	A Dot Product Example That Avoids Memory Bank Conflicts	129
4.5.3	Memory Bank Conflicts for Indexed Pointers	132
4.5.4	Memory Bank Conflict Algorithm	133
4.6	Memory Alias Disambiguation	133
4.6.1	How the Assembly Optimizer Handles Memory References (Default)	133
4.6.2	Using the --no_bad_aliases Option to Handle Memory References	133
4.6.3	Using the .no_mdep Directive	133
4.6.4	Using the .mdep Directive to Identify Specific Memory Dependencies	134
4.6.5	Memory Alias Examples	135
5	Linking C/C++ Code	137
5.1	Invoking the Linker Through the Compiler (-z Option)	138
5.1.1	Invoking the Linker Separately	138
5.1.2	Invoking the Linker as Part of the Compile Step	139
5.1.3	Disabling the Linker (--compile_only Compiler Option)	139
5.2	Linker Code Optimizations	140
5.2.1	Generating Function Subsections (--gen_func_subsections Compiler Option)	140
5.2.2	Conditional Linking	140
5.3	Controlling the Linking Process	141
5.3.1	Including the Run-Time-Support Library	141
5.3.2	Run-Time Initialization	142
5.3.3	Global Object Constructors	142
5.3.4	Specifying the Type of Global Variable Initialization	143
5.3.5	Specifying Where to Allocate Sections in Memory	143
5.3.6	A Sample Linker Command File	145
6	TMS320C6000C/C++ Language Implementation	146
6.1	Characteristics of TMS320C6000 C	147
6.2	Characteristics of TMS320C6000 C++	147
6.3	Using MISRA-C:2004	148
6.4	Data Types	149
6.5	Keywords	150
6.5.1	The const Keyword	150
6.5.2	The cregister Keyword	150
6.5.3	The interrupt Keyword	152
6.5.4	The near and far Keywords	152
6.5.5	The restrict Keyword	153

6.5.6	The volatile Keyword	154
6.6	C++ Exception Handling	155
6.7	Register Variables and Parameters	155
6.8	The asm Statement	156
6.9	Pragma Directives	157
6.9.1	The CHECK_MISRA Pragma	158
6.9.2	The CLINK Pragma	158
6.9.3	The CODE_SECTION Pragma	158
6.9.4	The DATA_ALIGN Pragma	159
6.9.5	The DATA_MEM_BANK Pragma	160
6.9.6	The DATA_SECTION Pragma	161
6.9.7	The Diagnostic Message Pragmas	162
6.9.8	The FUNC_ALWAYS_INLINE Pragma	162
6.9.9	The FUNC_CANNOT_INLINE Pragma	163
6.9.10	The FUNC_EXT_CALLED Pragma	163
6.9.11	The FUNC_INTERRUPT_THRESHOLD Pragma	164
6.9.12	The FUNC_IS_PURE Pragma	164
6.9.13	The FUNC_IS_SYSTEM Pragma	165
6.9.14	The FUNC_NEVER_RETURNS Pragma	165
6.9.15	The FUNC_NO_GLOBAL_ASG Pragma	165
6.9.16	The FUNC_NO_IND_ASG Pragma	166
6.9.17	The FUNCTION_OPTIONS Pragma	166
6.9.18	The INTERRUPT Pragma	166
6.9.19	The LOCATION Pragma	167
6.9.20	The MUST_ITERATE Pragma	167
6.9.21	The NMI_INTERRUPT Pragma	168
6.9.22	The NOINIT and PERSISTENT Pragmas	169
6.9.23	The NO_HOOKS Pragma	169
6.9.24	The PROB_ITERATE Pragma	170
6.9.25	The RESET_MISRA Pragma	170
6.9.26	The RETAIN Pragma	170
6.9.27	The SET_CODE_SECTION and SET_DATA_SECTION Pragmas	171
6.9.28	The STRUCT_ALIGN Pragma	172
6.9.29	The UNROLL Pragma	172
6.10	The _Pragma Operator	173
6.11	Application Binary Interface	174
6.11.1	COFF ABI	174
6.11.2	EABI	174
6.12	Object File Symbol Naming Conventions (Linknames)	174
6.13	Initializing Static and Global Variables in COFF ABI Mode	175
6.13.1	Initializing Static and Global Variables With the Linker	175
6.13.2	Initializing Static and Global Variables With the const Type Qualifier	175
6.14	Changing the ANSI/ISO C Language Mode	176
6.14.1	Compatibility With K&R C (--kr_compatible Option)	176
6.14.2	Enabling Strict ANSI/ISO Mode and Relaxed ANSI/ISO Mode (--strict_ansi and --relaxed_ansi Options)	177
6.14.3	Enabling Embedded C++ Mode (--embedded_cpp Option)	178
6.15	GNU Language Extensions	178
6.15.1	Extensions	178
6.15.2	Function Attributes	179
6.15.3	Variable Attributes	180
6.15.4	Type Attributes	180
6.15.5	Built-In Functions	181

7	Run-Time Environment	182
7.1	Memory Model	183
7.1.1	Sections	183
7.1.2	C/C++ System Stack	184
7.1.3	Dynamic Memory Allocation	185
7.1.4	Initialization of Variables in COFF ABI	185
7.1.5	Data Memory Models	185
7.1.6	Trampoline Generation for Function Calls	186
7.1.7	Position Independent Data	187
7.2	Object Representation	188
7.2.1	Data Type Storage	188
7.2.2	Bit Fields	194
7.2.3	Character String Constants	195
7.3	Register Conventions	196
7.4	Function Structure and Calling Conventions	197
7.4.1	How a Function Makes a Call	197
7.4.2	How a Called Function Responds	198
7.4.3	Accessing Arguments and Local Variables	199
7.5	Interfacing C and C++ With Assembly Language	200
7.5.1	Using Assembly Language Modules With C/C++ Code	200
7.5.2	Accessing Assembly Language Variables From C/C++	202
7.5.3	Sharing C/C++ Header Files With Assembly Source	203
7.5.4	Using Inline Assembly Language	204
7.5.5	Using Intrinsics to Access Assembly Language Statements	204
7.5.6	The <code>__x128_t</code> Container Type	223
7.5.7	The <code>__float2_t</code> Container Type	224
7.5.8	Using Intrinsics for Interrupt Control and Atomic Sections	225
7.5.9	Using Unaligned Data and 64-Bit Values	225
7.5.10	Using <code>MUST_ITERATE</code> and <code>_nassert</code> to Enable SIMD and Expand Compiler Knowledge of Loops	226
7.5.11	Methods to Align Data	227
7.5.12	SAT Bit Side Effects	229
7.5.13	IRP and AMR Conventions	230
7.5.14	Floating Point Control Register Side Effects	230
7.6	Interrupt Handling	230
7.6.1	Saving the SGIE Bit	230
7.6.2	Saving Registers During Interrupts	230
7.6.3	Using C/C++ Interrupt Routines	231
7.6.4	Using Assembly Language Interrupt Routines	232
7.7	Run-Time-Support Arithmetic Routines	232
7.8	System Initialization	234
7.8.1	COFF ABI Automatic Initialization of Variables	234
7.8.2	EABI Automatic Initialization of Variables	236
7.8.3	Initialization Tables	241
7.9	Support for Multi-Threaded Applications	245
7.9.1	Compiling with OpenMP	245
7.9.2	Multi-Threading Runtime Support	245
8	Using Run-Time-Support Functions and Building Libraries	247
8.1	C and C++ Run-Time Support Libraries	248
8.1.1	Linking Code With the Object Library	248
8.1.2	Header Files	249
8.1.3	Modifying a Library Function	249
8.1.4	Minimal Support for Internationalization	249

8.1.5	Allowable Number of Open Files	249
8.1.6	Library Naming Conventions	250
8.2	The C I/O Functions	250
8.2.1	High-Level I/O Functions	251
8.2.2	Overview of Low-Level I/O Implementation	252
8.2.3	Device-Driver Level I/O Functions	255
8.2.4	Adding a User-Defined Device Driver for C I/O	259
8.2.5	The device Prefix	260
8.3	Handling Reentrancy (_register_lock() and _register_unlock() Functions)	262
8.4	C6700 FastMath Library	263
8.5	Library-Build Process	264
8.5.1	Required Non-Texas Instruments Software	264
8.5.2	Using the Library-Build Process	264
8.5.3	Extending mklib	267
9	C++ Name Demangler	268
9.1	Invoking the C++ Name Demangler	269
9.2	C++ Name Demangler Options	269
9.3	Sample Usage of the C++ Name Demangler	270
A	Glossary	272

List of Figures

1-1.	TMS320C6000 Software Development Flow	15
3-1.	Software-Pipelined Loop.....	60
4-1.	4-Bank Interleaved Memory	127
4-2.	4-Bank Interleaved Memory With Two Memory Spaces.....	127
7-1.	Char and Short Data Storage Format	189
7-2.	32-Bit Data Storage Format	190
7-3.	Single-Precision Floating-Point Char Data Storage Format.....	190
7-4.	40-Bit Data Storage Format Signed <code>__int40_t</code> or 40-bit long.....	191
7-5.	Unsigned 40-bit <code>__int40_t</code> or long	191
7-6.	64-Bit Data Storage Format Signed 64-bit long.....	192
7-7.	Unsigned 64-bit long.....	192
7-8.	Double-Precision Floating-Point Data Storage Format	193
7-9.	Bit-Field Packing in Big-Endian and Little-Endian Formats	194
7-10.	Register Argument Conventions	198
7-11.	Autoinitialization at Run Time	235
7-12.	Initialization at Load Time.....	236
7-13.	Autoinitialization at Run Time in EABI Mode	237
7-14.	Initialization at Load Time in EABI Mode	241
7-15.	Constructor Table for EABI Mode	241
7-16.	Format of Initialization Records in the <code>.cinit</code> Section	242
7-17.	Format of Initialization Records in the <code>.pinit</code> Section	244

List of Tables

2-1.	Processor Options	20
2-2.	Optimization Options	20
2-3.	Debug Options	20
2-4.	Include Options	21
2-5.	Control Options	21
2-6.	Advanced Debug Options	21
2-7.	Language Options	21
2-8.	Parser Preprocessing Options	22
2-9.	Predefined Symbols Options	22
2-10.	Diagnostics Options	22
2-11.	Run-Time Model Options	23
2-12.	Advanced Optimization Options	24
2-13.	Entry/Exit Hook Options	24
2-14.	Feedback Options	24
2-15.	Library Function Assumptions Options	24
2-16.	Assembler Options	25
2-17.	File Type Specifier Options	25
2-18.	Directory Specifier Options	25
2-19.	Default File Extensions Options	26
2-20.	Dynamic Linking Support Compiler Options	26
2-21.	Command Files Options	26
2-22.	MISRA-C:2004 Options	26
2-23.	Performance Advisor Options	26
2-24.	Linker Basic Options	27
2-25.	File Search Path Options	27
2-26.	Command File Preprocessing Options	27
2-27.	Diagnostic Options	27
2-28.	Linker Output Options	27
2-29.	Symbol Management Options	29
2-30.	Run-Time Environment Options	29
2-31.	Link-Time Optimization Options	29
2-32.	Miscellaneous Options	29
2-33.	Dynamic Linking Support Options	30
2-34.	Compiler Options For Dynamic Linking	38
2-35.	Linker Options For Dynamic Linking	38
2-36.	Compiler Backwards-Compatibility Options Summary	39
2-37.	Predefined C6000 Macro Names	42
2-38.	Raw Listing File Identifiers	49
2-39.	Raw Listing File Diagnostic Identifiers	49
3-1.	Options That You Can Use With --opt_level=3	70
3-2.	Selecting a File-Level Optimization Option	70
3-3.	Selecting a Level for the --gen_opt_info Option	70
3-4.	Selecting a Level for the --call_assumptions Option	71
3-5.	Special Considerations When Using the --call_assumptions Option	72
4-1.	Options That Affect the Assembly Optimizer	107
4-2.	Assembly Optimizer Directives Summary	113
5-1.	Initialized Sections Created by the Compiler for COFFABI	143

5-2.	Initialized Sections Created by the Compiler for EABI	144
5-3.	Uninitialized Sections Created by the Compiler for Both ABIs.....	144
6-1.	TMS320C6000 C/C++ COFF ABI Data Types.....	149
6-2.	TMS320C6000 C/C++ EABI Data Types	149
6-3.	Valid Control Registers.....	150
6-4.	GCC Language Extensions.....	178
7-1.	Data Representation in Registers and Memory	188
7-2.	Register Usage	196
7-3.	C6000 C/C++ Intrinsic Support by Device	205
7-4.	TMS320C6000 C/C++ Compiler Intrinsic.....	212
7-5.	TMS320C6400, C6400+, C6740, and C6600 C/C++ Compiler Intrinsic	214
7-6.	TMS320C6400+, C6740, and C6600 C/C++ Compiler Intrinsic.....	216
7-7.	TMS320C6700, C6700+, C6740, and C6600 C/C++ Compiler Intrinsic	217
7-8.	TMS320C6600 C/C++ Compiler Intrinsic.....	218
7-9.	Vector-in-Scalar Support C/C++ Compiler v7.2 Intrinsic	224
7-10.	Summary of Run-Time-Support Arithmetic Functions	232
8-1.	The mklib Program Options	266

Read This First

About This Manual

The *TMS320C6000 Optimizing Compiler User's Guide* explains how to use these compiler tools:

- Compiler
- Assembly optimizer
- Library build utility
- C++ name demangler

The compiler accepts C and C++ code conforming to the International Organization for Standardization (ISO) standards for these languages. The compiler supports the 1989 version of the C language and the 1998 version of the C++ language.

This user's guide discusses the characteristics of the C/C++ compiler. It assumes that you already know how to write C/C++ programs. *The C Programming Language* (second edition), by Brian W. Kernighan and Dennis M. Ritchie, describes C based on the ISO C standard. You can use the Kernighan and Ritchie (hereafter referred to as K&R) book as a supplement to this manual. References to K&R C (as opposed to ISO C) in this manual refer to the C language as defined in the first edition of Kernighan and Ritchie's *The C Programming Language*.

Notational Conventions

This document uses the following conventions:

- Program listings, program examples, and interactive displays are shown in a special typeface. 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 of C code:

```
#include <stdio.h>
main()
{   printf("hello, cruel world\n");
}
```

- 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.
- 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 the **bold typeface**, do not enter the brackets themselves. The following is an example of a command that has an optional parameter:

```
cl6x [options] [filenames] [--run_linker [link_options] [object files]]
```

- Braces ({ and }) indicate that you must choose one of the parameters within the braces; you do not enter the braces themselves. This is an example of a command with braces that are not included in the actual syntax but indicate that you must specify either the --rom_model or --ram_model option:

```
cl6x --run_linker {--rom_model | --ram_model} filenames [--output_file= name.out]
--library= libraryname
```

- 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 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]
```

- Some directives can have a varying number of parameters. For example, the .byte directive. This syntax is shown as [, ..., *parameter*].
- The TMS320C6200™ core is referred to as C6200. The TMS320C6400 core is referred to as C6400. The TMS320C6700 core is referred to as C6700. TMS320C6000 and C6000 can refer to any of C6200, C6400, C6400+, C6700, C6700+, C6740, or C6600.

Related Documentation

You can use the following books to supplement this user's guide:

ANSI X3.159-1989, Programming Language - C (Alternate version of the 1989 C Standard), American National Standards Institute

ISO/IEC 9899:1989, International Standard - Programming Languages - C (The 1989 C Standard), International Organization for Standardization

ISO/IEC 9899:1999, International Standard - Programming Languages - C (The C Standard), International Organization for Standardization

ISO/IEC 14882-1998, International Standard - Programming Languages - C++ (The C++ Standard), International Organization for Standardization

The C Programming Language (second edition), by Brian W. Kernighan and Dennis M. Ritchie, published by Prentice-Hall, Englewood Cliffs, New Jersey, 1988

The Annotated C++ Reference Manual, Margaret A. Ellis and Bjarne Stroustrup, published by Addison-Wesley Publishing Company, Reading, Massachusetts, 1990

C: A Reference Manual (fourth edition), by Samuel P. Harbison, and Guy L. Steele Jr., published by Prentice Hall, Englewood Cliffs, New Jersey

Programming Embedded Systems in C and C++, by Michael Barr, Andy Oram (Editor), published by O'Reilly & Associates; ISBN: 1565923545, February 1999

Programming in C, Steve G. Kochan, Hayden Book Company

The C++ Programming Language (second edition), Bjarne Stroustrup, published by Addison-Wesley Publishing Company, Reading, Massachusetts, 1990

Tool Interface Standards (TIS) DWARF Debugging Information Format Specification Version 2.0, TIS Committee, 1995

DWARF Debugging Information Format Version 3, DWARF Debugging Information Format Workgroup, Free Standards Group, 2005 (<http://dwarfstd.org>)

Related Documentation From Texas Instruments

You can use the following books to supplement this user's guide:

[SPRAAB5](#)— *The Impact of DWARF on TI Object Files*. Describes the Texas Instruments extensions to the DWARF specification.

[SPRAB90](#)— *TMS320C6000 EABI Migration Guide Application Report*. Describes the changes which must be made to existing COFF ABI libraries and applications to add support for the new EABI.

[SPRU186](#)— *TMS320C6000 Assembly Language Tools User's Guide*. Describes the assembly language tools (assembler, linker, and other tools used to develop assembly language code), assembler directives, macros, common object file format, and symbolic debugging directives for the TMS320C6000 platform of devices (including the C64x+ and C67x+ generations).

[SPRU190](#)— *TMS320C6000 DSP Peripherals Overview Reference Guide*. Provides an overview and briefly describes the peripherals available on the TMS320C6000 family of digital signal processors (DSPs).

[SPRU198](#)— *TMS320C6000 Programmer's Guide*. Reference for programming the TMS320C6000 digital signal processors (DSPs). Before you use this manual, you should install your code generation and debugging tools. Includes a brief description of the C6000 DSP architecture and code development flow, includes C code examples and discusses optimization methods for the C code, describes the structure of assembly code and includes examples and discusses optimizations for the assembly code, and describes programming considerations for the C64x DSP.

[SPRU197](#)— *TMS320C6000 Technical Brief*. Provides an introduction to the TMS320C62x and TMS320C67x digital signal processors (DSPs) of the TMS320C6000 DSP family. Describes the CPU architecture, peripherals, development tools and third-party support for the C62x and C67x DSPs.

[SPRU423](#)— *TMS320 DSP/BIOS User's Guide*. DSP/BIOS gives developers of mainstream applications on Texas Instruments TMS320 digital signal processors (DSPs) the ability to develop embedded real-time software. DSP/BIOS provides a small firmware real-time library and easy-to-use tools for real-time tracing and analysis.

[SPRU731](#)— *TMS320C62x DSP CPU and Instruction Set Reference Guide*. Describes the CPU architecture, pipeline, instruction set, and interrupts for the TMS320C62x digital signal processors (DSPs) of the TMS320C6000 DSP family. The C62x DSP generation comprises fixed-point devices in the C6000 DSP platform.

[SPRU732](#)— *TMS320C64x/C64x+ DSP CPU and Instruction Set Reference Guide*. Describes the CPU architecture, pipeline, instruction set, and interrupts for the TMS320C64x and TMS320C64x+ digital signal processors (DSPs) of the TMS320C6000 DSP family. The C64x/C64x+ DSP generation comprises fixed-point devices in the C6000 DSP platform. The C64x+ DSP is an enhancement of the C64x DSP with added functionality and an expanded instruction set.

[SPRU733](#)— *TMS320C67x/C67x+ DSP CPU and Instruction Set Reference Guide*. Describes the CPU architecture, pipeline, instruction set, and interrupts for the TMS320C67x and TMS320C67x+ digital signal processors (DSPs) of the TMS320C6000 DSP platform. The C67x/C67x+ DSP generation comprises floating-point devices in the C6000 DSP platform. The C67x+ DSP is an enhancement of the C67x DSP with added functionality and an expanded instruction set.

[SPRUGH7](#)— *TMS320C66x CPU and Instruction Set Reference Guide*. Describes the CPU architecture, pipeline, instruction set, and interrupts for the TMS320C66x digital signal processors (DSPs) of the TMS320C6000 DSP platform. The C66x DSP generation comprises floating-point devices in the C6000 DSP platform.

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.

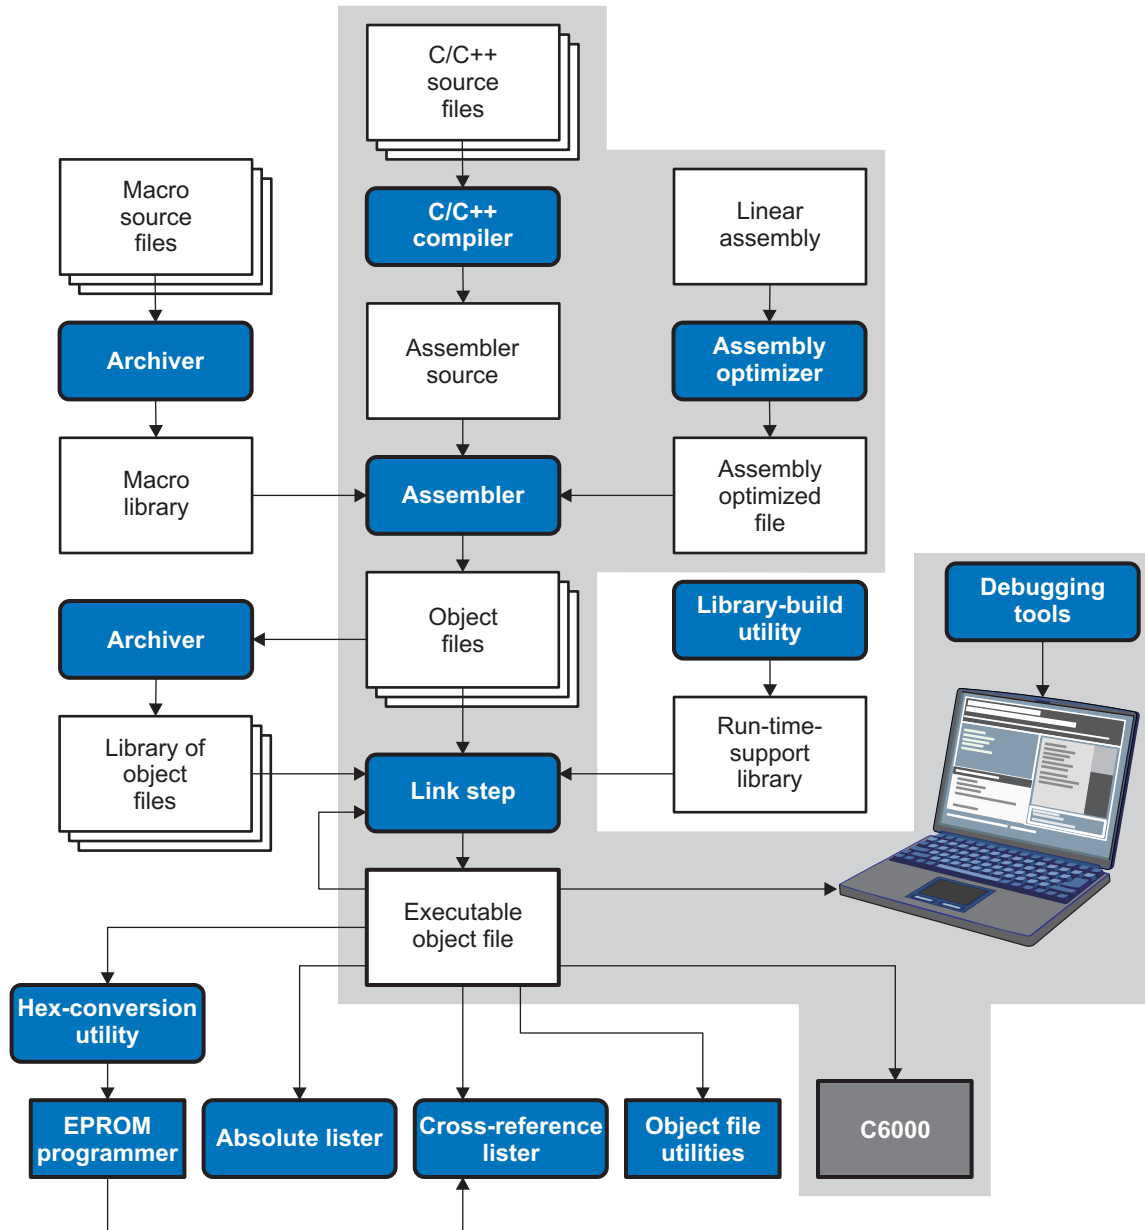
This chapter provides an overview of these tools and introduces the features of the optimizing C/C++ compiler. The assembly optimizer is discussed in [Chapter 4](#). The assembler and linker are discussed in detail in the *TMS320C6000 Assembly Language Tools User's Guide*.

Topic	Page
1.1 Software Development Tools Overview	15
1.2 C/C++ Compiler Overview	16

1.1 Software Development Tools Overview

Figure 1-1 illustrates the software development flow. The shaded portion of the figure highlights the most common path of software development for C language programs. The other portions are peripheral functions that enhance the development process.

Figure 1-1. TMS320C6000 Software Development Flow



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 accepts assembly code that has not been register-allocated and is unscheduled. The assembly optimizer assigns registers and uses loop optimization to turn linear assembly into highly parallel assembly that takes advantage of software pipelining. See Chapter 4.
- The **compiler** accepts C/C++ source code and produces C6000 assembly language source code. See Chapter 2.

- The **assembler** translates assembly language source files into machine language relocatable object files. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the assembler.
- The **linker** combines relocatable object files into a single absolute executable object file. As it creates the executable file, it performs relocation and resolves external references. The linker accepts relocatable object files and object libraries as input. See [Chapter 5](#). The *TMS320C6000 Assembly Language Tools User's Guide* provides a complete description of the linker.
- The **archiver** allows you to collect a group of files into a single archive file, called a *library*. Additionally, the archiver allows you to modify a library by deleting, replacing, extracting, or adding members. One of the most useful applications of the archiver is building a library of object files. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the archiver.
- The **run-time-support libraries** contain the standard ISO C and C++ library functions, compiler-utility functions, floating-point arithmetic functions, and C I/O functions that are supported by the compiler. See [Chapter 8](#).

You can use the **library-build utility** to build your own customized run-time-support library. See [Section 8.5](#). Source code for the standard run-time-support library functions for C and C++ are provided in the self-contained `rtsrc.zip` file.

- The **hex conversion utility** converts an object file into other object formats. You can download the converted file to an EPROM programmer. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the hex conversion utility and describes all supported formats.
- The **absolute lister** accepts linked object files as input and creates `.abs` files as output. You can assemble these `.abs` files to produce a listing that contains absolute, rather than relative, addresses. Without the absolute lister, producing such a listing would be tedious and would require many manual operations. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the absolute lister.
- The **cross-reference lister** uses object files to produce a cross-reference listing showing symbols, their definitions, and their references in the linked source files. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the cross-reference utility.
- The **C++ name demangler** is a debugging aid that converts names mangled by the compiler back to their original names as declared in the C++ source code. As shown in [Figure 1-1](#), you can use the C++ name demangler on the assembly file that is output by the compiler; you can also use this utility on the assembler listing file and the linker map file. See [Chapter 9](#).
- The **disassembler** decodes object files to show the assembly instructions that they represent. The *TMS320C6000 Assembly Language Tools User's Guide* explains how to use the disassembler.
- The main product of this development process is an executable object file 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-level and clock-accurate software simulator
 - An XDS emulator

1.2 C/C++ Compiler Overview

The following subsections describe the key features of the compiler.

1.2.1 ANSI/ISO Standard

The C and C++ language features in the compiler are implemented in conformance with these ISO standards:

- **ISO-standard C**

The C/C++ compiler conforms to the C Standard ISO/IEC 9889:1990. The ISO standard supercedes and is the same as the ANSI C standard. There is also a 1999 version of the ISO standard, but the TI compiler conforms to the 1990 standard, not the 1999 standard. The language is also described in the second edition of Kernighan and Ritchie's *The C Programming Language* (K&R).

- **ISO-standard C++**

The C/C++ compiler conforms to the C++ Standard ISO/IEC 14882:1998. The language is also described in Ellis and Stroustrup's *The Annotated C++ Reference Manual* (ARM), but this is not the

standard. The compiler also supports embedded C++. For a description of *unsupported C++* features, see [Section 6.2](#).

- **ISO-standard run-time support**

The compiler tools come with an extensive run-time library. All library functions conform to the ISO C/C++ library standard. The library includes functions for standard input and output, string manipulation, dynamic memory allocation, data conversion, timekeeping, trigonometry, and exponential and hyperbolic functions. Functions for signal handling are not included, because these are target-system specific. For more information, see [Chapter 8](#).

1.2.2 Output Files

These types of output files are created by the compiler:

- **COFF object files**

Common object file format (COFF) allows you to define your system's memory map at link time. This maximizes performance by enabling you to link C/C++ code and data objects into specific memory areas. COFF also supports source-level debugging.

- **ELF object files**

Executable and linking format (ELF) enables supporting modern language features like early template instantiation and exporting inline functions.

1.2.3 Compiler Interface

These features enable interfacing with the compiler:

- **Compiler program**

The compiler tools include a compiler program (cl6x) that you use to compile, optimize, assemble, and link programs in a single step. For more information, see [Section 2.1](#)

- **Flexible assembly language interface**

The compiler has straightforward calling conventions, so you can write assembly and C functions that call each other. For more information, see [Chapter 7](#).

1.2.4 Utilities

These features are compiler utilities:

- **Library-build utility**

The library-build utility lets you custom-build object libraries from source for any combination of run-time models. For more information, see [Section 8.5](#).

- **C++ name demangler**

The C++ name demangler (dem6x) is a debugging aid that translates each mangled name it detects in compiler-generated assembly code, disassembly output, or compiler diagnostic messages to its original name found in the C++ source code. For more information, see [Chapter 9](#).

- **Hex conversion utility**

For stand-alone embedded applications, the compiler has the ability to place all code and initialization data into ROM, allowing C/C++ code to run from reset. The COFF or ELF files output by the compiler can be converted to EPROM programmer data files by using the hex conversion utility, as described in the *TMS320C6000 Assembly Language Tools User's Guide*.

Using the C/C++ Compiler

The compiler translates your source program into machine language object code that the TMS320C6000 can execute. Source code must be compiled, assembled, and linked to create an executable object file. All of these steps are executed at once by using the compiler.

Topic	Page
2.1 About the Compiler	19
2.2 Invoking the C/C++ Compiler	19
2.3 Changing the Compiler's Behavior With Options	20
2.4 Controlling the Compiler Through Environment Variables	39
2.5 Precompiled Header Support	41
2.6 Controlling the Preprocessor	42
2.7 Understanding Diagnostic Messages	45
2.8 Other Messages	48
2.9 Generating Cross-Reference Listing Information (--gen_acp_xref Option)	48
2.10 Generating a Raw Listing File (--gen_acp_raw Option)	48
2.11 Using Inline Function Expansion	50
2.12 Interrupt Flexibility Options (--interrupt_threshold Option)	53
2.13 Linking C6400 Code With C6200/C6700/Older C6400 Object Code	54
2.14 Using Interlist	54
2.15 Generating and Using Performance Advice	55
2.16 Controlling Application Binary Interface	55
2.17 Enabling Entry Hook and Exit Hook Functions	57

2.1 About the Compiler

The compiler lets you compile, assemble, and optionally link in one step. The compiler performs the following steps on one or more source modules:

- The **compiler** accepts C/C++ source code and assembly code, and produces object code. You can compile C, C++, and assembly files in a single command. The compiler uses the filename extensions to distinguish between different file types. See [Section 2.3.9](#) for more information.
- The **linker** combines object files to create a static executable or dynamic object file. The linker is optional, so you can compile and assemble many modules independently and link them later. See [Chapter 5](#) for information about linking the files.

Invoking the Linker

NOTE: By default, the compiler does not invoke the linker. You can invoke the linker by using the `--run_linker` compiler option.

For a complete description of the assembler and the linker, see the *TMS320C6000 Assembly Language Tools User's Guide*.

2.2 Invoking the C/C++ Compiler

To invoke the compiler, enter:

```
cl6x [options] [filenames] [--run_linker [link_options] object files]
```

cl6x	Command that runs the compiler and the assembler.
<i>options</i>	Options that affect the way the compiler processes input files. The options are listed in Table 2-5 through Table 2-32 .
<i>filenames</i>	One or more C/C++ source files, assembly language source files, linear assembly files, or object files.
--run_linker	Option that invokes the linker. The <code>--run_linker</code> option's short form is <code>-z</code> . See Chapter 5 for more information.
<i>link_options</i>	Options that control the linking process.
<i>object files</i>	Name of the additional object files for the linking process.

The arguments to the compiler are of three types:

- Compiler options
- Link options
- Filenames

The `--run_linker` option indicates linking is to be performed. If the `--run_linker` option is used, any compiler options must precede the `--run_linker` option, and all link options must follow the `--run_linker` option.

Source code filenames must be placed before the `--run_linker` option. Additional object file filenames can be placed after the `--run_linker` option.

For example, if you want to compile two files named `symtab.c` and `file.c`, assemble a third file named `seek.asm`, and link to create an executable program called `myprogram.out`, you will enter:

```
cl6x symtab.c file.c seek.asm --run_linker --library=lnk.cmd
    --library=rts6200.lib --output_file=myprogram.out
```

2.3 Changing the Compiler's Behavior With Options

Options control the operation of the compiler. This section provides a description of option conventions and an option summary table. It also provides detailed descriptions of the most frequently used options, including options used for type-checking and assembling.

For a help screen summary of the options, enter **cl6x** with no parameters on the command line.

The following apply to the compiler options:

- Options are preceded by one or two hyphens.
- Options are case sensitive.
- Options are either single letters or sequences of characters.
- Individual options cannot be combined.
- An option with a *required* parameter should be specified with an equal sign before the parameter to clearly associate the parameter with the option. For example, the option to undefine a constant can be expressed as `--undefine=name`. Although not recommended, you can separate the option and the parameter with or without a space, as in `--undefine name` or `-undefinename`.
- An option with an *optional* parameter should be specified with an equal sign before the parameter to clearly associate the parameter with the option. For example, the option to specify the maximum amount of optimization can be expressed as `-O=3`. Although not recommended, you can specify the parameter directly after the option, as in `-O3`. No space is allowed between the option and the optional parameter, so `-O 3` is not accepted.
- Files and options except the `--run_linker` option can occur in any order. The `--run_linker` option must follow all other compile options and precede any link options.

You can define default options for the compiler by using the `C6X_C_OPTION` environment variable. For a detailed description of the environment variable, see [Section 2.4.1](#).

[Table 2-5](#) through [Table 2-32](#) summarize all options (including link options). Use the references in the tables for more complete descriptions of the options.

Table 2-1. Processor Options

Option	Alias	Effect	Section
<code>--silicon_version=id</code>	<code>-mv</code>	Selects target version. Defaults to 6200.	Section 2.3.4
<code>--big_endian</code>	<code>-me</code>	Produces object code in big-endian format	Section 2.13
<code>--abi={coffabi eabi}</code>		Selects application binary interface. Default is coffabi.	Section 2.16

Table 2-2. Optimization Options⁽¹⁾

Option	Alias	Effect	Section
<code>--opt_level=0</code>	<code>-O0</code>	Optimizes register usage	Section 3.1
<code>--opt_level=1</code>	<code>-O1</code>	Uses <code>-O0</code> optimizations and optimizes locally	Section 3.1
<code>--opt_level=2</code>	<code>-O2</code> or <code>-O</code>	Uses <code>-O1</code> optimizations and optimizes globally (default)	Section 3.1
<code>--opt_level=3</code>	<code>-O3</code>	Uses <code>-O2</code> optimizations and optimizes the file	Section 3.1 Section 3.6
<code>--opt_for_space=n</code>	<code>-ms</code>	Controls code size on four levels (0, 1, 2, and 3)	Section 3.5

⁽¹⁾ **Note:** Machine-specific options (see [Table 2-11](#)) can also affect optimization.

Table 2-3. Debug Options

Option	Alias	Effect	Section
<code>--symdebug:dwarf</code>	<code>-g</code>	Enables symbolic debugging	Section 2.3.5 Section 3.14.3.1
<code>--symdebug:coff</code>		Enables symbolic debugging using the alternate STABS debugging format. STABS format is not supported for C6400+ or C674x, or when using ELF.	Section 2.3.5 Section 3.14.3.1
<code>--symdebug:none</code>		Disables all symbolic debugging	Section 2.3.5

Table 2-3. Debug Options (continued)

Option	Alias	Effect	Section
--symdebug:profile_coff		Enables profiling using the alternate STABS debugging format. STABS format is not supported for C6400+ or C674x, or when using ELF.	Section 2.3.5
--symdebug:skeletal		Enables minimal symbolic debugging that does not hinder optimizations (default behavior)	Section 2.3.5
--disable_push_pop		Disables push-pop optimization	--
--optimize_with_debug	-mn	Reenables optimizations disabled with --symdebug:dwarf	Section 3.14.3.1
--symdebug:dwarf_version=2 3		Specifies the DWARF format version	Section 2.3.5

Table 2-4. Include Options

Option	Alias	Effect	Section
--include_path=directory	-I	Defines #include search path	Section 2.6.2.1
--preinclude=filename		Includes <i>filename</i> at the beginning of compilation	Section 2.3.2

Table 2-5. Control Options

Option	Alias	Effect	Section
--compile_only	-c	Disables linking (negates --run_linker)	Section 5.1.3
--help	-h	Prints (on the standard output device) a description of the options understood by the compiler.	Section 2.3.1
--run_linker	-z	Enables linking	Section 2.3.1
--skip_assembler	-n	Compiles or assembly optimizes only	Section 2.3.1

Table 2-6. Advanced Debug Options

Option	Alias	Effect	Section
--machine_regs		Displays reg operands as machine registers in assembly code	Section 2.3.11
--symdebug:keep_all_types		Keep unreferenced type information (default for ELF with debug)	Section 2.3.5

Table 2-7. Language Options

Option	Alias	Effect	Section
--cpp_default	-fg	Processes all source files with a C extension as C++ source files.	Section 2.3.7
--create_pch=filename		Creates a precompiled header file with the name specified	Section 2.5
--embedded_cpp	-pe	Enables embedded C++ mode	Section 6.14.3
--exceptions		Enables C++ exception handling	Section 6.6
--extern_c_can_throw		Allow extern C functions to propagate exceptions (EABI only)	--
--gcc		Enables support for GCC extensions	Section 6.15
--gen_acp_raw	-pl	Generates a raw listing file	Section 2.10
--gen_acp_xref	-px	Generates a cross-reference listing file	Section 2.9
--keep_unneeded_statics		Keeps unreferenced static variables.	Section 2.3.2
--kr_compatible	-pk	Allows K&R compatibility	Section 6.14.1
--multibyte_chars	-pc	Enables support for multibyte character sequences in comments, string literals and character constants.	--
--no_inlining	-pi	Disables definition-controlled inlining (but --opt_level=3 (or -O3) optimizations still perform automatic inlining)	Section 2.11
--no_intrinsics	-pn	Disables intrinsic functions. No predefinition of compiler-supplied intrinsic functions.	--
--pch		Creates or uses precompiled header files	Section 2.5

Table 2-7. Language Options (continued)

Option	Alias	Effect	Section
--pch_dir= <i>directory</i>		Specifies the path where the precompiled header file resides	Section 2.5.2
--pch_verbose		Displays a message for each precompiled header file that is considered but not used	Section 2.5.3
--program_level_compile	-pm	Combines source files to perform program-level optimization	Section 3.7
--relaxed_ansi	-pr	Enables relaxed mode; ignores strict ISO violations	Section 6.14.2
--rtti	-rtti	Enables run time type information (RTTI)	--
--static_template_instantiation		Instantiate all template entities with internal linkage	--
--strict_ansi	-ps	Enables strict ISO mode (for C/C++, not K&R C)	Section 6.14.2
--use_pch= <i>filename</i>		Specifies the precompiled header file to use for this compilation	Section 2.5.2

Table 2-8. Parser Preprocessing Options

Option	Alias	Effect	Section
--preproc_dependency[= <i>filename</i>]	-ppd	Performs preprocessing only, but instead of writing preprocessed output, writes a list of dependency lines suitable for input to a standard make utility	Section 2.6.7
--preproc_includes[= <i>filename</i>]	-ppi	Performs preprocessing only, but instead of writing preprocessed output, writes a list of files included with the #include directive	Section 2.6.8
--preproc_macros[= <i>filename</i>]	-ppm	Performs preprocessing only. Writes list of predefined and user-defined macros to a file with the same name as the input but with a .pp extension.	Section 2.6.9
--preproc_only	-ppo	Performs preprocessing only. Writes preprocessed output to a file with the same name as the input but with a .pp extension.	Section 2.6.3
--preproc_with_comment	-ppc	Performs preprocessing only. Writes preprocessed output, keeping the comments, to a file with the same name as the input but with a .pp extension.	Section 2.6.5
--preproc_with_compile	-ppa	Continues compilation after preprocessing	Section 2.6.4
--preproc_with_line	-ppl	Performs preprocessing only. Writes preprocessed output with line-control information (#line directives) to a file with the same name as the input but with a .pp extension.	Section 2.6.6

Table 2-9. Predefined Symbols Options

Option	Alias	Effect	Section
--define= <i>name</i> [= <i>def</i>]	-D	Predefines <i>name</i>	Section 2.3.1
--undefine= <i>name</i>	-U	Undefines <i>name</i>	Section 2.3.1

Table 2-10. Diagnostics Options

Option	Alias	Effect	Section
--compiler_revision		Prints out the compiler release revision and exits	--
--diag_error= <i>num</i>	-pdse	Categorizes the diagnostic identified by <i>num</i> as an error	Section 2.7.1
--diag_remark= <i>num</i>	-pdsr	Categorizes the diagnostic identified by <i>num</i> as a remark	Section 2.7.1
--diag_suppress= <i>num</i>	-pds	Suppresses the diagnostic identified by <i>num</i>	Section 2.7.1
--diag_warning= <i>num</i>	-pdsr	Categorizes the diagnostic identified by <i>num</i> as a warning	Section 2.7.1
--display_error_number	-pden	Displays a diagnostic's identifiers along with its text	Section 2.7.1
--emit_warnings_as_errors	-pdew	Treat warnings as errors	Section 2.7.1
--issue_remarks	-pdr	Issues remarks (nonserious warnings)	Section 2.7.1
--no_warnings	-pdw	Suppresses warning diagnostics (errors are still issued)	Section 2.7.1
--quiet	-q	Suppresses progress messages (quiet)	--
--set_error_limit= <i>num</i>	-pdel	Sets the error limit to <i>num</i> . The compiler abandons compiling after this number of errors. (The default is 100.)	Section 2.7.1

Table 2-10. Diagnostics Options (continued)

Option	Alias	Effect	Section
--super_quiet	-qq	Super quiet mode	--
--tool_version	-version	Displays version number for each tool	--
--verbose		Display banner and function progress information	--
--verbose_diagnostics	-pdv	Provides verbose diagnostics that display the original source with line-wrap	Section 2.7.1
--write_diagnostics_file	-pdf	Generates a diagnostics information file. Compiler only option.	Section 2.7.1

Table 2-11. Run-Time Model Options

Option	Alias	Effect	Section
--debug_software_pipeline	-mw	Produce verbose software pipelining report	Section 3.2.2
--disable_software_pipelining	-mu	Turns off software pipelining	Section 3.2.1
--dprel		Specifies that all non-const data is addressed using DP-relative addressing	Section 7.1.5.2
f--fp_not_associative	-mc	Prevents reordering of associative floating-point operations	Section 3.11
--fp_reassoc={on off}		Enables or disables the reassociation of floating-point arithmetic	Section 2.3.3
--gen_func_subsections={on off}	-mo	Puts each function in a separate subsection in the object file	Section 5.2.1
--interrupt_threshold[= <i>num</i>]	-mi	Specifies an interrupt threshold value	Section 2.12
--mem_model:const={ <i>far_aggregates</i> <i>far</i> <i>data</i> }		Allows const objects to be made far independently of the --mem_model:data option	Section 7.1.5.3
--mem_model:data={ <i>far_aggregates</i> <i>near</i> <i>far</i> }		Determines data access model	Section 7.1.5.1
--no_bad_aliases	-mt	Allows certain assumptions about aliasing and loops	Section 3.10.2 Section 4.6.2
--no_compress		Prevents compression on C6400+, C6740, and C6600	
--no_reload_errors		Turns off all reload-related loop buffer error messages for C6400+, C6740, and C6600	-
--profile:breakpt		Enables breakpoint-based profiling	Section 2.3.5 Section 3.14.3.2
--profile:power		Enables power profiling	Section 2.3.5 Section 3.14.3.2
--sat_reassoc={on off}		Enables or disables the reassociation of saturating arithmetic. Default is --sat_reassoc=off.	Section 2.3.2
--small_enum	--small-enum	Uses the smallest possible size for the enumeration type	Section 2.3.3
--speculate_loads= <i>n</i>	-mh	Specifies speculative load byte count threshold. Allows speculative execution of loads with bounded address ranges.	Section 3.2.3.1
--speculate_unknown_loads		Allows speculative execution of loads with unbounded addresses	Section 2.3.3
--target_compatibility_6200	-mb	Enables C62xx compatibility with C6400 code	Section 2.13
--use_const_for_alias_analysis	-ox	Uses const to disambiguate pointers	Section 2.3.3
--wchar_t={32 16}		Sets the size of the C/C++ type wchar_t. Default is 16 bits.	Section 2.3.3

Table 2-12. Advanced Optimization Options⁽¹⁾

Option	Alias	Effect	Section
--auto_inline=[size]	-oi	Sets automatic inlining size (--opt_level=3 only). If size is not specified, the default is 1.	Section 3.13
--call_assumptions=0	-op0	Specifies that the module contains functions and variables that are called or modified from outside the source code provided to the compiler	Section 3.7.1
--call_assumptions=1	-op1	Specifies that the module contains variables modified from outside the source code provided to the compiler but does not use functions called from outside the source code	Section 3.7.1
--call_assumptions=2	-op2	Specifies that the module contains no functions or variables that are called or modified from outside the source code provided to the compiler (default)	Section 3.7.1
--call_assumptions=3	-op3	Specifies that the module contains functions that are called from outside the source code provided to the compiler but does not use variables modified from outside the source code	Section 3.7.1
--gen_opt_info=0	-on0	Disables the optimization information file	Section 3.6.2
--gen_opt_info=1	-on1	Produces an optimization information file	Section 3.6.2
--gen_opt_info=2	-on2	Produces a verbose optimization information file	Section 3.6.2
--openmp	--omp	Enables support for OpenMP 3.0	Section 7.9.1
--opt_for_speed=[n]	-mf	Controls speed over space (0-5 range) (Default is 4.)	Section 3.14.4
--optimizer_interlist	-os	Interlists optimizer comments with assembly statements	Section 3.14.2
--remove_hooks_when_inlining		Removes entry/exit hooks for auto-inlined functions	Section 2.17
--single_inline		Inlines functions that are only called once	--
--aliased_variables	-ma	Assumes called functions create hidden aliases (rare)	Section 3.10.1

⁽¹⁾ **Note:** Machine-specific options (see [Table 2-11](#)) can also affect optimization.

Table 2-13. Entry/Exit Hook Options

Option	Alias	Effect	Section
--entry_hook[=name]		Enables entry hooks	Section 2.17
--entry_parm={none name address}		Specifies the parameters to the function to the --entry_hook option	Section 2.17
--exit_hook[=name]		Enables exit hooks	Section 2.17
--exit_parm={none name address}		Specifies the parameters to the function to the --exit_hook option	Section 2.17

Table 2-14. Feedback Options

Option	Alias	Effect	Section
--analyze={codecov callgraph}		Generate analysis info from profile data	Section 3.9.4.2
--analyze_only		Only generate analysis	Section 3.9.4.2
--gen_profile_info		Generates instrumentation code to collect profile information	Section 3.8.1.3
--use_profile_info=file1[, file2,...]		Specifies the profile information file(s)	Section 3.8.1.3

Table 2-15. Library Function Assumptions Options

Option	Alias	Effect	Section
--printf_support={nofloat full minimal}		Enables support for smaller, limited versions of the printf and sprintf run-time-support functions.	Section 2.3.2
--std_lib_func_defined	-ol1 or -oL1	Informs the optimizer that your file declares a standard library function	Section 3.6.1
--std_lib_func_not_defined	-ol2 or -oL2	Informs the optimizer that your file does not declare or alter library functions. Overrides the -ol0 and -ol1 options (default).	Section 3.6.1
--std_lib_func_redefined	-ol0 or -oL0	Informs the optimizer that your file alters a standard library function	Section 3.6.1

Table 2-16. Assembler Options

Option	Alias	Effect	Section
--keep_asm	-k	Keeps the assembly language (.asm) file	Section 2.3.11
--asm_listing	-al	Generates an assembly listing file	Section 2.3.11
--c_src_interlist	-ss	Interlists C source and assembly statements	Section 2.14 Section 3.14.2
--src_interlist	-s	Interlists optimizer comments (if available) and assembly source statements; otherwise interlists C and assembly source statements	Section 2.3.1
--absolute_listing	-aa	Enables absolute listing	Section 2.3.11
--asm_define=name[=def]	-ad	Sets the <i>name</i> symbol	Section 2.3.11
--asm_dependency	-apd	Performs preprocessing; lists only assembly dependencies	Section 2.3.11
--asm_includes	-api	Performs preprocessing; lists only included #include files	Section 2.3.11
--asm_undefine=name	-au	Undefines the predefined constant <i>name</i>	Section 2.3.11
--copy_file=filename	-ahc	Copies the specified file for the assembly module	Section 2.3.11
--cross_reference	-ax	Generates the cross-reference file	Section 2.3.11
--include_file=filename	-ahi	Includes the specified file for the assembly module	Section 2.3.11
--no_const_clink		Stops generation of .clink directives for const global arrays.	Section 2.3.2
--output_all_syms	-as	Puts labels in the symbol table	Section 2.3.11
--strip_coff_underscore		Aids in transitioning hand-coded assembly from COFF to EABI	
--syms_ignore_case	-ac	Makes case insignificant in assembly source files	Section 2.3.11

Table 2-17. File Type Specifier Options

Option	Alias	Effect	Section
--ap_file=filename	-fl	Identifies <i>filename</i> as a linear assembly source file regardless of its extension. By default, the compiler and assembly optimizer treat .sa files as linear assembly source files.	Section 2.3.7
--asm_file=filename	-fa	Identifies <i>filename</i> as an assembly source file regardless of its extension. By default, the compiler and assembler treat .asm files as assembly source files.	Section 2.3.7
--c_file=filename	-fc	Identifies <i>filename</i> as a C source file regardless of its extension. By default, the compiler treats .c files as C source files.	Section 2.3.7
--cpp_file=filename	-fp	Identifies <i>filename</i> as a C++ file, regardless of its extension. By default, the compiler treats .C, .cpp, .cc and .cxx files as a C++ files.	Section 2.3.7
--obj_file=filename	-fo	Identifies <i>filename</i> as an object code file regardless of its extension. By default, the compiler and linker treat .obj files as object code files.	Section 2.3.7

Table 2-18. Directory Specifier Options

Option	Alias	Effect	Section
--abs_directory=directory	-fb	Specifies an absolute listing file directory. By default, the compiler uses the .obj directory.	Section 2.3.10
--asm_directory=directory	-fs	Specifies an assembly file directory. By default, the compiler uses the current directory.	Section 2.3.10
--list_directory=directory	-ff	Specifies an assembly listing file and cross-reference listing file directory. By default, the compiler uses the .obj directory.	Section 2.3.10
--obj_directory=directory	-fr	Specifies an object file directory. By default, the compiler uses the current directory.	Section 2.3.10
--output_file=filename	-fe	Specifies a compilation output file name; can override --obj_directory.	Section 2.3.10
--pp_directory=dir		Specifies a preprocessor file directory. By default, the compiler uses the current directory.	Section 2.3.10
--temp_directory=directory	-ft	Specifies a temporary file directory. By default, the compiler uses the current directory.	Section 2.3.10

Table 2-19. Default File Extensions Options

Option	Alias	Effect	Section
--ap_extension=[.] <i>extension</i>	-el	Sets a default extension for linear assembly source files.	Section 2.3.9
--asm_extension=[.] <i>extension</i>	-ea	Sets a default extension for assembly source files	Section 2.3.9
--c_extension=[.] <i>extension</i>	-ec	Sets a default extension for C source files	Section 2.3.9
--cpp_extension=[.] <i>extension</i>	-ep	Sets a default extension for C++ source files	Section 2.3.9
--listing_extension=[.] <i>extension</i>	-es	Sets a default extension for listing files	Section 2.3.9
--obj_extension=[.] <i>extension</i>	-eo	Sets a default extension for object files	Section 2.3.9

Table 2-20. Dynamic Linking Support Compiler Options⁽¹⁾

Option	Alias	Description
--dsbt		Generates addressing via Dynamic Segment Base Table
--export_all_cpp_vtbl		Exports C++ virtual tables by default
--import_helper_functions		Treats compiler helper functions as imported references
--import_undef[={off on}]		Imports all undefined symbols. Default is on.
--inline_plt[={off on}]		Inlines the import function call stub. Default is on.
--linux		Generates code for Linux
--pic[={near far}]		Generates position independent addressing for a shared object. Default is near.
--visibility={hidden fhidden default protected}		Specifies visibility of global symbols

⁽¹⁾ See [Section 2.3.12](#) for more information.

Table 2-21. Command Files Options

Option	Alias	Effect	Section
--cmd_file= <i>filename</i>	-@	Interprets contents of a file as an extension to the command line. Multiple -@ instances can be used.	Section 2.3.1

Table 2-22. MISRA-C:2004 Options

Option	Alias	Effect	Section
--check_misra[={all required advisory none <i>rulespec</i> }]		Enables checking of the specified MISRA-C:2004 rules. Default is all.	Section 2.3.2
--misra_advisory={error warning remark suppress}		Sets the diagnostic severity for advisory MISRA-C:2004 rules	Section 2.3.2
--misra_required={error warning remark suppress}		Sets the diagnostic severity for required MISRA-C:2004 rules	Section 2.3.2

Table 2-23. Performance Advisor Options

Option	Alias	Effect	Section
--advice:performance[={all none}]		Generates compiler optimization advice. Default is all.	Section 2.15
--advice:performance_file={stdout stderr user_specified_filename}		Specifies that advice be written to stdout, stderr, or a file.	Section 2.15
--advice:performance_dir={user_specified_directory_name}		Specifies that advice file be created in the named directory.	Section 2.15

The following tables list the linker options. See the *TMS320C6000 Assembly Language Tools User's Guide* for details on these options.

Table 2-24. Linker Basic Options

Option	Alias	Description
--output_file= <i>file</i>	-o	Names the executable output file. The default filename is a.out.
--map_file= <i>file</i>	-m	Produces a map or listing of the input and output sections, including holes, and places the listing in <i>filename</i>
--stack_size= <i>size</i>	[-]-stack	Sets C system stack size to <i>size</i> bytes and defines a global symbol that specifies the stack size. Default = 1K bytes
--heap_size= <i>size</i>	[-]-heap	Sets heap size (for the dynamic memory allocation in C) to <i>size</i> bytes and defines a global symbol that specifies the heap size. Default = 1K bytes

Table 2-25. File Search Path Options

Option	Alias	Description
--library= <i>file</i>	-l	Names an archive library or link command <i>file</i> as linker input
--search_path= <i>pathname</i>	-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 --library option.
--priority	-priority	Satisfies unresolved references by the first library that contains a definition for that symbol
--reread_libs	-x	Forces rereading of libraries, which resolves back references
--disable_auto_rts		Disables the automatic selection of a run-time-support library

Table 2-26. Command File Preprocessing Options

Option	Alias	Description
--define= <i>name=value</i>		Predefines <i>name</i> as a preprocessor macro.
--undefine= <i>name</i>		Removes the preprocessor macro <i>name</i> .
--disable_pp		Disables preprocessing for command files

Table 2-27. Diagnostic Options

Option	Alias	Description
--diag_error= <i>num</i>		Categorizes the diagnostic identified by <i>num</i> as an error
--diag_remark= <i>num</i>		Categorizes the diagnostic identified by <i>num</i> as a remark
--diag_suppress= <i>num</i>		Suppresses the diagnostic identified by <i>num</i>
--diag_warning= <i>num</i>		Categorizes the diagnostic identified by <i>num</i> as a warning
--display_error_number		Displays a diagnostic's identifiers along with its text
--emit_warnings_as_errors	-pdew	Treat warnings as errors
--issue_remarks		Issues remarks (nonserious warnings)
--no_demangle		Disables demangling of symbol names in diagnostics
--no_warnings		Suppresses warning diagnostics (errors are still issued)
--set_error_limit= <i>count</i>		Sets the error limit to <i>count</i> . The linker abandons linking after this number of errors. (The default is 100.)
--verbose_diagnostics		Provides verbose diagnostics that display the original source with line-wrap
--warn_sections	-w	Displays a message when an undefined output section is created

Table 2-28. Linker Output Options

Option	Alias	Description
--absolute_exe	-a	Produces an absolute, executable object file. This is the default; if neither --absolute_exe nor --relocatable is specified, the linker acts as if --absolute_exe were specified.
--mapfile_contents= <i>attribute</i>		Controls the information that appears in the map file.
--relocatable	-r	Produces a nonexecutable, relocatable output object file
--rom		Creates a ROM object

Table 2-28. Linker Output Options (continued)

Option	Alias	Description
--run_abs	-abs	Produces an absolute listing file
--xml_link_info= <i>file</i>		Generates a well-formed XML <i>file</i> containing detailed information about the result of a link

Table 2-29. Symbol Management Options

Option	Alias	Description
--entry_point= <i>symbol</i>	-e	Defines a global symbol that specifies the primary entry point for the executable object file
--globalize= <i>pattern</i>		Changes the symbol linkage to global for symbols that match <i>pattern</i>
--hide= <i>pattern</i>		Hides symbols that match the specified <i>pattern</i>
--localize= <i>pattern</i>		Make the symbols that match the specified <i>pattern</i> local
--make_global= <i>symbol</i>	-g	Makes <i>symbol</i> global (overrides -h)
--make_static	-h	Makes all global symbols static
--no_sym_merge	-b	Disables merge of symbolic debugging information in COFF object files
--no_syntable	-s	Strips symbol table information and line number entries from the executable object file
--retain={ <i>symbol</i> <i>section specification</i> }		Specifies a symbol or section to be retained by the linker
--scan_libraries	-scanlibs	Scans all libraries for duplicate symbol definitions
--symbol_map= <i>refname</i> = <i>defname</i>		Specifies a symbol mapping; references to the <i>refname</i> symbol are replaced with references to the <i>defname</i> symbol
--undef_sym= <i>symbol</i>	-u	Adds <i>symbol</i> to the symbol table as an unresolved symbol
--unhide= <i>pattern</i>		Excludes symbols that match the specified <i>pattern</i> from being hidden

Table 2-30. Run-Time Environment Options

Option	Alias	Description
--arg_size= <i>size</i>	--args	Reserve <i>size</i> bytes for the argc/argv memory area
--fill_value= <i>value</i>	-f	Sets default fill value for holes within output sections
--ram_model	-cr	Initializes variables at load time
--rom_model	-c	Autoinitializes variables at run time
--trampolines[= <i>off</i> <i>on</i>]		Generates far call trampolines. Default is on.

Table 2-31. Link-Time Optimization Options

Option	Description
--cinit_compression[= <i>compression_kind</i>]	Specifies the type of compression to apply to the c auto initialization data. Default is <i>rlc</i> .
--compress_dwarf[= <i>off</i> <i>on</i>]	Aggressively reduces the size of DWARF information from input object files. Default is <i>on</i> .
--copy_compression[= <i>compression_kind</i>]	Compresses data copied by linker copy tables. Default is <i>rlc</i> .
--unused_section_elimination[= <i>off</i> <i>on</i>]	Eliminates sections that are not needed in the executable module. Default is <i>on</i> .

Table 2-32. Miscellaneous Options

Option	Alias	Description
--disable_clink	-j	Disables conditional linking of COFF object files
--linker_help	[-]help	Displays information about syntax and available options
--minimize_trampoline[= <i>off</i> <i>postorder</i>]		Places sections to minimize number of far trampolines required. Default is <i>postorder</i> .
--preferred_order= <i>function</i>		Prioritizes placement of functions
--strict_compatibility[= <i>off</i> <i>on</i>]		Performs more conservative and rigorous compatibility checking of input object files. Default is <i>on</i> .
--trampoline_min_spacing= <i>size</i>		When trampoline reservations are spaced more closely than the specified limit, tries to make them adjacent
--zero_init[= <i>off</i> <i>on</i>]		Controls preinitialization of uninitialized variables. Default is <i>on</i> .

Table 2-33. Dynamic Linking Support Options

Option	Description
<code>--bind_now</code>	Disables Linux default lazy binding mechanisms
<code>--dsbt_index=index</code>	Specifies the Data Segment Base Table (DSBT) index of this component
<code>--dsbt_size=size</code>	Specifies the size of the DSBT in words
<code>--dynamic[=exe lib]</code>	Generates dynamic executable or a dynamic library. Default is .exe.
<code>--export=symbol</code>	Specifies <i>symbol</i> exported by ELF object
<code>--fini=symbol</code>	Specifies <i>symbol</i> name of termination code
<code>--forced_static_binding[=off on]</code>	Forces all import references to bind during static linking; defaults to on
<code>--import=symbol</code>	Specifies <i>symbol</i> imported by ELF object
<code>--init=symbol</code>	Specifies <i>symbol</i> name of termination code
<code>--rpath=dir</code>	Adds directory to beginning of library search path
<code>--runpath=dir</code>	Adds directory to end of library search path
<code>--soname=soname</code>	Specifies ELF shared object file name

2.3.1 Frequently Used Options

Following are detailed descriptions of options that you will probably use frequently:

- c_src_interlist** Invokes the interlist feature, which interweaves original C/C++ source with compiler-generated assembly language. The interlisted C statements may appear to be out of sequence. You can use the interlist feature with the optimizer by combining the `--optimizer_interlist` and `--c_src_interlist` options. See [Section 3.14.2](#). The `--c_src_interlist` option can have a negative performance and/or code size impact.
- cmd_file=filename** Appends the contents of a file to the option set. You can use this option to avoid limitations on command line length or C style comments imposed by the host operating system. Use a `#` or `;` at the beginning of a line in the command file to include comments. You can also include comments by delimiting them with `/*` and `*/`. To specify options, surround hyphens with quotation marks. For example, `--"quiet`. You can use the `--cmd_file` option multiple times to specify multiple files. For instance, the following indicates that `file3` should be compiled as source and `file1` and `file2` are `--cmd_file` files:

```
c16x --cmd_file=file1 --cmd_file=file2 file3
```
- compile_only** Suppresses the linker and overrides the `--run_linker` option, which specifies linking. The `--compile_only` option's short form is `-c`. Use this option when you have `--run_linker` specified in the `C6X_C_OPTION` environment variable and you do not want to link. See [Section 5.1.3](#).
- define=name[=def]** Predefines the constant *name* for the preprocessor. This is equivalent to inserting `#define name def` at the top of each C source file. If the optional `[=def]` is omitted, the *name* is set to 1. The `--define` option's short form is `-D`.
 If you want to define a quoted string and keep the quotation marks, do one of the following:
- For Windows, use `--define=name=\"string def\"`. For example, `--define=car=\"sedan\"`
 - For UNIX, use `--define=name="string def"`. For example, `--define=car="sedan"`
 - For Code Composer Studio, enter the definition in a file and include that file with the `--cmd_file` option.

--help	Displays the syntax for invoking the compiler and lists available options. If the --help option is followed by another option or phrase, detailed information about the option or phrase is displayed. For example, to see information about debugging options use --help debug.
--include_path=directory	Adds <i>directory</i> to the list of directories that the compiler searches for #include files. The --include_path option's short form is -I. You can use this option several times to define several directories; be sure to separate the --include_path options with spaces. If you do not specify a directory name, the preprocessor ignores the --include_path option. See Section 2.6.2.1 .
--keep_asm	Retains the assembly language output from the compiler or assembly optimizer. Normally, the compiler deletes the output assembly language file after assembly is complete. The --keep_asm option's short form is -k.
--quiet	Suppresses banners and progress information from all the tools. Only source filenames and error messages are output. The --quiet option's short form is -q.
--run_linker	Runs the linker on the specified object files. The --run_linker option and its parameters follow all other options on the command line. All arguments that follow --run_linker are passed to the linker. The --run_linker option's short form is -z. See Section 5.1 .
--skip_assembler	Compiles only. The specified source files are compiled but not assembled or linked. The --skip_assembler option's short form is -n. This option overrides --run_linker. The output is assembly language output from the compiler.
--src_interlist	Invokes the interlist feature, which interweaves optimizer comments or C/C++ source with assembly source. If the optimizer is invoked (--opt_level= <i>n</i> option), optimizer comments are interlisted with the assembly language output of the compiler, which may rearrange code significantly. If the optimizer is not invoked, C/C++ source statements are interlisted with the assembly language output of the compiler, which allows you to inspect the code generated for each C/C++ statement. The --src_interlist option implies the --keep_asm option. The --src_interlist option's short form is -s.
--tool_version	Prints the version number for each tool in the compiler. No compiling occurs.
--undefine=name	Undefines the predefined constant <i>name</i> . This option overrides any --define options for the specified constant. The --undefine option's short form is -U.
--verbose	Displays progress information and toolset version while compiling. Resets the --quiet option.

2.3.2 Miscellaneous Useful Options

Following are detailed descriptions of miscellaneous options:

--check_misra={all required advisory none rulespec}	Displays the specified amount or type of MISRA-C documentation. The <i>rulespec</i> parameter is a comma-separated list of specifiers. See Section 6.3 for details.
--fp_reassoc={on off}	Enables or disables the reassociation of floating-point arithmetic. If --strict_ansi is set, --fp_reassoc=off is set since reassociation of floating-point arithmetic is an ANSI violation.

--keep_unneeded_statics	Does not delete unreferenced static variables. The parser by default remarks about and then removes any unreferenced static variables. The <code>--keep_unneeded_statics</code> option keeps the parser from deleting unreferenced static variables and any static functions that are referenced by these variable definitions. Unreferenced static functions will still be removed.
--no_const_clink	Tells the compiler to not generate <code>.clink</code> directives for const global arrays. By default, these arrays are placed in a <code>.const</code> subsection and conditionally linked.
--misra_advisory ={error warning remark suppress}	Sets the diagnostic severity for advisory MISRA-C:2004 rules.
--misra_required ={error warning remark suppress}	Sets the diagnostic severity for required MISRA-C:2004 rules.
--preinclude = <i>filename</i>	Includes the source code of <i>filename</i> at the beginning of the compilation. This can be used to establish standard macro definitions. The filename is searched for in the directories on the include search list. The files are processed in the order in which they were specified.
--printf_support ={full nofloat minimal}	Enables support for smaller, limited versions of the <code>printf</code> and <code>sprintf</code> run-time-support functions. The valid values are: <ul style="list-style-type: none"> • full: Supports all format specifiers. This is the default. • nofloat: Excludes support for printing and scanning floating-point values. Supports all format specifiers except <code>%f</code>, <code>%F</code>, <code>%g</code>, <code>%G</code>, <code>%e</code>, and <code>%E</code>. • minimal: Supports the printing and scanning of integer, char, or string values without width or precision flags. Specifically, only the <code>%%</code>, <code>%d</code>, <code>%o</code>, <code>%c</code>, <code>%s</code>, and <code>%x</code> format specifiers are supported There is no run-time error checking to detect if a format specifier is used for which support is not included. The <code>--printf_support</code> option precedes the <code>--run_linker</code> option, and must be used when performing the final link.
--sat_reassoc ={on off}	Enables or disables the reassociation of saturating arithmetic.

2.3.3 Run-Time Model Options

These options are specific to the TMS302C6000 toolset. See the referenced sections for more information. TMS320C6000-specific assembler options are listed in [Section 2.3.11](#).

--abi ={eabi coffabi}	Specifies application binary interface (ABI). Default support is for COFF ABI. See Section 2.16 . Also, refer to <i>The C6000 Embedded Application Binary Interface Application Report (SPRAB89)</i> . All code in an EABI application must be built for EABI. Make sure all your libraries are available in EABI mode before migrating your existing COFF ABI systems to C6000 EABI. See http://tiexpressdsp.com/index.php/EABI_Support_in_C6000_Compiler for full details.
--big_endian	Produces code in big-endian format. By default, little-endian code is produced.
--advice:performance	Generates compile-time optimization Advice. See Section 2.15 .
--debug_software_pipeline	Produces verbose software pipelining report. See Section 3.2.2 .
--disable_software_pipelining	Turns off software pipelining. See Section 3.2.1 .
--fp_not_associative	Compiler does not reorder floating-point operations. See Section 3.11 .

--interrupt_threshold=<i>n</i>	Specifies an interrupt threshold value <i>n</i> that sets the maximum cycles the compiler can disable interrupts. See Section 2.12 .
--mem_model:const=<i>type</i>	Allows const objects to be made far independently of the --mem_model:data option. The <i>type</i> can be data, far, or far_aggregates. See Section 7.1.5.3
--mem_model:data=<i>type</i>	Specifies data access model as <i>type</i> far, far_aggregates, or near. Default is far_aggregates. See Section 7.1.5.1 .
--silicon_version=<i>num</i>	Selects the target CPU version. See Section 2.3.4 .
--small_enum	By default, the C6000 compiler uses 32 bits for every enum. When you use the --small_enum option, the smallest possible byte size for the enumeration type is used. For example, enum example_enum {first = -128, second = 0, third = 127} uses only one byte instead of 32 bits when the --small_enum option is used. Similarly, enum a_short_enum {bottom = -32768, middle = 0, top = 32767} fits into two bytes instead of four. Do not link object files compiled with the --small_enum option with object files that have been compiled without it. If you use the --small_enum option, you must use it with all of your C/C++ files; otherwise, you will encounter errors that cannot be detected until run time.
--speculate_loads=<i>n</i>	Specifies speculative load byte count threshold. Allows speculative execution of loads with bounded addresses. See Section 3.2.3.1 .
--speculate_unknown_loads	Allows speculative execution of loads with unbounded addresses.
--target_compatibility_6200	Compiles C6400 code that is compatible with array alignment restrictions of version 4.0 tools or C6200/C6700 object code. This option is deprecated. See Section 2.13
--use_const_for_alias_analysis	Uses const to disambiguate pointers.
--wchar_t={32 16}	Sets the size (in bits) of the C/C++ type wchar_t. The --abi=eabi option is required when -wchar_t=32 is used. By default the compiler generates 16-bit wchar_t. In COFF ABI mode, a warning is generated and --wchar_t=32 is ignored. 16-bit wchar_t objects are not compatible with the 32-bit wchar_t objects; an error is generated if they are combined. When the --linux option is specified, it implies --wchar_t=32 since Linux uses 32-bit extended characters.

2.3.4 Selecting Target CPU Version (**--silicon_version** Option)

Select the target CPU version using the last four digits of the TMS320C6000 part number. This selection controls the use of target-specific instructions and alignment, such as --silicon_version=6701 or --silicon_version=6412. Alternatively, you can also specify the family of the part, for example, --silicon_version=6400 or --silicon_version=6700. If this option is not used, the compiler generates code for the C6200 parts. If the --silicon_version option is not specified, the code generated runs on all C6000 parts; however, the compiler does not take advantage of target-specific instructions or alignment. This option has the alias -mv. Common target CPU version options include:

- -mv6200
- -mv6700
- -mv6700+
- -mv6400+
- -mv6740
- -mv6600

2.3.5 Symbolic Debugging and Profiling Options

The following options are used to select symbolic debugging or profiling:

--profile:breakpt	Disables optimizations that would cause incorrect behavior when using a breakpoint-based profiler.
--profile:power	Enables power profiling support by inserting NOPs into the frame code. These NOPs can then be instrumented by the power profiling tooling to track the power usage of functions. If the power profiling tool is not used, this option increases the cycle count of each function because of the NOPs. The <code>--profile:power</code> option also disables optimizations that cannot be handled by the power-profiler.
--symdebug:coff	Enables symbolic debugging using the alternate STABS debugging format. This may be necessary to allow debugging with older debuggers or custom tools, which do not read the DWARF format. STABS format is not supported for C6400+ or ELF.
--symdebug:dwarf	Generates directives that are used by the C/C++ source-level debugger and enables assembly source debugging in the assembler. The <code>--symdebug:dwarf</code> option's short form is <code>-g</code> . The <code>--symdebug:dwarf</code> option disables many code generator optimizations, because they disrupt the debugger. You can use the <code>--symdebug:dwarf</code> option with the <code>--opt_level</code> (aliased as <code>-O</code>) option to maximize the amount of optimization that is compatible with debugging (see Section 3.14.3.1). For more information on the DWARF debug format, see <i>The DWARF Debugging Standard</i> .
--symdebug:keep_all_types	Effects the ability to view <i>unused</i> types in the debugger that are from a COFF executable. Use this option to view the details of a type that is defined but not used to define any symbols. Such unused types are not included in the debug information by default for COFF. However, in EABI mode, all types are included in the debug information and this option has no effect.
--symdebug:dwarf_ version={2 3}	Specifies the DWARF debugging format version (2 or 3) to be generated when <code>--symdebug:dwarf</code> or <code>--symdebug:skeletal</code> is specified. By default, the compiler generates DWARF version 3 debug information. For more information on TI extensions to the DWARF language, see <i>The Impact of DWARF on TI Object Files</i> (SPRAAB5).
--symdebug:none	Disables all symbolic debugging output. This option is not recommended; it prevents debugging and most performance analysis capabilities.
--symdebug:profile_coff	Adds the necessary debug directives to the object file which are needed by the profiler to allow function level profiling with minimal impact on optimization (when used). Using <code>--symdebug:coff</code> may hinder some optimizations to ensure that debug ability is maintained, while this option will not hinder optimization. STABS format is not supported for C6400+ or ELF. You can set breakpoints and profile on function-level boundaries in Code Composer Studio, but you cannot single-step through code as with full debug ability.
--symdebug:skeletal	Generates as much symbolic debugging information as possible without hindering optimization. Generally, this consists of global-scope information only. This option reflects the default behavior of the compiler.

See [Section 2.3.13](#) for a list of deprecated symbolic debugging options.

2.3.6 Specifying Filenames

The input files that you specify on the command line can be C source files, C++ source files, assembly source files, linear assembly files, or object files. The compiler uses filename extensions to determine the file type.

Extension	File Type
.asm, .abs, or .s* (extension begins with s)	Assembly source
.c	C source
.C	Depends on operating system
.cpp, .cxx, .cc	C++ source
.obj .o* .dll .so	Object
.sa	Linear assembly

NOTE: Case Sensitivity in Filename Extensions

Case sensitivity in filename extensions is determined by your operating system. If your operating system is not case sensitive, a file with a .C extension is interpreted as a C file. If your operating system is case sensitive, a file with a .C extension is interpreted as a C++ file.

For information about how you can alter the way that the compiler interprets individual filenames, see [Section 2.3.7](#). For information about how you can alter the way that the compiler interprets and names the extensions of assembly source and object files, see [Section 2.3.10](#).

You can use wildcard characters to compile or assemble multiple files. Wildcard specifications vary by system; use the appropriate form listed in your operating system manual. For example, to compile all of the files in a directory with the extension .cpp, enter the following:

```
c16x *.cpp
```

NOTE: No Default Extension for Source Files is Assumed

If you list a filename called example on the command line, the compiler assumes that the entire filename is example not example.c. No default extensions are added onto files that do not contain an extension.

2.3.7 Changing How the Compiler Interprets Filenames

You can use options to change how the compiler interprets your filenames. If the extensions that you use are different from those recognized by the compiler, you can use the filename options to specify the type of file. You can insert an optional space between the option and the filename. Select the appropriate option for the type of file you want to specify:

```
--ap_file=filename    for a linear assembly source file
--asm_file=filename    for an assembly language source file
--c_file=filename      for a C source file
--cpp_file=filename    for a C++ source file
--obj_file=filename    for an object file
```

For example, if you have a C source file called file.s and an assembly language source file called assy, use the --asm_file and --c_file options to force the correct interpretation:

```
c16x --c_file=file.s --asm_file=assy
```

You cannot use the filename options with wildcard specifications.

2.3.8 Changing How the Compiler Processes C Files

The `--cpp_default` option causes the compiler to process C files as C++ files. By default, the compiler treats files with a `.c` extension as C files. See [Section 2.3.9](#) for more information about filename extension conventions.

2.3.9 Changing How the Compiler Interprets and Names Extensions

You can use options to change how the compiler program interprets filename extensions and names the extensions of the files that it creates. The filename extension options must precede the filenames they apply to on the command line. You can use wildcard specifications with these options. An extension can be up to nine characters in length. Select the appropriate option for the type of extension you want to specify:

<code>--ap_extension=new extension</code>	for a linear assembly source file
<code>--asm_extension=new extension</code>	for an assembly language file
<code>--c_extension=new extension</code>	for a C source file
<code>--cpp_extension=new extension</code>	for a C++ source file
<code>--listing_extension=new extension</code>	sets default extension for listing files
<code>--obj_extension=new extension</code>	for an object file

The following example assembles the file `fit.rrr` and creates an object file named `fit.o`:

```
cl6x --asm_extension=.rrr --obj_extension=.o fit.rrr
```

The period (`.`) in the extension is optional. You can also write the example above as:

```
cl6x --asm_extension=rrr --obj_extension=o fit.rrr
```

2.3.10 Specifying Directories

By default, the compiler program places the object, assembly, and temporary files that it creates into the current directory. If you want the compiler program to place these files in different directories, use the following options:

<code>--abs_directory=directory</code>	Specifies the destination directory for absolute listing files. The default is to use the same directory as the object file directory. For example: <pre>cl6x --abs_directory=d:\abso_list</pre>
<code>--asm_directory=directory</code>	Specifies a directory for assembly files. For example: <pre>cl6x --asm_directory=d:\assembly</pre>
<code>--list_directory=directory</code>	Specifies the destination directory for assembly listing files and cross-reference listing files. The default is to use the same directory as the object file directory. For example: <pre>cl6x --list_directory=d:\listing</pre>
<code>--obj_directory=directory</code>	Specifies a directory for object files. For example: <pre>cl6x --obj_directory=d:\object</pre>
<code>--output_file=filename</code>	Specifies a compilation output file name; can override <code>--obj_directory</code> . For example: <pre>cl6x --output_file=transfer</pre>
<code>--pp_directory=directory</code>	Specifies a preprocessor file directory for object files (default is <code>.</code>). For example: <pre>cl6x --pp_directory=d:\preproc</pre>
<code>--temp_directory=directory</code>	Specifies a directory for temporary intermediate files. For example: <pre>cl6x --temp_directory=d:\temp</pre>

2.3.11 Assembler Options

Following are assembler options that you can use with the compiler. For more information, see the *TMS320C6000 Assembly Language Tools User's Guide*.

--absolute_listing	Generates a listing with absolute addresses rather than section-relative offsets.
--asm_define=name[=def]	<p>Predefines the constant <i>name</i> for the assembler; produces a <code>.set</code> directive for a constant or a <code>.arg</code> directive for a string. If the optional <code>[=def]</code> is omitted, the <i>name</i> is set to 1. If you want to define a quoted string and keep the quotation marks, do one of the following:</p> <ul style="list-style-type: none"> • For Windows, use <code>--asm_define=name="\string def"</code>. For example: <code>--asm_define=car="\sedan\ "</code> • For UNIX, use <code>--asm_define=name="string def"</code>. For example: <code>--asm_define=car=' "sedan" '</code> • For Code Composer Studio, enter the definition in a file and include that file with the <code>--cmd_file</code> option.
--asm_dependency	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 <code>.ppa</code> extension.
--asm_includes	Performs preprocessing for assembly files, but instead of writing preprocessed output, writes a list of files included with the <code>#include</code> directive. The list is written to a file with the same name as the source file but with a <code>.ppa</code> extension.
--asm_listing	Produces an assembly listing file.
--asm_undefine=name	Undefines the predefined constant <i>name</i> . This option overrides any <code>--asm_define</code> options for the specified name.
--copy_file=filename	Copies the specified file for the assembly module; acts like a <code>.copy</code> directive. The file is inserted before source file statements. The copied file appears in the assembly listing files.
--cross_reference	Produces a symbolic cross-reference in the listing file.
--include_file=filename	Includes the specified file for the assembly module; acts like a <code>.include</code> directive. The file is included before source file statements. The included file does not appear in the assembly listing files.
--machine_regs	Displays reg operands as machine registers in the assembly file for debugging purposes.
--no_compress	Prevents compression in the assembler. For C6400+, C6740, and C6600, compression is the changing of 32-bit instructions to 16-bit instructions, where possible/profitable.
--no_reload_errors	Turns off all reload-related loop buffer error messages in assembly code for C6400+, C6740, and C6600.
--output_all_syms	Puts labels in the symbol table. Label definitions are written to the COFF symbol table for use with symbolic debugging.
--strip_coff_underscore	Aids in transitioning hand-coded assembly from COFF to EABI.
--syms_ignore_case	Makes letter case insignificant in the assembly language source files. For example, <code>--syms_ignore_case</code> makes the symbols ABC and abc equivalent. <i>If you do not use this option, case is significant</i> (this is the default).

2.3.12 Dynamic Linking

The C6000 Code Generation Tools (CGT) (beginning with v7.3) support dynamic linking provided you build with EABI. For details on dynamic linking with the C6000 CGT, see the *TMS320C6000 Assembly Language Tools User's Guide* and http://processors.wiki.ti.com/index.php/C6000_Dynamic_Linking.

If you are not already familiar with the limitations of EABI support in the C6000 compiler, see http://processors.wiki.ti.com/index.php/EABI_Support_in_C6000_Compiler.

For more information about support for C6000 Linux ABI in the C6000 Code Generation Tools, see http://processors.wiki.ti.com/index.php/C6000_Linux_Support

Table 2-34 and Table 2-35 provide a brief summary of the compiler and linker options that are related to support for the Dynamic Linking Model in the C6000 CGT.

Table 2-34. Compiler Options For Dynamic Linking

Option	Description
<code>--abi=eabi</code>	Specifies that EABI run-time model is to be used.
<code>--dsbt</code>	Generates addressing via Dynamic Segment Base Table.
<code>--export_all_cpp_vtbl</code>	Exports C++ virtual tables by default.
<code>--import_undef[=off on]</code>	Specifies that all global symbol references that are not defined in a module are imported. Default is on.
<code>--import_helper_functions</code>	Specifies that all compiler generated calls to run-time-support functions are treated as calls to imported functions.
<code>--inline_plt[=off on]</code>	Inlines the import function call stub. Default is on.
<code>--linux</code>	Generates C6000 Linux ABI compliant code.
<code>--pic</code>	Generates position independent code suitable for a dynamic shared object.
<code>--visibility=={hidden fhidden default protected}</code>	Specifies a default visibility to be assumed for global symbols.
<code>-wchar_t</code>	Generates 32-bit <code>wchar_t</code> type when <code>--abi=eabi</code> is specified.

Table 2-35. Linker Options For Dynamic Linking

Option	Description
<code>--dsbt_index=int</code>	Requests a specific Data Segment Base Table (DSBT) index to be associated with the current output file. If the DSBT model is being used, and you do not request a specific DSBT index for the output file, then a DSBT index is assigned to the module at load time.
<code>--dsbt_size=int</code>	Specifies the size of the Data Segment Base Table (DSBT) for the current output file, in words. If the DSBT model is being used, this option can be used to override the default DSBT size (8 words).
<code>--dynamic[=exe]</code>	Specifies that the result of a link will be a lightweight dynamic executable.
<code>--dynamic=lib</code>	Specifies that the result of a link will be a dynamic library.
<code>--export=symbol</code>	Specifies that <i>symbol</i> is exported by the ELF object that is generated for this link.
<code>--fini=symbol</code>	Specifies the <i>symbol</i> name of the termination code for the output file currently being linked.
<code>--import=symbol</code>	Specifies that <i>symbol</i> is imported by the ELF object that is generated for this link.
<code>--init=symbol</code>	Specifies the <i>symbol</i> name of the initialization code for the output file currently being linked.
<code>--rpath=dir</code>	Adds a directory to the beginning of the dynamic library search path.
<code>--runpath=dir</code>	Adds a directory to the end of the dynamic library search path.
<code>--shared</code>	Generates an ELF dynamic shared object (DSO)
<code>--soname=string</code>	Specifies shared object name to be used to identify this ELF object to the any downstream ELF object consumers.
<code>--sysv</code>	Generates SysV ELF dynamic object module.

2.3.13 *Deprecated Options*

Several compiler options have been deprecated. The compiler continues to accept these options, but they are not recommended for use. Future releases of the tools will not support these options. [Table 2-36](#) lists the deprecated options and the options that have replaced them.

Table 2-36. Compiler Backwards-Compatibility Options Summary

Old Option	Effect	New Option
-gp	Allows function-level profiling of optimized code	--symdebug:dwarf or -g
-gt	Enables symbolic debugging using the alternate STABS debugging format	--symdebug:coff
-gw	Enables symbolic debugging using the DWARF debugging format	--symdebug:dwarf or -g

Additionally, the --symdebug:profile_coff option has been added to enable function-level profiling of optimized code with symbolic debugging using the STABS debugging format (the --symdebug:coff or -gt option).

Since C6400+, C6740, and C6600 produce only DWARF debug information, the -gp, -gt/--symdebug:coff, and --symdebug:profile_coff options are not supported for C6400+, C6740, and C6600.

2.4 Controlling the Compiler Through Environment Variables

An environment variable is a system symbol that you define and assign a string to. Setting environment variables is useful when you want to run the compiler repeatedly without re-entering options, input filenames, or pathnames.

NOTE: C_OPTION and C_DIR

The C_OPTION and C_DIR environment variables are deprecated. Use the device-specific environment variables instead.

2.4.1 *Setting Default Compiler Options (C6X_C_OPTION)*

You might find it useful to set the compiler, assembler, and linker default options using the C6X_C_OPTION environment variable. If you do this, the compiler uses the default options and/or input filenames that you name C6X_C_OPTION every time you run the compiler.

Setting the default options with these environment variables is useful when you want to run the compiler repeatedly with the same set of options and/or input files. After the compiler reads the command line and the input filenames, it looks for the C6X_C_OPTION environment variable and processes it.

The table below shows how to set the C6X_C_OPTION environment variable. Select the command for your operating system:

Operating System	Enter
UNIX (Bourne shell)	C6X_C_OPTION=" option₁ [option₂ . . .]"; export C6X_C_OPTION
Windows	set C6X_C_OPTION= option₁ [;option₂ . . .]

Environment variable options are specified in the same way and have the same meaning as they do on the command line. For example, if you want to always run quietly (the --quiet option), enable C/C++ source interlisting (the --src_interlist option), and link (the --run_linker option) for Windows, set up the C6X_C_OPTION environment variable as follows:

```
set C6X_C_OPTION=--quiet --src_interlist --run_linker
```

In the following examples, each time you run the compiler, it runs the linker. Any options following `--run_linker` on the command line or in `C6X_C_OPTION` are passed to the linker. Thus, you can use the `C6X_C_OPTION` environment variable to specify default compiler and linker options and then specify additional compiler and linker options on the command line. If you have set `--run_linker` in the environment variable and want to compile only, use the compiler `--compile_only` option. These additional examples assume `C6X_C_OPTION` is set as shown above:

```
cl6x *c                ; compiles and links
cl6x --compile_only *.c    ; only compiles
cl6x *.c --run_linker lnk.cmd ; compiles and links using a command file
cl6x --compile_only *.c --run_linker lnk.cmd
                        ; only compiles (--compile_only overrides --run_linker)
```

For details on compiler options, see [Section 2.3](#). For details on linker options, see the *Linker Description* chapter in the *TMS320C6000 Assembly Language Tools User's Guide*.

2.4.2 Naming an Alternate Directory (C6X_C_DIR)

The linker uses the `C6X_C_DIR` environment variable to name alternate directories that contain object libraries. The command syntaxes for assigning the environment variable are:

Operating System	Enter
UNIX (Bourne shell)	<code>C6X_C_DIR=" pathname₁ ; pathname₂ ;..."; export C6X_C_DIR</code>
Windows	<code>set C6X_C_DIR= pathname₁ ; pathname₂ ;...</code>

The *pathnames* are directories that contain input files. The pathnames must follow these constraints:

- Pathnames must be separated with a semicolon.
- Spaces or tabs at the beginning or end of a path are ignored. For example, the space before and after the semicolon in the following is ignored:

```
set C6X_C_DIR=c:\path\one\to\tools ; c:\path\two\to\tools
```

- Spaces and tabs are allowed within paths to accommodate Windows directories that contain spaces. For example, the pathnames in the following are valid:

```
set C6X_C_DIR=c:\first path\to\tools;d:\second path\to\tools
```

The environment variable remains set until you reboot the system or reset the variable by entering:

Operating System	Enter
UNIX (Bourne shell)	<code>unset C6X_C_DIR</code>
Windows	<code>set C6X_C_DIR=</code>

2.5 Precompiled Header Support

Precompiled header files may reduce the compile time for applications whose source files share a common set of headers, or a single file which has a large set of header files. Using precompiled headers, some recompilation is avoided thus saving compilation time.

There are two ways to use precompiled header files. One is the automatic precompiled header file processing and the other is called the manual precompiled header file processing.

2.5.1 Automatic Precompiled Header

The option to turn on automatic precompiled header processing is: `--pch`. Under this option, the compile step takes a snapshot of all the code prior to the header stop point, and dump it out to a file with suffix `.pch`. This snapshot does not have to be recompiled in the future compilations of this file or compilations of files with the same header files.

The stop point typically is the first token in the primary source file that does not belong to a preprocessing directive. For example, in the following the stopping point is before `int i`:

```
#include "x.h"  
#include "y.h"  
int i;
```

Carefully organizing the include directives across multiple files so that their header files maximize common usage can increase the compile time savings when using precompiled headers.

A precompiled header file is produced only if the header stop point and the code prior to it meet certain requirements.

2.5.2 Manual Precompiled Header

You can manually control the creation and use of precompiled headers by using several command line options. You specify a precompiled header file with a specific filename as follows:

`--create_pch=filename`

The `--use_pch=filename` option specifies that the indicated precompiled header file should be used for this compilation. If this precompiled header file is invalid, if its prefix does not match the prefix for the current primary source file for example, a warning is issued and the header file is not used.

If `--create_pch=filename` or `--use_pch=filename` is used with `--pch_dir`, the indicated filename, which can be a path name, is tacked on to the directory name, unless the filename is an absolute path name.

The `--create_pch`, `--use_pch`, and `--pch` options cannot be used together. If more than one of these options is specified, only the last one is applied. In manual mode, the header stop points are determined in the same way as in automatic mode. The precompiled header file applicability is determined in the same manner.

2.5.3 Additional Precompiled Header Options

The `--pch_verbose` option displays a message for each precompiled header file that is considered but not used. The `--pch_dir=pathname` option specifies the path where the precompiled header file resides.

2.6 Controlling the Preprocessor

This section describes specific features that control the preprocessor, which is part of the parser. A general description of C preprocessing is in section A12 of K&R. The C/C++ compiler includes standard C/C++ preprocessing functions, which are built into the first pass of the compiler. The preprocessor handles:

- Macro definitions and expansions
- #include files
- Conditional compilation
- Various preprocessor directives, specified in the source file as lines beginning with the # character

The preprocessor produces self-explanatory error messages. The line number and the filename where the error occurred are printed along with a diagnostic message.

2.6.1 Predefined Macro Names

The compiler maintains and recognizes the predefined macro names listed in [Table 2-37](#).

Table 2-37. Predefined C6000 Macro Names

Macro Name	Description
<code>__DATE__</code> ⁽¹⁾	Expands to the compilation date in the form <i>mmm dd yyyy</i>
<code>__FILE__</code> ⁽¹⁾	Expands to the current source filename
<code>__LINE__</code> ⁽¹⁾	Expands to the current line number
<code>__STDC__</code> ⁽¹⁾	Defined to indicate that compiler conforms to ISO C Standard. See Section 6.1 for exceptions to ISO C conformance.
<code>__STDC_VERSION__</code>	C standard macro
<code>__TI_32BIT_LONG__</code>	Defined to 1 if the EABI ABI is enabled (see Section 2.16); otherwise, it is undefined.
<code>__TI_40BIT_LONG__</code>	Defined to 1 if <code>__TI_32BIT_LONG__</code> is not defined; otherwise, it is undefined.
<code>__TI_COMPILER_VERSION__</code>	Defined to a 7-9 digit integer, depending on if X has 1, 2, or 3 digits. The number does not contain a decimal. For example, version 3.2.1 is represented as 3002001. The leading zeros are dropped to prevent the number being interpreted as an octal.
<code>__TI_EABI__</code>	Defined to 1 if the EABI is enabled (see Section 2.16); otherwise, it is undefined.
<code>__TI_GNU_ATTRIBUTE_SUPPORT__</code>	Defined if GCC extensions are enabled (the <code>--gcc</code> option is used); otherwise, it is undefined.
<code>__TI_STRICT_ANSI_MODE__</code>	Defined if strict ANSI/ISO mode is enabled (the <code>--strict_ansi</code> option is used); otherwise, it is undefined.
<code>__TI_WCHAR_T_BITS__</code>	
<code>__TIME__</code> ⁽¹⁾	Expands to the compilation time in the form " <i>hh:mm:ss</i> "
<code>__WCHAR_T_TYPE__</code>	
<code>_BIG_ENDIAN</code>	Defined if big-endian mode is selected (the <code>--big_endian</code> option is used); otherwise, it is undefined.
<code>_INLINE</code>	Expands to 1 if optimization is used (<code>--opt_level</code> or <code>-O</code> option); undefined otherwise. Regardless of any optimization, always undefined when <code>--no_inlining</code> is used.
<code>_LITTLE_ENDIAN</code>	Defined if little-endian mode is selected (the <code>--big_endian</code> option is not used); otherwise, it is undefined.
<code>_TMS320C6X</code>	Always defined
<code>_TMS320C6200</code>	Defined if target is C6200
<code>_TMS320C6400</code>	Defined if target is C6400, C6400+, C6740, or C6600
<code>_TMS320C6400_PLUS</code>	Defined if target is C6400+, C6740, or C6600
<code>_TMS320C6600</code>	Defined if target is C6600
<code>_TMS320C6700</code>	Defined if target is C6700, C6700+, C6740, or C6600
<code>_TMS320C6700_PLUS</code>	Defined if target is C6700+, C6740, or C6600
<code>_TMS320C6740</code>	Defined if target is C6740 or C6600
<code>__TMS320C6X__</code>	Always defined for use as alternate name for <code>_TMS320C6x</code>

⁽¹⁾ Specified by the ISO standard

You can use the names listed in [Table 2-37](#) in the same manner as any other defined name. For example,

```
printf ( "%s %s" , __TIME__ , __DATE__);
```

translates to a line such as:

```
printf ("%s %s" , "13:58:17" , "Jan 14 1997");
```

2.6.2 The Search Path for #include Files

The #include preprocessor directive tells the compiler to read source statements from another file. When specifying the file, you can enclose the filename in double quotes or in angle brackets. The filename can be a complete pathname, partial path information, or a filename with no path information.

- If you enclose the filename in double quotes (" "), the compiler searches for the file in the following directories in this order:
 1. The directory of the file that contains the #include directive and in the directories of any files that contain that file.
 2. Directories named with the --include_path option.
 3. Directories set with the C6X_C_DIR environment variable.
- If you enclose the filename in angle brackets (< >), the compiler searches for the file in the following directories in this order:
 1. Directories named with the --include_path option.
 2. Directories set with the C6X_C_DIR environment variable.

See [Section 2.6.2.1](#) for information on using the --include_path option. See [Section 2.4.2](#) for more information on input file directories.

2.6.2.1 Changing the #include File Search Path (--include_path Option)

The --include_path option names an alternate directory that contains #include files. The --include_path option's short form is -I. The format of the --include_path option is:

```
--include_path=directory1 [--include_path= directory2 ...]
```

There is no limit to the number of --include_path options per invocation of the compiler; each --include_path option names one *directory*. In C source, you can use the #include directive without specifying any directory information for the file; instead, you can specify the directory information with the --include_path option. For example, assume that a file called source.c is in the current directory. The file source.c contains the following directive statement:

```
#include "alt.h"
```

Assume that the complete pathname for alt.h is:

UNIX	/tools/files/alt.h
Windows	c:\tools\files\alt.h

The table below shows how to invoke the compiler. Select the command for your operating system:

Operating System	Enter
UNIX	c16x --include_path=/tools/files source.c
Windows	c16x --include_path=c:\tools\files source.c

NOTE: Specifying Path Information in Angle Brackets

If you specify the path information in angle brackets, the compiler applies that information relative to the path information specified with `--include_path` options and the `C6X_C_DIR` environment variable.

For example, if you set up `C6X_C_DIR` with the following command:

```
C6X_C_DIR "/usr/include:/usr/ucb"; export C6X_C_DIR
```

or invoke the compiler with the following command:

```
cl6x --include_path=/usr/include file.c
```

and `file.c` contains this line:

```
#include <sys/proc.h>
```

the result is that the included file is in the following path:

```
/usr/include/sys/proc.h
```

2.6.3 Generating a Preprocessed Listing File (`--preproc_only` Option)

The `--preproc_only` option allows you to generate a preprocessed version of your source file with an extension of `.pp`. The compiler's preprocessing functions perform the following operations on the source file:

- Each source line ending in a backslash (`\`) is joined with the following line.
- Trigraph sequences are expanded.
- Comments are removed.
- `#include` files are copied into the file.
- Macro definitions are processed.
- All macros are expanded.
- All other preprocessing directives, including `#line` directives and conditional compilation, are expanded.

2.6.4 Continuing Compilation After Preprocessing (`--preproc_with_compile` Option)

If you are preprocessing, the preprocessor performs preprocessing only; it does not compile your source code. To override this feature and continue to compile after your source code is preprocessed, use the `--preproc_with_compile` option along with the other preprocessing options. For example, use `--preproc_with_compile` with `--preproc_only` to perform preprocessing, write preprocessed output to a file with a `.pp` extension, and compile your source code.

2.6.5 Generating a Preprocessed Listing File With Comments (`--preproc_with_comment` Option)

The `--preproc_with_comment` option performs all of the preprocessing functions except removing comments and generates a preprocessed version of your source file with a `.pp` extension. Use the `--preproc_with_comment` option instead of the `--preproc_only` option if you want to keep the comments.

2.6.6 Generating a Preprocessed Listing File With Line-Control Information (`--preproc_with_line` Option)

By default, the preprocessed output file contains no preprocessor directives. To include the `#line` directives, use the `--preproc_with_line` option. The `--preproc_with_line` option performs preprocessing only and writes preprocessed output with line-control information (`#line` directives) to a file named as the source file but with a `.pp` extension.

2.6.7 **Generating Preprocessed Output for a Make Utility (--preproc_dependency Option)**

The `--preproc_dependency` option performs preprocessing only, but instead of writing preprocessed output, writes a list of dependency lines suitable for input to a standard make utility. If you do not supply an optional filename, the list is written to a file with the same name as the source file but with a `.pp` extension.

2.6.8 **Generating a List of Files Included With the #include Directive (--preproc_includes Option)**

The `--preproc_includes` option performs preprocessing only, but instead of writing preprocessed output, writes a list of files included with the `#include` directive. If you do not supply an optional filename, the list is written to a file with the same name as the source file but with a `.pp` extension.

2.6.9 **Generating a List of Macros in a File (--preproc_macros Option)**

The `--preproc_macros` option generates a list of all predefined and user-defined macros. If you do not supply an optional filename, the list is written to a file with the same name as the source file but with a `.pp` extension. Predefined macros are listed first and indicated by the comment `/* Predefined */`. User-defined macros are listed next and indicated by the source filename.

2.7 **Understanding Diagnostic Messages**

One of the compiler's primary functions is to report diagnostics for the source program. The new linker also reports diagnostics. When the compiler or linker detects a suspect condition, it displays a message in the following format:

"file.c", line n : diagnostic severity : diagnostic message

<i>"file.c"</i>	The name of the file involved
line n :	The line number where the diagnostic applies
<i>diagnostic severity</i>	The diagnostic message severity (severity category descriptions follow)
<i>diagnostic message</i>	The text that describes the problem

Diagnostic messages have an associated severity, as follows:

- A **fatal error** indicates a problem so severe that the compilation cannot continue. Examples of such problems include command-line errors, internal errors, and missing include files. If multiple source files are being compiled, any source files after the current one will not be compiled.
- An **error** indicates a violation of the syntax or semantic rules of the C/C++ language. Compilation continues, but object code is not generated.
- A **warning** indicates something that is valid but questionable. Compilation continues and object code is generated (if no errors are detected).
- A **remark** is less serious than a warning. It indicates something that is valid and probably intended, but may need to be checked. Compilation continues and object code is generated (if no errors are detected). By default, remarks are not issued. Use the `--issue_remarks` compiler option to enable remarks.

Diagnostics are written to standard error with a form like the following example:

```
"test.c", line 5: error: a break statement may only be used within a loop or switch
    break;
    ^
```

By default, the source line is omitted. Use the `--verbose_diagnostics` compiler option to enable the display of the source line and the error position. The above example makes use of this option.

The message identifies the file and line involved in the diagnostic, and the source line itself (with the position indicated by the `^` character) follows the message. If several diagnostics apply to one source line, each diagnostic has the form shown; the text of the source line is displayed several times, with an appropriate position indicated each time.

Long messages are wrapped to additional lines, when necessary.

You can use the `--display_error_number` command-line option to request that the diagnostic's numeric identifier be included in the diagnostic message. When displayed, the diagnostic identifier also indicates whether the diagnostic can have its severity overridden on the command line. If the severity can be overridden, the diagnostic identifier includes the suffix `-D` (for *discretionary*); otherwise, no suffix is present. For example:

```
"Test_name.c", line 7: error #64-D: declaration does not declare anything
    struct {};
    ^
"Test_name.c", line 9: error #77: this declaration has no storage class or type specifier
    xxxxx;
    ^
```

Because an error is determined to be discretionary based on the error severity associated with a specific context, an error can be discretionary in some cases and not in others. All warnings and remarks are discretionary.

For some messages, a list of entities (functions, local variables, source files, etc.) is useful; the entities are listed following the initial error message:

```
"test.c", line 4: error: more than one instance of overloaded function "f"
    matches the argument list:
    function "f(int)"
    function "f(float)"
    argument types are: (double)

    f(1.5);
    ^
```

In some cases, additional context information is provided. Specifically, the context information is useful when the front end issues a diagnostic while doing a template instantiation or while generating a constructor, destructor, or assignment operator function. For example:

```
"test.c", line 7: error: "A::A()" is inaccessible
    B x;
    ^
    detected during implicit generation of "B::B()" at line 7
```

Without the context information, it is difficult to determine to what the error refers.

2.7.1 Controlling Diagnostics

The C/C++ compiler provides diagnostic options to control compiler- and linker-generated diagnostics. The diagnostic options must be specified before the `--run_linker` option.

- `--diag_error=num`** Categorizes the diagnostic identified by *num* as an error. To determine the numeric identifier of a diagnostic message, use the `--display_error_number` option first in a separate compile. Then use `--diag_error=num` to recategorize the diagnostic as an error. You can only alter the severity of discretionary diagnostics.
- `--diag_remark=num`** Categorizes the diagnostic identified by *num* as a remark. To determine the numeric identifier of a diagnostic message, use the `--display_error_number` option first in a separate compile. Then use `--diag_remark=num` to recategorize the diagnostic as a remark. You can only alter the severity of discretionary diagnostics.
- `--diag_suppress=num`** Suppresses the diagnostic identified by *num*. To determine the numeric identifier of a diagnostic message, use the `--display_error_number` option first in a separate compile. Then use `--diag_suppress=num` to suppress the diagnostic. You can only suppress discretionary diagnostics.
- `--diag_warning=num`** Categorizes the diagnostic identified by *num* as a warning. To determine the numeric identifier of a diagnostic message, use the `--display_error_number` option first in a separate compile. Then use `--diag_warning=num` to recategorize the diagnostic as a warning. You can only alter the severity of discretionary diagnostics.

--display_error_number	Displays a diagnostic's numeric identifier along with its text. Use this option in determining which arguments you need to supply to the diagnostic suppression options (<code>--diag_suppress</code> , <code>--diag_error</code> , <code>--diag_remark</code> , and <code>--diag_warning</code>). This option also indicates whether a diagnostic is discretionary. A discretionary diagnostic is one whose severity can be overridden. A discretionary diagnostic includes the suffix <code>-D</code> ; otherwise, no suffix is present. See Section 2.7 .
--emit_warnings_as_errors	Treats all warnings as errors. This option cannot be used with the <code>--no_warnings</code> option. The <code>--diag_remark</code> option takes precedence over this option. This option takes precedence over the <code>--diag_warning</code> option.
--issue_remarks	Issues remarks (nonserious warnings), which are suppressed by default.
--no_warnings	Suppresses warning diagnostics (errors are still issued).
--set_error_limit=num	Sets the error limit to <i>num</i> , which can be any decimal value. The compiler abandons compiling after this number of errors. (The default is 100.)
--verbose_diagnostics	Provides verbose diagnostics that display the original source with line-wrap and indicate the position of the error in the source line
--write_diagnostics_file	Produces a diagnostics information file with the same source file name with an <code>.err</code> extension. (The <code>--write_diagnostics_file</code> option is not supported by the linker.)

2.7.2 How You Can Use Diagnostic Suppression Options

The following example demonstrates how you can control diagnostic messages issued by the compiler. You control the linker diagnostic messages in a similar manner.

```
int one();
int I;
int main()
{
    switch (I){
        case 1;
            return one ();
            break;
        default:
            return 0;
            break;
    }
}
```

If you invoke the compiler with the `--quiet` option, this is the result:

```
"err.c", line 9: warning: statement is unreachable
"err.c", line 12: warning: statement is unreachable
```

Because it is standard programming practice to include `break` statements at the end of each case arm to avoid the fall-through condition, these warnings can be ignored. Using the `--display_error_number` option, you can find out the diagnostic identifier for these warnings. Here is the result:

```
[err.c]
"err.c", line 9: warning #111-D: statement is unreachable
"err.c", line 12: warning #111-D: statement is unreachable
```

Next, you can use the diagnostic identifier of 111 as the argument to the `--diag_remark` option to treat this warning as a remark. This compilation now produces no diagnostic messages (because remarks are disabled by default).

Although this type of control is useful, it can also be extremely dangerous. The compiler often emits messages that indicate a less than obvious problem. Be careful to analyze all diagnostics emitted before using the suppression options.

2.8 Other Messages

Other error messages that are unrelated to the source, such as incorrect command-line syntax or inability to find specified files, are usually fatal. They are identified by the symbol >> preceding the message.

2.9 Generating Cross-Reference Listing Information (--gen_acp_xref Option)

The --gen_acp_xref option generates a cross-reference listing file that contains reference information for each identifier in the source file. (The --gen_acp_xref option is separate from --cross_reference, which is an assembler rather than a compiler option.) The cross-reference listing file has the same name as the source file with a .crl extension.

The information in the cross-reference listing file is displayed in the following format:

sym-id name X filename line number column number

<i>sym-id</i>	An integer uniquely assigned to each identifier
<i>name</i>	The identifier name
<i>X</i>	One of the following values:
	D Definition
	d Declaration (not a definition)
	M Modification
	A Address taken
	U Used
	C Changed (used and modified in a single operation)
	R Any other kind of reference
	E Error; reference is indeterminate
<i>filename</i>	The source file
<i>line number</i>	The line number in the source file
<i>column number</i>	The column number in the source file

2.10 Generating a Raw Listing File (--gen_acp_raw Option)

The --gen_acp_raw option generates a raw listing file that can help you understand how the compiler is preprocessing your source file. Whereas the preprocessed listing file (generated with the --preproc_only, --preproc_with_comment, --preproc_with_line, and --preproc_dependency preprocessor options) shows a preprocessed version of your source file, a raw listing file provides a comparison between the original source line and the preprocessed output. The raw listing file has the same name as the corresponding source file with an .rl extension.

The raw listing file contains the following information:

- Each original source line
- Transitions into and out of include files
- Diagnostics
- Preprocessed source line if nontrivial processing was performed (comment removal is considered trivial; other preprocessing is nontrivial)

Each source line in the raw listing file begins with one of the identifiers listed in [Table 2-38](#).

Table 2-38. Raw Listing File Identifiers

Identifier	Definition
N	Normal line of source
X	Expanded line of source. It appears immediately following the normal line of source if nontrivial preprocessing occurs.
S	Skipped source line (false #if clause)
L	Change in source position, given in the following format: L <i>line number filename key</i> Where <i>line number</i> is the line number in the source file. The <i>key</i> is present only when the change is due to entry/exit of an include file. Possible values of <i>key</i> are: 1 = entry into an include file 2 = exit from an include file

The `--gen_acp_raw` option also includes diagnostic identifiers as defined in [Table 2-39](#).

Table 2-39. Raw Listing File Diagnostic Identifiers

Diagnostic Identifier	Definition
E	Error
F	Fatal
R	Remark
W	Warning

Diagnostic raw listing information is displayed in the following format:

S filename line number column number diagnostic

S One of the identifiers in [Table 2-39](#) that indicates the severity of the diagnostic

filename The source file

line number The line number in the source file

column number The column number in the source file

diagnostic The message text for the diagnostic

Diagnostics after the end of file are indicated as the last line of the file with a column number of 0. When diagnostic message text requires more than one line, each subsequent line contains the same file, line, and column information but uses a lowercase version of the diagnostic identifier. For more information about diagnostic messages, see [Section 2.7](#).

2.11 Using Inline Function Expansion

When an inline function is called, the C/C++ source code for the function is inserted at the point of the call. This is known as inline function expansion. Inline function expansion is advantageous in short functions for the following reasons:

There are several types of inline function expansion:

- Inlining with intrinsic operators (intrinsic operators are always inlined)
- Automatic inlining
- Definition-controlled inlining with the unguarded inline keyword
- Definition-controlled inlining with the guarded inline keyword

NOTE: Function Inlining Can Greatly Increase Code Size

Expanding functions inline increases code size, especially inlining a function that is called in a number of places. Function inlining is optimal for functions that are called only from a small number of places and for small functions. If your code size seems too large, see [Section 3.5](#).

2.11.1 Inlining Intrinsic Operators

There are many intrinsic operators for the C6000. All of them are automatically inlined by the compiler. The inlining happens automatically whether or not you use the optimizer.

For details about intrinsic operators, and a list of the intrinsic operators, see [Section 7.5.5](#).

2.11.2 Automatic Inlining

When optimizing with the `--opt_level=3` or `--opt_level=2` option (aliased as `-O3` or `-O2`), the compiler automatically inlines certain functions. For more information, see [Section 3.13](#).

2.11.3 Unguarded Definition-Controlled Inlining

The inline keyword specifies that a function is expanded inline at the point at which it is called rather than by using standard calling procedures. The compiler performs inline expansion of functions declared with the inline keyword.

You must invoke the optimizer with any `--opt_level` option (`--opt_level=0`, `--opt_level=1`, `--opt_level=2`, or `--opt_level=3`) to turn on definition-controlled inlining. Automatic inlining is also turned on when using `--opt_level=3`.

The `--no_inlining` option turns off definition-controlled inlining. This option is useful when you need a certain level of optimization but do not want definition-controlled inlining.

[Example 2-1](#) shows usage of the inline keyword, where the function call is replaced by the code in the called function.

Example 2-1. Using the Inline Keyword

```
inline float volume_sphere(float r)
{
    return 4.0/3.0 * PI * r * r * r;
}
int foo(...)
{
    ...
    volume = volume_sphere(radius);
    ...
}
```

2.11.4 Guarded Inlining and the `_INLINE` Preprocessor Symbol

When declaring a function in a header file as static inline, you must follow additional procedures to avoid a potential code size increase when inlining is turned off with `--no_inlining` or the optimizer is not run.

To prevent a static inline function in a header file from causing an increase in code size when inlining gets turned off, use the following procedure. This allows external-linkage when inlining is turned off; thus, only one function definition will exist throughout the object files.

- Prototype a static inline version of the function. Then, prototype an alternative, nonstatic, externally-linked version of the function. Conditionally preprocess these two prototypes with the `_INLINE` preprocessor symbol, as shown in [Example 2-2](#).
- Create an identical version of the function definition in a `.c` or `.cpp` file, as shown in [Example 2-3](#).

In the following examples there are two definitions of the `strlen` function. The first ([Example 2-2](#)), in the header file, is an inline definition. This definition is enabled and the prototype is declared as static inline only if `_INLINE` is true (`_INLINE` is automatically defined for you when the optimizer is used and `--no_inlining` is not specified).

The second definition (see [Example 2-3](#)) for the library, ensures that the callable version of `strlen` exists when inlining is disabled. Since this is not an inline function, the `_INLINE` preprocessor symbol is undefined (`#undef`) before `string.h` is included to generate a noninline version of `strlen`'s prototype.

Example 2-2. Header File `string.h`

```

/*****
/* string.h vx.xx (Excerpted) */
/* Copyright (c) 1993-2011 Texas Instruments Incorporated */
/*****
#ifdef _INLINE
#define _IDECL static inline
#else
#define _IDECL extern _CODE_ACCESS
#endif

_IDECL size_t strlen(const char *_string);

#ifdef _INLINE

/*****
/* strlen */
/*****
static inline size_t strlen(const char *string)
{
    size_t    n = (size_t)-1;
    const char *s = string - 1;

    do n++; while (*++s);
    return n
}

#endif

```

Example 2-3. Library Definition File

```

/*****
/*  strlen
*****/
#undef _INLINE

#include <string.h>
{
_CODE_ACCESS size_t strlen(cont char * string)
    size_t      n = (size_t)-1;
    const char *s = string - 1;

    do n++; while (*++s);
    return n;
}

```

2.11.5 Inlining Restrictions

There are several restrictions on what functions can be inlined for both automatic inlining and definition-controlled inlining. Functions with local static variables or a variable number of arguments are not inlined, with the exception of functions declared as static inline. In functions declared as static inline, expansion occurs despite the presence of local static variables. In addition, a limit is placed on the depth of inlining for recursive or nonleaf functions. Furthermore, inlining should be used for small functions or functions that are called in a few places (though the compiler does not enforce this).

At a given call site, a function may be disqualified from inlining if it:

- Is not defined in the current compilation unit
- Never returns
- Is recursive
- Has a `FUNC_CANNOT_INLINE` pragma
- Has a variable length argument list
- Has a different number of arguments than the call site
- Has an argument whose type is incompatible with the corresponding call site argument
- Has a class, struct or union parameter
- Contains a volatile local variable or argument
- Is not declared inline and contains an `asm()` statement that is not a comment
- Is not declared inline and it is `main()`
- Is not declared inline and it is an interrupt function
- Is not declared inline and returns void but its return value is needed.
- Is not declared inline and will require too much stack space for local array or structure variables.

2.12 Interrupt Flexibility Options (--interrupt_threshold Option)

On the C6000 architecture, interrupts cannot be taken in the delay slots of a branch. In some instances the compiler can generate code that cannot be interrupted for a potentially large number of cycles. For a given real-time system, there may be a hard limit on how long interrupts can be disabled.

The `--interrupt_threshold=n` option specifies an interrupt threshold value *n*. The threshold value specifies the maximum number of cycles that the compiler can disable interrupts. If the *n* is omitted, the compiler assumes that the code is never interrupted. In Code Composer Studio, to specify that the code is never interrupted, select the Interrupt Threshold check box and leave the text box blank in the Build Options dialog box on the Compiler tab, Advanced category.

If the `--interrupt_threshold=n` option is not specified, then interrupts are only explicitly disabled around software pipelined loops. When using the `--interrupt_threshold=n` option, the compiler analyzes the loop structure and loop counter to determine the maximum number of cycles it takes to execute a loop. If it can determine that the maximum number of cycles is less than the threshold value, the compiler generates the fastest/optimal version of the loop. If the loop is smaller than six cycles, interrupts are not able to occur because the loop is always executing inside the delay slots of a branch. Otherwise, the compiler generates a loop that can be interrupted (and still generate correct results—single assignment code), which in most cases degrades the performance of the loop.

The `--interrupt_threshold=n` option does not comprehend the effects of the memory system. When determining the maximum number of execution cycles for a loop, the compiler does not compute the effects of using slow off-chip memory or memory bank conflicts. It is recommended that a conservative threshold value is used to adjust for the effects of the memory system.

See [Section 6.9.11](#) or the *TMS320C6000 Programmer's Guide* for more information.

RTS Library Files Are Not Built With the --interrupt_threshold Option

NOTE: The run-time-support library files provided with the compiler are not built with the interrupt flexibility option. Refer to the readme file to see how the run-time-support library files were built for your release. See [Section 8.5](#) to build your own run-time-support library files with the interrupt flexibility option.

Special Cases With the --interrupt_threshold Option

NOTE: The `--interrupt_threshold=0` option generates the same code to disable interrupts around software-pipelined loops as when the `--interrupt_threshold` option is not used.

The `--interrupt_threshold` option (the threshold value is omitted) means that no code is added to disable interrupts around software pipelined loops, which means that the code cannot be safely interrupted. Also, loop performance does not degrade because the compiler is not trying to make the loop interruptible by ensuring that there is at least one cycle in the loop kernel that is not in the delay slot of a branch instruction.

2.13 Linking C6400 Code With C6200/C6700/Older C6400 Object Code

In order to facilitate certain packed-data optimizations, the alignment of top-level arrays for the C6400 family was changed from 4 bytes to 8 bytes. (For C6200 and C6700 code, the alignment for top-level arrays is always 4 bytes.)

If you are linking C6400/C6400+/C6740/C6600 with C6200/C6700 code or older C6400 code, you may need to take steps to ensure compatibility. The following lists the potential alignment conflicts and possible solutions.

Potential alignment conflicts occur when:

- Linking new C6400/C6400+/C6740/C6600 code with any C6400 code already compiled with the 4.0 tools.
- Linking new C6400/C6400+/C6740/C6600 code with code already compiled with any version of the tools for the C6200 or C6700 family.

Solutions (pick one):

- Recompile the entire application with the `--silicon_version=6400` switch. This solution, if possible, is recommended because it can lead to better performance.
- Compile the new code with the `--target_compatibility_6200` option. When using this option, link with the 6200 RTS library to ensure proper compatibility. The `--target_compatibility_6200` option changes the alignment of top-level arrays to 4 bytes when the `--silicon_version=6400` or `--silicon_version=6400+` option is used.

This option is deprecated; it may be removed in the future.

The alignment of top-level arrays for the C6600 family is 16 bytes to facilitate compatibility with future C6600 family devices. This change in alignment does not have any compatibility issues with the C6400/C6400+/C6740 device code as the C6600 can safely accept top-level arrays aligned to an 8-byte boundary.

2.14 Using Interlist

The compiler tools include a feature that interlists C/C++ source statements into the assembly language output of the compiler. The interlist feature enables you to inspect the assembly code generated for each C statement. The interlist behaves differently, depending on whether or not the optimizer is used, and depending on which options you specify.

The easiest way to invoke the interlist feature is to use the `--c_src_interlist` option. To compile and run the interlist on a program called `function.c`, enter:

```
c16x --c_src_interlist function
```

The `--c_src_interlist` option prevents the compiler from deleting the interlisted assembly language output file. The output assembly file, `function.asm`, is assembled normally.

When you invoke the interlist feature without the optimizer, the interlist runs as a separate pass between the code generator and the assembler. It reads both the assembly and C/C++ source files, merges them, and writes the C/C++ statements into the assembly file as comments.

Using the `--c_src_interlist` option can cause performance and/or code size degradation.

[Example 2-4](#) shows a typical interlisted assembly file.

For more information about using the interlist feature with the optimizer, see [Section 3.14.2](#).

Example 2-4. An Interlisted Assembly Language File

```

_main:

        STW    .D2    B3,*SP--(12)
        STW    .D2    A10,*+SP(8)
;-----
; 5 | printf("Hello, world\n");
;-----
        B      .S1    _printf
        NOP
        MVKL   .S1    SL1+0,A0
        MVKH   .S1    SL1+0,A0
||      MVKL   .S2    RL0,B3
        STW    .D2    A0,*+SP(4)
||      MVKH   .S2    RL0,B3
RLO:    ; CALL OCCURS
;-----
; 6 | return 0;
;-----
        ZERO   .L1    A10
        MV     .L1    A10,A4
        LDW    .D2    *+SP(8),A10
        LDW    .D2    *++SP(12),B3
        NOP
        B      .S2    B3
        NOP
        ; BRANCH OCCURS

```

2.15 Generating and Using Performance Advice

The compiler can do better optimization in some cases, if the user aids the compiler by providing additional information in the code. The compiler can prompt you to take certain actions to improve performance, by emitting "Advice". To get this Advice, use the `--advice:performance` option:

```
cl6x --advice:performance -o3 filename.c
```

This Performance Advice is of 3 different types :

- Advice to use correct compiler options
- Advice to prevent software pipeline disqualification
- Advice to improve loop performance

For more details on Using Performance Advice to Optimize your Code, see [Section 3.14](#)

2.16 Controlling Application Binary Interface

An Application Binary Interface (ABI) defines the low level interface between object files, and between an executable and its execution environment. An ABI allows ABI-compliant object code to link together, regardless of its source, and allows the resulting executable to run on any system that supports that ABI

Object files conforming to different ABIs cannot be linked together. The linker detects this situation and generates an error.

The C6000 compiler supports two ABIs. The ABI is chosen through the `--abi` option as follows:

- **COFF ABI** (`--abi=coffabi`)
The COFF ABI is the original ABI format. There is no COFF to ELF conversion possible; recompile or reassemble assembly code.
- **C6000 EABI** (`--abi=eabi`)
Use this option to select the C6000 Embedded Application Binary Interface (EABI).
All code in an EABI application must be built for EABI. Make sure all your libraries are available in

EABI mode before migrating your existing COFF ABI systems to C6000 EABI. See http://tiexpressdsp.com/index.php/EABI_Support_in_C6000_Compiler for full details.

For more details on the different ABIs, see [Section 6.11](#).

2.17 Enabling Entry Hook and Exit Hook Functions

An entry hook is a routine that is called upon entry to each function in the program. An exit hook is a routine that is called upon exit of each function. Applications for hooks include debugging, trace, profiling, and stack overflow checking.

Entry and exit hooks are enabled using the following options:

- entry_hook[=*name*]** Enables entry hooks. If specified, the hook function is called *name*. Otherwise, the default entry hook function name is `__entry_hook`.
- entry_parm{=*name*|
address|none}** Specify the parameters to the hook function. The name parameter specifies that the name of the calling function is passed to the hook function as an argument. In this case the signature for the hook function is: `void hook(const char *name);`
The address parameter specifies that the address of the calling function is passed to the hook function. In this case the signature for the hook function is: `void hook(void (*addr)());`
The none parameter specifies that the hook is called with no parameters. This is the default. In this case the signature for the hook function is: `void hook(void);`
- exit_hook[=*name*]** Enables exit hooks. If specified, the hook function is called *name*. Otherwise, the default exit hook function name is `__exit_hook`.
- exit_parm{=*name*|
address|none}** Specify the parameters to the hook function. The name parameter specifies that the name of the calling function is passed to the hook function as an argument. In this case the signature for the hook function is: `void hook(const char *name);`
The address parameter specifies that the address of the calling function is passed to the hook function. In this case the signature for the hook function is: `void hook(void (*addr)());`
The none parameter specifies that the hook is called with no parameters. This is the default. In this case the signature for the hook function is: `void hook(void);`

The presence of the hook options creates an implicit declaration of the hook function with the given signature. If a declaration or definition of the hook function appears in the compilation unit compiled with the options, it must agree with the signatures listed above.

In C++, the hooks are declared extern "C". Thus you can define them in C (or assembly) without being concerned with name mangling.

Hooks can be declared inline, in which case the compiler tries to inline them using the same criteria as other inline functions.

Entry hooks and exit hooks are independent. You can enable one but not the other, or both. The same function can be used as both the entry and exit hook.

You must take care to avoid recursive calls to hook functions. The hook function should not call any function which itself has hook calls inserted. To help prevent this, hooks are not generated for inline functions, or for the hook functions themselves.

You can use the `--remove_hooks_when_inlining` option to remove entry/exit hooks for functions that are auto-inlined by the optimizer.

See [Section 6.9.23](#) for information about the `NO_HOOKS` pragma.

Optimizing Your Code

The compiler tools can perform many optimizations to improve the execution speed and reduce the size of C and C++ programs by simplifying loops, software pipelining, rearranging statements and expressions, and allocating variables into registers.

This chapter describes how to invoke different levels of optimization and describes which optimizations are performed at each level. This chapter also describes how you can use the Interlist feature when performing optimization and how you can profile or debug optimized code.

Topic	Page
3.1 Invoking Optimization	59
3.2 Optimizing Software Pipelining	60
3.3 Redundant Loops	67
3.4 Utilizing the Loop Buffer Using SPLOOP on C6400+, C6740, and C6600	68
3.5 Reducing Code Size (--opt_for_space (or -ms) Option)	68
3.6 Performing File-Level Optimization (--opt_level=3 option)	70
3.7 Performing Program-Level Optimization (--program_level_compile and --opt_level=3 options)	71
3.8 Using Feedback Directed Optimization	73
3.9 Using Profile Information to Get Better Program Cache Layout and Analyze Code Coverage	77
3.10 Indicating Whether Certain Aliasing Techniques Are Used	87
3.11 Prevent Reordering of Associative Floating-Point Operations	88
3.12 Use Caution With asm Statements in Optimized Code	89
3.13 Automatic Inline Expansion (--auto_inline Option)	89
3.14 Using Performance Advice to Optimize your Code	90

3.1 Invoking Optimization

The C/C++ compiler is able to perform various optimizations. High-level optimizations are performed in the optimizer and low-level, target-specific optimizations occur in the code generator. Use high-level optimization levels, such as `--opt_level=2` and `--opt_level=3`, to achieve optimal code.

The easiest way to invoke optimization is to use the compiler program, specifying the `--opt_level=n` option on the compiler command line. You can use `-On` to alias the `--opt_level` option. The *n* denotes the level of optimization (0, 1, 2, and 3), which controls the type and degree of optimization.

- **`--opt_level=0` or `-O0`**

- Performs control-flow-graph simplification
- Allocates variables to registers
- Performs loop rotation
- Eliminates unused code
- Simplifies expressions and statements
- Expands calls to functions declared inline

- **`--opt_level=1` or `-O1`**

Performs all `--opt_level=0` (`-O0`) optimizations, plus:

- Performs local copy/constant propagation
- Removes unused assignments
- Eliminates local common expressions

- **`--opt_level=2` or `-O2`**

Performs all `--opt_level=1` (`-O1`) optimizations, plus:

- Performs software pipelining (see [Section 3.2](#))
- Performs loop optimizations
- Eliminates global common subexpressions
- Eliminates global unused assignments
- Converts array references in loops to incremented pointer form
- Performs loop unrolling

The optimizer uses `--opt_level=2` (`-O2`) as the default if you use `--opt_level` (`-O`) without an optimization level.

- **`--opt_level=3` or `-O3`**

Performs all `--opt_level=2` (`-O2`) optimizations, plus:

- Removes all functions that are never called
- Simplifies functions with return values that are never used
- Inlines calls to small functions
- Reorders function declarations; the called functions attributes are known when the caller is optimized
- Propagates arguments into function bodies when all calls pass the same value in the same argument position
- Identifies file-level variable characteristics

If you use `--opt_level=3` (`-O3`), see [Section 3.6](#) and [Section 3.7](#) for more information.

The levels of optimizations described above are performed by the stand-alone optimization pass. The code generator performs several additional optimizations, particularly processor-specific optimizations. It does so regardless of whether you invoke the optimizer. These optimizations are always enabled, although they are more effective when the optimizer is used.

Do Not Lower the Optimization Level to Control Code Size

NOTE: To reduce code size, do not lower the level of optimization. Instead, use the `--opt_for_space` option to control the code size/performance tradeoff. Higher optimization levels (`--opt_level` or `-O`) combined with high `--opt_for_space` levels result in the smallest code size. For more information, see [Section 3.5](#).

The `--opt_level= n (-O n)` Option Applies to the Assembly Optimizer

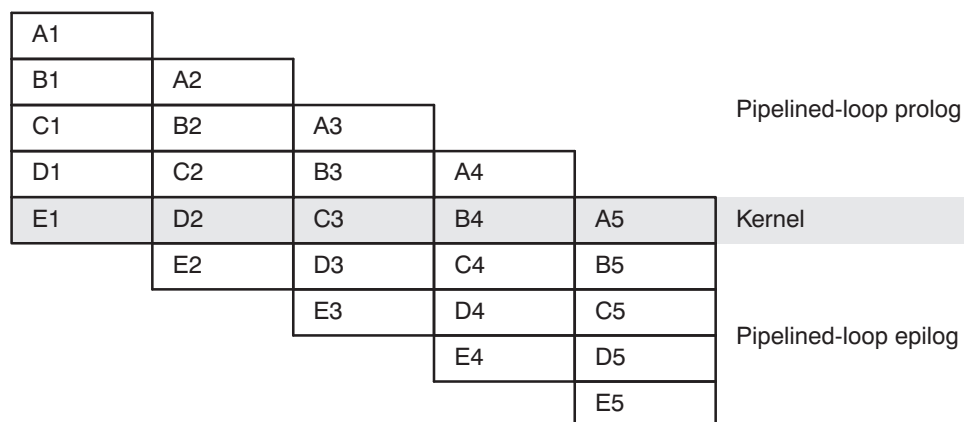
NOTE: The `--opt_level=n (-O n)` option should also be used with the assembly optimizer. Although the assembly optimizer does not perform all the optimizations described here, key optimizations such as software pipelining and loop unrolling require the `--opt_level (-O)` option.

3.2 Optimizing Software Pipelining

Software pipelining schedules instructions from a loop so that multiple iterations of the loop execute in parallel. At optimization levels `--opt_level=2` (or `-O2`) and `--opt_level=3` (or `-O3`), the compiler usually attempts to software pipeline your loops. The `--opt_for_space` option also affects the compiler's decision to attempt to software pipeline loops. In general, code size and performance are better when you use the `--opt_level=2` or `--opt_level=3` options. (See [Section 3.1](#).)

[Figure 3-1](#) illustrates a software-pipelined loop. The stages of the loop are represented by A, B, C, D, and E. In this figure, a maximum of five iterations of the loop can execute at one time. The shaded area represents the loop *kernel*. In the loop kernel, all five stages execute in parallel. The area above the kernel is known as the *pipelined loop prolog*, and the area below the kernel is known as the *pipelined loop epilog*.

Figure 3-1. Software-Pipelined Loop



If you enter comments on instructions in your linear assembly input file, the compiler moves the comments to the output file along with additional information. It attaches a 2-tuple `<x, y>` to the comments to specify the iteration and cycle of the loop an instruction is on in the software pipeline. The zero-based number `x` represents the iteration the instruction is on during the first execution of the loop kernel. The zero-based number `y` represents the cycle that the instruction is scheduled on within a single iteration of the loop.

For more information about software pipelining, see the *TMS320C6000 Programmer's Guide*.

3.2.1 Turn Off Software Pipelining (`--disable_software_pipelining` Option)

At optimization levels `--opt_level=2` (or `-O2`) and `-O3`, the compiler attempts to software pipeline your loops. You might not want your loops to be software pipelined for debugging reasons. Software-pipelined loops are sometimes difficult to debug because the code is not presented serially. The `--disable_software_pipelining` option affects both compiled C/C++ code and assembly optimized code.

Software Pipelining May Increase Code Size

NOTE: Software pipelining without the use of SPLOOP can lead to significant increases in code size. To control code size for loops that get software pipelined, it is preferable to use the `--opt_for_space` option rather than the `--disable_software_pipelining` option. The `--opt_for_space` option is capable of disabling non-SPLOOP software pipelining if necessary to achieve code size savings, but it does not affect the SPLOOP capability of C64x+ and C674x devices. SPLOOP does not significantly increase code size, but can greatly speed up a loop. Using `--disable_software_pipelining` disables all software pipelining including SPLOOP.

3.2.2 Software Pipelining Information

The compiler embeds software pipelined loop information in the `.asm` file. This information is used to optimize C/C++ code or linear assembly code.

The software pipelining information appears as a comment in the `.asm` file before a loop and for the assembly optimizer the information is displayed as the tool is running. [Example 3-1](#) illustrates the information that is generated for each loop.

The `--debug_software_pipeline` option adds additional information displaying the register usage at each cycle of the loop kernel and displays the instruction ordering of a single iteration of the software pipelined loop.

More Details on Software Pipelining Information

NOTE: Refer to the *TMS320C6000 Programmer's Guide* for details on the information and messages that can appear in the Software Pipelining Information comment block before each loop.

Example 3-1. Software Pipelining Information

```

*-----*
;*  SOFTWARE PIPELINE INFORMATION
;*
;*  Known Minimum Trip Count      : 2
;*  Known Maximum Trip Count     : 2
;*  Known Max Trip Count Factor  : 2
;*  Loop Carried Dependency Bound(^) : 4
;*  Unpartitioned Resource Bound  : 4
;*  Partitioned Resource Bound(*)  : 5
;*  Resource Partition:
;*
;*          A-side  B-side
;*  .L units      2      3
;*  .S units      4      4
;*  .D units      1      0
;*  .M units      0      0
;*  .X cross paths 1      3
;*  .T address paths 1      0
;*  Long read paths 0      0
;*  Long write paths 0      0
;*  Logical ops (.LS)      0      1      (.L or .S unit)
;*  Addition ops (.LSD)    6      3      (.L or .S or .D unit)
;*  Bound(.L .S .LS)      3      4
;*  Bound(.L .S .D .LS .LSD) 5*     4
;*
;*
;*  Searching for software pipeline schedule at ...
;*      ii = 5 Register is live too long
;*      ii = 6 Did not find schedule
;*      ii = 7 Schedule found with 3 iterations in parallel
;*  done
;*
;*  Epilog not entirely removed
;*  Collapsed epilog stages      : 1
;*
;*  Prolog not removed
;*  Collapsed prolog stages      : 0
;*
;*  Minimum required memory pad : 2 bytes
;*
;*  Minimum safe trip count     : 2
;*-----*
    
```

The terms defined below appear in the software pipelining information. For more information on each term, see the *TMS320C6000 Programmer's Guide*.

- **Loop unroll factor.** The number of times the loop was unrolled specifically to increase performance based on the resource bound constraint in a software pipelined loop.
- **Known minimum trip count.** The minimum number of times the loop will be executed.
- **Known maximum trip count.** The maximum number of times the loop will be executed.
- **Known max trip count factor.** Factor that would always evenly divide the loops trip count. This information can be used to possibly unroll the loop.
- **Loop label.** The label you specified for the loop in the linear assembly input file. This field is not present for C/C++ code.
- **Loop carried dependency bound.** The distance of the largest loop carry path. A loop carry path occurs when one iteration of a loop writes a value that must be read in a future iteration. Instructions that are part of the loop carry bound are marked with the ^ symbol.

- **Initiation interval** (ii). The number of cycles between the initiation of successive iterations of the loop. The smaller the initiation interval, the fewer cycles it takes to execute a loop.
- **Resource bound.** The most used resource constrains the minimum initiation interval. If four instructions require a .D unit, they require at least two cycles to execute (4 instructions/2 parallel .D units).
- **Unpartitioned resource bound.** The best possible resource bound values before the instructions in the loop are partitioned to a particular side.
- **Partitioned resource bound (*).** The resource bound values after the instructions are partitioned.
- **Resource partition.** This table summarizes how the instructions have been partitioned. This information can be used to help assign functional units when writing linear assembly. Each table entry has values for the A-side and B-side registers. An asterisk is used to mark those entries that determine the resource bound value. The table entries represent the following terms:
 - **.L units** is the total number of instructions that require .L units.
 - **.S units** is the total number of instructions that require .S units.
 - **.D units** is the total number of instructions that require .D units.
 - **.M units** is the total number of instructions that require .M units.
 - **.X cross paths** is the total number of .X cross paths.
 - **.T address paths** is the total number of address paths.
 - **Long read path** is the total number of long read port paths.
 - **Long write path** is the total number of long write port paths.
 - **Logical ops (.LS)** is the total number of instructions that can use either the .L or .S unit.
 - **Addition ops (.LSD)** is the total number of instructions that can use either the .L or .S or .D unit
- **Bound(.L .S .LS).** The resource bound value as determined by the number of instructions that use the .L and .S units. It is calculated with the following formula:

$$\text{Bound}(.L .S .LS) = \text{ceil}((.L + .S + .LS) / 2)$$
- **Bound(.L .S .D .LS .LSD).** The resource bound value as determined by the number of instructions that use the .D, .L, and .S units. It is calculated with the following formula:

$$\text{Bound}(.L .S .D .LS .LSD) = \text{ceil}((.L + .S + .D + .LS + .LSD) / 3)$$
- **Minimum required memory pad.** The number of bytes that are read if speculative execution is enabled. See [Section 3.2.3](#) for more information.

3.2.2.1 Loop Disqualified for Software Pipelining Messages

The following messages appear if the loop is completely disqualified for software pipelining:

- **Bad loop structure.** This error is very rare and can stem from the following:
 - An asm statement inserted in the C code inner loop
 - Parallel instructions being used as input to the Linear Assembly Optimizer
 - Complex control flow such as GOTO statements and breaks
- **Loop contains a call.** Sometimes the compiler may not be able to inline a function call that is in a loop. Because the compiler could not inline the function call, the loop could not be software pipelined.
- **Too many instructions.** There are too many instructions in the loop to software pipeline.
- **Software pipelining disabled.** Software pipelining has been disabled by a command-line option, such as when using the `--disable_software_pipelining` option, not using the `--opt_level=2` (or `-O2`) or `--opt_level=3` (or `-O3`) option, or using the `--opt_for_space=2` or `--opt_for_space=3` option.
- **Uninitialized trip counter.** The trip counter may not have been set to an initial value.
- **Suppressed to prevent code expansion.** Software pipelining may be suppressed because of the `--opt_for_space=1` option. When the `--opt_for_space=1` option is used, software pipelining is disabled in less promising cases to reduce code size. To enable pipelining, use `--opt_for_space=0` or omit the `--opt_for_space` option altogether.

- **Loop carried dependency bound too large.** If the loop has complex loop control, try `--speculate_loads` according to the recommendations in [Section 3.2.3.2](#).
- **Cannot identify trip counter.** The loop trip counter could not be identified or was used incorrectly in the loop body.

3.2.2.2 Pipeline Failure Messages

The following messages can appear when the compiler or assembly optimizer is processing a software pipeline and it fails:

- **Address increment is too large.** An address register's offset must be adjusted because the offset is out of range of the C6000's offset addressing mode. You must minimize address register offsets.
- **Cannot allocate machine registers.** A software pipeline schedule was found, but it cannot allocate machine registers for the schedule. Simplification of the loop may help.

The register usage for the schedule found at the given `ii` is displayed. This information can be used when writing linear assembly to balance register pressure on both sides of the register file. For example:

```
ii = 11 Cannot allocate machine registers
Regs Live Always : 3/0 (A/B-side)
Max Regs Live : 20/14
Max Condo Regs Live : 2/1
```

- **Regs Live Always.** The number of values that must be assigned a register for the duration of the whole loop body. This means that these values must always be allocated registers for any given schedule found for the loop.
- **Max Regs Live.** Maximum number of values live at any given cycle in the loop that must be allocated to a register. This indicates the maximum number of registers required by the schedule found.
- **Max Cond Regs Live.** Maximum number of registers live at any given cycle in the loop kernel that must be allocated to a condition register.
- **Cycle count too high. Never profitable.** With the schedule that the compiler found for the loop, it is more efficient to use a non-software-pipelined version.
- **Did not find schedule.** The compiler was unable to find a schedule for the software pipeline at the given `ii` (iteration interval). You should simplify the loop and/or eliminate loop carried dependencies.
- **Iterations in parallel > minimum or maximum trip count.** A software pipeline schedule was found, but the schedule has more iterations in parallel than the minimum or maximum loop trip count. You must enable redundant loops or communicate the trip information.
- **Speculative threshold exceeded.** It would be necessary to speculatively load beyond the threshold currently specified by the `--speculate_loads` option. You must increase the `--speculate_loads` threshold as recommended in the software-pipeline feedback located in the assembly file.
- **Register is live too long.** A register must have a value that exists (is live) for more than `ii` cycles. You may insert MV instructions to split register lifetimes that are too long.

If the assembly optimizer is being used, the `.sa` file line numbers of the instructions that define and use the registers that are live too long are listed after this failure message. For example:

```
ii = 9 Register is live too long
|10| -> |17|
```

This means that the instruction that defines the register value is on line 10 and the instruction that uses the register value is on line 17 in the `.sa` file.

- **Too many predicates live on one side.** The C6000 has predicate, or conditional, registers available for use with conditional instructions. There are five predicate registers on the C6200 and C6700, and six predicate registers on the C6400, C6400+, and C6700+. There are two or three on the A side and three on the B side. Sometimes the particular partition and schedule combination requires more than these available registers.
- **Schedule found with N iterations in parallel.** (This is not a failure message.) A software pipeline schedule was found with N iterations executing in parallel.

- **Too many reads of one register.** The same register can be read a maximum of four times per cycle with the C6200 or C6700 core. The C6400 core can read the same register any number of times per cycle.
- **Trip variable used in loop - Cannot adjust trip count.** The loop trip counter has a use in the loop other than as a loop trip counter.

3.2.2.3 Register Usage Table Generated by the `--debug_software_pipeline` Option

The `--debug_software_pipeline` option places additional software pipeline feedback in the generated assembly file. This information includes a single scheduled iteration view of the software pipelined loop.

If software pipelining succeeds for a given loop, and the `--debug_software_pipeline` option was used during the compilation process, a register usage table is added to the software pipelining information comment block in the generated assembly code.

The numbers on each row represent the cycle number within the loop kernel.

Each column represents one register on the TMS320C6000. The registers are labeled in the first three rows of the register usage table and should be read columnwise.

An * in a table entry indicates that the register indicated by the column header is live on the kernel execute packet indicated by the cycle number labeling each row.

An example of the register usage table follows:

```

;*      Searching for software pipeline schedule at
;*      ii = 15 Schedule found with 2 iterations in parallel
;*
;*      Register Usage Table:
;*
;*      +-----+
;*      |AAAAAAAAAAAAAAAA |BBBBBBBBBBBBBBBB |
;*      |0000000000111111 |0000000000111111 |
;*      |0123456789012345 |0123456789012345 |
;*      +-----+
;*      0:  ***  ****  | ***  *****
;*      1:  ****  ****  | ***  *****
;*      2:  ****  ****  | ***  *****
;*      3:  **   ***** | ***  *****
;*      4:  **   ***** | ***  *****
;*      5:  **   ***** | ***  *****
;*      6:  **   ***** | *****
;*      7:  ***  ***** | **  *****
;*      8:  ****  ***** | *****
;*      9:  *****      | **  *****
;*      10: *****      | **  *****
;*      11: *****      | **  *****
;*      12: *****      | *****
;*      13: ****  ***** | **  ***** *
;*      14: ***   ***** | ***  ***** *
;*
;*      +-----+

```

This example shows that on cycle 0 (first execute packet) of the loop kernel, registers A0, A1, A2, A6, A7, A8, A9, B0, B1, B2, B4, B5, B6, B7, B8, and B9 are all live during this cycle.

3.2.3 Collapsing Prologs and Epilogs for Improved Performance and Code Size

When a loop is software pipelined, a prolog and epilog are generally required. The prolog is used to pipe up the loop and epilog is used to pipe down the loop.

In general, a loop must execute a minimum number of iterations before the software-pipelined version can be safely executed. If the minimum known trip count is too small, either a redundant loop is added or software pipelining is disabled. Collapsing the prolog and epilog of a loop can reduce the minimum trip count necessary to safely execute the pipelined loop.

Collapsing can also substantially reduce code size. Some of this code size growth is due to the redundant loop. The remainder is due to the prolog and epilog.

The prolog and epilog of a software-pipelined loop consists of up to $p-1$ stages of length ii , where p is the number of iterations that are executed in parallel during the steady state and ii is the cycle time for the pipelined loop body. During prolog and epilog collapsing the compiler tries to collapse as many stages as possible. However, over-collapsing can have a negative performance impact. Thus, by default, the compiler attempts to collapse as many stages as possible without sacrificing performance. When the `--opt_for_space=0` or `--opt_for_space=1` options are invoked, the compiler increasingly favors code size over performance.

3.2.3.1 Speculative Execution

When prologs and epilogs are collapsed, instructions might be speculatively executed, thereby causing loads to addresses beyond either end of the range explicitly read within the loop. By default, the compiler cannot speculate loads because this could cause an illegal memory location to be read. Sometimes, the compiler can predicate these loads to prevent over execution. However, this can increase register pressure and might decrease the total amount collapsing which can be performed.

When the `--speculate_loads=n` option is used, the speculative threshold is increased from the default of 0 to n . When the threshold is n , the compiler can allow a load to be speculatively executed as the memory location it reads will be no more than n bytes before or after some location explicitly read within the loop. If the n is omitted, the compiler assumes the speculative threshold is unlimited. To specify this in Code Composer Studio, select the Speculate Threshold check box and leave the text box blank in the Build Options dialog box on the Compiler tab, Advanced category.

Collapsing can usually reduce the minimum safe trip count. If the minimum known trip count is less than the minimum safe trip count, a redundant loop is required. Otherwise, pipelining must be suppressed. Both these values can be found in the comment block preceding a software pipelined loop.

```

;*      Known Minimum Trip Count      : 1
.....
;*      Minimum safe trip count       : 7
    
```

If the minimum safe trip count is greater than the minimum known trip count, use of `--speculate_loads` is highly recommended, not only for code size, but for performance.

When using `--speculate_loads`, you must ensure that potentially speculated loads will not cause illegal reads. This can be done by padding the data sections and/or stack, as needed, by the required memory pad in both directions. The required memory pad for a given software-pipelined loop is also provided in the comment block for that loop.

```

;*      Minimum required memory pad   : 8 bytes
    
```

3.2.3.2 Selecting the Best Threshold Value

When a loop is software pipelined, the comment block preceding the loop provides the following information:

- Required memory pad for this loop
- The minimum value of n needed to achieve this software pipeline schedule and level of collapsing
- Suggestion for a larger value of n to use which might allow additional collapsing

This information shows up in the comment block as follows:

```

;*      Minimum required memory pad : 5 bytes
;*      Minimum threshold value     : --speculate_loads=7
;*
;*      For further improvement on this loop, try option --speculate_loads=14
    
```

For safety, the example loop requires that array data referenced within this loop be preceded and followed by a pad of at least 5 bytes. This pad can consist of other program data. The pad will not be modified. In many cases, the threshold value (namely, the minimum value of the argument to `--speculate_loads` that is needed to achieve a particular schedule and level of collapsing) is the same as the pad. However, when it is not, the comment block will also include the minimum threshold value. In the case of this loop, the threshold value must be at least 7 to achieve this level of collapsing.

Beginning with v7.4.0 of the C6000 Code Generation Tools, the compiler and linker can provide automatic load speculation via the *auto* argument to the `--speculate_loads` option (i.e. `--speculate_loads=auto` or `-mh=auto`). Use of the *auto* argument makes it easier to use and benefit from speculative load optimizations. This option can generate speculative loads of up to 256 bytes beyond memory that the compiler can prove to be allocated.

In addition, the compiler communicates information to the linker to help automatically ensure the required pre- and post-padding:

- If the symbol of the speculatively loaded buffer is known during compile time, the linker will ensure the object pointed to by the symbol has the required padding to let the speculative load access legal memory.
- If the symbol information is not known during compile time, the linker will ensure that the placement of data sections will allow legal accessing beyond the boundaries of the data sections. The linker does this by simply padding the start and end of the memory range(s) where the data sections are placed.

However, you can also specify the speculative loads threshold explicitly via the `--speculate_loads=n` option, where *n* is at least the minimum required pad (as explained earlier), but you also need to consider whether a larger threshold value would facilitate additional collapsing. This information is also provided, if applicable. For example, in the above comment block, a threshold value of 14 might facilitate further collapsing. If you choose the *auto* argument to `--speculate_loads`, the compiler will consider the larger threshold value automatically.

3.3 Redundant Loops

Every loop iterates some number of times before the loop terminates. The number of iterations is called the *trip count*. The variable used to count each iteration is the *trip counter*. When the trip counter reaches a limit equal to the trip count, the loop terminates. The C6000 tools use the trip count to determine whether or not a loop can be pipelined. The structure of a software pipelined loop requires the execution of a minimum number of loop iterations (a minimum trip count) in order to fill or prime the pipeline.

The minimum trip count for a software pipelined loop is determined by the number of iterations executing in parallel. In [Figure 3-1](#), the minimum trip count is five. In the following example A, B, and C are instructions in a software pipeline, so the minimum trip count for this single-cycle software pipelined loop is three.

```

A
B   A
C   B   A   ←Three iterations in parallel = minimum trip count
      C   B
          C

```

When the C6000 tools cannot determine the trip count for a loop, then by default two loops and control logic are generated. The first loop is not pipelined, and it executes if the run-time trip count is less than the loop's minimum trip count. The second loop is the software pipelined loop, and it executes when the run-time trip count is greater than or equal to the minimum trip count. At any given time, one of the loops is a *redundant loop*. For example:

```

foo(N) /* N is the trip count */
{
    for (I=0; I < N; I++) /* I is the trip counter */
}

```

After finding a software pipeline for the loop, the compiler transforms `foo()` as below, assuming the minimum trip count for the loop is 3. Two versions of the loop would be generated and the following comparison would be used to determine which version should be executed:

```

foo(N)
{
    if (N < 3)
    {
        for (I=0; I < N; I++) /* Unpipelined version */
    }
}

```



```

    else
    }
    for (I=0; I < N; I++) /* Pipelined version */
    }
}
foo(50); /* Execute software pipelined loop */
foo(2); /* Execute loop (unpipelined)*/

```

You may be able to help the compiler avoid producing redundant loops with the use of `--program_level_compile --opt_level=3` (see [Section 3.7](#)) or the use of the `MUST_ITERATE` pragma (see [Section 6.9.20](#)).

Turning Off Redundant Loops

NOTE: Specifying any `--opt_for_space` option turns off redundant loops.

3.4 Utilizing the Loop Buffer Using SPLOOP on C6400+, C6740, and C6600

The C6400+, C6740, and C6600 ISAs have a loop buffer which improves performance and reduces code size for software pipelined loops. The loop buffer provides the following benefits:

- Code size. A single iteration of the loop is stored in program memory.
- Interrupt latency. Loops executing out of the loop buffer are interruptible.
- Improves performance for loops with unknown trip counts and eliminates redundant loops.
- Reduces or eliminates the need for speculated loads.
- Reduces power usage.

You can tell that the compiler is using the loop buffer when you find `SPLOOP(D/W)` at the beginning of a software pipelined loop followed by an `SPKERNEL` at the end. Refer to the *TMS320C6400/C6400+ CPU and Instruction Set Reference Guide* for information on `SPLOOP`.

When the `--opt_for_space` option is not used, the compiler will not use the loop buffer if it can find a faster software pipelined loop without it. When using the `--opt_for_space` option, the compiler will use the loop buffer when it can.

The compiler does not generate code for the loop buffer (`SPLOOP(D/W)`) when any of the following occur:

- `ii` (initiation interval) > 14 cycles
- Dynamic length (of a single iteration) > 48 cycles
- The optimizer completely unrolls the loop
- Code contains elements that disqualify normal software pipelining (call in loop, complex control code in loop, etc.). See the *TMS320C6000 Programmer's Guide* for more information.

3.5 Reducing Code Size (`--opt_for_space` (or `-ms`) Option)

When using the `--opt_level=n` option (or `-On`), you are telling the compiler to optimize your code. The higher the value of `n`, the more effort the compiler invests in optimizing your code. However, you might still need to tell the compiler what your optimization priorities are. By default, when `--opt_level=2` or `--opt_level=3` is specified, the compiler optimizes primarily for performance. (Under lower optimization levels, the priorities are compilation time and debugging ease.) You can adjust the priorities between performance and code size by using the code size flag `--opt_for_space=n`. The `--opt_for_space=0`, `--opt_for_space=1`, `--opt_for_space=2` and `--opt_for_space=3` options increasingly favor code size over performance.

When you specify `--silicon_version=6400+` in conjunction with the `--opt_for_space` option, the code will be tailored for compression. That is, more instructions are tailored so they will more likely be converted from 32-bit to 16-bit instructions when assembled.

It is recommended that a code size flag not be used with the most performance-critical code. Using `--opt_for_space=0` or `--opt_for_space=1` is recommended for all but the most performance-critical code. Using `--opt_for_space=2` or `--opt_for_space=3` is recommended for seldom-executed code. Either `--opt_for_space=2` or `--opt_for_space=3` should be used if you need minimum code size. It is generally recommended that the code size flags be combined with `--opt_level=2` or `--opt_level=3`.

Disabling Code-Size Optimizations or Reducing the Optimization Level

NOTE: If you reduce optimization and/or do not use code size flags, you are disabling code-size optimizations and sacrificing performance.

The `--opt_for_space` Option is Equivalent to `--opt_for_space=0`

NOTE: If you use `--opt_for_space` with no code size level number specified, the option level defaults to `--opt_for_space=0`.

3.6 Performing File-Level Optimization (`--opt_level=3` option)

The `--opt_level=3` option (aliased as the `-O3` option) instructs the compiler to perform file-level optimization. You can use the `--opt_level=3` option alone to perform general file-level optimization, or you can combine it with other options to perform more specific optimizations. The options listed in [Table 3-1](#) work with `--opt_level=3` to perform the indicated optimization:

Table 3-1. Options That You Can Use With `--opt_level=3`

If You ...	Use this Option	See
Have files that redeclare standard library functions	<code>--std_lib_func_defined</code> <code>--std_lib_func_redefined</code>	Section 3.6.1
Want to create an optimization information file	<code>--gen_opt_level=n</code>	Section 3.6.2
Want to compile multiple source files	<code>--program_level_compile</code>	Section 3.7

Do Not Lower the Optimization Level to Control Code Size

NOTE: When trying to reduce code size, do not lower the level of optimization, as you might see an increase in code size. Instead, use the `--opt_for_space` option to control the code.

3.6.1 Controlling File-Level Optimization (`--std_lib_func_def` Options)

When you invoke the compiler with the `--opt_level=3` option, some of the optimizations use known properties of the standard library functions. If your file redeclares any of these standard library functions, these optimizations become ineffective. Use [Table 3-2](#) to select the appropriate file-level optimization option.

Table 3-2. Selecting a File-Level Optimization Option

If Your Source File...	Use this Option
Declares a function with the same name as a standard library function	<code>--std_lib_func_redefined</code>
Contains but does not alter functions declared in the standard library	<code>--std_lib_func_defined</code>
Does not alter standard library functions, but you used the <code>--std_lib_func_redefined</code> or <code>--std_lib_func_defined</code> option in a command file or an environment variable. The <code>--std_lib_func_not_defined</code> option restores the default behavior of the optimizer.	<code>--std_lib_func_not_defined</code>

3.6.2 Creating an Optimization Information File (`--gen_opt_info` Option)

When you invoke the compiler with the `--opt_level=3` option, you can use the `--gen_opt_info` option to create an optimization information file that you can read. The number following the option denotes the level (0, 1, or 2). The resulting file has an `.nfo` extension. Use [Table 3-3](#) to select the appropriate level to append to the option.

Table 3-3. Selecting a Level for the `--gen_opt_info` Option

If you...	Use this option
Do not want to produce an information file, but you used the <code>--gen_opt_level=1</code> or <code>--gen_opt_level=2</code> option in a command file or an environment variable. The <code>--gen_opt_level=0</code> option restores the default behavior of the optimizer.	<code>--gen_opt_info=0</code>
Want to produce an optimization information file	<code>--gen_opt_info=1</code>
Want to produce a verbose optimization information file	<code>--gen_opt_info=2</code>

3.7 Performing Program-Level Optimization (--program_level_compile and --opt_level=3 options)

You can specify program-level optimization by using the `--program_level_compile` option with the `--opt_level=3` option (aliased as `-O3`). With program-level optimization, all of your source files are compiled into one intermediate file called a *module*. The module moves to the optimization and code generation passes of the compiler. Because the compiler can see the entire program, it performs several optimizations that are rarely applied during file-level optimization:

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

To see which program-level optimizations the compiler is applying, use the `--gen_opt_level=2` option to generate an information file. See [Section 3.6.2](#) for more information.

In Code Composer Studio, when the `--program_level_compile` option is used, C and C++ files that have the same options are compiled together. However, if any file has a file-specific option that is not selected as a project-wide option, that file is compiled separately. For example, if every C and C++ file in your project has a different set of file-specific options, each is compiled separately, even though program-level optimization has been specified. To compile all C and C++ files together, make sure the files do not have file-specific options. Be aware that compiling C and C++ files together may not be safe if previously you used a file-specific option.

Compiling Files With the `--program_level_compile` and `--keep_asm` Options

NOTE: If you compile all files with the `--program_level_compile` and `--keep_asm` options, the compiler produces only one `.asm` file, not one for each corresponding source file.

3.7.1 Controlling Program-Level Optimization (--call_assumptions Option)

You can control program-level optimization, which you invoke with `--program_level_compile --opt_level=3`, by using the `--call_assumptions` option. Specifically, the `--call_assumptions` option indicates if functions in other modules can call a module's external functions or modify a module's external variables. The number following `--call_assumptions` indicates the level you set for the module that you are allowing to be called or modified. The `--opt_level=3` option combines this information with its own file-level analysis to decide whether to treat this module's external function and variable declarations as if they had been declared static. Use [Table 3-4](#) to select the appropriate level to append to the `--call_assumptions` option.

Table 3-4. Selecting a Level for the `--call_assumptions` Option

If Your Module ...	Use this Option
Has functions that are called from other modules and global variables that are modified in other modules	<code>--call_assumptions=0</code>
Does not have functions that are called by other modules but has global variables that are modified in other modules	<code>--call_assumptions=1</code>
Does not have functions that are called by other modules or global variables that are modified in other modules	<code>--call_assumptions=2</code>
Has functions that are called from other modules but does not have global variables that are modified in other modules	<code>--call_assumptions=3</code>

In certain circumstances, the compiler reverts to a different `--call_assumptions` level from the one you specified, or it might disable program-level optimization altogether. [Table 3-5](#) lists the combinations of `--call_assumptions` levels and conditions that cause the compiler to revert to other `--call_assumptions` levels.

Table 3-5. Special Considerations When Using the --call_assumptions Option

If Your Option is...	Under these Conditions...	Then the --call_assumptions Level...
Not specified	The --opt_level=3 optimization level was specified	Defaults to --call_assumptions=2
Not specified	The compiler sees calls to outside functions under the --opt_level=3 optimization level	Reverts to --call_assumptions=0
Not specified	Main is not defined	Reverts to --call_assumptions=0
--call_assumptions=1 or --call_assumptions=2	No function has main defined as an entry point and functions are not identified by the FUNC_EXT_CALLED pragma	Reverts to --call_assumptions=0
--call_assumptions=1 or --call_assumptions=2	No interrupt function is defined	Reverts to --call_assumptions=0
--call_assumptions=1 or --call_assumptions=2	Functions are identified by the FUNC_EXT_CALLED pragma	Remains --call_assumptions=1 or --call_assumptions=2
--call_assumptions=3	Any condition	Remains --call_assumptions=3

In some situations when you use --program_level_compile and --opt_level=3, you *must* use a --call_assumptions option or the FUNC_EXT_CALLED pragma. See [Section 3.7.2](#) for information about these situations.

3.7.2 Optimization Considerations When Mixing C/C++ and Assembly

If you have any assembly functions in your program, you need to exercise caution when using the --program_level_compile option. The compiler recognizes only the C/C++ source code and not any assembly code that might be present. Because the compiler does not recognize the assembly code calls and variable modifications to C/C++ functions, the --program_level_compile option optimizes out those C/C++ functions. To keep these functions, place the FUNC_EXT_CALLED pragma (see [Section 6.9.10](#)) before any declaration or reference to a function that you want to keep.

Another approach you can take when you use assembly functions in your program is to use the --call_assumptions=*n* option with the --program_level_compile and --opt_level=3 options (see [Section 3.7.1](#)).

In general, you achieve the best results through judicious use of the FUNC_EXT_CALLED pragma in combination with --program_level_compile --opt_level=3 and --call_assumptions=1 or --call_assumptions=2.

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

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

Solution — Compile with --program_level_compile --opt_level=3 --call_assumptions=2 to tell the compiler that outside functions do not call C/C++ functions or modify C/C++ variables. See [Section 3.7.1](#) for information about the --call_assumptions=2 option.

If you compile with the --program_level_compile --opt_level=3 options only, the compiler reverts from the default optimization level (--call_assumptions=2) to --call_assumptions=0. The compiler uses --call_assumptions=0, because it presumes that the calls to the assembly language functions that have a definition in C/C++ may call other C/C++ functions or modify C/C++ variables.

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

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

- Compile with --program_level_compile --opt_level=3 --call_assumptions=1.
- Add the volatile keyword to those variables that may be modified by the assembly functions and compile with --program_level_compile --opt_level=3 --call_assumptions=2.

See [Section 3.7.1](#) for information about the --call_assumptions=*n* option.

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

Solution — Add the volatile keyword to the C/C++ variables that may be modified by the interrupts. Then, you can optimize your code in one of these ways:

- You achieve the best optimization by applying the `FUNC_EXT_CALLED` pragma to all of the entry-point functions called from the assembly language interrupts, and then compiling with `--program_level_compile --opt_level=3 --call_assumptions=2`. *Be sure that you use the pragma with all of the entry-point functions.* If you do not, the compiler might remove the entry-point functions that are not preceded by the `FUNC_EXT_CALLED` pragma.
- Compile with `--program_level_compile --opt_level=3 --call_assumptions=3`. Because you do not use the `FUNC_EXT_CALLED` pragma, you must use the `--call_assumptions=3` option, which is less aggressive than the `--call_assumptions=2` option, and your optimization may not be as effective.

Keep in mind that if you use `--program_level_compile --opt_level=3` without additional options, the compiler removes the C functions that the assembly functions call. Use the `FUNC_EXT_CALLED` pragma to keep these functions.

3.8 Using Feedback Directed Optimization

Feedback directed optimization provides a method for finding frequently executed paths in an application using compiler-based instrumentation. This information is fed back to the compiler and is used to perform optimizations. This information is also used to provide you with information about application behavior.

3.8.1 Feedback Directed Optimization

Feedback directed optimization uses run-time feedback to identify and optimize frequently executed program paths. Feedback directed optimization is a two-phase process.

3.8.1.1 Phase 1: Collect Program Profile Information

In this phase the compiler is invoked with the option `--gen_profile_info`, which instructs the compiler to add instrumentation code to collect profile information. The compiler inserts a minimal amount of instrumentation code to determine control flow frequencies. Memory is allocated to store counter information.

The instrumented application program is executed on the target using representative input data sets. The input data sets should correlate closely with the way the program is expected to be used in the end product environment. When the program completes, a run-time-support function writes the collected information into a profile data file called a PDAT file. Multiple executions of the program using different input data sets can be performed and in such cases, the run-time-support function appends the collected information into the PDAT file. The resulting PDAT file is post-processed using a tool called the Profile Data Decoder or `pdd6x`. The `pdd6x` tool consolidates multiple data sets and formats the data into a feedback file (PRF file, see [Section 3.8.2](#)) for consumption by phase 2 of feedback directed optimization.

3.8.1.2 Phase 2: Use Application Profile Information for Optimization

In this phase, the compiler is invoked with the `--use_profile_info=file.prf` option, which reads the specified PRF file generated in phase 1. In phase 2, optimization decisions are made using the data generated during phase 1. The profile feedback file is used to guide program optimization. The compiler optimizes frequently executed program paths more aggressively.

The compiler uses data in the profile feedback file to guide certain optimizations of frequently executed program paths.

3.8.1.3 Generating and Using Profile Information

There are two options that control feedback directed optimization:

- gen_profile_info** tells the compiler to add instrumentation code to collect profile information. When the program executes the run-time-support `exit()` function, the profile data is written to a PDAT file. If the environment variable `TI_PROFDATA` on the host is set, the data is written into the specified file name. Otherwise, it uses the default filename: `pprofout.pdat`. The full pathname of the PDAT file (including the directory name) can be specified using the `TI_PROFDATA` host environment variable.
- By default, the RTS profile data output routine uses the C I/O mechanism to write data to the PDAT file. You can install a device handler for the PPHNDL device that enables you to re-direct the profile data to a custom device driver routine.
- Feedback directed optimization requires you to turn on at least skeletal debug information when using the `--gen_profile_info` option. This enables the compiler to output debug information that allows `pdd6x` to correlate compiled functions and their associated profile data.
- use_profile_info** specifies the profile information file(s) to use for performing phase 2 of feedback directed optimization. More than one profile information file can be specified on the command line; the compiler uses all input data from multiple information files. The syntax for the option is:
- use_profile_info==file1[, file2, ..., filen]**
- If no filename is specified, the compiler looks for a file named `pprofout.prf` in the directory where the compiler is invoked.

3.8.1.4 Example Use of Feedback Directed Optimization

These steps illustrate the creation and use of feedback directed optimization.

1. Generate profile information. (Skeletal debug is on by default.)

```
cl6x -mv6400+ --opt_level=2 --gen_profile_info foo.c --run_linker --output_file=foo.out
--library=lnk.cmd --library=rts64plus.lib
```

2. Execute the application.

The execution of the application creates a PDAT file named `pprofout.pdat` in the current (host) directory. The application can be run on a simulator or on target hardware connected to a host machine.

3. Process the profile data.

After running the application with multiple data-sets, run `pdd6x` on the PDAT files to create a profile information (PRF) file to be used with `--use_profile_info`.

```
pdd6x -e foo.out -o pprofout.prf pprofout.pdat
```

4. Re-compile using the profile feedback file. Skeletal debug is not required.

```
cl6x -mv6400+ --opt_level=2 --use_profile_info=pprofout.prf foo.c --run_linker
--output_file=foo.out --library=lnk.cmd --library=rts64plus.lib
```

3.8.1.5 The .ppdata Section

The profile information collected in phase 1 is stored in the `.ppdata` section, which must be allocated into target memory. The `.ppdata` section contains profiler counters for all functions compiled with `--gen_profile_info`. The default `lnk.cmd` file in code generation tools version 6.1 and later has directives to place the `.ppdata` section in data memory. If the link command file has no section directive for allocating `.ppdata` section, the link step places the `.ppdata` section in a writable memory range.

The `.ppdata` section must be allocated memory in multiples of 32 bytes. Please refer to the linker command file in the distribution for example usage.

3.8.1.6 Feedback Directed Optimization and Code Size Tune

Feedback directed optimization is different from the Code Size Tune feature in Code Composer Studio (CCS). The code size tune feature uses CCS profiling to select specific compilation options for each function in order to minimize code size while still maintaining a specific performance point. Code size tune is coarse-grained, since it is selecting an option set for the whole function. Feedback directed optimization selects different optimization goals along specific regions within a function.

3.8.1.7 Instrumented Program Execution Overhead

During profile collection, the execution time of the application may increase. The amount of increase depends on the size of the application and the number of files in the application compiled for profiling.

The profiling counters increase the code and data size of the application. Consider using the `--opt_for_space (-ms)` code size options when using profiling to mitigate the code size increase. This has no effect on the accuracy of the profile data being collected. Since profiling only counts execution frequency and not cycle counts, code size optimization flags do not affect profiler measurements.

3.8.1.8 Invalid Profile Data

When recompiling with `--use_profile_info`, the profile information is invalid in the following cases:

- The source file name changed between the generation of profile information (`gen-profile`) and the use of the profile information (`use-profile`).
- The source code was modified since `gen-profile`. In this case, profile information is invalid for the modified functions.
- Certain compiler options used with `gen-profile` are different from those with used with `use-profile`. In particular, options that affect parser behavior could invalidate profile data during `use-profile`. In general, using different optimization options during `use-profile` should not affect the validity of profile data.

3.8.2 Profile Data Decoder

The code generation tools include a new tool called the profile data decoder or `pdd6x`, which is used for post processing profile data (PDAT) files. The `pdd6x` tool generates a profile feedback (PRF) file. See [Section 3.8.1](#) for a discussion on where `pdd6x` fits in the profiling flow. The `pdd6x` tool is invoked with this syntax:

```
pdd6x -e exec.out -o application.prf filename .pdatt
```

-a	Computes the average of the data values in the data sets instead of accumulating data values
-e <i>exec.out</i>	Specifies <i>exec.out</i> is the name of the application executable.
-o <i>application.prf</i>	Specifies <i>application.prf</i> is the formatted profile feedback file that is used as the argument to <code>--use_profile_info</code> during recompilation. If no output file is specified, the default output filename is <code>pprofout.prf</code> .
<i>filename .pdatt</i>	Is the name of the profile data file generated by the run-time-support function. This is the default name and it can be overridden by using the host environment variable <code>TI_PROFDATA</code> .

The run-time-support function and pdd6x append to their respective output files and do not overwrite them. This enables collection of data sets from multiple runs of the application.

Profile Data Decoder Requirements

NOTE: Your application must be compiled with at least skeletal (dwarf) debug support to enable feedback directed optimization. When compiling for feedback directed optimization, the pdd6x tool relies on basic debug information about each function in generating the formatted .prf file.

The pprofout.pdat file generated by the run-time support is a raw data file of a fixed format understood only by pdd6x. You should not modify this file in any way.

3.8.3 Feedback Directed Optimization API

There are two user interfaces to the profiler mechanism. You can start and stop profiling in your application by using the following run-time-support calls.

- TI_start_pprof_collection()

This interface informs the run-time support that you wish to start profiling collection from this point on and causes the run-time support to clear all profiling counters in the application (that is, discard old counter values).

- TI_stop_pprof_collection()

This interface directs the run-time support to stop profiling collection and output profiling data into the output file (into the default file or one specified by the TI_PROFDATA host environment variable). The run-time support also disables any further output of profile data into the output file during exit(), unless you call TI_start_pprof_collection() again.

3.8.4 Feedback Directed Optimization Summary

Options

--gen_profile_info	Adds instrumentation to the compiled code. Execution of the code results in profile data being emitted to a PDAT file.
--use_profile_info= <i>file.prf</i>	Uses profile information for optimization and/or generating code coverage information.
--analyze=codecov	Generates a code coverage information file and continues with profile-based compilation. Must be used with --use_profile_info.
--analyze_only	Generates only a code coverage information file. Must be used with --use_profile_info. You must specify both --analyze=codecov and --analyze_only to do code coverage analysis of the instrumented application.

Host Environment Variables

TI_PROFDATA	Writes profile data into the specified file
TI_COVDIR	Creates code coverage files in the specified directory
TI_COVDATA	Writes code coverage data into the specified file

API

TI_start_pprof_collection()	Clears the profile counters to file
TI_stop_pprof_collection()	Writes out all profile counters to file
PPHDNL	Device driver handle for low-level C I/O based driver for writing out profile data from a target program.

Files Created

*.pdat	Profile data file, which is created by executing an instrumented program and used as input to the profile data decoder
*.prf	Profiling feedback file, which is created by the profile data decoder and used as input to the re-compilation step

3.9 Using Profile Information to Get Better Program Cache Layout and Analyze Code Coverage

There are two different types of analysis information you can get from the path profiler: code coverage information and call graph information.

The program cache layout tool helps you to develop better program instruction cache efficiency into your applications. Program cache layout is the process of controlling the relative placement of code sections into memory to minimize the occurrence of conflict misses in the program instruction cache.

3.9.1 Background and Motivation

Effective utilization of the program instruction cache is an important part of getting the best performance from a C6000. The dedicated program instruction cache (L1P) provides fast instruction fetches, but a *cache miss* can be very costly. Some applications (e.g. h264) can spend 30%+ of the processor's time stalling due to L1P cache misses. A cache miss occurs when a fetch fails to read an instruction from L1P and the process is required to access the instruction from the next level of memory. A request to L2 or external memory has a much higher latency than an access from L1P.

Careful placement of code sections can greatly reduce the number of cache misses. The C6000 L1P is especially sensitive to code placement because it is *direct-mapped*.

Many L1P cache misses are *conflict misses*. Conflict misses occur when the cache has recently evicted a block of code that is now needed again. In a program instruction cache this often occurs when two frequently executed blocks of code (usually from different functions) interleave their execution and are mapped to the same cache line.

For example, suppose there is a call to function B from inside a loop in function A. Suppose also that the code for function A's loop is mapped to the same cache line as a block of code from function B that is executed every time that B is called. Each time B is called from within this loop, the loop code in function A is evicted from the cache by the code in B that is mapped to the same cache line. Even worse, when B returns to A, the loop code in A evicts the code from function B that is mapped to the same cache line.

Every iteration through the loop will cause two program instruction cache conflict misses. If the loop is heavily traversed, then the number of processor cycles lost to program instruction cache stalls can become quite large.

Many program instruction cache conflict misses can be avoided with more intelligent placement of functions that are active at the same time. Program instruction cache efficiency can be significantly improved using code placement strategies that utilize dynamic profile information that is gathered during the run of an instrumented application.

The program cache layout tool (clt6x) takes dynamic profile information in the form of a weighted call graph and creates a preferred function order command file that can be used as input to the linker to guide the placement of function subsections.

You can use the program cache layout tool to help improve your program locality and reduce the number of L1P cache conflict misses that occur during the run of your application, thereby improving your application's performance.

3.9.2 Code Coverage

The information collected during feedback directed optimization can be used for generating code coverage reports. As with feedback directed optimization, the program must be compiled with the `--gen_profile_info` option.

Code coverage conveys the execution count of each line of source code in the file being compiled, using data collected during profiling.

3.9.2.1 Phase1: Collect Program Profile Information

In this phase the compiler is invoked with the option `--gen_profile_info`, which instructs the compiler to add instrumentation code to collect profile information. The compiler inserts a minimal amount of instrumentation code to determine control flow frequencies. Memory is allocated to store counter information.

The instrumented application program is executed on the target using representative input data sets. The input data sets should correlate closely with the way the program is expected to be used in the end product environment. When the program completes, a run-time-support function writes the collected information into a profile data file called a PDAT file. Multiple executions of the program using different input data sets can be performed and in such cases, the run-time-support function appends the collected information into the PDAT file. The resulting PDAT file is post-processed using a tool called the Profile Data Decoder or `pdd6x`. The `pdd6x` tool consolidates multiple data sets and formats the data into a feedback file (PRF file, see [Section 3.8.2](#)) for consumption by phase 2 of feedback directed optimization.

3.9.2.2 Phase 2: Generate Code Coverage Reports

In this phase, the compiler is invoked with the `--use_profile_info=file.prf` option, which indicates that the compiler should read the specified PRF file generated in phase 1. The application must also be compiled with either the `--codecov` or `--onlycodecov` option; the compiler generates a code-coverage info file. The `--codecov` option directs the compiler to continue compilation after generating code-coverage information, while the `--onlycodecov` option stops the compiler after generating code-coverage data. For example:

```
cl6x --opt_level=2 --use_profile_info=pprofout.prf --onlycodecov foo.c
```

You can specify two environment variables to control the destination of the code-coverage information file.

- The `TI_COVDIR` environment variable specifies the directory where the code-coverage file should be generated. The default is the directory where the compiler is invoked.
- The `TI_COVDATA` environment variable specifies the name of the code-coverage data file generated by the compiler. the default is `filename.csv` where `filename` is the base-name of the file being compiled. For example, if `foo.c` is being compiled, the default code-coverage data file name is `foo.csv`.

If the code-coverage data file already exists, the compiler appends the new dataset at the end of the file.

Code-coverage data is a comma-separated list of data items that can be conveniently handled by data-processing tools and scripting languages. The following is the format of code-coverage data:

```
"filename-with-full-path", "funcname", line#, column#, exec-frequency, "comments"
```

<code>"filename-with-full-path"</code>	Full pathname of the file corresponding to the entry
<code>"funcname"</code>	Name of the function
<code>line#</code>	Line number of the source line corresponding to frequency data
<code>column#</code>	Column number of the source line
<code>exec-frequency</code>	Execution frequency of the line
<code>"comments"</code>	Intermediate-level representation of the source-code generated by the parser

The full filename, function name, and comments appear within quotation marks ("). For example:

```
"/some_dir/zlib/c64p/deflate.c", "_deflateInit2_", 216, 5, 1, "( strm->zalloc )"
```

Other tools, such as a spreadsheet program, can be used to format and view the code coverage data.

3.9.3 What Performance Improvements Can You Expect to See?

If your application does not suffer from inefficient usage of the L1P cache, then the program cache layout capability will not have any effect on the performance of your application. Before applying the program cache layout tooling to your application, analyze the L1P cache performance in your application.

3.9.3.1 Evaluating L1P Cache Performance

Evaluating the L1P cache usage efficiency of your application will not only help you determine whether or not your application might benefit from using program cache layout, but it also gives you a rough estimate as to how much performance improvement you can reasonably expect from applying program cache layout.

There are several resources available to help you evaluate L1P cache usage in your application. One way of doing this is to use the Function Profiling capability in Code Composer Studio (CCS). This capability is available in the C6400+ Megamodule Cycle Accurate Simulator target configuration under CCS. You can find further information about using the CCS Function Profiling capabilities at <http://tiexpressdsp.com/index.php/Profiler>. You can find more information about how to use this capability in conjunction with the program cache layout tool at http://tiexpressdsp.com/index.php/Program_Cache_Layout.

The number of CPU stall cycles that occur due to L1P cache misses gives you a reasonable upper bound estimate of the number of CPU cycles that you may be able to recover with the use of the program cache layout tooling in your application. Please be aware that the performance impact due to program cache layout will tend to vary for the different data sets that are run through your application.

3.9.4 Program Cache Layout Related Features and Capabilities

Version 7.0 of the C6000 code generation tools introduce some features and capabilities that can be used in conjunction with the program cache layout tool, clt6x. The following is a summary:

3.9.4.1 Path Profiler

The C6000 tools include a path profiling utility, pprof6x, that is run from the compiler, cl6x. The pprof6x utility is invoked by the compiler when the `--gen_profile` or the `--use_profile` command is used from the compiler command line:

```
cl6x --gen_profile ... file.c
cl6x --use_profile ... file.c
```

For further information about profile-based optimization and a more detailed description of the profiling infrastructure within the C6000, see [Section 3.8](#).

3.9.4.2 Analysis Options

The path profiling utility, pprof6x, appends code coverage or weighted call graph analysis information to existing CSV (comma separated values) files that contain the same type of analysis information.

The utility checks to make sure that an existing CSV file contains analysis information that is consistent with the type of analysis information it is being asked to generate (whether it be code coverage or weighted call graph analysis). Attempts to mix code coverage and weighted call graph analysis information in the same output CSV file will be detected and pprof6x will emit a fatal error and abort.

--analyze=callgraph	Instructs the compiler to generate weighted call graph analysis information.
--analyze=codecov	Instructs the compiler to generate code coverage analysis information. This option replaces the previous <code>--codecov</code> option.
--analyze_only	Halts compilation after generation of analysis information is completed.

3.9.4.3 Environment Variables

To assist with the management of output CSV analysis files, pprof6x supports two new environment variables:

TI_WCGDATA	Allows you to specify a single output CSV file for all weighted call graph analysis information. New information is appended to the CSV file identified by this environment variable, if the file already exists.
TI_ANALYSIS_DIR	Specifies the directory in which the output analysis file will be generated. The same environment variable can be used for both code coverage information and weighted call graph information (all analysis files generated by pprof6x will be written to the directory specified by the TI_ANALYSIS_DIR environment variable).

TI_COVDIR Environment Variable

NOTE: The existing TI_COVDIR environment variable is still supported when generating code coverage analysis, but is overridden in the presence of a defined TI_ANALYSIS_DIR environment variable.

3.9.4.4 Program Cache Layout Tool, clt6x

The program cache layout tool creates a preferred function order command file from input weighted call graph (WCG) information. The syntax is:

```
clt6x CSV files with WCG info -o forder.cmd
```

3.9.4.5 Linker

The compiler prioritizes the placement of a function relative to others based on the order in which --preferred_order options are encountered during the linker invocation. The syntax is:

```
--preferred_order=function specification
```

3.9.4.6 Linker Command File Operator unordered()

The new linker command file keyword unordered relaxes placement constraints placed on an output section whose specification includes an explicit list of which input sections are contained in the output section. The syntax is:

```
unordered()
```

3.9.5 Program Instruction Cache Layout Development Flow

Once you have determined that your application is experiencing some inefficiencies in its usage of the program instruction cache, you may decide to include the program cache layout tooling in your development to attempt to recover some of the CPU cycles that are being lost to stalls due to program instruction cache conflict misses.

3.9.5.1 Gather Dynamic Profile Information

The program cache layout tool, clt6x, relies on the availability of dynamic profile information in the form of a weighted call graph in order to produce a preferred function order command file that can be used to guide function placement at link-time when your application is re-built.

There are several ways in which this dynamic profile information can be collected. For example, if you are running your application on hardware, you may have the capability to collect a PC discontinuity trace. The discontinuity trace can then be post-processed to construct weighted call graph input information for the clt6x.

The method for collecting dynamic profile information that is presented here relies on the path profiling capabilities in the C6000 code generation tools. Here is how it works:

1. Build an instrumented application using the `--gen_profile_info` option.

Using `--gen_profile_info` instructs the compiler to embed counters into the code along the execution paths of each function.

To compile only use:

```
cl6x options --gen_profile_info src_file(s)
```

The compile and link use:

```
cl6x options --gen_profile_info src_file(s) -run_linker --library lnk.cmd
```

2. Run an instrumented application to generate a .pdat file.

When the application runs, the counters embedded into the application by `--gen_profile_info` keep track of how many times a particular execution path through a function is traversed. The data collected in these counters is written out to a profile data file named `pprofout.pdat`.

The profile data file is automatically generated. For example, if you are using the C64+ simulator under CCS, you can load and run your instrumented program, and you will see that a new `pprofout.pdat` file is created in your working directory (where the instrumented application is loaded from).

3. Decode the profile data file.

Once you have a profile data file, the file is decoded by the profile data decoder tool, `pdd6x`, as follows:

```
pdd6x -e=instrumented app out file -o=pprofout.prf pprofout.pdat
```

Using `pdd6x` produces a .prf file which is then fed into the re-compile of the application that uses the profile information to generate weighted call graph input data.

4. Use decoded profile information to generate weighted call graph input.

The compiler now supports a new option, `--analyze`, which is used to tell the compiler to generate weighted call graph or code coverage analysis information. Its syntax are as follows:

--analyze=callgraph Instructs the compiler to generate weighted call graph information.
--analyze=codecov Instructs the compiler to generate code coverage information. This option replaces the previous `--codecov` option.

The compiler also supports a new `--analyze_only` option which instructs the compiler to halt compilation after the generation of analysis information has been completed. This option replaces the previous `--onlycodecov` option.

To make use of the dynamic profile information that you gathered, re-compile the source code for your application using the `--analyze=callgraph` option in combination with the `--use_profile_info` option:

```
cl6x options -mo --analyze=callgraph --use_profile_info=pprofout.prf src_file(s)
```

The use of `-mo` instructs the compiler to generate code for each function into its own subsection. This option provides the linker with the means to directly control the placement of the code for a given function.

The compiler generates a CSV file containing weighted call graph information for each source file that is specified on the command line. If such a CSV file already exists, then new call graph analysis information will be appended to the existing CSV file. These CSV files are then input to the cache layout tool (clt6x) to produce a preferred function order command file for your application.

For more details on the content of the CSV files (containing weighted call graph information) generated by the compiler, see [Section 3.9.6](#).

3.9.5.2 Generate Preferred Function Order from Dynamic Profile Information

At this point, the compiler has generated a CSV file for each C/C++ source file specified on the command line of the re-compile of the application. Each CSV file contains weighted call graph information about all of the call sites in each function defined in the C/C++ source file.

The program cache layout tool, clt6x, collects all of the weighted call graph information in these CSV files into a single, merged weighted call graph. The weighted call graph is processed to produce a preferred function order command file that is fed into the linker to guide the placement of the functions defined in your application source files. This is the syntax for clt6x:

```
clt6x *.csv -o forder.cmd
```

The output of clt6x is a text file containing a sequence of `--preferred_order=function specification` options. By default, the name of the output file is `forder.cmd`, but you can specify your own file name with the `-o` option. The order in which functions appear in this file is their preferred function order as determined by the clt6x.

In general, the proximity of one function to another in the preferred function order list is a reflection of how often the two functions call each other. If two functions are very close to each other in the list, then the linker interprets this as a suggestion that the two functions should be placed very near to one another. Functions that are placed close together are less likely to create a cache conflict miss at run time when both functions are active at the same time. The overall effect should be an improvement in program instruction cache efficiency and performance.

3.9.5.3 Utilize Preferred Function Order in Re-Build of Application

Finally, the preferred function order command file that is produced by the clt6x is fed into the linker during the re-build of the application, as follows:

```
cl6x options --run_linker *.obj forder.cmd -lInk.cmd
```

The preferred function order command file, `forder.cmd`, contains a list of `--preferred_order=function specification` options. The linker prioritizes the placement of functions relative to each other in the order that the `--preferred_order` options are encountered during the linker invocation.

Each `--preferred_order` option contains a function specification. A function specification can describe simply the name of the function for a global function, or it can provide the path name and source file name where the function is defined. A function specification that contains path and file name information is used to distinguish one static function from another that has the same function name.

The `--preferred_order` options are interpreted by the linker as suggestions to guide the placement of functions relative to each other. They are not explicit placement instructions. If an object file or input section is explicitly mentioned in a linker command file `SECTIONS` directive, then the placement instruction specified in the linker command file takes precedence over any suggestion from a `--preferred_order` option that is associated with a function that is defined in that object file or input section.

This precedence can be relaxed by applying the `unordered()` operator to an output specification as described in [Section 3.9.7](#).

3.9.6 Comma-Separated Values (CSV) Files with Weighted Call Graph (WCG) Information

The format of the CSV files generated by the compiler under the `--analyze=callgraph --use_profile_info` option combination is as follows:

```
"caller", "callee", "weight" [CR][LF]
caller spec, callee spec, call frequency [CR][LF]
caller spec, callee spec, call frequency [CR][LF]
caller spec, callee spec, call frequency [CR][LF]
...
```

Keep the following points in mind:

- Line 1 of the CSV file is the header line. It specifies the meaning of each field in each line of the remainder of the CSV file. In the case of CSV files that contain weighted call graph information, each line will have a caller function specification, followed by a callee function specification, followed by an unsigned integer that provides the number of times a call was executed during run time.
- There may be instances where the caller and callee function specifications are identical on multiple lines in the CSV file. This will happen when a caller function has multiple call sites to the callee function. In the merged weighted call graph that is created by the `clt6x`, the weights of each line that has the same caller and callee function specifications will be added together.
- The CSV file that is generated by the compiler using the path profiling instrumentation will not include information about indirect function calls or calls to runtime support helper functions (like `_remi` or `_divi`). However, you may be able to gather information about such calls with another method (like the PC discontinuity trace mentioned earlier).
- The format of these CSV files is in compliance with the RFC-4180 specification of Comma-Separated Values (CSV) files. For more details on this specification, please see <http://tools.ietf.org/html/rfc4180>.

3.9.7 Linker Command File Operator - `unordered()`

A new `unordered()` operator is now available for use in a linker command file. The effect of this operator is to relax the placement constraints placed on an output section specification in which the content of the output section is explicitly stated.

Consider an example output section specification:

```
SECTIONS
{
  .text:
  {
    file.obj(.text:func_a)
    file.obj(.text:func_b)
    file.obj(.text:func_c)
    file.obj(.text:func_d)
    file.obj(.text:func_e)
    file.obj(.text:func_f)
    file.obj(.text:func_g)
    file.obj(.text:func_h)

    *(.text)
  } > PMEM
  ...
}
```

In this `SECTIONS` directive, the specification of `.text` explicitly dictates the order in which functions are laid out in the output section. Thus by default, the linker will layout `func_a` through `func_h` in exactly the order that they are specified, regardless of any other placement priority criteria (such as a preferred function order list that is enumerated by `--preferred_order` options).

The `unordered()` operator can be used to relax this constraint on the placement of the functions in the `.text` output section so that placement can be guided by other placement priority criteria.

The `unordered()` operator can be applied to an output section as in [Example 3-2](#).

Example 3-2. Output Section for unordered() Operator

```

SECTIONS
{
  .text: unordered()
  {
    file.obj(.text:func_a)
    file.obj(.text:func_b)
    file.obj(.text:func_c)
    file.obj(.text:func_d)
    file.obj(.text:func_e)
    file.obj(.text:func_f)
    file.obj(.text:func_g)
    file.obj(.text:func_h)

    *(.text)
  } > PMEM
  ...
}
    
```

So that, given this list of `--preferred_order` options:

- `--preferred_order="func_g"`
- `--preferred_order="func_b"`
- `--preferred_order="func_d"`
- `--preferred_order="func_a"`
- `--preferred_order="func_c"`
- `--preferred_order="func_f"`
- `--preferred_order="func_h"`
- `--preferred_order="func_e"`

The placement of the functions in the `.text` output section is guided by this preferred function order list. This placement will be reflected in a linker generated map file, as follows:

Example 3-3. Generated Linker Map File for Example 3-2

```

SECTION ALLOCATION MAP

output
section  page      origin      length      attributes/
-----  -
.text    0       00000020   00000120
          00000020   00000020   file.obj (.text:func_g:func_g)
          00000040   00000020   file.obj (.text:func_b:func_b)
          00000060   00000020   file.obj (.text:func_d:func_d)
          00000080   00000020   file.obj (.text:func_a:func_a)
          000000a0   00000020   file.obj (.text:func_c:func_c)
          000000c0   00000020   file.obj (.text:func_f:func_f)
          000000e0   00000020   file.obj (.text:func_h:func_h)
          00000100   00000020   file.obj (.text:func_e:func_e)
    
```

3.9.7.1 About Dot (.) Expressions in the Presence of unordered()

Another aspect of the `unordered()` operator that should be taken into consideration is that even though the operator causes the linker to relax constraints imposed by the explicit specification of an output section's contents, the `unordered()` operator will still respect the position of a dot (.) expression within such a specification.

Consider the output section specification in [Example 3-4](#).

Example 3-4. Respecting Position of a . Expression

```
SECTIONS
{
  .text: unordered()
  {
    file.obj(.text:func_a)
    file.obj(.text:func_b)
    file.obj(.text:func_c)
    file.obj(.text:func_d)

    . += 0x100;

    file.obj(.text:func_e)
    file.obj(.text:func_f)
    file.obj(.text:func_g)
    file.obj(.text:func_h)

    *(.text)
  } > PMEM
  ...
}
```

In [Example 3-4](#), a dot (.) expression, ". += 0x100;", separates the explicit specification of two groups of functions in the output section. In this case, the linker will honor the specified position of the dot (.) expression with respect to the functions on either side of the expression. That is, the `unordered()` operator will allow the preferred function order list to guide the placement of `func_a` through `func_d` relative to each other, but none of those functions will be placed after the hole that is created by the dot (.) expression. Likewise, the `unordered()` operator allows the preferred function order list to influence the placement of `func_e` through `func_h` relative to each other, but none of those functions will be placed before the hole that is created by the dot (.) expression.

3.9.7.2 GROUPs and UNIONs

The `unordered()` operator can only be applied to an output section. This includes members of a `GROUP` or `UNION` directive.

Example 3-5. Applying `unordered()` to `GROUPs`

```
SECTIONS
{
  GROUP
  {
    .grp1:
    {
      file.obj(.grp1:func_a)
      file.obj(.grp1:func_b)
      file.obj(.grp1:func_c)
      file.obj(.grp1:func_d)
    } unordered()

    .grp2:
    {
      file.obj(.grp2:func_e)
      file.obj(.grp2:func_f)
      file.obj(.grp2:func_g)
      file.obj(.grp2:func_h)
    }
  }
}
```

Example 3-5. Applying unordered() to GROUPs (continued)

```

        .text: { *(.text) }

    } > PMEM
    ...
}
    
```

The SECTIONS directive in [Example 3-5](#) applies the unordered() operator to the first member of the GROUP. The .grp1 output section layout can then be influenced by other placement priority criteria (like the preferred function order list), whereas the .grp2 output section will be laid out as explicitly specified.

The unordered() operator cannot be applied to an entire GROUP or UNION. Attempts to do so will result in a linker command file syntax error and the link will be aborted.

3.9.8 Things To Be Aware Of

There are some behavioral characteristics and limitations of the program cache layout development flow that you should bear in mind:

- Generation of Path Profiling Data File (.pdat)

When running an application that has been instrumented to collect path-profiling data (using --gen_profile_info compiler option during build), the application will use functions in the run-time-support library to write out information to the path profiling data file (pprofout.pdat in above tutorial). If there is a path profiling data file already in existence when the application starts to run, then any new path profiling data generated will be appended to the existing file.

To prevent combining path profiling data from separate runs of an application, you need to either rename the path profiling data file from the previous run of the application or remove it before running the application again.

- Indirect Calls Not Recognized by Path Profiling Mechanisms

When using available path profiling mechanisms to collect weighted call graph information from the path profiling data, pprof6x does not recognize indirect calls. An indirect call site will not be represented in the CSV output file that is generated by pprof6x.

You can work around this limitation by introducing your own information about indirect call sites into the relevant CSV file(s). If you take this approach, please be sure to follow the format of the callgraph analysis CSV file ("caller", "callee", "call frequency").

If you are able to get weighted call graph information from a PC trace into a callgraph analysis CSV, this limitation will no longer apply (as the PC trace can always identify the callee of an indirect call).

- Multiple --preferred_order Options Associated With Single Function

There may be cases in which you might want to input more than one preferred function order command file to the linker during the link of an application. For example, you may have developed or received a separate preferred function order command file for one or more of the object libraries that are used by your application.

In such cases, it is possible that one function may be specified in multiple preferred function order command files. If this happens, the linker will honor only the first instance of the --preferred_order option in which the function is specified.

3.10 Indicating Whether Certain Aliasing Techniques Are Used

Aliasing occurs when you can access a single object in more than one way, such as when two pointers point to the same object or when a pointer points to a named object. Aliasing can disrupt optimization, because any indirect reference can refer to another object. The compiler analyzes the code to determine where aliasing can and cannot occur, then optimizes as much as possible while preserving the correctness of the program. The compiler behaves conservatively.

The following sections describe some aliasing techniques that may be used in your code. These techniques are valid according to the ISO C standard and are accepted by the C6000 compiler; however, they prevent the optimizer from fully optimizing your code.

3.10.1 Use the `--aliased_variables` Option When Certain Aliases are Used

The compiler, when invoked with optimization, assumes that any variable whose address is passed as an argument to a function is not subsequently modified by an alias set up in the called function. Examples include:

- Returning the address from a function
- Assigning the address to a global variable

If you use aliases like this in your code, you must use the `--aliased_variables` option when you are optimizing your code. For example, if your code is similar to this, use the `--aliased_variables` option:

```
int *glob_ptr;

g()
{
    int x = 1;
    int *p = f(&x);

    *p      = 5;    /* p aliases x */
    *glob_ptr = 10; /* glob_ptr aliases x */

    h(x);
}

int *f(int *arg)
{
    glob_ptr = arg;
    return arg;
}
```

3.10.2 Use the `--no_bad_aliases` Option to Indicate That These Techniques Are Not Used

The `--no_bad_aliases` option informs the compiler that it can make certain assumptions about how aliases are used in your code. These assumptions allow the compiler to improve optimization. The `--no_bad_aliases` option also specifies that loop-invariant counter increments and decrements are non-zero. Loop invariant means the value of an expression does not change within the loop.

- The `--no_bad_aliases` option indicates that your code does not use the aliasing technique described in [Section 3.10.1](#). If your code uses that technique, do *not* use the `--no_bad_aliases` option. You must compile with the `--aliased_variables` option.

Do *not* use the `--aliased_variables` option with the `--no_bad_aliases` option. If you do, the `--no_bad_aliases` option overrides the `--aliased_variables` option.

- The `--no_bad_aliases` option indicates that a pointer to a character type does *not* alias (point to) an object of another type. That is, the special exception to the general aliasing rule for these types given in section 3.3 of the ISO specification is ignored. If you have code similar to the following example, do *not* use the `--no_bad_aliases` option:

```

{
    long l;
    char *p = (char *) &l;

    p[2] = 5;
}
    
```

- The `--no_bad_aliases` option indicates that indirect references on two pointers, P and Q, are not aliases if P and Q are distinct parameters of the same function activated by the same call at run time. If you have code similar to the following example, do *not* use the `--no_bad_aliases` option:

```

g(int j)
{
    int a[20];

    f(&a, &a)          /* Bad */
    f(&a+42, &a+j)     /* Also Bad */
}

f(int *ptr1, int *ptr2)
{
    ...
}
    
```

- The `--no_bad_aliases` option indicates that each subscript expression in an array reference `A[E1]..[En]` evaluates to a nonnegative value that is less than the corresponding declared array bound. Do *not* use `--no_bad_aliases` if you have code similar to the following example:

```

static int ary[20][20];

int g()
{
    return f(5, -4); /* -4 is a negative index */
    return f(0, 96); /* 96 exceeds 20 as an index */
    return f(4, 16); /* This one is OK */
}

int f(int I, int j)
{
    return ary[i][j];
}
    
```

In this example, `ary[5][-4]`, `ary[0][96]`, and `ary[4][16]` access the same memory location. Only the reference `ary[4][16]` is acceptable with the `--no_bad_aliases` option because both of its indices are within the bounds (0..19).

- The `--no_bad_aliases` option indicates that loop-invariant counter increments and decrements of loop counters are non-zero. Loop invariant means a value of an expression does not change within the loop.

If your code does *not* contain any of the aliasing techniques described above, you should use the `--no_bad_aliases` option to improve the optimization of your code. However, you must use discretion with the `--no_bad_aliases` option; unexpected results may occur if these aliasing techniques appear in your code and the `--no_bad_aliases` option is used.

3.10.3 Using the `--no_bad_aliases` Option With the Assembly Optimizer

The `--no_bad_aliases` option allows the assembly optimizer to assume there are no memory aliases in your linear assembly; i.e., no memory references ever depend on each other. However, the assembly optimizer still recognizes any memory dependencies you point out with the `.mdep` directive. For more information about the `.mdep` directive, see and [Section 4.6.4](#).

3.11 Prevent Reordering of Associative Floating-Point Operations

The compiler freely reorders associative floating-point operations. If you do not wish to have the compiler reorder associative floating point operations, use the `--fp_not_associative` option. Specifying the `--fp_not_associative` option may decrease performance.

3.12 Use Caution With asm Statements in Optimized Code

You must be extremely careful when using asm (inline assembly) statements in optimized code. The compiler rearranges code segments, uses registers freely, and can completely remove variables or expressions. Although the compiler never optimizes out an asm statement (except when it is unreachable), the surrounding environment where the assembly code is inserted can differ significantly from the original C/C++ source code.

It is usually safe to use asm statements to manipulate hardware controls such as interrupt masks, but asm statements that attempt to interface with the C/C++ environment or access C/C++ variables can have unexpected results. After compilation, check the assembly output to make sure your asm statements are correct and maintain the integrity of the program.

3.13 Automatic Inline Expansion (--auto_inline Option)

When optimizing with the --opt_level=3 option or --opt_level=2 option (aliased as -O3 and -O2, respectively), the compiler automatically inlines small functions. A command-line option, --auto_inline=size, specifies the size threshold for automatic inlining. This option controls only the inlining of functions that are not explicitly declared as inline.

When the --auto_inline option is not used, the compiler sets the size limit based on the optimization level and the optimization goal (performance versus code size). If the -auto_inline size parameter is set to 0, automatic inline expansion is disabled. If the --auto_inline size parameter is set to a non-zero integer, the compiler automatically inlines any function smaller than size. (This is a change from previous releases, which inlined functions for which the product of the function size and the number of calls to it was less than size. The new scheme is simpler, but will usually lead to more inlining for a given value of size.)

The compiler measures the size of a function in arbitrary units; however the optimizer information file (created with the --gen_opt_info=1 or --gen_opt_info=2 option) reports the size of each function in the same units that the --auto_inline option uses. When --auto_inline is used, the compiler does not attempt to prevent inlining that causes excessive growth in compile time or size; use with care.

When --auto_inline option is not used, the decision to inline a function at a particular call-site is based on an algorithm that attempts to optimize benefit and cost. The compiler inlines eligible functions at call-sites until a limit on size or compilation time is reached.

When deciding what to inline, the compiler collects all eligible call-sites in the module being compiled and sorts them by the estimated benefit over cost. Functions declared static inline are ordered first, then leaf functions, then all others eligible. Functions that are too big are not included.

Inlining behavior varies, depending on which compile-time options are specified:

- The code size limit is smaller when compiling for code size rather than performance. The --auto_inline option overrides this size limit.
- At --opt_level=3, the compiler auto-inlines aggressively if compiling for performance.
- At --opt_level=2, the compiler only automatically inlines small functions.

Some Functions Cannot Be Inlined

NOTE: For a call-site to be considered for inlining, it must be legal to inline the function and inlining must not be disabled in some way. See the inlining restrictions in [Section 2.11.5](#).

Optimization Level 3 or 2 and Inlining

NOTE: In order to turn on automatic inlining, you must use the --opt_level=3 option or --opt_level=2 option. At --opt_level=2, only small functions are auto-inlined. If you desire the --opt_level=3 or 2 optimizations, but not automatic inlining, use --auto_inline=0 with the --opt_level=3 or 2 option.

Inlining and Code Size

NOTE: Expanding functions inline increases code size, especially inlining a function that is called in a number of places. Function inlining is optimal for functions that are called only from a small number of places and for small functions. To prevent increases in code size because of inlining, use the `--auto_inline=0` and `--no_inlining` options. These options, used together, cause the compiler to inline intrinsics only.

3.14 Using Performance Advice to Optimize your Code

You can use compiler generated Performance Advice to optimize your code. To get Performance Advice, compile with with the following options:

<code>--advice:performance</code>	Instructs the compiler to emit advice to stdout (default).
<code>--advice:performance_file</code>	Instructs the compiler to emit advice into a file.
<code>--advice:performance_dir</code>	Instructs the compiler to emit advice into a file in a specific directory.

If an Advice file is requested, but there is no advice, the advice file will not be created; rather the compiler prints a message to stdout :

```
"filename.c": advice #27004: No Performance Advice is generated.
```

Example 3-6. Output Advice to stdout (default)

```
cl6x -mv6400 -o2 -k --advice:performance func.c
"func.c", line 10: advice #30006: Loop at line 8 cannot be scheduled efficiently
as it contains a function call ("_init"). Try making "_init" an inline
function.
"func.c", line 12: advice #30000: Loop at line 8 cannot be scheduled efficiently
as it contains a function call ("_calculate"). Try to inline call or
consider rewriting loop.
```

Note that Advice to prevent Software Pipeline Disqualification (such as that presented above) will also be printed in the .asm file. So, func.asm will contain :

```

;-----*
;*  SOFTWARE PIPELINE INFORMATION
;*  Disqualified loop: Loop contains a call
;*      Loop at line 8 cannot be scheduled efficiently as it contains a
;*      function call ("_init"). Try making "_init" an inline function.
;*  Disqualified loop: Loop contains non-pipelineable instructions
;*  Disqualified loop: Loop contains a call
;*      Loop at line 8 cannot be scheduled efficiently as it contains a
;*      function call ("_calculate"). Try to inline call or consider
;*      rewriting loop.
;*  Disqualified loop: Loop contains non-pipelineable instructions
;-----*

```

Example 3-7. Output Advice to a file filename.advice

```
cl6x -mv6400 --advice:performance --advice:performance_file=filename.advice func.c
;*****
;* TMS320C6x C/C++ Codegen      Unix v7.5.0P12047 (a0322878 - Feb 16 2012) *
;* Date/Time created: Thu Feb 16 10:26:02 2012 *
;* * * * *
;* This file contains Performance Advice generated by the compiler under the *
```

Example 3-7. Output Advice to a file filename.advise (continued)

```

;* --advice:performance option. Additional information on the specific piece *
;* of advice can be found at: *
;* *
;* http://processors.wiki.ti.com/index.php/Compiler/diagnostic_messages/C6000 *
;* *
;* or by searching for `Advice #advice_id`. *
;* *
;* Warning: This file is auto generated by the compiler and can be *
;* overwritten during the next compile. *
;* *
;*****
;* User Options: --silicon_version=6400

"func.c": advice #27000: Detecting compilation without optimization. Use
    optimization option -o2 or higher.

```

Example 3-8. Output advice to a specific file "myfile.adv" in directory "mydir"

You can do this by using the `--advice:performance_file` and `--advice:performance_dir` option :

```
cl6x -mv6400 -o2 -k --advice:performance_file=myfile.adv --advice:performance_dir=mydir basicloop.c
```

Or by specifying the full path name to the `--advice:performance_file` option :

```
cl6x -mv6400 -o2 -k --advice:performance_file=mydir/myfile.adv basicloop.c
```

If `--advice_dir` option and full pathname are specified together, `--advice:performance_dir` option is ignored, and the advice is generated in the full pathname advice file. Also, note that directory "mydir" must already exist for an advice file to be created in there.

3.14.1 Types of Advice

3.14.1.1 Advice #27000

```
advice #27000: Detecting compilation without optimization. Use optimization option -o2 or higher.
```

Your compilation is being done without any optimization options (-o0 and above). This prevents the compiler from using its most powerful optimization techniques, since the -o (--opt_level) options are the foundations for most other optimizations. You could get substantially better performance using -o2 (or above) optimization. For C6000, optimization option -o2 is required for the software pipelining loop optimization, which is crucial to getting good performance.

The C/C++ compiler is able to perform various optimizations, but you need to specify optimization options on the command line so that these optimizations are performed. The easiest way to invoke optimization is to specify the `--opt_level=n` option on the compiler command line. You can use -On to alias the `--opt_level` option. The n denotes the level of optimization (0, 1, 2, and 3), which controls the type and degree of optimization. See "Invoking Optimization" in [Section 3.1](#) for more information on Optimization Options.

3.14.1.2 Advice #27001: Increase Optimization Level

```
advice #27001: Detecting compilation with low optimization level.
    Use optimization option -o2 or higher.
```

Your compilation uses low-level optimization options (-o1 and below), which prevents the compiler from using its most powerful optimization techniques.

The C/C++ compiler is able to perform various optimizations, but you can control the level of these optimizations. High-level optimizations are performed in the optimizer and low-level, target-specific optimizations occur in the code generator. You must use high-level optimizations to achieve optimal code. You can invoke optimization by specifying the `--opt_level=n` option on the compiler command line.

See "Invoking Optimization" in [Section 3.1](#) for more information on Optimization Options. Also see information for Advice #27000 in [Section 3.14.1.1](#).

3.14.1.3 Advice #27002: Do not turn off software pipelining

```
advice #27002: Detecting compilation with "-mu" which turns off
               software pipelining. To optimize, turn off this option.
```

Your compilation is being done using `-mu`, which turns off software-pipelining. Software-pipelining is a key optimization for achieving good performance. This Advice is issued to alert you to NOT use compiler option `-mu`. `-mu` is a good option for debugging, but it is recommended that this option not be used for production code because of the negative performance implications.

In general, to achieve maximal performance, avoid using the following in production code :

- `-g`: Support full symbolic debug, and is a good option for debugging. This inhibits code reordering across source line boundaries and limits optimizations around function boundaries. This results in less parallelism, more nops and generally less efficient schedules. Using this option can cause a 30-50% performance degradation for control code, generally somewhat less but still significant degradation for performance critical code.
- `-ss`: Interlist source code into assembly file. As with `-g`, this option can negatively impact performance.
- `-mu`: Turns off software-pipelining, which is a key optimization for achieving good performance. This is a good option for debugging, but is not recommended for use in production code due to negative performance implications.

3.14.1.4 Advice #27003: Avoid compiling with debug options

```
advice #27003: Detecting compilation with debug option "-g",
               which hinders optimization. To optimize, remove -g or
               compile with --optimize_with_debug_option.
```

The compiler detects that your compilation is being done using `-g`, which limits software-pipelining. Software-pipelining is a key optimization for achieving good performance. This Advice is issued to alert you to NOT use compiler option `-g`. `-g` is a good option for debugging, but it is recommended that this option not be used for production code because of the negative performance implications. Also see Advice #27002 in [Section 3.14.1.3](#).

3.14.1.5 Advice #27004: No Performance Advice generated

```
advice #27004: No Performance Advice is generated.
```

The compiler detects that your compilation is being done using `--advice:performance` option, but the compiler has no Advice to report. This Advice is issued to alert you to the fact that no Advice is being emitted, and an Advice file will not be created (if one was requested).

3.14.1.6 Advice #30000: Prevent Loop Disqualification due to call

```
advice #30000: Loop at line 10 cannot be scheduled efficiently,
               as it contains a function call ("function_name").
               Try to inline call or consider rewriting loop.
```

The compiler attempts to perform the software pipeline loop optimization at optimization levels `--opt_level=2` (or `-O2`) and `-O3`. If there is a call in the loop, the compiler will attempt to completely inline the called function, but sometimes this is not possible. If the compiler cannot inline the called function, software pipelining cannot be performed. This can severely reduce the performance of the loop.

In the testcase below, the call to the function "func2" prevents software pipelining. Inlining function "func2" or rewriting the loop to avoid a function call can avoid pipeline disqualification. If the loop pipelines successfully you may see performance improvement.

```
void func1(int *p, int *q, int n)
```



```

{
    unsigned int i;

    for (i = 0; i < n; i++)
    {
        int t = func2(i);

        p[i] = q[i] + t;
    }
}

```

3.14.1.7 Advice #30001: Prevent Loop Disqualification due to rts-call

advice #30001: Loop at line 18 cannot be scheduled efficiently, as it contains conversion from "type-a" to "type-b". Try to use "suggested" type.

The compiler can insert calls to special functions in the run-time support library (RTS) to support operations that are not natively supported by the ISA. For instance, while C6000 floating-point ISAs support instructions to convert between floating-point and signed integer values, they don't support conversion between floating-point and unsigned integer values. If you use unsigned variables in floating point expressions, the compiler will generate a call to an RTS routine to carry out this function. Such a call will disable software pipelining.

You can change the unsigned variables in your code to signed variables and prevent this from happening. The compiler will then be able to use the native hardware instead of adding the special function call, so you may get better performance.

3.14.1.8 Advice #30002: Prevent Loop Disqualification due to asm statement

advice #30002: Loop at line 8 cannot be scheduled efficiently, as it contains an asm() statement. Try to replace the asm() statement with C or intrinsic statement.

An asm statement inserted in a C code loop will disqualify the loop for software pipelining. Software-pipelining is a key optimization for achieving good performance. You may see reduced performance without software pipelining.

Replace the asm() statement with native C, or an intrinsic function call to prevent this from happening.

3.14.1.9 Advice #30003: Prevent Loop Disqualification due to complex condition

advice #30003: Loop at line 8 cannot be scheduled efficiently, as it contains complex conditional expression. Try to simplify condition.

Your code contains a complex conditional expression, possibly a large "if" clause, within a loop, which is preventing optimization. The C6000 compiler will optimize small "if" statements ("if" statements with "if" and "else" blocks that are short or empty). The compiler will not optimize large "if" statements, and such large if statements within the loop body will disqualify the loop for software pipelining. Software-pipelining is a key optimization; you may see reduced performance without it.

In the examples below, Example 1 will pipeline, but Example 2 won't :

Example 1:

```

for (i=0; i < N; i++)
{
    if (!flag)
    {
        //statements
    }
    else
    {
        x[i] = y[i];
    }
}

```

Example 2:

```

for (i = 0; i < n; i++)
{
    if (!flag)
    {
        //statements
    }
    else
    {
        if (flag == 1) x[i] = y[i];
    }
}

```

Example 1 will have significantly better performance than Example 2 because it pipelines successfully. But Example 2 can be pipelined if the code is modified to eliminate the nested "if" :

```

for (i = 0; i < n; i++)
{
    if (!flag)
    {
        //statements
    }
    else
    {
        p = (flag == 1);
        x[i] = !p * x[i] + p * y[i] ;
    }
}

```

3.14.1.10 Advice #30004: Prevent Loop Disqualification due to switch statement

advice #30004: Loop at line 257 cannot be scheduled efficiently, as it contains a switch statement. Try to rewrite loop.

There is a switch statement within the loop. A switch statement in a loop will disqualify the loop for software pipelining. Software-pipelining is a key optimization; you may see reduced performance without it.

Try and rewrite the loop without a switch statement.

3.14.1.11 Advice #30005: Prevent Loop Disqualification due to arithmetic operation

advice #30005: Loop at line 5 cannot be scheduled efficiently, as it contains a "division" operation. Rewrite using simpler operations if possible.

The compiler can insert calls to special functions in the run-time support library (RTS) to support operations that are not natively supported by the ISA. For example, the compiler calls `__c6xabi_divi()` (`_divi()` in COFF) function to perform 32-bit integer divide operation. Such functions are called compiler helper functions, and result in a function call within the loop body. In the example below, the compiler will accomplish the division operation by calling the compiler helper function `"_divi"` :

```

void func(float *p, float n)
{
    int i;

    for (i = 1; i < 1000; i++)
    {
        p[i] /= n;
    }
}

```

However if we modify this loop, like below, the loop pipelines :

```

void func_adjusted(float *p, float n)
{
    int i;

```

```

float inv = 1/n;

for (i = 1; i < 1000; i++)
{
    p[i] *= inv;
}

```

3.14.1.12 Advice #30006: Prevent Loop Disqualification due to call(2)

advice #30006: Loop at line 22 cannot be scheduled efficiently, as it contains a function call ("function_name"). Try making "function_name" an inline function.

For improved performance, at optimization levels `--opt_level=2 (-O2)` and `--opt_level=3 (-O3)`, the compiler attempts to software pipeline your loops. Sometimes the compiler may not be able to inline a function call that is in a loop. Because the compiler could not inline the function call, the loop could not be software pipelined, and the loop could not be efficiently scheduled.

For example, in the testcase below, call to function "func2" prevents software pipelining:

```

void func1(int *p, int *q, int n)
{
    unsigned int i;

    for (i = 0; i < n; i++)
    {
        int t = func2(i);

        ; other operations
    }
}
int function func2() { . . . }

```

However if function func2 is inlined, it saves the overhead of a function call. The compiler is free to optimize the function in context with surrounding code. Automatic inlining is controlled by the "inline" keyword; use it to allow inlining of specific functions :

```
inline int function func2() { . . . }
```

Also see #Advice 30000 in [Section 3.14.1.6](#).

3.14.1.13 Advice #30007: Prevent Loop Disqualification due to rts-call(2)

advice #30007: Attempting to use floating-point operation "__mpyd" on fixed-point device, at line 5 (there may be other instances of this). Such calls reduce loop performance; use fixed point operation if possible.

The compiler inserts calls to special functions in the run-time support library (RTS) to support operations that are not natively supported by the instruction set architecture (ISA). For example, C6000 fixed point ISAs do not support floating-point instructions and the compiler will generate a call to an RTS routine to carry out the floating point operation. In the testcase below, the floating-point multiplication is unavailable for a fixed-point device such as C6200:

```

void func(float *p, float *q, int n)
{
    unsigned int i;

    for (i = 1; i < n; i++)
    {
        p[i] = (q[i] * 12.4) / p[i - 1];
    }
}

```

If compiled for C6200 (compiler option `-mv6200`) the compiler will use an RTS call to carry out the operation. Such a call will disable software pipelining. You can rewrite the operation, or use a fixed point operation to prevent this.

Also see Advice #30001 in [Section 3.14.1.7](#).

3.14.1.14 Advice #30008: Improve Loop; Qualify with restrict

advice #30008: Consider adding the restrict qualifier to the definition of `inp1`, `inp2` if they don't access the same memory location.

To help the compiler determine memory dependencies, you can qualify a pointer, reference, or array with the `restrict` keyword. The `restrict` keyword is a type qualifier that can be applied to pointers, references, and arrays. Its use represents a guarantee by you, the programmer, that within the scope of the pointer declaration the object pointed to can be accessed only by that pointer. Any violation of this guarantee renders the program undefined.

To see more information on using `restrict`, refer to [Section 6.5.5](#)

3.14.1.15 Advice #30009: Improve Loop; Add MUST_ITERATE pragma

advice #30009: If you know that this loop will always execute at a multiple of `<2>` and at least `<2>` times, try adding `"#pragma MUST_ITERATE(2, ,2)"` just before the loop.

The C6000 architecture is partitioned into two nearly symmetric halves. The resource breakdown displayed in the software pipelining information in the asm file, is computed after the compiler has partitioned instructions to either the A-side or the B-side. If the resources are imbalanced (i.e.; some resources on one side are used more than resources on the other) software pipelining is resource-bound, and the loop cannot be efficiently scheduled. If the compiler has information about the trip-count for the loop, it can unroll the loop to balance resource usage, and get better pipelining. You can give loop trip-count information to the compiler using the "MUST_ITERATE" pragma.

To see more information on using the `MUST_ITERATE` pragma, refer to [Section 6.9.20](#)

3.14.1.16 Advice #30010: Improve Loop; Add MUST_ITERATE pragma(2)

advice #30010: If you know that this loop will always execute at least `<2>` times, try adding `"#pragma MUST_ITERATE(2)"` just before the loop.

See Advice #30009 at [Section 3.14.1.15](#).

3.14.1.17 Advice #30011: Improve Loop; Add _nassert()

advice #30011: Consider adding assertions to indicate n-byte alignment of variables `input1`, `input2`, `output` if they are actually n-byte aligned: `_nassert((int)(input1) % 8 == 0)`.

Most loops have memory access instructions. The compiler attempts to use wider load instructions, and aligned memory accesses instead of non-aligned memory accesses to reduce/balance out resources used for the memory access instructions. One of the ways to let the compiler know that it is safe to use "wider" loads is to use the keyword "`_nassert`".

To find out more on using the `_nassert` keyword, see [Section 7.5.10](#).

3.14.2 Using the Interlist Feature With Optimization

You control the output of the interlist feature when compiling with optimization (the `--opt_level=n` or `-On` option) with the `--optimizer_interlist` and `--c_src_interlist` options.

- The `--optimizer_interlist` option interlists compiler comments with assembly source statements.
- The `--c_src_interlist` and `--optimizer_interlist` options together interlist the compiler comments and the original C/C++ source with the assembly code.

When you use the `--optimizer_interlist` option with optimization, the interlist feature does *not* run as a separate pass. Instead, the compiler inserts comments into the code, indicating how the compiler has rearranged and optimized the code. These comments appear in the assembly language file as comments starting with `;`. The C/C++ source code is not interlisted, unless you use the `--c_src_interlist` option also.

The interlist feature can affect optimized code because it might prevent some optimization from crossing C/C++ statement boundaries. Optimization makes normal source interlisting impractical, because the compiler extensively rearranges your program. Therefore, when you use the `--optimizer_interlist` option, the compiler writes reconstructed C/C++ statements.

[Example 3-9](#) shows a function that has been compiled with optimization (`--opt_level=2`) and the `--optimizer_interlist` option. The assembly file contains compiler comments interlisted with assembly code.

Impact on Performance and Code Size

NOTE: The `--c_src_interlist` option can have a negative effect on performance and code size.

When you use the `--c_src_interlist` and `--optimizer_interlist` options with optimization, the compiler inserts its comments and the interlist feature runs before the assembler, merging the original C/C++ source into the assembly file.

[Example 3-10](#) shows the function from [Example 3-9](#) compiled with the optimization (`--opt_level=2`) and the `--c_src_interlist` and `--optimizer_interlist` options. The assembly file contains compiler comments and C source interlisted with assembly code.

Example 3-9. The Function From Example 2-4 Compiled With the -O2 and --optimizer_interlist Options

```

_main:
;** 5 ----- printf("Hello, world\n");
;** 6 ----- return 0;
        STW      .D2      B3,*SP--(12)
        .line    3
        B        .S1      _printf
        NOP      2
        MVKL     .S1      SL1+0,A0
        MVKH     .S1      SL1+0,A0
||       MVKL     .S2      RL0,B3
        STW      .D2      A0,*+SP(4)
||       MVKH     .S2      RL0,B3
RL0:    ; CALL OCCURS
        .line    4
        ZERO     .L1      A4
        .line    5
        LDW      .D2      *++SP(12),B3
        NOP      4
        B        .S2      B3
        NOP      5
        ; BRANCH OCCURS
    
```

Example 3-10. The Function From Example 2-4 Compiled with the --opt_level=2, --optimizer_interlist, and --c_src_interlist Options

```

_main:
;** 5 ----- printf("Hello, world\n");
;** 6 ----- return 0;
        STW      .D2      B3,*SP--(12)
;-----
; 5 | printf("Hello, world\n");
;-----
        B        .S1      _printf
        NOP      2
        MVKL     .S1      SL1+0,A0
        MVKH     .S1      SL1+0,A0
||       MVKL     .S2      RL0,B3
        STW      .D2      A0,*+SP(4)
||       MVKH     .S2      RL0,B3
RL0:    ; CALL OCCURS
;-----
; 6 | return 0;
;-----
        ZERO     .L1      A4
        LDW      .D2      *++SP(12),B3
        NOP      4
        B        .S2      B3
        NOP      5
        ; BRANCH OCCURS
    
```

3.14.3 Debugging and Profiling Optimized Code

Debugging fully optimized code is not recommended, because the compiler's extensive rearrangement of code and the many-to-many allocation of variables to registers often make it difficult to correlate source code with object code. Profiling code that has been built with the `--symdebug:dwarf` (aliased as `-g`) option or the `--symdebug:coff` option (STABS debug) is not recommended either, because these options can significantly degrade performance. To remedy these problems, you can use the options described in the following sections to optimize your code in such a way that you can still debug or profile the code.

3.14.3.1 Debugging Optimized Code (`--symdebug:dwarf`, `--symdebug:coff`, and `--opt_level` Options)

To debug optimized code, use the `--opt_level` (aliased as `-O`) option in conjunction with one of the symbolic debugging options (`--symdebug:dwarf` or `--symdebug:coff`). The symbolic debugging options generate directives that are used by the C/C++ source-level debugger, but they disable many compiler optimizations. When you use the `--opt_level` option (which invokes optimization) with the `--symdebug:dwarf` or `--symdebug:coff` option, you turn on the maximum amount of optimization that is compatible with debugging.

If you want to use symbolic debugging and still generate fully optimized code, use the `--optimize_with_debug` option. This option reenables the optimizations disabled by `--symdebug:dwarf` or `--symdebug:coff`. However, if you use the `--optimize_with_debug` option, portions of the debugger's functionality will be unreliable.

If you are having trouble debugging loops in your code, you can use the `--disable_software_pipelining` option to turn off software pipelining. See [Section 3.2.1](#) for more information.

Symbolic Debugging Options Affect Performance and Code Size

NOTE: Using the `--symdebug:dwarf` or `--symdebug:coff` option can cause a significant performance and code size degradation of your code. Use these options for debugging only. Using `--symdebug:dwarf` or `--symdebug:coff` when profiling is not recommended.

C6400+, C6740, and C6600 Support Only DWARF Debugging

NOTE: Since C6400+, C6740, and C6600 produce only DWARF debug information, the `--symdebug:coff` option is not supported when compiling with `-mv6400+`, `-mv6740`, or `-mv6600`.

3.14.3.2 Profiling Optimized Code

To profile optimized code, use optimization (`--opt_level=0` through `--opt_level=3`) without any debug option. By default, the compiler generates a minimal amount of debug information without affecting optimizations, code size, or performance.

If you have a breakpoint-based profiler, use the `--profile:breakpt` option with the `--opt_level` option. The `--profile:breakpt` option disables optimizations that would cause incorrect behavior when using a breakpoint-based profiler.

If you have a power profiler, use the `--profile:power` option with the `--opt_level` option. The `--profile:power` option produces instrument code for the power profiler.

If you need to profile code at a finer grain than the function level in Code Composer Studio, you can use the `--symdebug:dwarf` or `--symdebug:coff` option, although this is not recommended. You might see a significant performance degradation because the compiler cannot use all optimizations with `--symdebug:dwarf` or `--symdebug:coff`. It is recommended that outside of Code Composer Studio, you use the `clock()` function.

Profile Points

NOTE: In Code Composer Studio, when symbolic debugging is not used, profile points can only be set at the beginning and end of functions.

3.14.4 Controlling Code Size Versus Speed

The latest mechanism for controlling the goal of optimizations in the compiler is represented by the `--opt_for_speed=num` option. The *num* denotes the level of optimization (0-5), which controls the type and degree of code size or code speed optimization:

- `--opt_for_speed=0`
Enables optimizations geared towards improving the code size with a *high* risk of worsening or impacting performance.
- `--opt_for_speed=1`
Enables optimizations geared towards improving the code size with a *medium* risk of worsening or impacting performance.
- `--opt_for_speed=2`
Enables optimizations geared towards improving the code size with a *low* risk of worsening or impacting performance.
- `--opt_for_speed=3`
Enables optimizations geared towards improving the code performance/speed with a *low* risk of worsening or impacting code size.
- `--opt_for_speed=4`
Enables optimizations geared towards improving the code performance/speed with a *medium* risk of worsening or impacting code size.
- `--opt_for_speed=5`
Enables optimizations geared towards improving the code performance/speed with a *high* risk of worsening or impacting code size.

If you specify the option without a parameter, the default setting is `--opt_for_speed=4`.

The initial mechanism for controlling code space, the `--opt_for_space` option, has the following equivalences with the `--opt_for_speed` option:

<code>--opt_for_space</code>	<code>--opt_for_speed</code>
none	=4
=0	=3
=1	=2
=2	=1
=3	=0

3.14.5 What Kind of Optimization Is Being Performed?

The TMS320C6000 C/C++ compiler uses a variety of optimization techniques to improve the execution speed of your C/C++ programs and to reduce their size.

Following are some of the optimizations performed by the compiler:

Optimization	See
Cost-based register allocation	Section 3.14.5.1
Alias disambiguation	Section 3.14.5.1
Branch optimizations and control-flow simplification	Section 3.14.5.3
Data flow optimizations <ul style="list-style-type: none"> • Copy propagation • Common subexpression elimination • Redundant assignment elimination 	Section 3.14.5.4
Expression simplification	Section 3.14.5.5
Inline expansion of functions	Section 3.14.5.6
Function Symbol Aliasing	Section 3.14.5.7
Induction variable optimizations and strength reduction	Section 3.14.5.8
Loop-invariant code motion	Section 3.14.5.9
Loop rotation	Section 3.14.5.10
Instruction scheduling	Section 3.14.5.11

C6000-Specific Optimization	See
Register variables	Section 3.14.5.12
Register tracking/targeting	Section 3.14.5.13
Software pipelining	Section 3.14.5.14

3.14.5.1 Cost-Based Register Allocation

The compiler, when optimization is enabled, allocates registers to user variables and compiler temporary values according to their type, use, and frequency. Variables used within loops are weighted to have priority over others, and those variables whose uses do not overlap can be allocated to the same register.

Induction variable elimination and loop test replacement allow the compiler to recognize the loop as a simple counting loop and software pipeline, unroll, or eliminate the loop. Strength reduction turns the array references into efficient pointer references with autoincrements.

3.14.5.2 Alias Disambiguation

C and C++ programs generally use many pointer variables. Frequently, compilers are unable to determine whether or not two or more l values (lowercase L: symbols, pointer references, or structure references) refer to the same memory location. This aliasing of memory locations often prevents the compiler from retaining values in registers because it cannot be sure that the register and memory continue to hold the same values over time.

Alias disambiguation is a technique that determines when two pointer expressions cannot point to the same location, allowing the compiler to freely optimize such expressions.

3.14.5.3 Branch Optimizations and Control-Flow Simplification

The compiler analyzes the branching behavior of a program and rearranges the linear sequences of operations (basic blocks) to remove branches or redundant conditions. Unreachable code is deleted, branches to branches are bypassed, and conditional branches over unconditional branches are simplified to a single conditional branch.

When the value of a condition is determined at compile time (through copy propagation or other data flow analysis), the compiler can delete a conditional branch. Switch case lists are analyzed in the same way as conditional branches and are sometimes eliminated entirely. Some simple control flow constructs are reduced to conditional instructions, totally eliminating the need for branches.

3.14.5.4 Data Flow Optimizations

Collectively, the following data flow optimizations replace expressions with less costly ones, detect and remove unnecessary assignments, and avoid operations that produce values that are already computed. The compiler with optimization enabled performs these data flow optimizations both locally (within basic blocks) and globally (across entire functions).

- **Copy propagation.** Following an assignment to a variable, the compiler replaces references to the variable with its value. The value can be another variable, a constant, or a common subexpression. This can result in increased opportunities for constant folding, common subexpression elimination, or even total elimination of the variable.
- **Common subexpression elimination.** When two or more expressions produce the same value, the compiler computes the value once, saves it, and reuses it.
- **Redundant assignment elimination.** Often, copy propagation and common subexpression elimination optimizations result in unnecessary assignments to variables (variables with no subsequent reference before another assignment or before the end of the function). The compiler removes these dead assignments.

3.14.5.5 Expression Simplification

For optimal evaluation, the compiler simplifies expressions into equivalent forms, requiring fewer instructions or registers. Operations between constants are folded into single constants. For example, $a = (b + 4) - (c + 1)$ becomes $a = b - c + 3$.

3.14.5.6 Inline Expansion of Functions

The compiler replaces calls to small functions with inline code, saving the overhead associated with a function call as well as providing increased opportunities to apply other optimizations.

3.14.5.7 Function Symbol Aliasing

The compiler recognizes a function whose definition contains only a call to another function. If the two functions have the same signature (same return value and same number of parameters with the same type, in the same order), then the compiler can make the calling function an alias of the called function.

For example, consider the following:

```
int bbb(int arg1, char *arg2);

int aaa(int n, char *str)
{
    return bbb(n, str);
}
```

For this example, the compiler makes `aaa` an alias of `bbb`, so that at link time all calls to function `aaa` should be redirected to `bbb`. If the linker can successfully redirect all references to `aaa`, then the body of function `aaa` can be removed and the symbol `aaa` is defined at the same address as `bbb`.

3.14.5.8 Induction Variables and Strength Reduction

Induction variables are variables whose value within a loop is directly related to the number of executions of the loop. Array indices and control variables for loops are often induction variables.

Strength reduction is the process of replacing inefficient expressions involving induction variables with more efficient expressions. For example, code that indexes into a sequence of array elements is replaced with code that increments a pointer through the array.

Induction variable analysis and strength reduction together often remove all references to your loop-control variable, allowing its elimination.

3.14.5.9 Loop-Invariant Code Motion

This optimization identifies expressions within loops that always compute to the same value. The computation is moved in front of the loop, and each occurrence of the expression in the loop is replaced by a reference to the precomputed value.

3.14.5.10 Loop Rotation

The compiler evaluates loop conditionals at the bottom of loops, saving an extra branch out of the loop. In many cases, the initial entry conditional check and the branch are optimized out.

3.14.5.11 Instruction Scheduling

The compiler performs instruction scheduling, which is the rearranging of machine instructions in such a way that improves performance while maintaining the semantics of the original order. Instruction scheduling is used to improve instruction parallelism and hide pipeline latencies. It can also be used to reduce code size.

3.14.5.12 Register Variables

The compiler helps maximize the use of registers for storing local variables, parameters, and temporary values. Accessing variables stored in registers is more efficient than accessing variables in memory. Register variables are particularly effective for pointers.

3.14.5.13 Register Tracking/Targeting

The compiler tracks the contents of registers to avoid reloading values if they are used again soon. Variables, constants, and structure references such as (a.b) are tracked through straight-line code. Register targeting also computes expressions directly into specific registers when required, as in the case of assigning to register variables or returning values from functions.

3.14.5.14 Software Pipelining

Software pipelining is a technique used to schedule from a loop so that multiple iterations of a loop execute in parallel. See [Section 3.2](#) for more information.

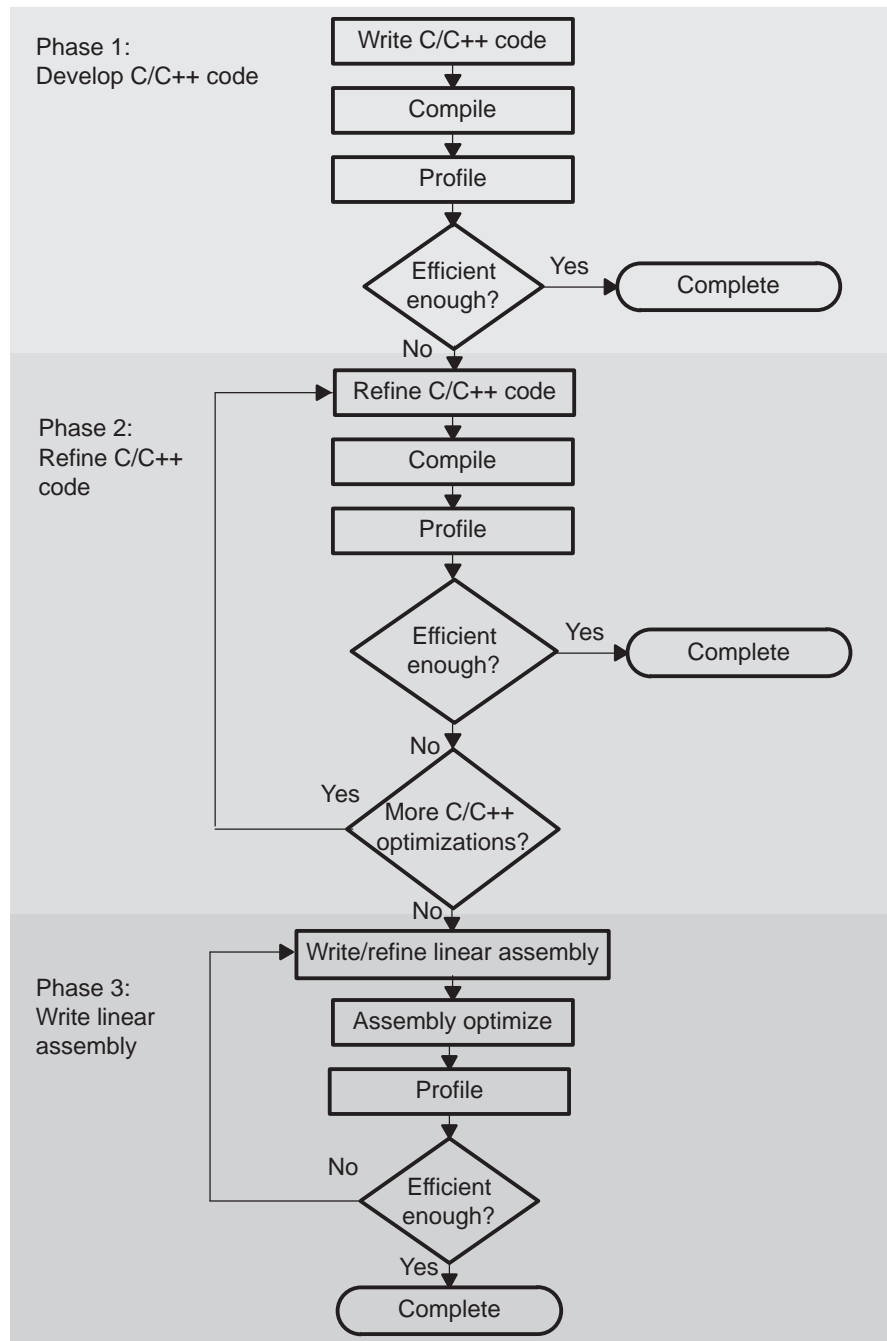
Using the Assembly Optimizer

The assembly optimizer allows you to write assembly code without being concerned with the pipeline structure of the C6000 or assigning registers. It accepts *linear assembly code*, which is assembly code that may have had register-allocation performed and is unscheduled. The assembly optimizer assigns registers and uses loop optimizations to turn linear assembly into highly parallel assembly.

Topic	Page
4.1 Code Development Flow to Increase Performance	105
4.2 About the Assembly Optimizer	106
4.3 What You Need to Know to Write Linear Assembly	107
4.4 Assembly Optimizer Directives	113
4.5 Avoiding Memory Bank Conflicts With the Assembly Optimizer	127
4.6 Memory Alias Disambiguation	133

4.1 Code Development Flow to Increase Performance

You can achieve the best performance from your C6000 code if you follow this flow when you are writing and debugging your code:



There are three phases of code development for the C6000:

- **Phase 1: write in C**

You can develop your C/C++ code for phase 1 without any knowledge of the C6000. Use a simulator after compiling with the `--opt_level=3` option without any `--debug` option to identify any inefficient areas in your C/C++ code. See [Section 3.14.3](#) for more information about debugging and profiling optimized code. To improve the performance of your code, proceed to phase 2.

- **Phase 2: refine your C/C++ code**

In phase 2, use the intrinsics and compiler options that are described in this book to improve your C/C++ code. Use a simulator to check the performance of your altered code. Refer to the *TMS320C6000 Programmer's Guide* for hints on refining C/C++ code. If your code is still not as efficient as you would like it to be, proceed to phase 3.

- **Phase 3: write linear assembly**

In this phase, you extract the time-critical areas from your C/C++ code and rewrite the code in linear assembly. You can use the assembly optimizer to optimize this code. When you are writing your first pass of linear assembly, you should not be concerned with the pipeline structure or with assigning registers. Later, when you are refining your linear assembly code, you might want to add more details to your code, such as partitioning registers.

Improving performance in this stage takes more time than in phase 2, so try to refine your code as much as possible before using phase 3. Then, you should have smaller sections of code to work on in this phase.

4.2 About the Assembly Optimizer

If you are not satisfied with the performance of your C/C++ code after you have used all of the C/C++ optimizations that are available, you can use the assembly optimizer to make it easier to write assembly code for the C6000.

The assembly optimizer performs several tasks including the following:

- Optionally, partitions instructions and/or registers
- Schedules instructions to maximize performance using the instruction-level parallelism of the C6000
- Ensures that the instructions conform to the C6000 latency requirements
- Optionally, allocates registers for your source code

Like the C/C++ compiler, the assembly optimizer performs software pipelining. *Software pipelining* is a technique used to schedule instructions from a loop so that multiple iterations of the loop execute in parallel. The code generation tools attempt to software pipeline your code with inputs from you and with information that it gathers from your program. For more information, see [Section 3.2](#).

To invoke the assembly optimizer, use the compiler program (cl6x). The assembly optimizer is automatically invoked by the compiler program if one of your input files has a .sa extension. You can specify C/C++ source files along with your linear assembly files. For more information about the compiler program, see [Chapter 2](#).

4.3 What You Need to Know to Write Linear Assembly

By using the C6000 profiling tools, you can identify the time-critical sections of your code that need to be rewritten as linear assembly. The source code that you write for the assembly optimizer is similar to assembly source code. However, linear assembly code does not need to be partitioned, scheduled, or register allocated. The intention is for you to let the assembly optimizer determine this information for you. When you are writing linear assembly code, you need to know about these items:

- **Assembly optimizer directives**

Your linear assembly file can be a combination of linear assembly code segments and regular assembly source. Use the assembly optimizer directives to differentiate the assembly optimizer code from the regular assembly code and to provide the assembly optimizer with additional information about your code. The assembly optimizer directives are described in [Section 4.4](#).

- **Options that affect what the assembly optimizer does**

The compiler options in [Table 4-1](#) affect the behavior of the assembly optimizer.

Table 4-1. Options That Affect the Assembly Optimizer

Option	Effect	See
--ap_extension	Changes the default extension for assembly optimizer source files	Section 2.3.9
--ap_file	Changes how assembly optimizer source files are identified	Section 2.3.7
--disable_software_pipelining	Turns off software pipelining	Section 3.2.1
--debug_software_pipeline	Generates verbose software pipelining information	Section 3.2.2
--interrupt_threshold= <i>n</i>	Specifies an interrupt threshold value	Section 2.12
--keep_asm	Keeps the assembly language (.asm) file	Section 2.3.1
--no_bad_aliases	Presumes no memory aliasing	Section 3.10.3
--opt_for_space= <i>n</i>	Controls code size on four levels (<i>n</i> =0, 1, 2, or 3)	Section 3.5
--opt_level= <i>n</i>	Increases level of optimization (<i>n</i> =0, 1, 2, or 3)	Section 3.1
--quiet	Suppresses progress messages	Section 2.3.1
--silicon_version= <i>n</i>	Select target version	Section 2.3.4
--skip_assembler	Compiles or assembly optimizes only (does not assemble)	Section 2.3.1
--speculate_loads= <i>n</i>	Allows speculative execution of loads with bounded address ranges	Section 3.2.3

- **TMS320C6000 instructions**

When you are writing your linear assembly, your code does *not* need to indicate the following:

- Pipeline latency
- Register usage
- Which unit is being used

As with other code generation tools, you might need to modify your linear assembly code until you are satisfied with its performance. When you do this, you will probably want to add more detail to your linear assembly. For example, you might want to partition or assign some registers.

Do Not Use Scheduled Assembly Code as Source

NOTE: The assembly optimizer assumes that the instructions in the input file are placed in the logical order in which you would like them to occur (that is, linear assembly code). Parallel instructions are illegal.

If the compiler cannot make your instructions linear (non-parallel), it produces an error message. The compiler assumes instructions occur in the order the instructions appear in the file. Scheduled code is illegal (even non-parallel scheduled code). Scheduled code may not be detected by the compiler but the resulting output may not be what you intended.

- **Linear assembly source statement syntax**

The linear assembly source programs consist of source statements that can contain assembly optimizer directives, assembly language instructions, and comments. See [Section 4.3.1](#) for more information on the elements of a source statement.

- **Specifying registers or register sides**

Registers can be assigned explicitly to user symbols. Alternatively, symbols can be assigned to the A-side or B-side leaving the compiler to do the actual register allocation. See [Section 4.3.2](#) for information on specifying registers.

- **Specifying the functional unit**

The functional unit specifier is optional in linear assembly code. Data path information is respected; unit information is ignored.

- **Source comments**

The assembly optimizer attaches the comments on instructions from the input linear assembly to the output file. It attaches the 2-tuple $\langle x, y \rangle$ to the comments to specify which iteration and cycle of the loop an instruction is on in the software pipeline. The zero-based number x represents the iteration the instruction is on during the first execution of the kernel. The zero-based number y represents the cycle the instruction is scheduled on within a single iteration of the loop. See [Section 4.3.4](#), for an illustration of the use of source comments and the resulting assembly optimizer output.

4.3.1 Linear Assembly Source Statement Format

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:

<i>label[:]</i>	Labels are optional for all assembly language instructions and for most (but not all) assembly optimizer directives. When used, a label must begin in column 1 of a source statement. A label can be followed by a colon.
[<i>register</i>]	Square brackets ([]) enclose conditional instructions. The machine-instruction mnemonic is executed based on the value of the register within the brackets; valid register names are A0 for C6400, C6400+, C6740, and C6600 only; A1, A2, B0, B1, B2, or symbolic.
<i>mnemonic</i>	The mnemonic is a machine-instruction (such as ADDK, MVKH, B) or assembly optimizer directive (such as .proc, .trip)
<i>unit specifier</i>	The optional unit specifier enables you to specify the functional unit operand. Only the specified unit side is used; other specifications are ignored. The preferred method is specifying register sides.
<i>operand list</i>	The operand list is not required for all instructions or directives. The operands can be symbols, constants, or expressions and must be separated by commas.
<i>comment</i>	Comments are optional. Comments that begin in column 1 must begin with a semicolon or an asterisk; comments that begin in any other column must begin with a semicolon.

The C6000 assembly optimizer 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 character limit, but the truncated portion is not included in the .asm file.

Follow these guidelines in writing linear assembly code:

- 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 characters are interpreted as 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.
- If you set up a conditional instruction, the register must be surrounded by square brackets.
- A mnemonic cannot begin in column 1 or it is interpreted as a label.

Refer to the *TMS320C6000 Assembly Language Tools User's Guide* for information on the syntax of C6000 instructions, including conditional instructions, labels, and operands.

4.3.2 Register Specification for Linear Assembly

There are only two cross paths in the C6000. This limits the C6000 to one source read from each data path's opposite register file per cycle. The compiler must select a side for each instruction; this is called partitioning.

It is recommended that you do not initially partition the linear assembly source code by hand. This allows the compiler more freedom to partition and optimize your code. If the compiler does not find an optimal partition in a software pipelined loop, then you can partition enough instructions by hand to force optimal partitioning by partitioning registers.

The assembly optimizer chooses a register for you such that its use agrees with the functional units chosen for the instructions that operate on the value.

Registers can be directly partitioned through two directives. The `.rega` directive is used to constrain a symbolic name to A-side registers. The `.regb` directive is used to constrain a symbolic name to B-side registers. See [the .rega/.regb topic](#) for further details on these directives. The `.reg` directive allows you to use descriptive names for values that are stored in registers. See [the .reg topic](#) for further details and examples of the `.reg` directive.

[Example 4-1](#) is a hand-coded linear assembly program that computes a dot product; compare to [Example 4-2](#), which illustrates C code.

Example 4-1. Linear Assembly Code for Computing a Dot Product

```

_dotp: .cproc a_0, b_0

    .rega    a_4, tmp0, sum0, prod1, prod2
    .regb    b_4, tmp1, sum1, prod3, prod4
    .reg     cnt, sum
    .reg     val0, val1

    ADD     4, a_0, a_4
    ADD     4, b_0, b_4
    MVK     100, cnt
    ZERO    sum0
    ZERO    sum1

loop:  .trip   25
    LDW     *a_0++[2], val0      ; load a[0-1]
    LDW     *b_0++[2], val1      ; load b[0-1]
    MPY     val0, val1, prod1     ; a[0] * b[0]
    MPYH    val0, val1, prod2     ; a[1] * b[1]
    ADD     prod1, prod2, tmp0    ; sum0 += (a[0]*b[0]) +
    ADD     tmp0, sum0, sum0      ;         (a[1]*b[1])

    LDW     *a_4++[2], val0      ; load a[2-3]
    LDW     *b_4++[2], val1      ; load b[2-3]
    MPY     val0, val1, prod3     ; a[2] * b[2]
    MPYH    val0, val1, prod4     ; a[3] * b[3]
    ADD     prod3, prod4, tmp1    ; sum1 += (a[2]*b[2]) +
    ADD     tmp1, sum1, sum1      ;         (a[3]*b[3])

[cnt] SUB   cnt, 4, cnt          ; cnt -= 4
[cnt] B     loop                ; if (cnt!=0) goto loop

    ADD     sum0, sum1, sum      ; compute final result

    .return sum
    .endproc

```

[Example 4-2](#) is refined C code for computing a dot product.

Example 4-2. C Code for Computing a Dot Product

```
int dotp(short a[], shortb[])
{
    int sum0 = 0;
    int sum1 = 0;

    int sum, I;

    for (I = 0; I < 100/4; I +=4)
    {
        sum0 += a[i] * b[i];
        sum0 += a[i+1] * b[i+1];

        sum1 += a[i+2] * b[i+2];
        sum1 += a[i+3] * b[i+3];
    }
    return
}
```

The old method of partitioning registers indirectly by partitioning instructions can still be used. Side and functional unit specifiers can still be used on instructions. However, functional unit specifiers (.L/.S/.D/.M) are ignored. Side specifiers are translated into partitioning constraints on the corresponding symbolic names, if any. For example:

```
MV .L    x, y    ; translated to .REGA y
LDW .D2T2 *u, v:w ; translated to .REGB u, v, w
```

In the linear assembler, you can also specify register pairs using the .cproc and/or .reg directive as in [Example 4-3](#):

Example 4-3. Specifying a Register Pair

```
.global foopair
foopair: .cproc q1:q0,s0
    .reg r1:r0
    ADD q1:q0, s0, r1:r0
    .return r1:r0
    .endproc
```

In [Example 4-3](#), the expression "q1:q0" means that the first argument into the linear assembly function is a register pair. By the C calling conventions, the pair "q1:q0" symbols are mapped to register pair "a5:a4". When a register pair syntax is used as the argument to a .reg directive (as shown), it means that the two register symbols are constrained to be an aligned register pair when the compiler processes the linear assembler source and allocates actual registers that the register pair symbols map to "r1:r0" as shown.

The 7.2. Beta compiler supports a register quad syntax (C6600 only), in order to specify 128-bit operands of 128-bit capable instructions in linear assembly and assembly source code. [Example 4-4](#) illustrates how you can specify register quads:

Example 4-4. Specifying a Register Quad (C6600 Only)

```
.global fooquad
fooquad: .cproc q3:q2:q1:q0, s3:s2:s1:s0
    .reg r3:r2:r1:r0
    QMPY32 s3:s2:s1:s0, q3:q2:q1:q0, r3:r2:r1:r0
    .return r3:r2:r1:r0
    .endproc
```

In [Example 4-4](#), the expression "q3:q2:q1:q0" means that the first argument into the linear assembly function is a register quad. By the C calling conventions, the quad "q3:q2:q1:q0" symbols are mapped to register quad "a7:a6:a5:a4". When a register quad syntax is used as the argument to a .reg directive (as shown), it means that the four register symbols are constrained to be an aligned register quad when the compiler processes the linear assembler source and allocates actual registers that the register quad symbols map to "r3:r2:r1:r0" as shown.

4.3.3 Functional Unit Specification for Linear Assembly

Specifying functional units has been deprecated by the ability to partition registers directly. (See [Section 4.3.2](#) for details.) While you can use the unit specifier field in linear assembly, only the register side information is used by the compiler.

You specify a functional unit by following the assembler instruction with a period (.) and a functional unit specifier. One instruction can be assigned to each functional unit in a single instruction cycle. There are eight functional units, two of each functional type, and two address paths. The two of each functional type are differentiated by the data path each uses, A or B.

.D1 and .D2	Data/addition/subtraction operations
.L1 and .L2	Arithmetic logic unit (ALU)/compares/long data arithmetic
.M1 and .M2	Multiply operations
.S1 and .S2	Shift/ALU/branch/field operations
.T1 and .T2	Address paths

There are several ways to enter the unit specifier filed in linear assembly. Of these, only the specific register side information is recognized and used:

- You can specify the particular functional unit (for example, .D1).
- You can specify the .D1 or .D2 functional unit followed by T1 or T2 to specify that the nonmemory operand is on a specific register side. T1 specifies side A and T2 specifies side B. For example:

```
LDW .D1T2 *A3[A4], B3
LDW .D1T2 *src, dst
```

- You can specify only the data path (for example, .1), and the assembly optimizer assigns the functional type (for example, .L1).

For more information on functional units refer to the *TMS320C6000 CPU and Instruction Set Reference Guide*.

4.3.4 Using Linear Assembly Source Comments

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

The assembly optimizer schedules instructions; that is, it rearranges instructions. Stand-alone comments are moved to the top of a block of instructions. Comments at the end of an instruction statement remain in place with the instruction.

[Example 4-5](#) shows code for a function called Lmac that contains comments.

Example 4-5. Lmac Function Code Showing Comments

```

Lmac:  .cproc   A4,B4

       .reg    t0,t1,p,i,sh:s1

       MVK    100,i
       ZERO   sh
       ZERO   sl

loop:  .trip   100

       LDH    *a4++, t0      ; t0 = a[i]
       LDH    *b4++, t1      ; t1 = b[i]
       MPY    t0,t1,p        ; prod = t0 * t1
       ADD    p,sh:s1,sh:s1  ; sum += prod
[I]     ADD    -1,i,i         ; --I
[I]     B     loop           ; if (I) goto loop

       .return sh:s1

       .endproc

```

4.3.5 Assembly File Retains Your Symbolic Register Names

In the output assembly file, register operands contain your symbolic name. This aids you in debugging your linear assembly files and in gluing snippets of linear assembly output into assembly files.

A `.map` directive (see [the .map topic](#)) at the beginning of an assembly function associates the symbolic name with the actual register. In other words, the symbolic name becomes an alias for the actual register. The `.map` directive can be used in assembly and linear assembly code.

When the compiler splits a user symbol into two symbols and each is mapped to distinct machine register, a suffix is appended to instances of the symbolic name to generate unique names so that each unique name is associated with one machine register.

For example, if the compiler associated the symbolic name `y` with `A5` in some instructions and `B6` in some others, the output assembly code might look like:

```

.MAP y/A5
.MAP y'/B6
...
ADD .S2X y, 4, y' ; Equivalent to add A5, 4, B6

```

To disable this format with symbolic names and display assembly instructions with actual registers instead, compile with the `--machine_regs` option.

4.4 Assembly Optimizer Directives

Assembly optimizer directives supply data for and control the assembly optimization process. The assembly optimizer optimizes linear assembly code that is contained within procedures; that is, code within the `.proc` and `.endproc` directives or within the `.cproc` and `.endcproc` directives. If you do not use `.cproc`/`.proc` directives in your linear assembly file, your code will not be optimized by the assembly optimizer. This section describes these directives and others that you can use with the assembly optimizer.

[Table 4-2](#) summarizes the assembly optimizer directives. It provides the syntax for each directive, a description of each directive, and any restrictions that you should keep in mind. See the specific directive topic for more detail.

In [Table 4-2](#) and the detailed directive topics, the following terms for parameters are used:

argument— Symbolic variable name or machine register

memref— Symbol used for a memory reference (not a register)

register— Machine (hardware) register

symbol— Symbolic user name or symbolic register name

variable— Symbolic variable name or machine register

Table 4-2. Assembly Optimizer Directives Summary

Syntax	Description	Restrictions
<code>.call</code> [<i>ret_reg</i> =] <i>func_name</i> (<i>argument</i> ₁ , <i>argument</i> ₂ , ...)	Calls a function	Valid only within procedures
<code>.circ</code> <i>symbol</i> ₁ / <i>register</i> ₁ [, <i>symbol</i> ₂ / <i>register</i> ₂]	Declares circular addressing	Must manually insert setup/teardown code for circular addressing. Valid only within procedures
<i>label</i> <code>.cproc</code> [<i>argument</i> ₁ [, <i>argument</i> ₂ , ...]]	Start a C/C++ callable procedure	Must use with <code>.endproc</code>
<code>.endproc</code>	End a C/C++ callable procedure	Must use with <code>.cproc</code>
<code>.endproc</code> [<i>variable</i> ₁ [, <i>variable</i> ₂ ,...]]	End a procedure	Must use with <code>.proc</code>
<code>.map</code> <i>symbol</i> ₁ / <i>register</i> ₁ [, <i>symbol</i> ₂ / <i>register</i> ₂]	Assigns a symbol to a register	Must use an actual machine register
<code>.mdep</code> [<i>memref</i> ₁ [, <i>memref</i> ₂]]	Indicates a memory dependence	Valid only within procedures
<code>.mptr</code> { <i>variable</i> <i>memref</i> }, <i>base</i> [+ <i>offset</i>] [, <i>stride</i>]	Avoid memory bank conflicts	Valid only within procedures
<code>.no_mdep</code>	No memory aliases in the function	Valid only within procedures
<code>.pref</code> <i>symbol</i> / <i>register</i> ₁ [/ <i>register</i> ₂ /...]	Assigns a symbol to a register in a set	Must use actual machine registers
<i>label</i> <code>.proc</code> [<i>variable</i> ₁ [, <i>variable</i> ₂ , ...]]	Start a procedure	Must use with <code>.endproc</code>
<code>.reg</code> <i>symbol</i> ₁ [, <i>symbol</i> ₂ ,...]	Declare variables	Valid only within procedures
<code>.rega</code> <i>symbol</i> ₁ [, <i>symbol</i> ₂ ,...]	Partition symbol to A-side register	Valid only within procedures
<code>.regb</code> <i>symbol</i> ₁ [, <i>symbol</i> ₂ ,...]	Partition symbol to B-side register	Valid only within procedures
<code>.reserve</code> [<i>register</i> ₁ [, <i>register</i> ₂ ,...]]	Prevents the compiler from allocating a register	Valid only within procedures
<code>.return</code> [<i>argument</i>]	Return a value to a procedure	Valid only within <code>.cproc</code> procedures
<i>label</i> <code>.trip</code> <i>min</i>	Specify trip count value	Valid only within procedures
<code>.volatile</code> <i>memref</i> ₁ [, <i>memref</i> ₂ ,...]	Designate memory reference volatile	Use <code>--interrupt_threshold=1</code> if reference may be modified during an interrupt

.call *Calls a Function*
Syntax

```
.call [ret_reg=] func_name ([argument1, argument2,...])
```

Description

Use the **.call** directive to call a function. Optionally, you can specify a register that is assigned the result of the call. The register can be a symbolic or machine register. The **.call** directive adheres to the same register and function calling conventions as the C/C++ compiler. For information, see [Section 7.3](#) and [Section 7.4](#). There is no support for alternative register or function calling conventions.

You cannot call a function that has a variable number of arguments, such as `printf`. No error checking is performed to ensure the correct number and/or type of arguments is passed. You cannot pass or return structures through the **.call** directive.

Following is a description of the **.call** directive parameters:

<i>ret_reg</i>	(Optional) Symbolic/machine register that is assigned the result of the call. If not specified, the assembly optimizer presumes the call overwrites the registers A5 and A4 with a result.
<i>func_name</i>	The name of the function to call, or the name of the symbolic/machine register for indirect calls. A register pair is not allowed. The label of the called function must be defined in the file. If the code for the function is not in the file, the label must be defined with the .global or .ref directive (refer to the <i>TMS320C6000 Assembly Language Tools User's Guide</i> for details). If you are calling a C/C++ function, you must use the appropriate linkname of that function. See Section 6.12 for more information.
<i>arguments</i>	(Optional) Symbolic/machine registers passed as an argument. The arguments are passed in this order and cannot be a constant, memory reference, or other expression.

By default, the compiler generates near calls and the linker utilizes trampolines if the near call will not reach its destination. To force a far call, you must explicitly load the address of the function into a register, and then issue an indirect call. For example:

```
MVK    func,reg
MVKH   func,reg
.call  reg(op1)           ; forcing a far call
```

If you want to use ***** for indirection, you must abide by C/C++ syntax rules, and use the following alternate syntax:

```
.call [ret_reg =] (* ireg)([arg1, arg2,...])
```

For example:

```
.call  (*driver)(op1, op2) ; indirect call

.reg   driver
.call  driver(op1, op2)   ; also an indirect call
```

Here are other valid examples that use the **.call** syntax.

```
.call  fir(x, h, y)           ; void function

.call  minimal( )           ; no arguments

.call  sum = vecsum(a, b)    ; returns an int

.call  hi:lo = _atol(string) ; returns a long
```

Since you can use machine register names anywhere you can use symbolic registers, it may appear you can change the function calling convention. For example:

```
.call  A6 = compute()
```


It appears that the result is returned in A6 instead of A4. This is incorrect. Using machine registers does not override the calling convention. After returning from the *compute* function with the returned result in A4, a MV instruction transfers the result to A6.

Example

Here is a complete .call example:

```
.global _main
.global _puts, _rand, _ltoa
.sect ".const"
string1: .string "The random value returned is ", 0
string2: .string " ", 10, 0 ; '10' == newline
.bss charbuf, 20
.text
_main: .cproc
.reg random_value, bufptr, ran_val_hi:ran_val_lo
.call random_value = _rand() ; get a random value

MVKL string1, bufptr ; load address of string1
MVKH string1, bufptr
.call _puts(bufptr) ; print out string1
MV random_value, ran_val_lo
SHR ran_val_lo, 31, ran_val_hi ; sign extend random value
.call _ltoa(ran_val_hi:ran_val_lo, bufptr) ; convert it to a string
.call _puts(bufptr) ; print out the random value
MVKL string2, bufptr ; load address of string2
MVKH string2, bufptr
.call _puts(bufptr) ; print out a newline
.endproc
```

.circ***Declare Circular Registers*****Syntax**

.circ *symbol*₁ / *register*₁ [, *symbol*₂ / *register*₂ , ...]

Description

The **.circ** directive assigns a symbolic register name to a machine register and declares the symbolic register as available for circular addressing. The compiler then assigns the variable to the register and ensures that all code transformations are safe in this situation. You must insert setup/teardown code for circular addressing.

<i>symbol</i>	A valid symbol name to be assigned to the register. The variable is up to 128 characters long and must begin with a letter. Remaining characters of the variable can be a combination of alphanumeric characters, the underscore (_), and the dollar sign (\$).
<i>register</i>	Name of the actual register to be assigned a variable.

The compiler assumes that it is safe to speculate any load using an explicitly declared circular addressing variable as the address pointer and may exploit this assumption to perform optimizations.

When a symbol is declared with the **.circ** directive, it is not necessary to declare that symbol with the **.reg** directive.

The **.circ** directive is equivalent to using **.map** with a circular declaration.

Example

Here the symbolic name Ri is assigned to actual machine register Mi and Ri is declared as potentially being used for circular addressing.

```
.CIRC R1/M1, R2/M2 ...
```

.cproc/endproc **Define a C Callable Procedure**

Syntax `label .cproc [argument1 [, argument2 , ...]]
 .endproc`

Description Use the **.cproc/endproc** directive pair to delimit a section of your code that you want the assembly optimizer to optimize and treat as a C/C++ callable function. This section is called a procedure. The `.cproc` directive is similar to the `.proc` directive in that you use `.cproc` at the beginning of a section and `.endproc` at the end of a section. In this way, you can set off sections of your assembly code that you want to be optimized, like functions. The directives must be used in pairs; do not use `.cproc` without the corresponding `.endproc`. Specify a label with the `.cproc` directive. You can have multiple procedures in a linear assembly file.

The `.cproc` directive differs from the `.proc` directive in that the compiler treats the `.cproc` region as a C/C++ callable function. The assembly optimizer performs some operations automatically in a `.cproc` region in order to make the function conform to the C/C++ calling conventions and to C/C++ register usage conventions.

These operations include the following:

- When you use save-on-entry registers (A10 to A15 and B10 to B15), the assembly optimizer saves the registers on the stack and restores their original values at the end of the procedure.
- If the compiler cannot allocate machine registers to symbolic register names specified with the `.reg` directive (see [the .reg topic](#)) it uses local temporary stack variables. With `.cproc`, the compiler manages the stack pointer and ensures that space is allocated on the stack for these variables.

For more information, see [Section 7.3](#) and [Section 7.4](#).

Use the optional *argument* to represent function parameters. The argument entries are very similar to parameters declared in a C/C++ function. The arguments to the `.cproc` directive can be of the following types:

- **Machine-register names.** If you specify a machine-register name, its position in the argument list must correspond to the argument passing conventions for C (see [Section 7.4](#)). For example, the C/C++ compiler passes the first argument to a function in register A4. This means that the first argument in a `.cproc` directive must be A4 or a symbolic name. Up to ten arguments can be used with the `.cproc` directive.
- **Variable names.** If you specify a variable name, then the assembly optimizer ensures that either the variable name is allocated to the appropriate argument passing register or the argument passing register is copied to the register allocated for the variable name. For example, the first argument in a C/C++ call is passed in register A4, so if you specify the following `.cproc` directive:

```
frame .cproc arg1
```

The assembly optimizer either allocates `arg1` to A4, or `arg1` is allocated to a different register (such as B7) and an `MV A4, B7` is automatically generated.

- **Register pairs.** A register pair is specified as `arghi:arglo` and represents a 40-bit argument or a 64-bit type double argument.

For example, the `.cproc` defined as follows:

```
_fcn: .cproc arg1, arg2hi:arg2lo, arg3, B6, arg5, B9:B8
      ...
      .return res
      ...
      .endproc
```

corresponds to a C function declared as:

```
int fcn(int arg1, long arg2, int arg3, int arg4, int arg5, long arg6);
```

In this example, the fourth argument of `.cproc` is register B6. This is allowed since the fourth argument in the C/C++ calling conventions is passed in B6. The sixth argument of `.cproc` is the actual register pair B9:B8. This is allowed since the sixth argument in the C/C++ calling conventions is passed in B8 or B9:B8 for longs.

- **Register quads** (C6600 only). A register quad is specified as `r3:r2:r1:r0` and represents a 128-bit type, `__x128_t`. See [Example 4-4](#).

If you are calling a procedure from C++ source, you must use the appropriate linkname for the procedure label. Otherwise, you can force C naming conventions by using the extern C declaration. See [Section 6.12](#) and [Section 7.5](#) for more information.

When `.endproc` is used with a `.cproc` directive, it cannot have arguments. The *live out* set for a `.cproc` region is determined by any `.return` directives that appear in the `.cproc` region. (A value is *live out* if it has been defined before or within the procedure and is used as an output from the procedure.) Returning a value from a `.cproc` region is handled by the `.return` directive. The return branch is automatically generated in a `.cproc` region. See [the .return topic](#) for more information.

Only code within procedures is optimized. The assembly optimizer copies any code that is outside of procedures to the output file and does not modify it. See [Section 4.4.1](#) for a list of instruction types that cannot appear in a `.cproc` region.

Example

Here is an example in which `.cproc` and `.endproc` are used:

```
_if_then: .cproc  a, cword, mask, theta

        .reg    cond, if, ai, sum, cntr

        MVK     32,cntr          ; cntr = 32
        ZERO   sum              ; sum = 0

LOOP:
        AND     cword,mask,cond  ; cond = codeword & mask
[cond] MVK     1,cond            ; !(!(cond))
        CMPEQ  theta,cond,if    ; (theta == !(!(cond)))
        LDH    *a++,ai          ; a[i]
[if]   ADD     sum,ai,sum        ; sum += a[i]
[!if]  SUB     sum,ai,sum        ; sum -= a[i]
        SHL    mask,1,mask      ; mask = mask << 1
[cntr] ADD     -1,cntr,cntr      ; decrement counter
[cntr] B       LOOP            ; for LOOP

        .return sum

        .endproc
```

.map *Assign a Variable to a Register*

Syntax `.map symbol1 / register1 [, symbol2 / register2 , ...]`

Description The **.map** directive assigns symbol names to machine registers. Symbols are stored in the substitution symbol table. The association between symbolic names and actual registers is wiped out at the beginning and end of each linear assembly function. The **.map** directive can be used in assembly and linear assembly files.

variable A valid symbol name to be assigned to the register. The substitution symbol is up to 128 characters long and must begin with a letter. Remaining characters of the variable can be a combination of alphanumeric characters, the underscore (`_`), and the dollar sign (`$`).

register Name of the actual register to be assigned a variable.

When a symbol is declared with the **.map** directive, it is not necessary to declare that symbol with the **.reg** directive.

Example Here the **.map** directive is used to assign `x` to register `A6` and `y` to register `B7`. The symbols are used with a move statement.

```
.map x/A6, y/B7
MV    x, y           ; equivalent to MV A6, B7
```

.mdep *Indicates a Memory Dependence*

Syntax `.mdep memref1 , memref2`

Description The **.mdep** directive identifies a specific memory dependence. Following is a description of the **.mdep** directive parameters:

memref The symbol parameter is the name of the memory reference.

The symbol used to name a memory reference has the same syntax restrictions as any assembly symbol. (For more information about symbols, refer to the *TMS320C6000 Assembly Language Tools User's Guide*.) It is in the same space as the symbolic registers. You cannot use the same name for a symbolic register and annotating a memory reference.

The **.mdep** directive tells the assembly optimizer that there is a dependence between two memory references.

The **.mdep** directive is valid only within procedures; that is, within occurrences of the **.proc** and **.endproc** directive pair or the **.cproc** and **.endproc** directive pair.

Example Here is an example in which **.mdep** is used to indicate a dependence between two memory references.

```
.mdep ld1, st1

LDW   *p1++{ld1}, inp1 ;memory reference "ld1"
;other code ...
STW   outp2, *p2++{st1} ;memory reference "st1"
```

.mptr *Avoid Memory Bank Conflicts*

Syntax `.mptr {variable | memref}, base [+ offset] [, stride]`

Description The **.mptr** directive associates a register with the information that allows the assembly optimizer to determine automatically whether two memory operations have a memory bank conflict. If the assembly optimizer determines that two memory operations have a memory bank conflict, then it does not schedule them in parallel.

A memory bank conflict occurs when two accesses to a single memory bank in a given cycle result in a memory stall that halts all pipeline operation for one cycle while the second value is read from memory. For more information on memory bank conflicts, including how to use the **.mptr** directive to prevent them, see [Section 4.5](#).

Following are descriptions of the **.mptr** directive parameters:

<i>variable memref</i>	The name of the register symbol or memory reference used to identify a load or store involved in a dependence.
<i>base</i>	A symbolic address that associates related memory accesses
<i>offset</i>	The offset in bytes from the starting base symbol. The offset is an optional parameter and defaults to 0.
<i>stride</i>	The register loop increment in bytes. The stride is an optional parameter and defaults to 0.

The **.mptr** directive tells the assembly optimizer that when the *symbol* or *memref* is used as a memory pointer in an LD(B/BU)(H/HU)(W) or ST(B/H/W) instruction, it is initialized to point to *base + offset* and is incremented by *stride* each time through the loop.

The **.mptr** directive is valid within procedures only; that is, within occurrences of the **.proc** and **.endproc** directive pair or the **.cproc** and **.endproc** directive pair.

The *symbolic addresses* used for base symbol names are in a name space separate from all other labels. This means that a symbolic register or assembly label can have the same name as a memory bank base name. For example:

```
.mptr Darray,Darray
```

Example

Here is an example in which **.mptr** is used to avoid memory bank conflicts.

```
_blkcp: .cproc I

        .reg    ptr1, ptr2, tmp1, tmp2

        MVK    0x0, ptr1           ; ptr1 = address 0
        MVK    0x8, ptr2           ; ptr2 = address 8

loop:   .trip   50

        .mptr  ptr1, a+0, 4
        .mptr  foo, a+8, 4

        LDW   *ptr1++, tmp1       ; potential conflict
        STW   tmp1, *ptr2++{foo}  ; load *0, bank 0
                                       ; store *8, bank 0

        [I]   ADD   -1,i,i         ; I--
        [I]   B     loop           ; if (!0) goto loop

        .endproc
```


A value is *live in* if it has been defined before the procedure and is used as an input to the procedure. A value is *live out* if it has been defined before or within the procedure and is used as an output from the procedure. If you do not specify any registers with the `.endproc` directive, it is assumed that no registers are live out.

Only code within procedures is optimized. The assembly optimizer copies any code that is outside of procedures to the output file and does not modify it.

See [Section 4.4.1](#) for a list of instruction types that cannot appear in a `.proc` region.

Example

Here is a block move example in which `.proc` and `.endproc` are used:

```
move    .proc A4, B4, B0
        .no_mdep

loop:
    LDW    *B4++, A1
    MV     A1, B1
    STW    B1, *A4++
    ADD    -4, B0, B0
[B0] B    loop
        .endproc
```

.reg
Declare Symbolic Registers
Syntax

```
.reg symbol1 [, symbol2 , ...]
```

Description

The **.reg** directive allows you to use descriptive names for values that are stored in registers. The assembly optimizer chooses a register for you such that its use agrees with the functional units chosen for the instructions that operate on the value.

The `.reg` directive is valid within procedures only; that is, within occurrences of the `.proc` and `.endproc` directive pair or the `.cproc` and `.endproc` directive pair.

Declaring register pairs (or register quads for C6600) explicitly is optional. Doing so is only necessary if the registers should be allocated as a pair, but they are not used that way. It is a best practice to declare register pairs and register quads with the pair/quad syntax. Here is an example of declaring a register pair:

```
.reg    A7:A6
```

Example 1

This example uses the same code as the block move example shown for `.proc/.endproc` but the `.reg` directive is used:

```
move    .cproc dst, src, cnt

        .reg tmp1, tmp2

loop:
    LDW    *src++, tmp1
    MV     tmp1, tmp2
    STW    tmp2, *dst++
    ADD    -4, cnt, cnt
[cnt] B    loop
```

Notice how this example differs from the `.proc` example: symbolic registers declared with `.reg` are allocated as machine registers.

Example 2

The code in the following example is invalid, because a variable defined by the `.reg` directive cannot be used outside of the defined procedure:

```
move    .proc A4
        .reg tmp
    LDW    *A4++, top
    MV     top, B5
        .endproc
    MV top, B6    ; WRONG: top is invalid outside of the procedure
```

.rega/.regb **Partition Registers Directly**

Syntax

```
.rega symbol1 [, symbol2 , ...]
.regb symbol1 [, symbol2 , ...]
```

Description Registers can be directly partitioned through two directives. The **.rega** directive is used to constrain a symbol name to A-side registers. The **.regb** directive is used to constrain a symbol name to B-side registers. For example:

```
.REGA y
.REGB u, v, w
MV    x, y
LDW   *u, v:w
```

The **.rega** and **.regb** directives are valid within procedures only; that is, within occurrences of the **.proc** and **.endproc** directive pair or the **.cproc** and **.endproc** directive pair.

When a symbol is declared with the **.rega** or **.regb** directive, it is not necessary to declare that symbol with the **.reg** directive.

The old method of partitioning registers indirectly by partitioning instructions can still be used. Side and functional unit specifiers can still be used on instructions. However, functional unit specifiers (**.L/S/D/M**) and crosspath information are ignored. Side specifiers are translated into partitioning constraints on the corresponding symbol names, if any. For example:

```
MV .lX    z, y    ; translated to .REGA y
LDW .D2T2 *u, v:w ; translated to .REGB u, v, w
```

.reserve **Reserve a Register**

Syntax

```
.reserve [register1 [, register2 , ...]]
```

Description The **.reserve** directive prevents the assembly optimizer from using the specified *register* in a **.proc** or **.cproc** region.

If a **.reserved** register is explicitly assigned in a **.proc** or **.cproc** region, then the assembly optimizer can also use that register. For example, the variable **tmp1** can be allocated to register A7, even though it is in the **.reserve** list, since A7 was explicitly defined in the **ADD** instruction:

```
.cproc
.reserve a7
.reg     tmp1
....
ADD     a6, b4, a7
....
.endproc
```

Reserving Registers A4 and A5

NOTE: When inside of a **.cproc** region that contains a **.call** statement, A4 and A5 cannot be specified in a **.reserve** statement. The calling convention mandates that A4 and A5 are used as the return registers for a **.call** statement.

Example 1

The `.reserve` in this example guarantees that the assembly optimizer does not use A10 to A13 or B10 to B13 for the variables `tmp1` to `tmp5`:

```
test .proc    a4, b4
     .reg     tmp1, tmp2, tmp3, tmp4, tmp5
     .reserve a10, a11, a12, a13, b10, b11, b12, b13
     .....
     .endproc a4
```

Example 2

The assembly optimizer may generate less efficient code if the available register pool is overly restricted. In addition, it is possible that the available register pool is constrained such that allocation is not possible and an error message is generated. For example, the following code generates an error since all of the conditional registers have been reserved, but a conditional register is required for the variable `tmp`:

```
.cproc ...
.reserve a1,a2,b0,b1,b2
.reg tmp
....
[tmp] ....
....
.endproc
```

.return
Return a Value to a C callable Procedure
Syntax

.return [*argument*]

Description

The `.return` directive function is equivalent to the `return` statement in C/C++ code. It places the optional argument in the appropriate register for a return value as per the C/C++ calling conventions (see [Section 7.4](#)).

The optional *argument* can have the following meanings:

- Zero arguments implies a `.cproc` region that has no return value, similar to a void function in C/C++ code.
- An argument implies a `.cproc` region that has a 32-bit return value, similar to an `int` function in C/C++ code.
- A register pair of the format `hi:lo` implies a `.cproc` region that has a 40-bit long, a 64-bit long long, or a 64-bit type double return value; similar to a long/long long/double function in C/C++ code.

Arguments to the `.return` directive can be either symbolic register names or machine-register names.

All return statements in a `.cproc` region must be consistent in the type of the return value. It is not legal to mix a `.return arg` with a `.return hi:lo` in the same `.cproc` region.

The `.return` directive is unconditional. To perform a conditional `.return`, simply use a conditional branch around a `.return`. The assembly optimizer removes the branch and generates the appropriate conditional code. For example, to return if condition `cc` is true, return as:

```
[!cc] B around
     .return
around:
```

Example

This example uses a symbolic register, `tmp`, and a machine-register, `A5`, as `.return` arguments:

```
.cproc ...
.reg tmp
...
.return tmp    = legal symbolic name
...
.return a5    = legal actual name
```

.trip *Specify Trip Count Values*

Syntax *label .trip minimum value [, maximum value[, factor]]*

Description The **.trip** directive specifies the value of the trip count. The *trip count* indicates how many times a loop iterates. The **.trip** directive is valid within procedures only. Following are descriptions of the **.trip** directive parameters:

label The label represents the beginning of the loop. This is a required parameter.

minimum value The minimum number of times that the loop can iterate. This is a required parameter. The default is 1.

maximum value The maximum number of times that the loop can iterate. The maximum value is an optional parameter.

factor The factor used, along with *minimum value* and *maximum value*, to determine the number of times that the loop can iterate. In the following example, the loop executes some multiple of 8, between 8 and 48, times:

```
loop: .trip 8, 48, 8
```

A *factor* of 2 states that your loop always executes an even number of times allowing the compiler to unroll once; this can result in a performance increase.

The factor is optional when the maximum value is specified.

If the assembly optimizer cannot ensure that the trip count is large enough to pipeline a loop for maximum performance, a pipelined version and an unpipelined version of the same loop are generated. This makes one of the loops a *redundant loop*. The pipelined or the unpipelined loop is executed based on a comparison between the trip count and the number of iterations of the loop that can execute in parallel. If the trip count is greater or equal to the number of parallel iterations, the pipelined loop is executed; otherwise, the unpipelined loop is executed. For more information about redundant loops, see [Section 3.3](#).

You are not required to specify a **.trip** directive with every loop; however, you should use **.trip** if you know that a loop iterates some number of times. This generally means that redundant loops are not generated (unless the minimum value is really small) saving code size and execution time.

If you know that a loop always executes the same number of times whenever it is called, define maximum value (where maximum value equals minimum value) as well. The compiler may now be able to unroll your loop thereby increasing performance.

When you are compiling with the interrupt flexibility option (`--interrupt_threshold=n`), using a **.trip** maximum value allows the compiler to determine the maximum number of cycles that the loop can execute. Then, the compiler compares that value to the threshold value given by the `--interrupt_threshold` option. See [Section 2.12](#) for more information.

Example

The `.trip` directive states that the loop will execute 16, 24, 32, 40 or 48 times when the `w_vecsum` routine is called.

```
w_vecsum:  .cproc ptr_a, ptr_b, ptr_c, weight, cnt
           .reg   ai, bi, prod, scaled_prod, ci
           .no_mdep

loop:     .trip 16, 48, 8
         ldh   *ptr_a++, ai
         ldh   *ptr_b++, bi
         mpy   weight, ai, prod
         shr   prod, 15, scaled_prod
         add   scaled_prod, bi, ci
         sth   ci, *ptr_c++
        [cnt] sub   cnt, 1, cnt
        [cnt] b    loop
           .endproc
```

.volatile
Declare Memory References as Volatile

Syntax

.volatile *memref*₁ [, *memref*₂ , ...]

Description

The **.volatile** directive allows you to designate memory references as volatile. Volatile loads and stores are not deleted. Volatile loads and stores are not reordered with respect to other volatile loads and stores.

If the `.volatile` directive references a memory location that may be modified during an interrupt, compile with the `--interrupt_threshold=1` option to ensure all code referencing the volatile memory location can be interrupted.

Example

The `st` and `ld` memory references are designated as volatile.

```
.volatile st, ld

STW W, *X{st}      ; volatile store
STW U, *V
LDW *Y{ld}, Z      ; volatile load
```

4.4.1 Instructions That Are Not Allowed in Procedures

These types of instructions are not allowed in `.cproc` or `.proc topic` regions:

- The stack pointer (register B15) can be read, but it cannot be written to. Instructions that write to B15 are not allowed in a `.proc` or `.cproc` region. Stack space can be allocated by the assembly optimizer in a `.proc` or `.cproc` region for storage of temporary values. To allocate this storage area, the stack pointer is decremented on entry to the region and incremented on exit from the region. Since the stack pointer can change value on entry to the region, the assembly optimizer does not allow code that changes the stack pointer register.
- Indirect branches are not allowed in a `.proc` or `.cproc` region so that the `.proc` or `.cproc` region exit protocols cannot be bypassed. Here is an example of an indirect branch:

```
B B4    <= illegal
```

- Direct branches to labels not defined in the `.proc` or `.cproc` region are not allowed so that the `.proc` or `.cproc` region exit protocols cannot be bypassed. Here is an example of a direct branch outside of a `.proc` region:

```
.proc
...
B outside = illegal
.endproc
outside:
```

- Direct branches to the label associated with a `.proc` directive are not allowed. If you require a branch back to the start of the linear assembly function, then use the `.call` directive. Here is an example of a direct branch to the label of a `.proc` directive:

```
_func: .proc
...
B _func <= illegal
...
.endproc
```

- An `.if/.endif` loop must be entirely inside or outside of a `proc` or `.cproc` region. It is not allowed to have part of an `.if/.endif` loop inside of a `.proc` or `.cproc` region and the other part of the `.if/.endif` loop outside of the `.proc` or `.cproc` region. Here are two examples of legal `.if/.endif` loops. The first loop is outside a `.cproc` region, the second loop is inside a `.proc` region:

```
.if
.cproc
...
.endproc
.endif
```

```
.proc
.if
...
.endif
.endproc
```

Here are two examples of `.if/.endif` loops that are partly inside and partly outside of a `.cproc` or `.proc` region:

```
.if
.cproc
.endif
.endproc
```

```
.proc
.if
...
.else
.endproc
.endif
```

- The following assembly instructions cannot be used from linear assembly:
 - EFI
 - SPLOOP, SPLOOPD and SPLOOPW and all other loop-buffer related instructions

- C6700+ instructions
- ADDKSP and DP-relative addressing


4.5 Avoiding Memory Bank Conflicts With the Assembly Optimizer

The internal memory of the C6000 family varies from device to device. See the appropriate device data sheet to determine the memory spaces in your particular device. This section discusses how to write code to avoid memory bank conflicts.

Most C6000 devices use an interleaved memory bank scheme, as shown in [Figure 4-1](#). Each number in the diagram represents a byte address. A load byte (LDB) instruction from address 0 loads byte 0 in bank 0. A load halfword (LDH) from address 0 loads the halfword value in bytes 0 and 1, which are also in bank 0. A load word (LDW) from address 0 loads bytes 0 through 3 in banks 0 and 1.


Because each bank is single-ported memory, only one access to each bank is allowed per cycle. Two accesses to a single bank in a given cycle result in a memory stall that halts all pipeline operation for one cycle while the second value is read from memory. Two memory operations per cycle are allowed without any stall, as long as they do not access the same bank.

Figure 4-1. 4-Bank Interleaved Memory



For devices that have more than one memory space ([Figure 4-2](#)), an access to bank 0 in one memory space does not interfere with an access to bank 0 in another memory space, and no pipeline stall occurs.

Figure 4-2. 4-Bank Interleaved Memory With Two Memory Spaces



4.5.1 Preventing Memory Bank Conflicts

The assembly optimizer uses the assumptions that memory operations do not have bank conflicts. If it determines that two memory operations have a bank conflict on any loop iteration it does *not* schedule the operations in parallel. The assembly optimizer checks for memory bank conflicts only for those loops that it is trying to software pipeline.

The information required for memory bank analysis indicates a base, an offset, a stride, a width, and an iteration delta. The width is implicitly determined by the type of memory access (byte, halfword, word, or double word for the C6400 and C6700). The iteration delta is determined by the assembly optimizer as it constructs the schedule for the software pipeline. The base, offset, and stride are supplied by the load and store instructions and/or by the `.mptr` directive.

An LD(B/BU)(H/HU)(W) or ST(B/H/W) operation in linear assembly can have memory bank information associated with it implicitly, by using the `.mptr` directive. The `.mptr` directive associates a register with the information that allows the assembly optimizer to determine automatically whether two memory operations have a bank conflict. If the assembly optimizer determines that two memory operations have a memory bank conflict, then it does not schedule them in parallel within a software pipelined loop. The syntax is:

`.mptr variable , base + offset , stride`

For example:

```
.mptr a_0,a+0,16
.mptr a_4,a+4,16
LDW *a_0++[4], val1 ; base=a, offset=0, stride=16
LDW *a_4++[4], val2 ; base=a, offset=4, stride=16
.mptr dptr,D+0,8
LDH *dptr++, d0 ; base=D, offset=0, stride=8
LDH *dptr++, d1 ; base=D, offset=2, stride=8
LDH *dptr++, d2 ; base=D, offset=4, stride=8
LDH *dptr++, d3 ; base=D, offset=6, stride=8
```

In this example, the offset for `dptr` is updated after every memory access. The offset is updated only when the pointer is modified by a constant. This occurs for the pre/post increment/decrement addressing modes.

See [the `.mptr` topic](#) for more information.

[Example 4-6](#) shows loads and stores extracted from a loop that is being software pipelined.

Example 4-6. Load and Store Instructions That Specify Memory Bank Information

```
.mptr Ain,IN,-16
.mptr Bin,IN-4,-16

.mptr Aco,COEF,16
.mptr Bco,COEF+4,16

.mptr Aout,optr+0,4
.mptr Bout,optr+2,4

LDW *Ain--[2],Ain12 ; IN(k-I) & IN(k-I+1)
LDW *Bin--[2],Bin23 ; IN(k-I-2) & IN(k-I-1)
LDW *Ain--[2],Ain34 ; IN(k-I-4) & IN(k-I-3)
LDW *Bin--[2],Bin56 ; IN(k-I-6) & IN(k-I-5)

LDW *Bco++[2],Bco12 ; COEF(I) & COEF(I+1)
LDW *Aco++[2],Aco23 ; COEF(I+2) & COEF(I+3)
LDW *Bco++[2],Bin34 ; COEF(I+4) & COEF(I+5)
LDW *Aco++[2],Ain56 ; COEF(I+6) & COEF(I+7)

STH Assum,*Aout++[2] ; *oPtr++ = (r >> 15)
STH Bssum,*Bout++[2] ; *oPtr++ = (I >> 15)
```


4.5.2 A Dot Product Example That Avoids Memory Bank Conflicts

The C code in [Example 4-7](#) implements a dot product function. The inner loop is unrolled once to take advantage of the C6000's ability to operate on two 16-bit data items in a single 32-bit register. LDW instructions are used to load two consecutive short values. The linear assembly instructions in [Example 4-8](#) implement the dotp loop kernel. [Example 4-9](#) shows the loop kernel determined by the assembly optimizer.

For this loop kernel, there are two restrictions associated with the arrays a[] and b[]:

- Because LDW is being used, the arrays must be aligned to start on word boundaries.
- To avoid a memory bank conflict, one array must start in bank 0 and the other array in bank 2. If they start in the same bank, then a memory bank conflict occurs every cycle and the loop computes a result every two cycles instead of every cycle, due to a memory bank stall. For example:

Bank conflict:

```

    MVK    0, A0
||  MVK    8, B0
    LDW   *A0, A1

```

No bank conflict:

```

    MVK    0, A0
||  MVK    4, B0
    LDW   *A0, A1
||  LDW   *B0, B1

```

Example 4-7. C Code for Dot Product

```

int dot(short a[], short b[])
{
    int sum0 = 0, sum1 = 0, sum, I;

    for (I = 0; I < 100/2; I+= 2)
    {
        sum0 += a[i] * b[i];
        sum1 += a[i + 1] * b[i + 1];
    }
    return sum0 + sum1;
}

```

Example 4-8. Linear Assembly for Dot Product

```

_dot: .cproc a, b
      .reg sum0, sum1, I
      .reg val1, val2, prod1, prod2

      MVK    50,i ; I = 100/2
      ZERO   sum0 ; multiply result = 0
      ZERO   sum1 ; multiply result = 0

loop: .trip 50
      LDW    *a++,val1      ; load a[0-1] bank0
      LDW    *b++,val2      ; load b[0-1] bank2
      MPY    val1,val2,prod1 ; a[0] * b[0]
      MPYH   val1,val2,prod2 ; a[1] * b[1]
      ADD    prod1,sum0,sum0 ; sum0 += a[0] * b[0]
      ADD    prod2,sum1,sum1 ; sum1 += a[1] * b[1]

      [I] ADD    -1,i,i      ; I--
      [I] B      loop       ; if (!I) goto loop

      ADD    sum0,sum1,A4    ; compute final result
      .return A4
      .endproc
    
```

Example 4-9. Dot Product Software-Pipelined Kernel

```

L2:      ; PIPED LOOP KERNEL

      ADD    .L2    B7,B4,B4      ; |14| <0,7> sum0 += a[0]*b[0]
||      ADD    .L1    A5,A0,A0     ; |15| <0,7> sum1 += a[1]*b[1]
||      MPY    .M2X   B6,A4,B7     ; |12| <2,5> a[0] * b[0]
||      MPYH   .M1X   B6,A4,A5     ; |13| <2,5> a[1] * b[1]
|| [ B0] B      .S1    L2          ; |18| <5,2> if (!I) goto loop
|| [ B0] ADD    .S2    0xffffffff,B0,B0 ; |17| <6,1> I--
||      LDW    .D2T2  *B5++,B6     ; |10| <7,0> load a[0-1] bank0
||      LDW    .D1T1  *A3++,A4     ; |11| <7,0> load b[0-1] bank2
|| LDW    *B0, B1
    
```

It is not always possible to control fully how arrays and other memory objects are aligned. This is especially true when a pointer is passed into a function and that pointer may have different alignments each time the function is called. A solution to this problem is to write a dot product routine that cannot have memory hits. This would eliminate the need for the arrays to use different memory banks.

If the dot product loop kernel is unrolled once, then four LDW instructions execute in the loop kernel. Assuming that nothing is known about the bank alignment of arrays a and b (except that they are word aligned), the only safe assumptions that can be made about the array accesses are that a[0-1] cannot conflict with a[2-3] and that b[0-1] cannot conflict with b[2-3]. [Example 4-10](#) shows the unrolled loop kernel.

Example 4-10. Dot Product From [Example 4-8](#) Unrolled to Prevent Memory Bank Conflicts

```

_dotp2: .cproc   a_0, b_0
        .reg    a_4, b_4, sum0, sum1, I
        .reg    val1, val2, prod1, prod2

        ADD     4,a_0,a_4
        ADD     4,b_0,b_4
        MVK     25,i      ; I = 100/4
        ZERO    sum0     ; multiply result = 0
        ZERO    sum1     ; multiply result = 0

        .mptr   a_0,a+0,8
        .mptr   a_4,a+4,8
        .mptr   b_0,b+0,8
        .mptr   b_4,b+4,8

loop:   .trip    25

        LDW     *a_0++[2],val1 ; load a[0-1] bankx
        LDW     *b_0++[2],val2 ; load b[0-1] banky
        MPY     val1,val2,prod1 ; a[0] * b[0]
        MPYH    val1,val2,prod2 ; a[1] * b[1]
        ADD     prod1,sum0,sum0 ; sum0 += a[0] * b[0]
        ADD     prod2,sum1,sum1 ; sum1 += a[1] * b[1]

        LDW     *a_4++[2],val1 ; load a[2-3] bankx+2
        LDW     *b_4++[2],val2 ; load b[2-3] banky+2
        MPY     val1,val2,prod1 ; a[2] * b[2]
        MPYH    val1,val2,prod2 ; a[3] * b[3]
        ADD     prod1,sum0,sum0 ; sum0 += a[2] * b[2]
        ADD     prod2,sum1,sum1 ; sum1 += a[3] * b[3]

        [I] ADD     -1,i,i      ; I--
        [I] B      loop        ; if (!0) goto loop

        ADD     sum0,sum1,A4    ; compute final result
        .return A4
        .endproc

```

The goal is to find a software pipeline in which the following instructions are in parallel:

```
LDW *a0++[2],val1 ; load a[0-1] bankx
|| LDW *a2++[2],val2 ; load a[2-3] bankx+2
LDW *b0++[2],val1 ; load b[0-1] banky
|| LDW *b2++[2],val2 ; load b[2-3] banky+2
```

Example 4-11. Unrolled Dot Product Kernel From Example 4-9

```
L2:      ; PIPED LOOP KERNEL

[ B1] SUB  .S2   B1,1,B1      ; <0,8>
||      ADD  .L2   B9,B5,B9      ; |21| <0,8> ^ sum0 += a[0] * b[0]
||      ADD  .L1   A6,A0,A0      ; |22| <0,8> ^ sum1 += a[1] * b[1]
||      MPY  .M2X  B8,A4,B9      ; |19| <1,6> a[0] * b[0]
||      MPYH .M1X  B8,A4,A6      ; |20| <1,6> a[1] * b[1]
|| [ B0] B    .S1   L2          ; |32| <2,4> if (!I) goto loop
|| [ B1] LDW  .D1T1 *A3++(8),A4   ; |24| <3,2> load a[2-3] bankx+2
|| [ A1] LDW  .D2T2 *B6++(8),B8   ; |17| <4,0> load a[0-1] bankx

[ A1] SUB  .S1   A1,1,A1      ; <0,9>
||      ADD  .L2   B5,B9,B5      ; |28| <0,9> ^ sum0 += a[2] * b[2]
||      ADD  .L1   A6,A0,A0      ; |29| <0,9> ^ sum1 += a[3] * b[3]
||      MPY  .M2X  A4,B7,B5      ; |26| <1,7> a[2] * b[2]
||      MPYH .M1X  A4,B7,A6      ; |27| <1,7> a[3] * b[3]
|| [ B0] ADD  .S2   -1,B0,B0      ; |31| <3,3> I--
|| [ A1] LDW  .D2T2 *B4++(8),B7   ; |25| <4,1> load b[2-3] banky+2
|| [ A1] LDW  .D1T1 *A5++(8),A4   ; |18| <4,1> load b[0-1] banky
```

Without the `.mptr` directives in Example 4-10, the loads of `a[0-1]` and `b[0-1]` are scheduled in parallel, and the loads of `a[2-3]` and `b[2-3]` might be scheduled in parallel. This results in a 50% chance that a memory conflict will occur on every cycle. However, the loop kernel shown in Example 4-11 can never have a memory bank conflict.

In Example 4-8, if `.mptr` directives had been used to specify that `a` and `b` point to different bases, then the assembly optimizer would never find a schedule for a 1-cycle loop kernel, because there would always be a memory bank conflict. However, it would find a schedule for a 2-cycle loop kernel.

4.5.3 Memory Bank Conflicts for Indexed Pointers

When determining memory bank conflicts for indexed memory accesses, it is sometimes necessary to specify that a pair of memory accesses always conflict, or that they never conflict. This can be accomplished by using the `.mptr` directive with a stride of 0.

A stride of 0 indicates that there is a constant relation between the memory accesses regardless of the iteration delta. Essentially, only the base, offset, and width are used by the assembly optimizer to determine a memory bank conflict. Recall that the stride is optional and defaults to 0.

In Example 4-12, the `.mptr` directive is used to specify which memory accesses conflict and which never conflict.

Example 4-12. Using `.mptr` for Indexed Pointers

```
.mptr a,RS
.mptr b,RS
.mptr c,XY
.mptr d,XY+2
LDW  *a++[i0a],A0 ; a and b always conflict with each other
LDW  *b++[i0b],B0 ;
STH  A1,*c++[i1a] ; c and d never conflict with each other
STH  B2,*d++[i1b] ;
```

4.5.4 Memory Bank Conflict Algorithm

The assembly optimizer uses the following process to determine if two memory access instructions might have a memory bank conflict:

1. If either access does not have memory bank information, then they do not conflict.
2. If both accesses do not have the same base, then they conflict.
3. The offset, stride, access width, and iteration delta are used to determine if a memory bank conflict will occur. The assembly optimizer uses a straightforward analysis of the access patterns and determines if they ever access the same relative bank. The stride and offset values are always expressed in bytes.

The iteration delta is the difference in the loop iterations of the memory references being scheduled in the software pipeline. For example, given three instructions A, B, C and a software pipeline with a single-cycle kernel, then A and C have an iteration delta of 2:

```

A
B  A
C  B  A
   C  B
    C
  
```

4.6 Memory Alias Disambiguation

Memory aliasing occurs when two instructions can access the same memory location. Such memory references are called ambiguous. Memory alias disambiguation is the process of determining when such ambiguity is not possible. When you cannot determine whether two memory references are ambiguous, you presume they are ambiguous. This is the same as saying the two instructions have a memory dependence between them.

Dependencies between instructions constrain the instruction schedule, including the software pipeline schedule. In general, the fewer the Dependencies, the greater freedom you have in choosing a schedule and the better the final schedule performs.

4.6.1 How the Assembly Optimizer Handles Memory References (Default)

The assembly optimizer assumes memory references are aliased, unless it can prove otherwise.

Because alias analysis is very limited in the assembly optimizer, this presumption is often overly conservative. In such cases, the extra instruction Dependencies, due to the presumed memory aliases, can cause the assembly optimizer to emit instruction schedules that have less parallelism and do not perform well. To handle these cases, the assembly optimizer provides one option and two directives.

4.6.2 Using the `--no_bad_aliases` Option to Handle Memory References

In the assembly optimizer, the `--no_bad_aliases` option means no memory references ever depend on each other. The `--no_bad_aliases` option does not mean the same thing to the C/C++ compiler. The C/C++ compiler interprets the `--no_bad_aliases` switch to indicate several specific cases of memory aliasing are guaranteed not to occur. For more information about using the `--no_bad_aliases` option, see [Section 3.10.2](#).

4.6.3 Using the `.no_mdep` Directive

You can specify the `.no_mdep` directive anywhere in a `.(c)proc` function. Whenever it is used, you guarantee that no memory Dependencies occur within that function.

Memory Dependency Exception

NOTE: For both of these methods, `--no_bad_aliases` and `.no_mdep`, the assembly optimizer recognizes any memory Dependencies you point out with the `.mdep` directive.

4.6.4 Using the .mdep Directive to Identify Specific Memory Dependencies

You can use the .mdep directive to identify specific memory Dependencies by annotating each memory reference with a name, and using those names with the .mdep directive to indicate the actual dependence. Annotating a memory reference requires adding information right next to the memory reference in the assembly stream. Include the following immediately after a memory reference:

```
{ memref }
```

The *memref* has the same syntax restrictions as any assembly symbol. (For more information about symbols, refer to the *TMS320C6000 Assembly Language Tools User's Guide*.) It is in the same name space as the symbolic registers. You cannot use the same name for a symbolic register and annotating a memory reference.

Example 4-13. Annotating a Memory Reference

```
LDW    *p1++ {ld1}, inp1 ;name memory reference "ld1"
;other code ...
STW    outp2, *p2++ {st1} ;name memory reference "st1"

*<The directive to indicate...:

.mdep ld1, st1 <<bold>>
```

The directive to indicate a specific memory dependence in the previous example is as follows:

```
.mdep ld1, st1
```

This means that whenever ld1 accesses memory at location X, some later time in code execution, st1 may also access location X. This is equivalent to adding a dependence between these two instructions. In terms of the software pipeline, these two instructions must remain in the same order. The ld1 reference must always occur before the st1 reference; the instructions cannot even be scheduled in parallel.

It is important to note the directional sense of the directive from ld1 to st1. The opposite, from st1 to ld1, is not implied. In terms of the software pipeline, while every ld1 must occur before every st1, it is still legal to schedule the ld1 from iteration n+1 before the st1 from iteration n.

[Example 4-14](#) is a picture of the software pipeline with the instructions from two different iterations in different columns. In the actual instruction sequence, instructions on the same horizontal line are in parallel.

Example 4-14. Software Pipeline Using .mdep ld1, st1

```
iteration n          iteration n+1
-----
LDW { ld1 }
...
STW { st1 }

LDW { ld1 }
...
STW { st1 }

*<If that schedule...>

.mdep st1, ld1
```

If that schedule does not work because the iteration n st1 might write a value the iteration n+1 ld1 should read, then you must note a dependence relationship from st1 to ld1.

```
.mdep st1, ld1
```

Both directives together force the software pipeline shown in [Example 4-15](#).

Example 4-15. Software Pipeline Using .mdep st1, ld1 and .mdep ld1, st1

iteration n	iteration n+1
-----	-----
LDW { ld1 }	
...	
STW { st1 }	
	LDW { ld1 }
	...
	STW { st1 }
 <Indexed addressing,...>	
.mdep ld1, st1	
.mdep st1, ld1	

Indexed addressing, `*+base[index]`, is a good example of an addressing mode where you typically do not know anything about the relative sequence of the memory accesses, except they sometimes access the same location. To correctly model this case, you need to note the dependence relation in both directions, and you need to use both directives.

```
.mdep ld1, st1 .mdep st1, ld1
```

4.6.5 Memory Alias Examples

Following are memory alias examples that use the `.mdep` and `.no_mdep` directives.

- **Example 1**

The `.mdep r1, r2` directive declares that LDW must be before STW. In this case, `src` and `dst` might point to the same array.

```
fn:      .cproc      dst, src, cnt
        .reg        tmp
        .no_mdep
        .mdep       r1, r2

        LDW        *src{r1}, tmp
        STW        cnt, *dst{r2}

        .return    tmp
        .endproc
```

- **Example 2**

Here, `.mdep r2, r1` indicates that STW must occur before LDW. Since STW is after LDW in the code, the dependence relation is across loop iterations. The STW instruction writes a value that may be read by the LDW instruction on the next iteration. In this case, a 6-cycle recurrence is created.

```
fn:      .cproc      dst, src, cnt
        .reg        tmp
        .no_mdep
        .mdep       r2, r1

LOOP:    .trip      100
        LDW        *src++{r1}, tmp
        STW        tmp, *dst++{r2}
[ cnt ] SUB        cnt, 1, cnt
[ cnt ] B          LOOP

        .endproc
```

Memory Dependence/Bank Conflict

NOTE: Do not confuse memory alias disambiguation with the handling of memory bank conflicts. These may seem similar because they each deal with memory references and the effect of those memory references on the instruction schedule. Alias disambiguation is a correctness issue, bank conflicts are a performance issue. A memory dependence has a much broader impact on the instruction schedule than a bank conflict. It is best to keep these two topics separate.

Volatile References

NOTE: For volatile references, use `.volatile` rather than `.mdep`.

Linking C/C++ Code

The C/C++ compiler and assembly language tools provide two methods for linking your programs:

- You can compile individual modules and link them together. This method is especially useful when you have multiple source files.
- You can compile and link in one step. This method is useful when you have a single source module.

This chapter describes how to invoke the linker with each method. It also discusses special requirements of linking C/C++ code, including the run-time-support libraries, specifying the type of initialization, and allocating the program into memory. For a complete description of the linker, see the *TMS320C6000 Assembly Language Tools User's Guide*.

Topic	Page
5.1 Invoking the Linker Through the Compiler (-z Option)	138
5.2 Linker Code Optimizations	140
5.3 Controlling the Linking Process	141

5.1 Invoking the Linker Through the Compiler (-z Option)

This section explains how to invoke the linker after you have compiled and assembled your programs: as a separate step or as part of the compile step.

5.1.1 Invoking the Linker Separately

This is the general syntax for linking C/C++ programs as a separate step:

```
cl6x --run_linker {--rom_model | --ram_model} filenames
      [options] [--output_file= name.out] --library= library [Ink.cmd]
```

cl6x --run_linker	The command that invokes the linker.
--rom_model --ram_model	Options that tell the linker to use special conventions defined by the C/C++ environment. When you use <code>cl6x --run_linker</code> , you must use --rom_model or --ram_model . The <code>--rom_model</code> option uses automatic variable initialization at run time; the <code>--ram_model</code> option uses variable initialization at load time.
<i>filenames</i>	Names of object files, linker command files, or archive libraries. The default extension for all input files is <code>.obj</code> ; 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 <code>a.out</code> , unless you use the <code>--output_file</code> option to name the output file.
<i>options</i>	Options affect how the linker handles your object files. Linker options can only appear after the --run_linker option on the command line, but otherwise may be in any order. (Options are discussed in detail in the <i>TMS320C6000 Assembly Language Tools User's Guide</i> .)
--output_file= name.out	Names the output file.
--library= library	Identifies the appropriate archive library containing C/C++ run-time-support and floating-point math functions, or linker command files. If you are linking C/C++ code, you must use a run-time-support library. You can use the libraries included with the compiler, or you can create your own run-time-support library. If you have specified a run-time-support library in a linker command file, you do not need this parameter. The <code>--library</code> option's short form is <code>-l</code> .
<i>Ink.cmd</i>	Contains options, filenames, directives, or commands for the linker.

When you specify a library as linker input, the linker includes and links only those library members that resolve undefined references. The linker uses a default allocation algorithm to allocate your program into memory. You can use the `MEMORY` and `SECTIONS` directives in the linker command file to customize the allocation process. For information, see the *TMS320C6000 Assembly Language Tools User's Guide*.

You can link a C/C++ program consisting of object files `prog1.obj`, `prog2.obj`, and `prog3.obj`, with an executable object file filename of `prog.out` with the command:

```
cl6x --run_linker --rom_model prog1 prog2 prog3 --output_file=prog.out
      --library=rts6200.lib
```

5.1.2 Invoking the Linker as Part of the Compile Step

This is the general syntax for linking C/C++ programs as part of the compile step:

```
cl6xfilenames [options] --run_linker {--rom_model | --ram_model} filenames
      [options] [--output_file= name.out] --library= library [lnk.cmd]
```

The `--run_linker` option divides the command line into the compiler options (the options before `--run_linker`) and the linker options (the options following `--run_linker`). The `--run_linker` option must follow all source files and compiler options on the command line.

All arguments that follow `--run_linker` on the command line are passed to the linker. These arguments can be linker command files, additional object files, linker options, or libraries. These arguments are the same as described in [Section 5.1.1](#).

All arguments that precede `--run_linker` on the command line are compiler arguments. These arguments can be C/C++ source files, assembly files, linear assembly files, or compiler options. These arguments are described in [Section 2.2](#).

You can compile and link a C/C++ program consisting of object files `prog1.c`, `prog2.c`, and `prog3.c`, with an executable object file filename of `prog.out` with the command:

```
cl6x prog1.c prog2.c prog3.c --run_linker --rom_model --output_file=prog.out --library=rts6200.lib
```

NOTE: Order of Processing Arguments in the Linker

The order in which the linker processes arguments is important. The compiler passes arguments to the linker in the following order:

1. Object filenames from the command line
 2. Arguments following the `--run_linker` option on the command line
 3. Arguments following the `--run_linker` option from the `C6X_C_OPTION` environment variable
-

5.1.3 Disabling the Linker (--compile_only Compiler Option)

You can override the `--run_linker` option by using the `--compile_only` compiler option. The `--run_linker` option's short form is `-z` and the `--compile_only` option's short form is `-c`.

The `--compile_only` option is especially helpful if you specify the `--run_linker` option in the `C6X_C_OPTION` environment variable and want to selectively disable linking with the `--compile_only` option on the command line.

5.2 Linker Code Optimizations

5.2.1 Generating Function Subsections (`--gen_func_subsections` Compiler Option)

When the linker places code into an executable file, it allocates all the functions in a single source file as a group. This means that if any function in a file needs to be linked into an executable, then all the functions in the file are linked in. This can be undesirable if a file contains many functions and only a few are required for an executable.

This situation may exist in libraries where a single file contains multiple functions, but the application only needs a subset of those functions. An example is a library .obj file that contains a signed divide routine and an unsigned divide routine. If the application requires only signed division, then only the signed divide routine is required for linking. By default, both the signed and unsigned routines are linked in since they exist in the same .obj file.

The `--gen_func_subsections` compiler option remedies this problem by placing each function in a file in its own subsection. Thus, only the functions that are referenced in the application are linked into the final executable. This can result in an overall code size reduction.

However, be aware that using the `--gen_func_subsections` compiler option can result in overall code size growth if all or nearly all functions are being referenced. This is because any section containing code must be aligned to a 32-byte boundary to support the C6000 branching mechanism. When the `--gen_func_subsections` option is not used, all functions in a source file are usually placed in a common section which is aligned. When `--gen_func_subsections` is used, each function defined in a source file is placed in a unique section. Each of the unique sections requires alignment. If all the functions in the file are required for linking, code size may increase due to the additional alignment padding for the individual subsections.

Thus, the `--gen_func_subsections` compiler option is advantageous for use with libraries where normally only a limited number of the functions in a file are used in any one executable.

The alternative to the `--gen_func_subsections` option is to place each function in its own file.

5.2.2 Conditional Linking

The conditional linking paradigm is different under COFF compared to ELF. In COFF, you must mark a section with the `.clink` directive to make it eligible for removal during conditional linking. In ELF, all sections are considered eligible for removal through conditional linking. Sections are not removed if they are referenced or if they are marked with the `.retain` directive.

Under COFF, when you compile with the `-gen_func_subsections` option, in addition to placing each function in a separate subsection, the compiler also annotates that subsection with the conditional linking directive, `.clink`. This directive marks the section as a candidate to be removed if it is not referenced by any other section in the program. The compiler does not place a `.clink` directive in a subsection for a trap or interrupt function, as these may be needed by a program even though there is no symbolic reference to them anywhere in the program.

Under COFF, if a section that has been marked for conditional linking is never referenced by any other section in the program, that section is removed from the program. Under ELF, a section that is never referenced by any other section in the program is removed from the program automatically, unless it is marked with `.retain`. Conditional linking is disabled when performing a partial link or when relocation information is kept with the output of the link. Conditional linking can also be disabled with the `--disable_clink` linker option.

5.3 Controlling the Linking Process

Regardless of the method you choose for invoking the linker, special requirements apply when linking C/C++ programs. You must:

- Include the compiler's run-time-support library
- Specify the type of boot-time initialization
- Determine how you want to allocate your program into memory

This section discusses how these factors are controlled and provides an example of the standard default linker command file.

For more information about how to operate the linker, see the linker description in the *TMS320C6000 Assembly Language Tools User's Guide*

5.3.1 Including the Run-Time-Support Library

You must link all C/C++ programs with a run-time-support library. The library contains standard C/C++ functions as well as functions used by the compiler to manage the C/C++ environment. The following sections describe two methods for including the run-time-support library.

5.3.1.1 Automatic Run-Time-Support Library Selection

If the `--rom_model` or `--ram_model` option is specified during the linker and the entry point for the program (normally `c_int00`) is not resolved by any specified object file or library, the linker attempts to automatically include the best compatible run-time-support library for your program. The chosen run-time-support library is linked in after any other libraries specified with the `--library` option on the command line. Alternatively, you can force the linker to choose an appropriate run-time-support library by specifying "libc.a" as an argument to the `--library` option, or when specifying the run-time-support library name explicitly in a linker command file.

The automatic selection of a run-time-support library can be disabled with the `--disable_auto_rts` option.

If the `--issue_remarks` option is specified before the `--run_linker` option during the linker, a remark is generated indicating which run-time support library was linked in. If a different run-time-support library is desired, you must specify the name of the desired run-time-support library using the `--library` option and in your linker command files when necessary.

Example 5-1. Using the `--issue_remarks` Option

```
cl6x --silicon_version=6400+ --issue_remarks main.c --run_linker --rom_model
<Linking>
remark: linking in "libc.a"
remark: linking in "rts64plus.lib" in place of "libc.a"
```

5.3.1.2 Manual Run-Time-Support Library Selection

You should use the `--library` linker option to specify which C6000 run-time-support library to use. The `--library` option also tells the linker to look at the `--search_path` options and then the `C6X_C_DIR` environment variable to find an archive path or object file. To use the `--library` linker option, type on the command line:

```
cl6x --run_linker {--rom_model | --ram_model} filenames --library= libraryname
```

5.3.1.3 Library Order for Searching for Symbols

Generally, you should specify the run-time-support library as the last name on the command line because the linker searches libraries for unresolved references in the order that files are specified on the command line. If any object files follow a library, references from those object files to that library are not resolved. You can use the `--reread_libs` option to force the linker to reread all libraries until references are resolved. Whenever you specify a library as linker input, the linker includes and links only those library members that resolve undefined references.

By default, if a library introduces an unresolved reference and multiple libraries have a definition for it, then the definition from the same library that introduced the unresolved reference is used. Use the `--priority` option if you want the linker to use the definition from the first library on the command line that contains the definition.

5.3.2 Run-Time Initialization

C/C++ programs require initialization of the run-time environment before execution of the program itself may begin. This initialization is performed by a *bootstrap routine*. This routine is responsible for creating the stack, initializing global variables, and calling the function `main`. The bootstrap routine should be the entry point for the program, and it typically should be the RESET interrupt handler. The bootstrap routine is responsible for the following tasks:

1. Set up the stack by initializing SP
2. Set up the data page pointer DP (for architectures that have one)
3. Set configuration registers
4. Process the `.cinit` table to autoinitialize global variables (when using the `--rom_model` option)
5. Process the `.pinit` table to construct global C++ objects.
6. Call the function `main` with appropriate arguments
7. Call `exit` when `main` returns

When you compile a C program and use `--rom_model` or `--ram_model`, the linker looks for a bootstrap routine named `_c_int00`. The run-time support library provides a sample `_c_int00` in `boot.obj`, which performs the required tasks. If you use the run-time support's bootstrap routine, you should set `_c_int00` as the entry point.

A sample bootstrap routine is `_c_int00`, provided in `boot.obj` in the run-time support object libraries. The entry point is usually set to the starting address of the bootstrap routine.

NOTE: The `_c_int00` Symbol

If you use the `--ram_model` or `--rom_model` link option, `_c_int00` is automatically defined as the entry point for the program.

5.3.3 Global Object Constructors

Global C++ variables that have constructors and destructors require their constructors to be called during program initialization and their destructors to be called during program termination. The C++ compiler produces a table of constructors to be called at startup.

Constructors for global objects from a single module are invoked in the order declared in the source code, but the relative order of objects from different object files is unspecified.

Global constructors are called after initialization of other global variables and before the function `main` is called. Global destructors are invoked during the function exit, similar to functions registered through `atexit`.

[Section 7.8.1.3](#) discusses the format of the global constructor table for COFFABI mode and [Section 7.8.1.3](#) for EABI mode..

5.3.4 Specifying the Type of Global Variable Initialization

The C/C++ compiler produces data tables for initializing global variables. [Section 7.8.3](#) discusses the format of these initialization tables for COFFABI. [Section 7.8.2.4](#) discusses the format of these initialization tables for EABI. The initialization tables are used in one of the following ways:

- Global variables are initialized at *run time*. Use the `--rom_model` linker option (see [Section 7.8.1.1](#)).
- Global variables are initialized at *load time*. Use the `--ram_model` linker option (see [Section 7.8.1.2](#)).

When you link a C/C++ program, you must use either the `--rom_model` or `--ram_model` option. These options tell the linker to select initialization at run time or load time.

When you compile and link programs, the `--rom_model` option is the default. If used, the `--rom_model` option must follow the `--run_linker` option (see [Section 5.1](#)). The following list outlines the linking conventions for COFFABI used with `--rom_model` or `--ram_model`:

- The symbol `_c_int00` is defined as the program entry point; it identifies the beginning of the C/C++ boot routine in `boot.obj`. When you use `--rom_model` or `--ram_model`, `_c_int00` is automatically referenced, ensuring that `boot.obj` is automatically linked in from the run-time-support library.
- The initialization output section is padded with a termination record so that the loader (load-time initialization) or the boot routine (run-time initialization) knows when to stop reading the initialization tables.
- When initializing at load time (the `--ram_model` option), the following occur:
 - The linker sets the initialization table symbol to `-1`. This indicates that the initialization tables are not in memory, so no initialization is performed at run time.
 - The `STYP_COPY` flag is set in the initialization table section header. `STYP_COPY` is the special attribute that tells the loader to perform autoinitialization directly and not to load the initialization table into memory. The linker does not allocate space in memory for the initialization table.
- When autoinitializing at run time (`--rom_model` option), the linker defines the initialization table symbol as the starting address of the initialization table. The boot routine uses this symbol as the starting point for autoinitialization.

For details on linking conventions for EABI used with `--rom_model` and `--ram_model`, see [Section 7.8.2.3](#) and [Section 7.8.2.5](#), respectively.

5.3.5 Specifying Where to Allocate Sections in Memory

The compiler produces relocatable blocks of code and data. These blocks, called *sections*, are allocated in memory in a variety of ways to conform to a variety of system configurations.

The compiler creates two basic kinds of sections: initialized and uninitialized. [Table 5-1](#) summarizes the initialized sections created under the COFF ABI mode. [Table 5-2](#) summarizes the initialized sections created under the EABI mode. [Table 5-3](#) summarizes the uninitialized sections. Be aware that the COFF ABI `.cinit` and `.pinit` (`.init_array` in EABI) tables have different formats in EABI.

Table 5-1. Initialized Sections Created by the Compiler for COFFABI

Name	Contents
<code>.args</code>	Command argument for host-based loader; read-only (see the <code>--arg_size</code> option)
<code>.cinit</code>	Tables for explicitly initialized global and static variables
<code>.const</code>	Global and static const variables that are explicitly initialized and contain string literals
<code>.pinit</code>	Table of constructors to be called at startup
<code>.ppdata</code>	Data tables for compiler-based profiling (see the <code>--gen_profile_info</code> option)
<code>.ppinfo</code>	Correlation tables for compiler-based profiling (see the <code>--gen_profile_info</code> option)
<code>.switch</code>	Jump tables for large switch statements
<code>.text</code>	Executable code and constants

Table 5-2. Initialized Sections Created by the Compiler for EABI

Name	Contents
.args	Command argument for host-based loader; read-only (see the --arg_size option)
.binit	Boot time copy tables (See the <i>TMS320C6000 Assembly Language Tools User's Guide</i> for information on BINIT in linker command files.)
.cinit	In EABI mode, the compiler does not generate a .cinit section. However, when the --rom_mode linker option is specified, the linker creates this section, which contains tables for explicitly initialized global and static variables.
.const	Far, const global and static variables, and string constants
.c6xabi.exidx	Index table for exception handling; read-only (see --exceptions option)
.c6xabi.exstab	Unwinded instructions for exception handling; read-only (see --exceptions option)
.fardata	Far non-const global and static variables that are explicitly initialized
.init_array	Table of constructors to be called at startup
.name.load	Compressed image of section <i>name</i> ; read-only (See the <i>TMS320C6000 Assembly Language Tools User's Guide</i> for information on copy tables.)
.neardata	Near non-const global and static variables that are explicitly initialized
.ppdata	Data tables for compiler-based profiling (see the --gen_profile_info option)
.ppinfo	Correlation tables for compiler-based profiling (see the --gen_profile_info option)
.rodata	Global and static variables that have near and const qualifiers
.switch	Jump tables for large switch statements
.text	Executable code and constants

Table 5-3. Uninitialized Sections Created by the Compiler for Both ABIs

Name	Contents
.bss	Global and static variables
.far	Global and static variables declared far
.stack	Stack
.systemem	Memory for malloc functions (heap)

When you link your program, you must specify where to allocate the sections in memory. In general, initialized sections are linked into ROM or RAM; uninitialized sections are linked into RAM. With the exception of code sections, the initialized and uninitialized sections created by the compiler cannot be allocated into internal program memory. See [Section 7.1.1](#) for a complete description of how the compiler uses these sections.

The linker provides MEMORY and SECTIONS directives for allocating sections. For more information about allocating sections into memory, see the *TMS320C6000 Assembly Language Tools User's Guide*.

5.3.6 A Sample Linker Command File

[Example 5-2](#) shows a typical linker command file that links a C program. The command file in this example is named `Ink.cmd` and lists several linker options:

--rom_model	Tells the linker to use autoinitialization at run time.
--heap_size	Tells the linker to set the C heap size at 0x2000 bytes.
--stack_size	Tells the linker to set the stack size to 0x0100 bytes.
--library	Tells the linker to use an archive library file, <code>rts6200.lib</code> , for input.

To link the program, use the following syntax:

```
cl6x --run_linker object_file(s) --output_file= outfile --map_file= mapfile Ink.cmd
```

The `MEMORY` and possibly the `SECTIONS` directives, might require modification to work with your system. See the *TMS320C6000 Assembly Language Tools User's Guide* for more information on these directives.

Example 5-2. Linker Command File

```
--rom_model
--heap_size=0x2000
--stack_size=0x0100
--library=rts6200.lib

MEMORY
{
    VECS:   o = 0x00000000      l = 0x000000400 /* reset & interrupt vectors */
    PMEM:   o = 0x00000400     l = 0x00000FC00 /* intended for initialization */
    BMEM:   o = 0x80000000     l = 0x000010000 /* .bss, .system, .stack, .cinit */
}

SECTIONS
{
    vectors      >      VECS
    .text        >      PMEM
    .data        >      BMEM
    .stack       >      BMEM
    .bss         >      BMEM
    .systemem    >      BMEM
    .cinit       >      BMEM
    .const       >      BMEM
    .cio         >      BMEM
    .far         >      BMEM
}
```

TMS320C6000C/C++ Language Implementation

The C/C++ compiler supports the C/C++ language standard that was developed by a committee of the American National Standards Institute (ANSI) and subsequently adopted by the International Standards Organization (ISO).

The C++ language supported by the C6000 is defined by the ANSI/ISO/IEC 14882:1998 standard with certain exceptions.

Topic	Page
6.1 Characteristics of TMS320C6000 C	147
6.2 Characteristics of TMS320C6000 C++	147
6.3 Using MISRA-C:2004	148
6.4 Data Types	149
6.5 Keywords	150
6.6 C++ Exception Handling	155
6.7 Register Variables and Parameters	155
6.8 The asm Statement	156
6.9 Pragma Directives	157
6.10 The _Pragma Operator	173
6.11 Application Binary Interface	174
6.12 Object File Symbol Naming Conventions (Linknames)	174
6.13 Initializing Static and Global Variables in COFF ABI Mode	175
6.14 Changing the ANSI/ISO C Language Mode	176
6.15 GNU Language Extensions	178

6.1 Characteristics of TMS320C6000 C

The compiler supports the C language as defined by ISO/IEC 9899:1990, which is equivalent to American National Standard for Information Systems-Programming Language C X3.159-1989 standard, commonly referred to as C89, published by the American National Standards Institute. The compiler can also accept many of the language extensions found in the GNU C compiler (see [Section 6.15](#)). The compiler does not support C99.

The ANSI/ISO standard identifies some features of the C language that are affected by characteristics of the target processor, run-time environment, or host environment. For reasons of efficiency or practicality, this set of features can differ among standard compilers.

Unsupported features of the C library are:

- The run-time library has minimal support for wide and multi-byte characters. The type `wchar_t` is implemented as `int`. The wide character set is equivalent to the set of values of type `char`. The library includes the header files `<wchar.h>` and `<wctype.h>`, but does not include all the functions specified in the standard. So-called multi-byte characters are limited to single characters. There are no shift states. The mapping between multi-byte characters and wide characters is simple equivalence; that is, each wide character maps to and from exactly a single multi-byte character having the same value.
- The run-time library includes the header file `<locale.h>`, but with a minimal implementation. The only supported locale is the C locale. That is, library behavior that is specified to vary by locale is hard-coded to the behavior of the C locale, and attempting to install a different locale by way of a call to `setlocale()` will return `NULL`.

6.2 Characteristics of TMS320C6000 C++

The C6000 compiler supports C++ as defined in the ANSI/ISO/IEC 14882:1998 standard, including these features:

- Complete C++ standard library support, with exceptions noted below.
- Templates
- Exceptions, which are enabled with the `--exceptions` option; see [Section 6.6](#).
- Run-time type information (RTTI), which can be enabled with the `--rtti` compiler option.

The *exceptions* to the standard are as follows:

- The compiler does not support embedded C++ run-time-support libraries.
- The library supports wide chars (`wchar_t`), in that template functions and classes that are defined for `char` are also available for `wchar_t`. For example, wide char stream classes `wios`, `wiostream`, `wstreambuf` and so on (corresponding to char classes `ios`, `iostream`, `streambuf`) are implemented. However, there is no low-level file I/O for wide chars. Also, the C library interface to wide char support (through the C++ headers `<wchar>` and `<wctype>`) is limited as described above in the C library.
- For COFF ABI only: If the definition of an inline function contains a static variable, and it appears in multiple compilation units (usually because it's a member function of a class defined in a header file), the compiler generates multiple copies of the static variable rather than resolving them to a single definition. The compiler emits a warning (`#1369`) in such cases.
- No support for `bad_cast` or `bad_type_id` is included in the `typeinfo` header.
- Two-phase name binding in templates, as described in `[tesp.res]` and `[temp.dep]` of the standard, is not implemented.
- The `export` keyword for templates is not implemented.
- A `typedef` of a function type cannot include member function cv-qualifiers.
- A partial specialization of a class member template cannot be added outside of the class definition.

6.3 Using MISRA-C:2004

You can alter your code to work with the MISRA-C:2004 rules. The following enable/disable the rules:

- The `--check_misra` option enables checking of the specified MISRA-C:2004 rules.
- The `CHECK_MISRA` pragma enables/disables MISRA-C:2004 rules at the source level. This pragma is equivalent to using the `--check_misra` option. See [Section 6.9.1](#).
- `RESET_MISRA` pragma resets the specified MISRA-C:2004 rules to the state they were before any `CHECK_MISRA` pragmas were processed. See [Section 6.9.25](#).

The syntax of the option and pragmas is:

```
--check_misra={all|required|advisory|none|rulespec}
#pragma CHECK_MISRA ("{all|required|advisory|none|rulespec}");
#pragma RESET_MISRA ("{all|required|advisory|rulespec}");
```

The *rulespec* parameter is a comma-separated list of these specifiers:

- `[-]X` Enable (or disable) all rules in topic X.
- `[-]X-Z` Enable (or disable) all rules in topics X through Z.
- `[-]X.A` Enable (or disable) rule A in topic X.
- `[-]X.A-C` Enable (or disable) rules A through C in topic X.

Example: `--check_misra=1-5,-1.1,7.2-4`

- Checks topics 1 through 5
- Disables rule 1.1 (all other rules from topic 1 remain enabled)
- Checks rules 2 through 4 in topic 7

Two options control the severity of certain MISRA-C:2004 rules:

- The `--misra_required` option sets the diagnostic severity for required MISRA-C:2004 rules.
- The `--misra_advisory` option sets the diagnostic severity for advisory MISRA-C:2004 rules.

The syntax for these options is:

```
--misra_advisory={error|warning|remark|suppress}
--misra_required={error|warning|remark|suppress}
```

6.4 Data Types

Table 6-1 lists the size, representation, and range of each scalar data type for the C6000 compiler for COFF ABI. See Table 6-2 for the EABI data types. Many of the range values are available as standard macros in the header file limits.h.

Table 6-1. TMS320C6000 C/C++ COFF ABI Data Types

Type	Size	Representation	Range	
			Minimum	Maximum
char, signed char	8 bits	ASCII	-128	127
unsigned char	8 bits	ASCII	0	255
short	16 bits	2s complement	-32 768	32 767
unsigned short	16 bits	Binary	0	65 535
int, signed int	32 bits	2s complement	-2 147 483 648	2 147 483 647
unsigned int	32 bits	Binary	0	4 294 967 295
long, signed long	40 bits	2s complement	-549 755 813 888	549 755 813 887
unsigned long	40 bits	Binary	0	1 099 511 627 775
__int40_t	40 bits	2s complement	-549 755 813 888	549 755 813 887
unsigned __int40_t	40 bits	Binary	0	1 099 511 627 775
long long, signed long long	64 bits	2s complement	-9 223 372 036 854 775 808	9 223 372 036 854 775 807
unsigned long long	64 bits	Binary	0	18 446 744 073 709 551 615
enum	32 bits	2s complement	-2 147 483 648	2 147 483 647
float	32 bits	IEEE 32-bit	1.175 494e-38 ⁽¹⁾	3.40 282 346e+38
double	64 bits	IEEE 64-bit	2.22 507 385e-308 ⁽¹⁾	1.79 769 313e+308
long double	64 bits	IEEE 64-bit	2.22 507 385e-308 ⁽¹⁾	1.79 769 313e+308
pointers, references, pointer to data members	32 bits	Binary	0	0xFFFFFFFF

⁽¹⁾ Figures are minimum precision.

Table 6-2. TMS320C6000 C/C++ EABI Data Types

Type	Size	Representation	Range	
			Minimum	Maximum
char, signed char	8 bits	ASCII	-128	127
unsigned char	8 bits	ASCII	0	255
short	16 bits	2s complement	-32 768	32 767
unsigned short	16 bits	Binary	0	65 535
int, signed int	32 bits	2s complement	-2 147 483 648	2 147 483 647
unsigned int	32 bits	Binary	0	4 294 967 295
long, signed long	32 bits	2s complement	-2 147 483 648	2 147 483 647
unsigned long	32 bits	Binary	0	4 294 967 295
__int40_t	40 bits	2s complement	-549 755 813 888	549 755 813 887
unsigned __int40_t	40 bits	Binary	0	1 099 511 627 775
long long, signed long long	64 bits	2s complement	-9 223 372 036 854 775 808	9 223 372 036 854 775 807
unsigned long long	64 bits	Binary	0	18 446 744 073 709 551 615
enum	32 bits	2s complement	-2 147 483 648	2 147 483 647
float	32 bits	IEEE 32-bit	1.175 494e-38 ⁽¹⁾	3.40 282 346e+38
double	64 bits	IEEE 64-bit	2.22 507 385e-308 ⁽¹⁾	1.79 769 313e+308
long double	64 bits	IEEE 64-bit	2.22 507 385e-308 ⁽¹⁾	1.79 769 313e+308
pointers, references, pointer to data members	32 bits	Binary	0	0xFFFFFFFF

⁽¹⁾ Figures are minimum precision.

6.5 Keywords

The C6000 C/C++ compiler supports the standard `const`, `register`, `restrict`, and `volatile` keywords. In addition, the C/C++ compiler extends the C/C++ language through the support of the `cregister`, `interrupt`, `near`, and `far` keywords.

6.5.1 The `const` Keyword

The C/C++ compiler supports the ANSI/ISO standard keyword `const`. This keyword gives you greater optimization and control over allocation of storage for certain data objects. You can apply the `const` qualifier to the definition of any variable or array to ensure that its value is not altered.

If you define an object as `far const`, the `.const` section allocates storage for the object. The `const` data storage allocation rule has two exceptions:

- If the keyword `volatile` is also specified in the definition of an object (for example, `volatile const int x`). Volatile keywords are assumed to be allocated to RAM. (The program does not modify a `const volatile` object, but something external to the program might.)
- If the object has automatic storage (allocated on the stack).

In both cases, the storage for the object is the same as if the `const` keyword were not used.

The placement of the `const` keyword within a definition is important. For example, the first statement below defines a constant pointer `p` to a variable `int`. The second statement defines a variable pointer `q` to a constant `int`:

```
int * const p = &x;
const int * q = &x;
```

Using the `const` keyword, you can define large constant tables and allocate them into system ROM. For example, to allocate a ROM table, you could use the following definition:

```
far const int digits[] = {0,1,2,3,4,5,6,7,8,9};
```

6.5.2 The `cregister` Keyword

The compiler extends the C/C++ language by adding the `cregister` keyword to allow high level language access to control registers.

When you use the `cregister` keyword on an object, the compiler compares the name of the object to a list of standard control registers for the C6000 (see [Table 6-3](#)). If the name matches, the compiler generates the code to reference the control register. If the name does not match, the compiler issues an error.

Table 6-3. Valid Control Registers

Register	Description
AMR	Addressing mode register
CSR	Control status register
DESR	(C6700+ only) dMAX event status register
DETR	(C6700+ only) dMAX event trigger register
DNUM	(C6400+ only) DSP core number register
ECR	(C6400+ only) Exception clear register
EFR	(C6400+ only) Exception flag 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
GPLYA	(C6400+ only) GMPY A-side polynomial register
CPLYB	(C6400+ only) GMPY B-side polynomial register
ICR	Interrupt clear register
IER	Interrupt enable register
IERR	(C6400+ only) Internal exception report register

Table 6-3. Valid Control Registers (continued)

Register	Description
IFR	Interrupt flag register. (IFR is read only.)
ILC	(C6400+ only) Inner loop count register
IRP	Interrupt return pointer
ISR	Interrupt set register
ISTP	Interrupt service table pointer
ITSR	(C6400+ only) Interrupt task state register
NRP	Nonmaskable interrupt return pointer
NTSR	(C6400+ only) NMI/exception task state register
REP	(C6400+ only) Restricted entry point address register
RILC	(C6400+ only) Reload inner loop count register
SSR	(C6400+ only) Saturation status register
TSCH	(C6400+ only) Time-stamp counter (high 32) register
TSCCL	(C6400+ only) Time-stamp counter (low 32) register
TSR	(C6400+ only) Task state register

The `cregister` keyword can be used only in file scope. The `cregister` keyword is not allowed on any declaration within the boundaries of a function. It can only be used on objects of type integer or pointer. The `cregister` keyword is not allowed on objects of any floating-point type or on any structure or union objects.

The `cregister` keyword does not imply that the object is volatile. If the control register being referenced is volatile (that is, can be modified by some external control), then the object must be declared with the `volatile` keyword also.

To use the control registers in [Table 6-3](#), you must declare each register as follows. The `c6x.h` include file defines all the control registers through this syntax:

```
extern cregister volatile unsigned int register ;
```

Once you have declared the register, you can use the register name directly. See the *TMS320C62x DSP CPU and Instruction Set Reference Guide*, *TMS320C64x/C64x+ DSP CPU and Instruction Set Reference Guide*, the *TMS320C67x/C67x+ DSP CPU and Instruction Set Reference Guide*, or *TMS320C66x+ DSP CPU and Instruction Set Reference Guide* for detailed information on the control registers.

See [Example 6-1](#) for an example that declares and uses control registers.

Example 6-1. Define and Use Control Registers

```
extern cregister volatile unsigned int AMR;
extern cregister volatile unsigned int CSR;
extern cregister volatile unsigned int IFR;
extern cregister volatile unsigned int ISR;
extern cregister volatile unsigned int ICR;
extern cregister volatile unsigned int IER;
extern cregister volatile unsigned int FADCR;
extern cregister volatile unsigned int FAUCR;
extern cregister volatile unsigned int FMCR;
main()
{
    printf("AMR = %x\n", AMR);
}
```

6.5.3 The interrupt Keyword

The compiler extends the C/C++ language by adding the interrupt keyword, which specifies that a function is treated as an interrupt function.

Functions that handle interrupts follow special register-saving rules and a special return sequence. The implementation stresses safety. The interrupt routine does not assume that the C run-time conventions for the various CPU register and status bits are in effect; instead, it re-establishes any values assumed by the run-time environment. When C/C++ code is interrupted, the interrupt routine must preserve the contents of all machine registers that are used by the routine or by any function called by the routine. When you use the interrupt keyword with the definition of the function, the compiler generates register saves based on the rules for interrupt functions and the special return sequence for interrupts.

You can only use the interrupt keyword with a function that is defined to return void and that has no parameters. The body of the interrupt function can have local variables and is free to use the stack or global variables. For example:

```
interrupt void int_handler()
{
    unsigned int flags;
    ...
}
```

The name `c_int00` is the C/C++ entry point. This name is reserved for the system reset interrupt. This special interrupt routine initializes the system and calls the function `main`. Because it has no caller, `c_int00` does not save any registers.

Use the alternate keyword, `__interrupt`, if you are writing code for strict ANSI/ISO mode (using the `--strict_ansi` compiler option).

HWI Objects and the interrupt Keyword

NOTE: The interrupt keyword must not be used when BIOS HWI objects are used in conjunction with C functions. The `HWI_enter`/`HWI_exit` macros and the HWI dispatcher contain this functionality, and the use of the C modifier can cause negative results.

6.5.4 The near and far Keywords

The C6000 C/C++ compiler extends the C/C++ language with the near and far keywords to specify how global and static variables are accessed and how functions are called.

Syntactically, the near and far keywords are treated as storage class modifiers. They can appear before, after, or in between the storage class specifiers and types. With the exception of near and far, two storage class modifiers cannot be used together in a single declaration. The following examples are legal combinations of near and far with other storage class modifiers:

```
far static int x;
static near int x;
static int far x;
far int foo();
static far int foo();
```

6.5.4.1 near and far Data Objects

Global and static data objects can be accessed in the following two ways:

near keyword The compiler assumes that the data item can be accessed relative to the data page pointer. For example:

```
LDW    *+dp(_address), a0
```

far keyword The compiler cannot access the data item via the DP. This can be required if the total amount of program data is larger than the offset allowed (32K) from the DP. For example:

```
MVKL  _address, a1
MVKH  _address, a1
LDW   *a1, a0
```


Once a variable has been defined to be far, all external references to this variable in other C files or headers must also contain the far keyword. This is also true of the near keyword. However, you will get compiler or linker errors when the far keyword is not used everywhere. Not using the near keyword everywhere only leads to slower data access times.

If you use the `DATA_SECTION` pragma, the object is indicated as a far variable, and this cannot be overridden. If you reference this object in another file, then you need to use *extern far* when declaring this object in the other source file. This ensures access to the variable, since the variable might not be in the `.bss` section. For details, see [Section 6.9.6](#).

NOTE: Defining Global Variables in Assembly Code

If you also define a global variable in assembly code with the `.usect` directive (where the variable is not assigned in the `.bss` section) or you allocate a variable into separate section using a `#pragma DATA_SECTION` directive; and you want to reference that variable in C code, you must declare the variable as *extern far*. This ensures the compiler does not try to generate an illegal access of the variable by way of the data page pointer.

When data objects do not have the near or far keyword specified, the compiler will use far accesses to aggregate data and near accesses to non-aggregate data. For more information on the data memory model and ways to control accesses to data, see [Section 7.1.5.1](#).

6.5.4.2 Near and far Function Calls

Function calls can be invoked in one of two ways:

near keyword	The compiler assumes that destination of the call is within ± 1 M word of the caller. Here the compiler uses the PC-relative branch instruction. <pre>B _func</pre>
far keyword	The compiler is told by you that the call is not within ± 1 M word. <pre>MVKL _func, a1 MVKH _func, a1 B _func</pre>

By default, the compiler generates small-memory model code, which means that every function call is handled as if it were declared near, unless it is actually declared far.

For more information on function calls, see [Section 7.1.6](#).

6.5.5 The restrict Keyword

To help the compiler determine memory dependencies, you can qualify a pointer, reference, or array with the restrict keyword. The restrict keyword is a type qualifier that can be applied to pointers, references, and arrays. Its use represents a guarantee by you, the programmer, that within the scope of the pointer declaration the object pointed to can be accessed only by that pointer. Any violation of this guarantee renders the program undefined. This practice helps the compiler optimize certain sections of code because aliasing information can be more easily determined.

In [Example 6-2](#), the restrict keyword is used to tell the compiler that the function `func1` is never called with the pointers `a` and `b` pointing to objects that overlap in memory. You are promising that accesses through `a` and `b` will never conflict; therefore, a write through one pointer cannot affect a read from any other pointers. The precise semantics of the restrict keyword are described in the 1999 version of the ANSI/ISO C Standard.

Example 6-2. Use of the restrict Type Qualifier With Pointers

```
void func1(int * restrict a, int * restrict b)
{
    /* func1's code here */
}
```

[Example 6-3](#) illustrates using the restrict keyword when passing arrays to a function. Here, the arrays c and d should not overlap, nor should c and d point to the same array.

Example 6-3. Use of the restrict Type Qualifier With Arrays

```
void func2(int c[restrict], int d[restrict])
{
    int i;

    for(i = 0; i < 64; i++)
    {
        c[i] += d[i];
        d[i] += 1;
    }
}
```

6.5.6 The volatile Keyword

The compiler eliminates redundant memory accesses whenever possible, using data flow analysis to figure out when it is legal. However, some memory accesses may be special in some way that the compiler cannot see, and in such cases you must use the volatile keyword to prevent the compiler from optimizing away something important. The compiler does not optimize out any accesses to variables declared volatile. The number and order of accesses of a volatile variable are exactly as they appear in the C/C++ code, no more and no less.

There are different ways to understand how volatile works, but fundamentally it is a hint to the compiler that something it cannot understand is going on, and so the compiler should not try to be over-clever.

Any variable which might be modified by something external to the obvious control flow of the program (such as an interrupt service routine) must be declared volatile. This tells the compiler that an interrupt function might modify the value at any time, so the compiler should not perform optimizations which will change the number or order of accesses of that variable. This is the primary purpose of the volatile keyword.

In the following example, the loop intends to wait for a location to be read as 0xFF:

```
unsigned int *ctrl;
while (*ctrl !=0xFF);
```

However, in this example, *ctrl is a loop-invariant expression, so the loop is optimized down to a single-memory read. To get the desired result, define ctrl as:

```
volatile unsigned int *ctrl;
```

Here the *ctrl pointer is intended to reference a hardware location, such as an interrupt flag.

Volatile must also be used when accessing memory locations that represent memory-mapped peripheral devices. Such memory locations might change value in ways that the compiler cannot predict. These locations might change if accessed, or when some other memory location is accessed, or when some signal occurs.

Volatile must also be used for local variables in a function which calls setjmp, if the value of the local variables needs to remain valid if a longjmp occurs.

Example 6-4. Volatile for Local Variables With setjmp

```

#include <stdlib.h>
jmp_buf context;
void function()
{
    volatile int x = 3;
    switch(setjmp(context))
    {
        case 0: setup(); break;
        default:
        {
            printf("x == %d\n", x); /* We can only reach here if longjmp
                                     has occurred; because x's lifetime
                                     begins before the setjmp and lasts
                                     through the longjmp, the C standard
                                     requires x be declared "volatile" */

            break;
        }
    }
}

```

The `--interrupt_threshold=1` option should be used when compiling with volatiles.

6.6 C++ Exception Handling

The compiler supports all the C++ exception handling features as defined by the ANSI/ISO 14882 C++ Standard. More details are discussed in *The C++ Programming Language, Third Edition* by Bjarne Stroustrup.

The compiler `--exceptions` option enables exception handling. The compiler's default is no exception handling support.

For exceptions to work correctly, all C++ files in the application must be compiled with the `--exceptions` option, regardless of whether exceptions occur in a particular file. Mixing exception-enabled object files and libraries with object files and libraries that do not have exceptions enabled can lead to undefined behavior.

Exception handling requires support in the run-time-support library, which come in exception-enabled and exception-disabled forms; you must link with the correct form. When using automatic library selection (the default), the linker automatically selects the correct library [Section 5.3.1.1](#). If you select the library manually, you must use run-time-support libraries whose name contains `_eh` if you enable exceptions.

Using `--exceptions` causes the compiler to insert exception handling code. This code will increase the code size of the program, particularly for COFF ABI. In addition, COFF ABI will increase the execution time, even if an exception is never thrown. EABI will not increase code size as much, and has a minimal execution time cost if exceptions are never thrown, but will slightly increase the data size for the exception-handling tables.

See [Section 8.1](#) for details on the run-time libraries.

6.7 Register Variables and Parameters

The C/C++ compiler treats register variables (variables defined with the `register` keyword) differently, depending on whether you use the `--opt_level (-O)` option.

- **Compiling with optimization**

The compiler ignores any register definitions and allocates registers to variables and temporary values by using an algorithm that makes the most efficient use of registers.

- **Compiling without optimization**

If you use the `register` keyword, you can suggest variables as candidates for allocation into registers. The compiler uses the same set of registers for allocating temporary expression results as it uses for

allocating register variables.

The compiler attempts to honor all register definitions. If the compiler runs out of appropriate registers, it frees a register by moving its contents to memory. If you define too many objects as register variables, you limit the number of registers the compiler has for temporary expression results. This limit causes excessive movement of register contents to memory.

Any object with a scalar type (integral, floating point, or pointer) can be defined as a register variable. The register designator is ignored for objects of other types, such as arrays.

The register storage class is meaningful for parameters as well as local variables. Normally, in a function, some of the parameters are copied to a location on the stack where they are referenced during the function body. The compiler copies a register parameter to a register instead of the stack, which speeds access to the parameter within the function.

For more information about register conventions, see [Section 7.3](#).

6.8 The asm Statement

The C/C++ compiler can embed assembly language instructions or directives directly into the assembly language output of the compiler. This capability is an extension to the C/C++ language—the *asm* statement. The *asm* (or `__asm`) statement provides access to hardware features that C/C++ cannot provide. The *asm* statement is syntactically like a call to a function named *asm*, with one string constant argument:

```
asm(" assembler text ");
```

The compiler copies the argument string directly into your output file. The assembler text must be enclosed in double quotes. All the usual character string escape codes retain their definitions. For example, you can insert a `.byte` directive that contains quotes as follows:

```
asm("STR: .byte \"abc\"");
```

The inserted code must be a legal assembly language statement. Like all assembly language statements, the line of code inside the quotes must begin with a label, a blank, a tab, or a comment (asterisk or semicolon). The compiler performs no checking on the string; if there is an error, the assembler detects it. For more information about the assembly language statements, see the *TMS320C6000 Assembly Language Tools User's Guide*.

The *asm* statements do not follow the syntactic restrictions of normal C/C++ statements. Each can appear as a statement or a declaration, even outside of blocks. This is useful for inserting directives at the very beginning of a compiled module.

Use the alternate statement `__asm("assembler text")` if you are writing code for strict ANSI/ISO C mode (using the `--strict_ansi` option).

NOTE: Avoid Disrupting the C/C++ Environment With *asm* Statements

Be careful not to disrupt the C/C++ environment with *asm* statements. The compiler does not check the inserted instructions. Inserting jumps and labels into C/C++ code can cause unpredictable results in variables manipulated in or around the inserted code. Directives that change sections or otherwise affect the assembly environment can also be troublesome.

Be especially careful when you use optimization with *asm* statements. Although the compiler cannot remove *asm* statements, it can significantly rearrange the code order near them and cause undesired results.

6.9 Pragma Directives

Pragma directives tell the compiler how to treat a certain function, object, or section of code. The C6000 C/C++ compiler supports the following pragmas:

- CHECK_MISRA (See [Section 6.9.1](#))
- CLINK (See [Section 6.9.2](#))
- CODE_SECTION (See [Section 6.9.3](#))
- DATA_ALIGN (See [Section 6.9.4](#))
- DATA_MEM_BANK (See [Section 6.9.5](#))
- DATA_SECTION (See [Section 6.9.6](#))
- DIAG_SUPPRESS, DIAG_REMARK, DIAG_WARNING, DIAG_ERROR, and DIAG_DEFAULT (See [Section 6.9.7](#))
- FUNC_ALWAYS_INLINE (See [Section 6.9.8](#))
- FUNC_CANNOT_INLINE (See [Section 6.9.9](#))
- FUNC_EXT_CALLED (See [Section 6.9.10](#))
- FUNC_INTERRUPT_THRESHOLD (See [Section 6.9.11](#))
- FUNC_IS_PURE (See [Section 6.9.12](#))
- FUNC_IS_SYSTEM (See [Section 6.9.13](#))
- FUNC_NEVER_RETURNS (See [Section 6.9.14](#))
- FUNC_NO_GLOBAL_ASG (See [Section 6.9.15](#))
- FUNC_NO_IND_ASG (See [Section 6.9.16](#))
- FUNCTION_OPTIONS (See [Section 6.9.17](#))
- INTERRUPT (See [Section 6.9.18](#))
- LOCATION (See [Section 6.9.19](#))
- MUST_ITERATE (See [Section 6.9.20](#))
- NMI_INTERRUPT (See [Section 6.9.21](#))
- NOINIT (See [Section 6.9.22](#))
- NO_HOOKS (See [Section 6.9.23](#))
- PERSISTENT (See [Section 6.9.22](#))
- PROB_ITERATE (See [Section 6.9.24](#))
- RESET_MISRA (See [Section 6.9.25](#))
- RETAIN (See [Section 6.9.26](#))
- SET_CODE_SECTION (See [Section 6.9.27](#))
- SET_DATA_SECTION (See [Section 6.9.27](#))
- STRUCT_ALIGN (See [Section 6.9.28](#))
- UNROLL (See [Section 6.9.29](#))

Most of these pragmas apply to functions. Except for the DATA_MEM_BANK pragma, the arguments *func* and *symbol* cannot be defined or declared inside the body of a function. You must specify the pragma outside the body of a function; and the pragma specification must occur before any declaration, definition, or reference to the func or symbol argument. If you do not follow these rules, the compiler issues a warning and may ignore the pragma.

For the pragmas that apply to functions or symbols (except CLINK and RETAIN), the syntax for the pragmas differs between C and C++. In C, you must supply the name of the object or function to which you are applying the pragma as the first argument. In C++, the name is omitted; the pragma applies to the declaration of the object or function that follows it.

6.9.1 The CHECK_MISRA Pragma

The CHECK_MISRA pragma enables/disables MISRA-C:2004 rules at the source level. This pragma is equivalent to using the `--check_misra` option.

The syntax of the pragma in C is:

```
#pragma CHECK_MISRA (" {all|required|advisory|none|rulespec} ");
```

The *rulespec* parameter is a comma-separated list of specifiers. See [Section 6.3](#) for details.

The RESET_MISRA pragma can be used to reset any CHECK_MISRA pragmas; see [Section 6.9.25](#).

6.9.2 The CLINK Pragma

The CLINK pragma can be applied to a code or data symbol. It causes a `.clink` directive to be generated into the section that contains the definition of the symbol. The `.clink` directive indicates to the linker that the section is eligible for removal during conditional linking. Therefore, if the section is not referenced by any other section in the application that is being compiled and linked, it will not be included in the output file result of the link.

The syntax of the pragma in C/C++ is:

```
#pragma CLINK (symbol)
```

The RETAIN pragma has the opposite effect of the CLINK pragma. See [Section 6.9.26](#) for more details.

6.9.3 The CODE_SECTION Pragma

The CODE_SECTION pragma allocates space for the *symbol* in C, or the next symbol declared in C++, in a section named *section name*.

The syntax of the pragma in C is:

```
#pragma CODE_SECTION (symbol, "section name")
```

The syntax of the pragma in C++ is:

```
#pragma CODE_SECTION ("section name")
```

The CODE_SECTION pragma is useful if you have code objects that you want to link into an area separate from the `.text` section.

The following examples demonstrate the use of the CODE_SECTION pragma.

Example 6-5. Using the CODE_SECTION Pragma C Source File

```
#pragma CODE_SECTION(fn, "my_sect")

int fn(int x)
{
    return x;
}
```

Example 6-6. Generated Assembly Code From [Example 6-5](#)

```

.sect    "my_sect"
.global _fn

;*****
;* FUNCTION NAME: _fn                                     *
;*                                                         *
;*   Regs Modified   : SP                               *
;*   Regs Used       : A4,B3,SP                         *
;*   Local Frame Size : 0 Args + 4 Auto + 0 Save = 4 byte *
;*****
_fn:
;*- - - - - - - - - - - - - - - - - - - - - - - - - - - - *
      RET     .S2    B3                ; |6|
      SUB     .D2    SP,8,SP           ; |4|
      STW     .D2T1  A4,**+SP(4)      ; |4|
      ADD     .S2    8,SP,SP          ; |6|
      NOP                                ;
      ; BRANCH OCCURS                ; |6|

```

6.9.4 The DATA_ALIGN Pragma

The DATA_ALIGN pragma aligns the *symbol* in C, or the next symbol declared in C++, to an alignment boundary. The alignment boundary is the maximum of the symbol's default alignment value or the value of the *constant* in bytes. The constant must be a power of 2.

The syntax of the pragma in C is:

```
#pragma DATA_ALIGN ( symbol , constant );
```

The syntax of the pragma in C++ is:

```
#pragma DATA_ALIGN ( constant );
```

6.9.5 The DATA_MEM_BANK Pragma

The DATA_MEM_BANK pragma aligns a symbol or variable to a specified C6000 internal data memory bank boundary. The *constant* specifies a specific memory bank to start your variables on. (See [Figure 4-1](#) for a graphic representation of memory banks.) The value of *constant* depends on the C6000 device:

C6200	The C6200 devices contain four memory banks (0, 1, 2, and 3); <i>constant</i> can be 0 or 2.
C6400	The C6400 devices contain 8 memory banks; <i>constant</i> can be 0, 2, 4, or 6.
C6400+	The C6400+ devices contain 8 memory banks; <i>constant</i> can be 0, 2, 4, or 6.
C6600	The C6600 devices contain 8 memory banks; <i>constant</i> can be 0, 2, 4, or 6.
C6700	The C6700 devices contain 8 memory banks; <i>constant</i> can be 0, 2, 4, or 6.
C6740	The C6740 devices contain 8 memory banks; <i>constant</i> can be 0, 2, 4, or 6.

The syntax of the pragma in C is:

```
#pragma DATA_MEM_BANK ( symbol , constant );
```

The syntax of the pragma in C++ is:

```
#pragma DATA_MEM_BANK ( constant );
```

Both global and local variables can be aligned with the DATA_MEM_BANK pragma. The DATA_MEM_BANK pragma must reside inside the function that contains the local variable being aligned. The *symbol* can also be used as a parameter in the DATA_SECTION pragma.

When optimization is enabled, the tools may or may not use the stack to store the values of local variables.

The DATA_MEM_BANK pragma allows you to align data on any data memory bank that can hold data of the type size of the *symbol*. This is useful if you need to align data in a particular way to avoid memory bank conflicts in your hand-coded assembly code versus padding with zeros and having to account for the padding in your code.

This pragma increases the amount of space used in data memory by a small amount as padding is used to align data onto the correct bank.

For C6200, the code in [Example 6-7](#) guarantees that array x begins at an address ending in 4 or c (in hexadecimal), and that array y begins at an address ending in 4 or c. The alignment for array y affects its stack placement. Array z is placed in the .z_sect section, and begins at an address ending in 0 or 8.

Example 6-7. Using the DATA_MEM_BANK Pragma

```
#pragma DATA_MEM_BANK (x, 2);
short x[100];

#pragma DATA_MEM_BANK (z, 0);
#pragma DATA_SECTION (z, ".z_sect");
short z[100];

void main()
{
    #pragma DATA_MEM_BANK (y, 2);
    short y[100];
    ...
}
```


6.9.6 The DATA_SECTION Pragma

The DATA_SECTION pragma allocates space for the *symbol* in C, or the next symbol declared in C++, in a section named *section name*.

The syntax of the pragma in C is:

```
#pragma DATA_SECTION ( symbol , " section name " );
```

The syntax of the pragma in C++ is:

```
#pragma DATA_SECTION (" section name " );
```

The DATA_SECTION pragma is useful if you have data objects that you want to link into an area separate from the .bss section. If you allocate a global variable using a DATA_SECTION pragma and you want to reference the variable in C code, you must declare the variable as extern far.

[Example 6-8](#) through [Example 6-10](#) demonstrate the use of the DATA_SECTION pragma.

Example 6-8. Using the DATA_SECTION Pragma C Source File

```
#pragma DATA_SECTION(bufferB, "my_sect")
char bufferA[512];
char bufferB[512];
```

Example 6-9. Using the DATA_SECTION Pragma C++ Source File

```
char bufferA[512];
#pragma DATA_SECTION("my_sect")
char bufferB[512];
```

Example 6-10. Using the DATA_SECTION Pragma Assembly Source File

```
.global _bufferA
.bss _bufferA,512,4
.global _bufferB
```

6.9.7 The Diagnostic Message Pragas

The following pragmas can be used to control diagnostic messages in the same ways as the corresponding command line options:

Pragma	Option	Description
DIAG_SUPPRESS <i>num</i>	-pds= <i>num</i> [, <i>num</i> ₂ , <i>num</i> ₃ ...]	Suppress diagnostic <i>num</i>
DIAG_REMARK <i>num</i>	-pdsr= <i>num</i> [, <i>num</i> ₂ , <i>num</i> ₃ ...]	Treat diagnostic <i>num</i> as a remark
DIAG_WARNING <i>num</i>	-pdsw= <i>num</i> [, <i>num</i> ₂ , <i>num</i> ₃ ...]	Treat diagnostic <i>num</i> as a warning
DIAG_ERROR <i>num</i>	-pdse= <i>num</i> [, <i>num</i> ₂ , <i>num</i> ₃ ...]	Treat diagnostic <i>num</i> as an error
DIAG_DEFAULT <i>num</i>	n/a	Use default severity of the diagnostic

The syntax of the pragmas in C is:

```
#pragma DIAG_XXX [=]num[, num2, num3...]
```

The diagnostic affected (*num*) is specified using either an error number or an error tag name. The equal sign (=) is optional. Any diagnostic can be overridden to be an error, but only diagnostics with a severity of discretionary error or below can have their severity reduced to a warning or below, or be suppressed. The `diag_default` pragma is used to return the severity of a diagnostic to the one that was in effect before any pragmas were issued (i.e., the normal severity of the message as modified by any command-line options).

The diagnostic identifier number is output along with the message when the `-pden` command line option is specified.

6.9.8 The FUNC_ALWAYS_INLINE Pragma

The `FUNC_ALWAYS_INLINE` pragma instructs the compiler to always inline the named function. The compiler only inlines the function if it is legal to inline the function and the compiler is invoked with any level of optimization (`--opt_level=0`).

The pragma must appear before any declaration or reference to the function that you want to inline. In C, the argument *func* is the name of the function that will be inlined. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_ALWAYS_INLINE ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_ALWAYS_INLINE;
```

Use Caution with the FUNC_ALWAYS_INLINE Pragma

NOTE: The `FUNC_ALWAYS_INLINE` pragma overrides the compiler's inlining decisions. Overuse of the pragma could result in increased compilation times or memory usage, potentially enough to consume all available memory and result in compilation tool failures.

6.9.9 The `FUNC_CANNOT_INLINE` Pragma

The `FUNC_CANNOT_INLINE` pragma instructs the compiler that the named function cannot be expanded inline. Any function named with this pragma overrides any inlining you designate in any other way, such as using the `inline` keyword. Automatic inlining is also overridden with this pragma; see [Section 2.11](#).

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function that cannot be inlined. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_CANNOT_INLINE ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_CANNOT_INLINE;
```

6.9.10 The `FUNC_EXT_CALLED` Pragma

When you use the `--program_level_compile` option, the compiler uses program-level optimization. When you use this type of optimization, the compiler removes any function that is not called, directly or indirectly, by `main`. You might have C/C++ functions that are called by hand-coded assembly instead of `main`.

The `FUNC_EXT_CALLED` pragma specifies to the optimizer to keep these C functions or any other functions that these C/C++ functions call. These functions act as entry points into C/C++.

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function that you do not want removed. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_EXT_CALLED ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_EXT_CALLED;
```

Except for `_c_int00`, which is the name reserved for the system reset interrupt for C/C++ programs, the name of the interrupt (the *func* argument) does not need to conform to a naming convention.

When you use program-level optimization, you may need to use the `FUNC_EXT_CALLED` pragma with certain options. See [Section 3.7.2](#).

6.9.11 The `FUNC_INTERRUPT_THRESHOLD` Pragma

The compiler allows interrupts to be disabled around software pipelined loops for threshold cycles within the function. This implements the `--interrupt_threshold` option for a single function (see [Section 2.12](#)). The `FUNC_INTERRUPT_THRESHOLD` pragma always overrides the `--interrupt_threshold=n` command line option. A threshold value less than 0 assumes that the function is never interrupted, which is equivalent to an interrupt threshold of infinity.

The syntax of the pragma in C is:

```
#pragma FUNC_INTERRUPT_THRESHOLD ( func , threshold );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_INTERRUPT_THRESHOLD ( threshold );
```

The following examples demonstrate the use of different thresholds:

- The function `foo()` must be interruptible at least every 2,000 cycles:

```
#pragma FUNC_INTERRUPT_THRESHOLD (foo, 2000)
```
- The function `foo()` must always be interruptible.

```
#pragma FUNC_INTERRUPT_THRESHOLD (foo, 1)
```
- The function `foo()` is never interrupted.

```
#pragma FUNC_INTERRUPT_THRESHOLD (foo, -1)
```

6.9.12 The `FUNC_IS_PURE` Pragma

The `FUNC_IS_PURE` pragma specifies to the compiler that the named function has no side effects. This allows the compiler to do the following:

- Delete the call to the function if the function's value is not needed
- Delete duplicate functions

The pragma must appear before any declaration or reference to the function. In C, the argument `func` is the name of a function. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_IS_PURE ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_IS_PURE;
```

6.9.13 The **FUNC_IS_SYSTEM** Pragma

The **FUNC_IS_SYSTEM** pragma specifies to the compiler that the named function has the behavior defined by the ANSI/ISO standard for a function with that name.

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function to treat as an ANSI/ISO standard function. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_IS_SYSTEM ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_IS_SYSTEM;
```

6.9.14 The **FUNC_NEVER_RETURNS** Pragma

The **FUNC_NEVER_RETURNS** pragma specifies to the compiler that the function never returns to its caller.

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function that does not return. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_NEVER_RETURNS ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_NEVER_RETURNS;
```

6.9.15 The **FUNC_NO_GLOBAL_ASG** Pragma

The **FUNC_NO_GLOBAL_ASG** pragma specifies to the compiler that the function makes no assignments to named global variables and contains no asm statements.

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function that makes no assignments. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_NO_GLOBAL_ASG ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_NO_GLOBAL_ASG;
```

6.9.16 The **FUNC_NO_IND_ASG** Pragma

The **FUNC_NO_IND_ASG** pragma specifies to the compiler that the function makes no assignments through pointers and contains no asm statements.

The pragma must appear before any declaration or reference to the function that you want to keep. In C, the argument *func* is the name of the function that makes no assignments. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma FUNC_NO_IND_ASG ( func );
```

The syntax of the pragma in C++ is:

```
#pragma FUNC_NO_IND_ASG;
```

6.9.17 The **FUNCTION_OPTIONS** Pragma

The **FUNCTION_OPTIONS** pragma allows you to compile a specific function in a C or C++ file with additional command-line compiler options. The affected function will be compiled as if the specified list of options appeared on the command line after all other compiler options. In C, the pragma is applied to the function specified. In C++, the pragma is applied to the next function.

The syntax of the pragma in C is:

```
#pragma FUNCTION_OPTIONS ( func, "additional options" );
```

The syntax of the pragma in C++ is:

```
#pragma FUNCTION_OPTIONS( "additional options" );
```

6.9.18 The **INTERRUPT** Pragma

The **INTERRUPT** pragma enables you to handle interrupts directly with C code. In C, the argument *func* is the name of a function. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma INTERRUPT ( func );
```

The syntax of the pragma in C++ is:

```
#pragma INTERRUPT ;
```

The code for the function will return via the IRP (interrupt return pointer).

Except for `_c_int00`, which is the name reserved for the system reset interrupt for C programs, the name of the interrupt (the *func* argument) does not need to conform to a naming convention.

HWI Objects and the **INTERRUPT** Pragma

NOTE: The **INTERRUPT** pragma must not be used when BIOS HWI objects are used in conjunction with C functions. The `HWI_enter`/`HWI_exit` macros and the HWI dispatcher contain this functionality, and the use of the C modifier can cause negative results.

6.9.19 The *LOCATION* Pragma

The compiler supports the ability to specify the run-time address of a variable at the source level. This can be accomplished with the *LOCATION* pragma or attribute. Location support is only available in EABI.

The syntax of the pragma in C is:

```
#pragma LOCATION( x , address );
int x;
```

The syntax of the pragmas in C++ is:

```
#pragma LOCATION(address );
int x;
```

The syntax of the pragma for IAR is:

```
#pragma location=address
int x;
```

The syntax of the GCC attribute is:

```
int x __attribute__((location(address)));
```

The *noinit* pragma may be used in conjunction with the location pragma to map variables to special memory locations; see [Section 6.9.22](#).

6.9.20 The *MUST_ITERATE* Pragma

The *MUST_ITERATE* pragma specifies to the compiler certain properties of a loop. You guarantee that these properties are always true. Through the use of the *MUST_ITERATE* pragma, you can guarantee that a loop executes a specific number of times. Anytime the *UNROLL* pragma is applied to a loop, *MUST_ITERATE* should be applied to the same loop. For loops the *MUST_ITERATE* pragma's third argument, *multiple*, is the most important and should always be specified.

Furthermore, the *MUST_ITERATE* pragma should be applied to any other loops as often as possible. This is because the information provided via the pragma (especially the minimum number of iterations) aids the compiler in choosing the best loops and loop transformations (that is, software pipelining and nested loop transformations). It also helps the compiler reduce code size.

No statements are allowed between the *MUST_ITERATE* pragma and the *for*, *while*, or *do-while* loop to which it applies. However, other pragmas, such as *UNROLL* and *PROB_ITERATE*, can appear between the *MUST_ITERATE* pragma and the loop.

6.9.20.1 The *MUST_ITERATE* Pragma Syntax

The syntax of the pragma for C and C++ is:

```
#pragma MUST_ITERATE ( min, max, multiple );
```

The arguments *min* and *max* are programmer-guaranteed minimum and maximum trip counts. The trip count is the number of times a loop iterates. The trip count of the loop must be evenly divisible by *multiple*. All arguments are optional. For example, if the trip count could be 5 or greater, you can specify the argument list as follows:

```
#pragma MUST_ITERATE(5);
```

However, if the trip count could be any nonzero multiple of 5, the pragma would look like this:

```
#pragma MUST_ITERATE(5, , 5); /* Note the blank field for max */
```

It is sometimes necessary for you to provide min and multiple in order for the compiler to perform unrolling. This is especially the case when the compiler cannot easily determine how many iterations the loop will perform (that is, the loop has a complex exit condition).

When specifying a multiple via the `MUST_ITERATE` pragma, results of the program are undefined if the trip count is not evenly divisible by multiple. Also, results of the program are undefined if the trip count is less than the minimum or greater than the maximum specified.

If no min is specified, zero is used. If no max is specified, the largest possible number is used. If multiple `MUST_ITERATE` pragmas are specified for the same loop, the smallest max and largest min are used.

6.9.20.2 Using `MUST_ITERATE` to Expand Compiler Knowledge of Loops

Through the use of the `MUST_ITERATE` pragma, you can guarantee that a loop executes a certain number of times. The example below tells the compiler that the loop is guaranteed to run exactly 10 times:

```
#pragma MUST_ITERATE(10,10);

for(i = 0; i < trip_count; i++) { ...
```

In this example, the compiler attempts to generate a software pipelined loop even without the pragma. However, if `MUST_ITERATE` is not specified for a loop such as this, the compiler generates code to bypass the loop, to account for the possibility of 0 iterations. With the pragma specification, the compiler knows that the loop iterates at least once and can eliminate the loop-bypassing code.

`MUST_ITERATE` can specify a range for the trip count as well as a factor of the trip count. For example:

```
pragma MUST_ITERATE(8, 48, 8);

for(i = 0; i < trip_count; i++) { ...
```

This example tells the compiler that the loop executes between 8 and 48 times and that the `trip_count` variable is a multiple of 8 (8, 16, 24, 32, 40, 48). The multiple argument allows the compiler to unroll the loop.

You should also consider using `MUST_ITERATE` for loops with complicated bounds. In the following example:

```
for(i2 = ipos[2]; i2 < 40; i2 += 5) { ...
```

The compiler would have to generate a divide function call to determine, at run time, the exact number of iterations performed. The compiler will not do this. In this case, using `MUST_ITERATE` to specify that the loop always executes eight times allows the compiler to attempt to generate a software pipelined loop:

```
#pragma MUST_ITERATE(8, 8);

for(i2 = ipos[2]; i2 < 40; i2 += 5) { ...
```

6.9.21 The `NMI_INTERRUPT` Pragma

The `NMI_INTERRUPT` pragma enables you to handle non-maskable interrupts directly with C code. In C, the argument *func* is the name of a function. In C++, the pragma applies to the next function declared.

The syntax of the pragma in C is:

```
#pragma NMI_INTERRUPT( func );
```

The syntax of the pragma in C++ is:

```
#pragma NMI_INTERRUPT;
```

The code generated for the function will return via the NRP versus the IRP as for a function declared with the interrupt keyword or `INTERRUPT` pragma.

Except for `_c_int00`, which is the name reserved for the system reset interrupt for C programs, the name of the interrupt (function) does not need to conform to a naming convention.

6.9.22 The **NOINIT** and **PERSISTENT** Pragmas

When using EABI, global and static variables will be zero-initialized. However, in applications using non-volatile memory, it may be desirable to have variables that are not initialized. Noinit variables are global or static variables that are not zero-initialized at startup or reset.

The **NOINIT** pragma may be used in conjunction with the location pragma to map variables to special memory locations, like memory-mapped registers, without generating unwanted writes.

Noinit variables cannot have an initializer. However, statically-initialized variables may be declared "persistent" to disable startup initialization. Persistent variables are given an initial value when the code is loaded, but are never again initialized. Persistent and noinit variables behave identically with the exception of whether or not they are initialized at load time.

The syntax of the pragmas in C is:

```
#pragma NOINIT (x);
int x;
```

```
#pragma PERSISTENT (x);
int x=10;
```

The syntax of the pragmas in C++ is:

```
#pragma NOINIT;
int x;
```

```
#pragma PERSISTENT;
int x=10;
```

The syntax of the GCC attributes is:

```
int x __attribute__((noinit));
```

```
int x =0 __attribute__((persistent));
```

6.9.23 The **NO_HOOKS** Pragma

The **NO_HOOKS** pragma prevents entry and exit hook calls from being generated for a function.

The syntax of the pragma in C is:

```
#pragma NO_HOOKS ( func );
```

The syntax of the pragma in C++ is:

```
#pragma NO_HOOKS;
```

See [Section 2.17](#) for details on entry and exit hooks.

6.9.24 The **PROB_ITERATE** Pragma

The **PROB_ITERATE** pragma specifies to the compiler certain properties of a loop. You assert that these properties are true in the common case. The **PROB_ITERATE** pragma aids the compiler in choosing the best loops and loop transformations (that is, software pipelining and nested loop transformations). **PROB_ITERATE** is useful only when the **MUST_ITERATE** pragma is not used or the **PROB_ITERATE** parameters are more constraining than the **MUST_ITERATE** parameters.

No statements are allowed between the **PROB_ITERATE** pragma and the for, while, or do-while loop to which it applies. However, other pragmas, such as **UNROLL** and **MUST_ITERATE**, may appear between the **PROB_ITERATE** pragma and the loop.

The syntax of the pragma for C and C++ is:

```
#pragma PROB_ITERATE( min , max )
```

Where *min* and *max* are the minimum and maximum trip counts of the loop in the common case. The trip count is the number of times a loop iterates. Both arguments are optional.

For example, **PROB_ITERATE** could be applied to a loop that executes for eight iterations in the majority of cases (but sometimes may execute more or less than eight iterations):

```
#pragma PROB_ITERATE(8, 8);
```

If only the minimum expected trip count is known (say it is 5), the pragma would look like this:

```
#pragma PROB_ITERATE(5);
```

If only the maximum expected trip count is known (say it is 10), the pragma would look like this:

```
#pragma PROB_ITERATE(, 10); /* Note the blank field for min */
```

6.9.25 The **RESET_MISRA** Pragma

The **RESET_MISRA** pragma resets the specified MISRA-C:2004 rules to the state they were before any **CHECK_MISRA** pragmas (see [Section 6.9.1](#)) were processed. For instance, if a rule was enabled on the command line but disabled in the source, the **RESET_MISRA** pragma resets it to enabled. This pragma accepts the same format as the `--check_misra` option, except for the "none" keyword.

The syntax of the pragma in C is:

```
#pragma RESET_MISRA (" {all|required|advisory|rulespec} ")
```

The *rulespec* parameter is a comma-separated list of specifiers. See [Section 6.3](#) for details.

6.9.26 The **RETAIN** Pragma

The **RETAIN** pragma can be applied to a code or data symbol. It causes a `.retain` directive to be generated into the section that contains the definition of the symbol. The `.retain` directive indicates to the linker that the section is ineligible for removal during conditional linking. Therefore, regardless whether or not the section is referenced by another section in the application that is being compiled and linked, it will be included in the output file result of the link.

The syntax of the pragma in C/C++ is:

```
#pragma RETAIN ( symbol )
```

The **CLINK** pragma has the opposite effect of the **RETAIN** pragma. See [Section 6.9.2](#) for more details.

6.9.27 The `SET_CODE_SECTION` and `SET_DATA_SECTION` Pragmas

These pragmas can be used to set the section for all declarations below the pragma.

The syntax of the pragmas in C/C++ is:

```
#pragma SET_CODE_SECTION ("section name")
```

```
#pragma SET_DATA_SECTION ("section name")
```

In [Example 6-11](#) `x` and `y` are put in the section `mydata`. To reset the current section to the default used by the compiler, a blank parameter should be passed to the pragma. An easy way to think of the pragma is that it is like applying the `CODE_SECTION` or `DATA_SECTION` pragma to all symbols below it.

Example 6-11. Setting Section With `SET_DATA_SECTION` Pragma

```
#pragma SET_DATA_SECTION("mydata")
int x;
int y;
#pragma SET_DATA_SECTION()
```

The pragmas apply to both declarations and definitions. If applied to a declaration and not the definition, the pragma that is active at the declaration is used to set the section for that symbol. Here is an example:

Example 6-12. Setting a Section With `SET_CODE_SECTION` Pragma

```
#pragma SET_CODE_SECTION("func1")
extern void func1();
#pragma SET_CODE_SECTION()
...
void func1() { ... }
```

In [Example 6-12](#) `func1` is placed in section `func1`. If conflicting sections are specified at the declaration and definition, a diagnostic is issued.

The current `CODE_SECTION` and `DATA_SECTION` pragmas and GCC attributes can be used to override the `SET_CODE_SECTION` and `SET_DATA_SECTION` pragmas. For example:

Example 6-13. Overriding `SET_DATA_SECTION` Setting

```
#pragma DATA_SECTION(x, "x_data")
#pragma SET_DATA_SECTION("mydata")
int x;
int y;
#pragma SET_DATA_SECTION()
```

In [Example 6-13](#) `x` is placed in `x_data` and `y` is placed in `mydata`. No diagnostic is issued for this case.

The pragmas work for both C and C++. In C++, the pragmas are ignored for templates and for implicitly created objects, such as implicit constructors and virtual function tables.

6.9.28 The **STRUCT_ALIGN** Pragma

The **STRUCT_ALIGN** pragma is similar to **DATA_ALIGN**, but it can be applied to a structure, union type, or typedef and is inherited by any symbol created from that type. The **STRUCT_ALIGN** pragma is supported only in C.

The syntax of the pragma is:

```
#pragma STRUCT_ALIGN( type , constant expression )
```

This pragma guarantees that the alignment of the named type or the base type of the named typedef is at least equal to that of the expression. (The alignment may be greater as required by the compiler.) The alignment must be a power of 2. The *type* must be a type or a typedef name. If a type, it must be either a structure tag or a union tag. If a typedef, its base type must be either a structure tag or a union tag.

Since ANSI/ISO C declares that a typedef is simply an alias for a type (i.e. a struct) this pragma can be applied to the struct, the typedef of the struct, or any typedef derived from them, and affects all aliases of the base type.

This example aligns any `st_tag` structure variables on a page boundary:

```
typedef struct st_tag
{
    int    a;
    short b;
} st_typedef;

#pragma STRUCT_ALIGN (st_tag, 128);
#pragma STRUCT_ALIGN (st_typedef, 128);
```

Any use of **STRUCT_ALIGN** with a basic type (int, short, float) or a variable results in an error.

6.9.29 The **UNROLL** Pragma

The **UNROLL** pragma specifies to the compiler how many times a loop should be unrolled. The **UNROLL** pragma is useful for helping the compiler utilize SIMD instructions on the C6400 family. It is also useful in cases where better utilization of software pipeline resources are needed over a non-unrolled loop.

The optimizer must be invoked (use `--opt_level=[1|2|3]` or `-O1`, `-O2`, or `-O3`) in order for pragma-specified loop unrolling to take place. The compiler has the option of ignoring this pragma.

No statements are allowed between the **UNROLL** pragma and the `for`, `while`, or `do-while` loop to which it applies. However, other pragmas, such as **MUST_ITERATE** and **PROB_ITERATE**, can appear between the **UNROLL** pragma and the loop.

The syntax of the pragma for C and C++ is:

```
#pragma UNROLL( n );
```

If possible, the compiler unrolls the loop so there are *n* copies of the original loop. The compiler only unrolls if it can determine that unrolling by a factor of *n* is safe. In order to increase the chances the loop is unrolled, the compiler needs to know certain properties:

- The loop iterates a multiple of *n* times. This information can be specified to the compiler via the multiple argument in the **MUST_ITERATE** pragma.
- The smallest possible number of iterations of the loop
- The largest possible number of iterations of the loop

The compiler can sometimes obtain this information itself by analyzing the code. However, sometimes the compiler can be overly conservative in its assumptions and therefore generates more code than is necessary when unrolling. This can also lead to not unrolling at all.

Furthermore, if the mechanism that determines when the loop should exit is complex, the compiler may not be able to determine these properties of the loop. In these cases, you must tell the compiler the properties of the loop by using the **MUST_ITERATE** pragma.

Specifying `#pragma UNROLL(1)`; asks that the loop not be unrolled. Automatic loop unrolling also is not performed in this case.

If multiple `UNROLL` pragmas are specified for the same loop, it is undefined which pragma is used, if any.

6.10 The _Pragma Operator

The C6000 C/C++ compiler supports the C99 preprocessor `_Pragma()` operator. This preprocessor operator is similar to `#pragma` directives. However, `_Pragma` can be used in preprocessing macros (`#defines`).

The syntax of the operator is:

```
_Pragma (" string_literal ");
```

The argument *string_literal* is interpreted in the same way the tokens following a `#pragma` directive are processed. The *string_literal* must be enclosed in quotes. A quotation mark that is part of the *string_literal* must be preceded by a backward slash.

You can use the `_Pragma` operator to express `#pragma` directives in macros. For example, the `DATA_SECTION` syntax:

```
#pragma DATA_SECTION( func , " section ");
```

Is represented by the `_Pragma()` operator syntax:

```
_Pragma ("DATA_SECTION( func , \" section \")")
```

The following code illustrates using `_Pragma` to specify the `DATA_SECTION` pragma in a macro:

```
...

#define EMIT_PRAGMA(x) _Pragma(#x)
#define COLLECT_DATA(var) EMIT_PRAGMA(DATA_SECTION(var, "mysection"))

COLLECT_DATA(x)
int x;

...
```

The `EMIT_PRAGMA` macro is needed to properly expand the quotes that are required to surround the section argument to the `DATA_SECTION` pragma.

6.11 Application Binary Interface

Selecting one of the two ABIs supported by the C6000 compiler is discussed in [Section 2.16](#).

An ABI should define how functions that are written separately, and compiled or assembled separately can work together. This involves standardizing the data type representation, register conventions, and function structure and calling conventions. It should define linkname generation from C symbol names. It should define the object file format and the debug format. It should document how the system is initialized. In the case of C++ it defines C++ name mangling and exception handling support.

An application must be only one of COFF ABI and EABI; these ABIs are not compatible.

6.11.1 COFF ABI

COFF ABI is the only ABI supported by older compilers. To generate object files compatible with older COFF ABI object files, you must use COFF ABI (`--abi=coffabi`, the default). This option must also be used when assembly hand-coded assembly source files intended to be used in a COFF ABI application.

6.11.2 EABI

EABI requires the ELF object file format which enables supporting modern language features like early template instantiation and export inline functions support.

TI-specific information on EABI mode is described in [Section 7.8.2](#).

To generate object files compatible with EABI, you must use C6000 compiler version 7.2 or greater; see [Section 2.16](#). The `__TI_EABI__` predefined symbol is defined and set to 1 if compiling for EABI and is not defined otherwise.

6.12 Object File Symbol Naming Conventions (Linknames)

Each externally visible identifier is assigned a unique symbol name to be used in the object file, a so-called *linkname*. This name is assigned by the compiler according to an algorithm which depends on the name, type, and source language of the symbol. This algorithm may add a prefix to the identifier (typically an underscore), and it may *mangle* the name. This algorithm may *mangle* the name.

In COFF ABI, the linkname for all objects and functions is the same as the name in the C source with an added underscore prefix. This prevents any C identifier from colliding with any identifier in the assembly code namespace, such as an assembler keyword.

In EABI, no prefix is used. If a C identifier would collide with an assembler keyword, the compiler will escape the identifier with double parallel bars, which instructs the assembler not to treat the identifier as a keyword. You are responsible for making sure that C identifiers do not collide with user-defined assembly code identifiers.

Name mangling encodes the types of the parameters of a function in the linkname for a function. Name mangling only occurs for C++ functions which are not declared 'extern "C"'. Mangling allows function overloading, operator overloading, and type-safe linking. Be aware that the return value of the function is not encoded in the mangled name, as C++ functions cannot be overloaded based on the return value.

For COFF ABI, the mangling algorithm used closely follows that described in The Annotated Reference Manual (ARM).

For example, the general form of a C++ linkname for a function named `func` is:

`__func__F parmcodes`

Where `parmcodes` is a sequence of letters that encodes the parameter types of `func`.

For this simple C++ source file:

```
int foo(int i){ } //global C++ function
```

This is the resulting assembly code:

```
__foo__Fi
```

The linkname of foo is `_foo__Fi`, indicating that foo is a function that takes a single argument of type int. To aid inspection and debugging, a name demangling utility is provided that demangles names into those found in the original C++ source. See [Chapter 9](#) for more information.

For EABI, the mangling algorithm follows that described in the Itanium C++ ABI (<http://www.codesourcery.com/cxx-abi/abi.html>).

`int foo(int i) { }` would be mangled "`_Z3foo`"

6.13 Initializing Static and Global Variables in COFF ABI Mode

The ANSI/ISO C standard specifies that global (extern) and static variables without explicit initializations must be initialized to 0 before the program begins running. This task is typically done when the program is loaded. Because the loading process is heavily dependent on the specific environment of the target application system, in COFF ABI mode the compiler itself makes no provision for initializing to 0 otherwise uninitialized static storage class variables at run time. It is up to your application to fulfill this requirement.

Initialize Global Objects

NOTE: You should explicitly initialize all global objects which you expected the compiler would set to zero by default.

In C6000 EABI mode the uninitialized variables are zero initialized automatically.

6.13.1 Initializing Static and Global Variables With the Linker

If your loader does not preinitialize variables, you can use the linker to preinitialize the variables to 0 in the object file. For example, in the linker command file, use a fill value of 0 in the `.bss` section:

```
SECTIONS
{
...
.bss: {} = 0x00;
...
}
```

Because the linker writes a complete load image of the zeroed `.bss` section into the output COFF file, this method can have the unwanted effect of significantly increasing the size of the output file (but not the program).

If you burn your application into ROM, you should explicitly initialize variables that require initialization. The preceding method initializes `.bss` to 0 only at load time, not at system reset or power up. To make these variables 0 at run time, explicitly define them in your code.

For more information about linker command files and the `SECTIONS` directive, see the linker description information in the *TMS320C6000 Assembly Language Tools User's Guide*.

6.13.2 Initializing Static and Global Variables With the `const` Type Qualifier

Static and global variables of type `const` without explicit initializations are similar to other static and global variables because they might not be preinitialized to 0 (for the same reasons discussed in [Section 6.13](#)). For example:

```
const int zero; /* may not be initialized to 0 */
```

However, the initialization of `const` global and static variables is different because these variables are declared and initialized in a section called `.const`. For example:

```
const int zero = 0 /* guaranteed to be 0 */
```

This corresponds to an entry in the `.const` section:

```
.sect .const
_zero
.word 0
```


This feature is particularly useful for declaring a large table of constants, because neither time nor space is wasted at system startup to initialize the table. Additionally, the linker can be used to place the `.const` section in ROM.

You can use the `DATA_SECTION` pragma to put the variable in a section other than `.const`. For example, the following C code:

```
#pragma DATA_SECTION (var, ".mysect");
const int zero=0;
```

is compiled into this assembly code:

```
.sect .mysect
_zero
.word 0
```

6.14 Changing the ANSI/ISO C Language Mode

The `--kr_compatible`, `--relaxed_ansi`, and `--strict_ansi` options let you specify how the C/C++ compiler interprets your source code. You can compile your source code in the following modes:

- Normal ANSI/ISO mode
- K&R C mode
- Relaxed ANSI/ISO mode
- Strict ANSI/ISO mode

The default is normal ANSI/ISO mode. Under normal ANSI/ISO mode, most ANSI/ISO violations are emitted as errors. Strict ANSI/ISO violations (those idioms and allowances commonly accepted by C/C++ compilers, although violations with a strict interpretation of ANSI/ISO), however, are emitted as warnings. Language extensions, even those that conflict with ANSI/ISO C, are enabled.

K&R C mode does not apply to C++ code.

6.14.1 Compatibility With K&R C (`--kr_compatible` Option)

The ANSI/ISO C/C++ language is a superset of the de facto C standard defined in Kernighan and Ritchie's *The C Programming Language*. Most programs written for other non-ANSI/ISO compilers correctly compile and run without modification.

There are subtle changes, however, in the language that can affect existing code. Appendix C in *The C Programming Language* (second edition, referred to in this manual as K&R) summarizes the differences between ANSI/ISO C and the first edition's C standard (the first edition is referred to in this manual as K&R C).

To simplify the process of compiling existing C programs with the ANSI/ISO C/C++ compiler, the compiler has a K&R option (`--kr_compatible`) that modifies some semantic rules of the language for compatibility with older code. In general, the `--kr_compatible` option relaxes requirements that are stricter for ANSI/ISO C than for K&R C. The `--kr_compatible` option does not disable any new features of the language such as function prototypes, enumerations, initializations, or preprocessor constructs. Instead, `--kr_compatible` simply liberalizes the ANSI/ISO rules without revoking any of the features.

The specific differences between the ANSI/ISO version of C and the K&R version of C are as follows:

- The integral promotion rules have changed regarding promoting an unsigned type to a wider signed type. Under K&R C, the result type was an unsigned version of the wider type; under ANSI/ISO, the result type is a signed version of the wider type. This affects operations that perform differently when applied to signed or unsigned operands; namely, comparisons, division (and mod), and right shift:

```
unsigned short u;
int i;
if (u < i)          /* SIGNED comparison, unless --kr_compatible used */
```

- ANSI/ISO prohibits combining two pointers to different types in an operation. In most K&R compilers, this situation produces only a warning. Such cases are still diagnosed when `--kr_compatible` is used, but with less severity:

```
int *p;
char *q = p;       /* error without --kr_compatible, warning with --kr_compatible */
```


- External declarations with no type or storage class (only an identifier) are illegal in ANSI/ISO but legal in K&R:

```
a; /* illegal unless --kr_compatible used */
```

- ANSI/ISO interprets file scope definitions that have no initializers as *tentative definitions*. In a single module, multiple definitions of this form are fused together into a single definition. Under K&R, each definition is treated as a separate definition, resulting in multiple definitions of the same object and usually an error. For example:

```
int a;
int a; /* illegal if --kr_compatible used, OK if not */
```

Under ANSI/ISO, the result of these two definitions is a single definition for the object a. For most K&R compilers, this sequence is illegal, because int a is defined twice.

- ANSI/ISO prohibits, but K&R allows objects with external linkage to be redeclared as static:

```
extern int a;
static int a; /* illegal unless --kr_compatible used */
```

- Unrecognized escape sequences in string and character constants are explicitly illegal under ANSI/ISO but ignored under K&R:

```
char c = '\q'; /* same as 'q' if --kr_compatible used, error if not */
```

- ANSI/ISO specifies that bit fields must be of type int or unsigned. With --kr_compatible, bit fields can be legally defined with any integral type. For example:

```
struct s
{
    short f : 2; /* illegal unless --kr_compatible used */
};
```

- K&R syntax allows a trailing comma in enumerator lists:

```
enum { a, b, c, }; /* illegal unless --kr_compatible used */
```

- K&R syntax allows trailing tokens on preprocessor directives:

```
#endif NAME /* illegal unless --kr_compatible used */
```

6.14.2 Enabling Strict ANSI/ISO Mode and Relaxed ANSI/ISO Mode (--strict_ansi and --relaxed_ansi Options)

Use the --strict_ansi option when you want to compile under strict ANSI/ISO mode. In this mode, error messages are provided when non-ANSI/ISO features are used, and language extensions that could invalidate a strictly conforming program are disabled. Examples of such extensions are the inline and asm keywords.

Use the --relaxed_ansi option when you want the compiler to ignore strict ANSI/ISO violations rather than emit a warning (as occurs in normal ANSI/ISO mode) or an error message (as occurs in strict ANSI/ISO mode). In relaxed ANSI/ISO mode, the compiler accepts extensions to the ANSI/ISO C standard, even when they conflict with ANSI/ISO C. The GCC language extensions described in [Section 6.15](#) are available in relaxed ANSI/ISO mode.

6.14.3 Enabling Embedded C++ Mode (`--embedded_cpp` Option)

The compiler supports the compilation of embedded C++. In this mode, some features of C++ are removed that are of less value or too expensive to support in an embedded system. When compiling for embedded C++, the compiler generates diagnostics for the use of omitted features.

Embedded C++ is enabled by compiling with the `--embedded_cpp` option.

Embedded C++ omits these C++ features:

- Templates
- Exception handling
- Run-time type information
- The new cast syntax
- The keyword `mutable`
- Multiple inheritance
- Virtual inheritance

Under the standard definition of embedded C++, namespaces and using-declarations are not supported. The C6000 compiler nevertheless allows these features under embedded C++ because the C++ run-time-support library makes use of them. Furthermore, these features impose no run-time penalty.

The compiler does not support embedded C++ run-time-support libraries.

6.15 GNU Language Extensions

The GNU compiler collection (GCC) defines a number of language features not found in the ANSI/ISO C and C++ standards. The definition and examples of these extensions (for GCC version 3.4) can be found at the GNU web site, <http://gcc.gnu.org/onlinedocs/gcc-3.4.6/gcc/C-Extensions.html>.

Most of these extensions are also available for C++ source code.

6.15.1 Extensions

Most of the GCC language extensions are available in the TI compiler when compiling in relaxed ANSI mode (`--relaxed_ansi`) or if the `--gcc` option is used.

The extensions that the TI compiler supports are listed in [Table 6-4](#), which is based on the list of extensions found at the GNU web site. The shaded rows describe extensions that are not supported.

Table 6-4. GCC Language Extensions

Extensions	Descriptions
Statement expressions	Putting statements and declarations inside expressions (useful for creating smart 'safe' macros)
Local labels	Labels local to a statement expression
Labels as values	Pointers to labels and computed gotos
Nested functions	As in Algol and Pascal, lexical scoping of functions
Constructing calls	Dispatching a call to another function
Naming types ⁽¹⁾	Giving a name to the type of an expression
<code>typeof</code> operator	<code>typeof</code> referring to the type of an expression
Generalized lvalues	Using question mark (?) and comma (,) and casts in lvalues
Conditionals	Omitting the middle operand of a ?: expression
<code>long long</code>	Double long word integers and long long int type
Hex floats	Hexadecimal floating-point constants
Complex	Data types for complex numbers
Zero length	Zero-length arrays
Variadic macros	Macros with a variable number of arguments
Variable length	Arrays whose length is computed at run time

⁽¹⁾ Feature defined for GCC 3.0; definition and examples at <http://gcc.gnu.org/onlinedocs/gcc-3.0.4/gcc/C-Extensions.html>

Table 6-4. GCC Language Extensions (continued)

Extensions	Descriptions
Empty structures	Structures with no members
Subscripting	Any array can be subscripted, even if it is not an lvalue.
Escaped newlines	Slightly looser rules for escaped newlines
Multi-line strings ⁽¹⁾	String literals with embedded newlines
Pointer arithmetic	Arithmetic on void pointers and function pointers
Initializers	Non-constant initializers
Compound literals	Compound literals give structures, unions, or arrays as values
Designated initializers	Labeling elements of initializers
Cast to union	Casting to union type from any member of the union
Case ranges	'Case 1 ... 9' and such
Mixed declarations	Mixing declarations and code
Function attributes	Declaring that functions have no side effects, or that they can never return
Attribute syntax	Formal syntax for attributes
Function prototypes	Prototype declarations and old-style definitions
C++ comments	C++ comments are recognized.
Dollar signs	A dollar sign is allowed in identifiers.
Character escapes	The character ESC is represented as <code>\e</code>
Variable attributes	Specifying the attributes of variables
Type attributes	Specifying the attributes of types
Alignment	Inquiring about the alignment of a type or variable
Inline	Defining inline functions (as fast as macros)
Assembly labels	Specifying the assembler name to use for a C symbol
Extended asm	Assembler instructions with C operands
Constraints	Constraints for asm operands
Alternate keywords	Header files can use <code>__const__</code> , <code>__asm__</code> , etc
Explicit reg vars	Defining variables residing in specified registers
Incomplete enum types	Define an enum tag without specifying its possible values
Function names	Printable strings which are the name of the current function
Return address	Getting the return or frame address of a function (limited support)
Other built-ins	Other built-in functions (see Section 6.15.5)
Vector extensions	Using vector instructions through built-in functions
Target built-ins	Built-in functions specific to particular targets
Pragmas	Pragmas accepted by GCC
Unnamed fields	Unnamed struct/union fields within structs/unions
Thread-local	Per-thread variables

6.15.2 Function Attributes

The following GCC function attributes are supported: `always_inline`, `const`, `constructor`, `deprecated`, `format`, `format_arg`, `malloc`, `noinline`, `noreturn`, `pure`, `section`, `unused`, `used` and `warn_unused_result`.

In addition, the visibility function attribute is supported for EABI mode (`--abi=eabi`).

The `format` attribute is applied to the declarations of `printf`, `fprintf`, `sprintf`, `snprintf`, `vprintf`, `vfprintf`, `vsprintf`, `vsnprintf`, `scanf`, `fscanf`, `vfscanf`, `vscanf`, `vsscanf`, and `sscanf` in `stdio.h`. Thus when GCC extensions are enabled, the data arguments of these functions are type checked against the format specifiers in the format string argument and warnings are issued when there is a mismatch. These warnings can be suppressed in the usual ways if they are not desired.

The `malloc` attribute is applied to the declarations of `malloc`, `calloc`, `realloc` and `memalign` in `stdlib.h`.

6.15.3 Variable Attributes

The following variable attributes are supported: aligned, deprecated, mode, packed, section, transparent_union, unused, and used.

The used attribute is defined in GCC 4.2 (see <http://gcc.gnu.org/onlinedocs/gcc-4.2.4/gcc/Variable-Attributes.html#Variable-Attributes>).

The packed attribute for structure and union types is available only when there is hardware support for unaligned accesses. For C6000 this means C6400+, C6400, C6740, and C6600.

In addition, the weak variable attribute is supported for EABI mode (--abi=eabi).

6.15.4 Type Attributes

The following type attributes are supported: aligned, deprecated, packed, transparent_union, and unused.

In addition, the visibility type attribute is supported for EABI mode (--abi=eabi).

The packed attribute on struct and union types is available only for target architectures that have hardware support for unaligned access (such as C64x+, C64x).

Members of a packed structure are stored as closely to each other as possible, omitting additional bytes of padding usually added to preserve word-alignment. For example, assuming a word-size of 4 bytes ordinarily has 3 bytes of padding between members c1 and i, and another 3 bytes of trailing padding after member c2, leading to a total size of 12 bytes:

```
struct unpacked_struct { char c1; int i; char c2;};
```

However, the members of a packed struct are byte-aligned. Thus the following does not have any bytes of padding between or after members and totals 6 bytes:

```
struct __attribute__((__packed__)) packed_struct { char c1; int i; char c2;};
```

Subsequently, packed structures in an array are packed together without trailing padding between array elements.

Bit fields of a packed structure are bit-aligned. The byte alignment of adjacent struct members that are not bit fields does not change. However, there are no bits of padding between adjacent bit fields.

The packed attribute can only be applied to the original definition of a structure or union type. It cannot be applied with a typedef to a non-packed structure that has already been defined, nor can it be applied to the declaration of a struct or union object. Therefore, any given structure or union type can only be packed or non-packed, and all objects of that type will inherit its packed or non-packed attribute.

The packed attribute is not applied recursively to structure types that are contained within a packed structure. Thus, in the following example the member s retains the same internal layout as in the first example above. There is no padding between c and s, so s falls on an unaligned boundary:

```
struct __attribute__((__packed__)) outer_packed_struct { char c; struct unpacked_struct s;};
```

It is illegal to implicitly or explicitly cast the address of a packed struct member as a pointer to any non-packed type except an unsigned char. In the following example, p1, p2, and the call to foo are all illegal.

```
void foo(int *param);
struct packed_struct ps;

int *p1 = &ps.i;
int *p2 = (int *)&ps.i;
foo(&ps.i);
```

However, it is legal to explicitly cast the address of a packed struct member as a pointer to an unsigned char:

```
unsigned char *pc = (unsigned char *)&ps.i;
```

The TI compiler also supports an unpacked attribute for an enumeration type to allow you to indicate that the representation is to be an integer type that is no smaller than int; in other words, it is not *packed*.

6.15.5 Built-In Functions

The following builtin functions are supported: `__builtin_abs`, `__builtin_classify_type`, `__builtin_constant_p`, `__builtin_expect`, `__builtin_fabs`, `__builtin_fabsf`, `__builtin_frame_address`, `__builtin_labs`, `__builtin_memcpy`, and `__builtin_return_address`.

The `__builtin_frame_address` function returns zero unless the argument is a constant zero.

The `__builtin_return_address` function always returns zero.

Run-Time Environment

This chapter describes the TMS320C6000 C/C++ run-time environment. To ensure successful execution of C/C++ programs, it is critical that all run-time code maintain this environment. It is also important to follow the guidelines in this chapter if you write assembly language functions that interface with C/C++ code.

Topic	Page
7.1 Memory Model	183
7.2 Object Representation	188
7.3 Register Conventions	196
7.4 Function Structure and Calling Conventions	197
7.5 Interfacing C and C++ With Assembly Language	200
7.6 Interrupt Handling	230
7.7 Run-Time-Support Arithmetic Routines	232
7.8 System Initialization	234
7.9 Support for Multi-Threaded Applications	245

7.1 Memory Model

The C6000 compiler treats memory as a single linear block that is partitioned into subblocks of code and data. Each subblock of code or data generated by a C program is placed in its own continuous memory space. The compiler assumes that a full 32-bit address space is available in target memory.

The Linker Defines the Memory Map

NOTE: The linker, not the compiler, defines the memory map and allocates code and data into target memory. The compiler assumes nothing about the types of memory available, about any locations not available for code or data (holes), or about any locations reserved for I/O or control purposes. The compiler produces relocatable code that allows the linker to allocate code and data into the appropriate memory spaces.

For example, you can use the linker to allocate global variables into on-chip RAM or to allocate executable code into external ROM. You can allocate each block of code or data individually into memory, but this is not a general practice (an exception to this is memory-mapped I/O, although you can access physical memory locations with C/C++ pointer types).

7.1.1 Sections

The compiler produces relocatable blocks of code and data called *sections*. The sections are allocated into memory in a variety of ways to conform to a variety of system configurations. For more information about sections and allocating them, see the introductory object file information in the *TMS320C6000 Assembly Language Tools User's Guide*.

There are two basic types of sections:

- **Initialized sections** contain data or executable code. The C/C++ compiler creates the following initialized sections:
 - The **.args section** contains the command argument for a host-based loader. This section is read-only. See the `--arg_size` option for details.
 - For EABI only, the **.binit section** contains boot time copy tables. For details on BINIT, see the *TMS320C6000 Assembly Language Tools User's Guide* for linker command file information.
 - For COFF ABI only, the **.cinit section** contains tables for initializing variables and constants.
 - The **.pinit section** for COFF ABI, or the **.init_array section** for EABI, contains the table for calling global constructor tables.
 - For EABI only, the **.c6xabi.exidx section** contains the index table for exception handling. The **.c6xabi.extab section** contains un-winded instructions for exception handling. These sections are read-only. See the `--exceptions` option for details.
 - The **.name.load section** contains the compressed image of section *name*. This section is read-only. See the *TMS320C6000 Assembly Language Tools User's Guide* for information on copy tables.
 - The **.ppinfo section** contains correlation tables and the **.ppdata section** contains data tables for compiler-based profiling. See the `--gen_profile_info` option for details.
 - The **.const section** contains string literals, floating-point constants, and data defined with the C/C++ qualifiers *far* and *const* (provided the constant is not also defined as *volatile*).
 - For EABI only, the **.fardata section** reserves space for non-const, initialized far global and static variables.
 - For EABI only, the **.neardata section** reserves space for non-const, initialized near global and static variables.
 - For EABI only, the **.rodata section** reserves space for const near global and static variables.
 - The **.switch section** contains jump tables for large switch statements.
 - The **.text section** contains all the executable code.

- **Uninitialized sections** reserve space in memory (usually RAM). A program can use this space at run time to create and store variables. The compiler creates the following uninitialized sections:
 - For COFF ABI only, the **.bss section** reserves space for global and static variables. When you specify the `--rom_model` linker option, at program startup, the C boot routine copies data out of the `.cinit` section (which can be in ROM) and stores it in the `.bss` section. The compiler defines the global symbol `$bss` and assigns `$bss` the value of the starting address of the `.bss` section.
 - For EABI only, the **.bss section** reserves space for uninitialized global and static variables.
 - The **.far section** reserves space for global and static variables that are declared far.
 - The **.stack section** reserves memory for the system stack.
 - The **.systemem section** reserves space for dynamic memory allocation. The reserved space is used by dynamic memory allocation routines, such as `malloc`, `calloc`, `realloc`, or `new`. If a C/C++ program does not use these functions, the compiler does not create the `.systemem` section.

Use Only Code in Program Memory

NOTE: With the exception of code sections, the initialized and uninitialized sections cannot be allocated into internal program memory.

The assembler creates the default sections `.text`, `.bss`, and `.data`. The C/C++ compiler, however, does not use the `.data` section. You can instruct the compiler to create additional sections by using the `CODE_SECTION` and `DATA_SECTION` pragmas (see [Section 6.9.3](#) and [Section 6.9.6](#)).

7.1.2 C/C++ System Stack

The C/C++ compiler uses a stack to:

- Save function return addresses
- Allocate local variables
- Pass arguments to functions
- Save temporary results

The run-time stack grows from the high addresses to the low addresses. The compiler uses the B15 register to manage this stack. B15 is the *stack pointer* (SP), which points to the next unused location on the stack.

The linker sets the stack size, creates a global symbol, `__TI_STACK_SIZE`, and assigns it a value equal to the stack size in bytes. The default stack size is 1K bytes. You can change the stack size at link time by using the `--stack_size` option with the linker command. For more information on the `--stack_size` option, see the linker description chapter in the *TMS320C6000 Assembly Language Tools User's Guide*.

At system initialization, SP is set to the first 8-byte aligned address before the end (highest numerical address) of the `.stack` section. For C6600, SP is set to the first 16-byte aligned address. Since the position of the stack depends on where the `.stack` section is allocated, the actual address of the stack is determined at link time.

The C/C++ environment automatically decrements SP at the entry to a function to reserve all the space necessary for the execution of that function. The stack pointer is incremented at the exit of the function to restore the stack to the state before the function was entered. If you interface assembly language routines to C/C++ programs, be sure to restore the stack pointer to the same state it was in before the function was entered.

For more information about the stack and stack pointer, see [Section 7.4](#).

Unaligned SP Can Cause Application Crash

NOTE: The HWI dispatcher uses SP during an interrupt call regardless of SP alignment. Therefore, SP must never be misaligned, even for 1 cycle.

Stack Overflow

NOTE: The compiler provides no means to check for stack overflow during compilation or at run time. A stack overflow disrupts the run-time environment, causing your program to fail. Be sure to allow enough space for the stack to grow. You can use the `--entry_hook` option to add code to the beginning of each function to check for stack overflow; see [Section 2.17](#).

7.1.3 Dynamic Memory Allocation

The run-time-support library supplied with the C6000 compiler contains several functions (such as `malloc`, `calloc`, and `realloc`) that allow you to allocate memory dynamically for variables at run time.

Memory is allocated from a global pool, or heap, that is defined in the `.system` section. You can set the size of the `.system` section by using the `--heap_size=size` option with the linker command. The linker also creates a global symbol, `__TI_SYSMEM_SIZE`, and assigns it a value equal to the size of the heap in bytes. The default size is 1K bytes. For more information on the `--heap_size` option, see the linker description chapter in the *TMS320C6000 Assembly Language Tools User's Guide*.

Dynamically allocated objects are not addressed directly (they are always accessed with pointers) and the memory pool is in a separate section (`.system`); therefore, the dynamic memory pool can have a size limited only by the amount of available memory in your system. To conserve space in the `.bss` section, you can allocate large arrays from the heap instead of defining them as global or static. For example, instead of a definition such as:

```
struct big table[100];
```

Use a pointer and call the `malloc` function:

```
struct big *table
table = (struct big *)malloc(100*sizeof(struct big));
```

7.1.4 Initialization of Variables in COFF ABI

The C/C++ compiler produces code that is suitable for use as firmware in a ROM-based system. In such a system, the initialization tables in the `.cinit` section are stored in ROM. At system initialization time, the C/C++ boot routine copies data from these tables (in ROM) to the initialized variables in `.bss` (RAM).

In situations where a program is loaded directly from an object file into memory and run, you can avoid having the `.cinit` section occupy space in memory. A loader can read the initialization tables directly from the object file (instead of from ROM) and perform the initialization directly at load time instead of at run time. You can specify this to the linker by using the `--ram_model` link option. For more information, see [Section 7.8](#).

7.1.5 Data Memory Models

Several options extend the C6000 data addressing model.

7.1.5.1 Determining the Data Address Model

As of the 5.1.0 version of the compiler tools, if a near or far keyword is not specified for an object, the compiler generates far accesses to aggregate data and near accesses to all other data. This means that structures, unions, C++ classes, and arrays are not accessed through the data-page (DP) pointer.

Non-aggregate data, by default, is placed in the `.bss` section and is accessed using relative-offset addressing from the data page pointer (DP, which is B14). DP points to the beginning of the `.bss` section. Accessing data via the data page pointer is generally faster and uses fewer instructions than the mechanism used for far data accesses.

If you want to use near accesses to aggregate data, you must specify the `--mem_model:data=near` option, or declare your data with the `near` keyword.

If you have too much static and extern data to fit within a 15-bit scaled offset from the beginning of the `.bss` section, you cannot use `--mem_model:data=near`. The linker will issue an error message if there is a DP-relative data access that will not reach.

The `--mem_model:data=type` option controls how data is accessed:

<code>--mem_model:data=near</code>	Data accesses default to near
<code>--mem_model:data=far</code>	Data accesses default to far
<code>--mem_model:data=far_aggregates</code>	Data accesses to aggregate data default to far, data accesses to non-aggregate data default to near. This is the default behavior.

The `--mem_model:data` options do not affect the access to objects explicitly declared with the `near` or `far` keyword.

By default, all run-time-support data is defined as `far`.

For more information on near and far accesses to data, see [Section 6.5.4](#).

7.1.5.2 Using DP-Relative Addressing

The default behavior of the compiler is to use DP-relative addressing for near (`.bss`) data, and absolute addressing for all other (`far`) data. The `--dprel` option specifies that all data, including const data and far data, is addressed using DP-relative addressing.

The purpose of the `--dprel` option is to support a shared object model so multiple applications running simultaneously can share code, but each have their own copy of the data.

The `--dprel` option is supported for ELF only.

7.1.5.3 Const Objects as Far

The `--mem_model:const` option allows const objects to be made far independently of the `--mem_model:data` option. This enables an application with a small amount of non-const data but a large amount of const data to move the const data out of `.bss`. Also, since consts can be shared, but `.bss` cannot, it saves memory by moving the const data into `.const`.

The `--mem_model:const=type` option has the following values:

<code>--mem_model:const=data</code>	Const objects are placed according to the <code>--mem_model:data</code> option. This is the default behavior.
<code>--mem_model:const=far</code>	Const objects default to far independent of the <code>--mem_model:data</code> option.
<code>--mem_model:const=far_aggregates</code>	Const aggregate objects default to far, scalar consts default to near.

Consts that are declared far, either explicitly through the `far` keyword or implicitly using `--mem_model:const` are always placed in the `.const` section.

7.1.6 Trampoline Generation for Function Calls

The C6000 compiler generates trampolines by default. Trampolines are a method for modifying function calls at link time to reach destinations that would normally be too far away. When a function call is more than +/- 1M instructions away from its destination, the linker will generate an indirect branch (or trampoline) to that destination, and will redirect the function call to point to the trampoline. The end result is that these function calls branch to the trampoline, and then the trampoline branches to the final destination. With trampolines, you no longer need to specify memory model options to generate far calls.

7.1.7 Position Independent Data

Near global and static data are stored in the .bss section. All near data for a program must fit within 32K bytes of memory. This limit comes from the addressing mode used to access near data, which is limited to a 15-bit unsigned offset from DP (B14), which is the data page pointer.

For some applications, it may be desirable to have multiple data pages with separate instances of near data. For example, a multi-channel application may have multiple copies of the same program running with different data pages. The functionality is supported by the C6000 compiler's memory model, and is referred to as position independent data.

Position independent data means that all near data accesses are relative to the data page (DP) pointer, allowing for the DP to be changed at run time. There are three areas where position independent data is implemented by the compiler:

- Near direct memory access

```
STW  B4,*DP(_a)
.global _a
.bss  _a,4,4
```

All near direct accesses are relative to the DP.

- Near indirect memory access

```
MVK  (_a - $bss),A0
ADD  DP,A0,A0
```

The expression `(_a - $bss)` calculates the offset of the symbol `_a` from the start of the .bss section. The compiler defines the global `$bss` in generated assembly code. The value of `$bss` is the starting address of the .bss section.

- Initialized near pointers

The .cinit record for an initialized near pointer value is stored as an offset from the beginning of the .bss section. During the autoinitialization of global variables, the data page pointer is added to these offsets. (See [Section 7.8.3](#).)

7.2 Object Representation

This section explains how various data objects are sized, aligned, and accessed.

7.2.1 Data Type Storage

[Table 7-1](#) lists register and memory storage for various data types:

Table 7-1. Data Representation in Registers and Memory

Data Type	Register Storage	Memory Storage
char	Bits 0-7 of register	8 bits aligned to 8-bit boundary
unsigned char	Bits 0-7 of register	8 bits aligned to 8-bit boundary
short	Bits 0-15 of register	16 bits aligned to 16-bit boundary
unsigned short	Bits 0-15 of register	16 bits aligned to 16-bit boundary
int	Entire register	32 bits aligned to 32-bit boundary
unsigned int	Entire register	32 bits aligned to 32-bit boundary
enum	Entire register	32 bits aligned to 32-bit boundary
float	Entire register	32 bits aligned to 32-bit boundary
long (EABI)	Entire register	32 bits aligned to 32-bit boundary
unsigned long (EABI)	Entire register	32 bits aligned to 32-bit boundary
long (COFF ABI)	Bits 0-39 of even/odd register pair	64 bits aligned to 64-bit boundary
unsigned long (COFF ABI)	Bits 0-39 of even/odd register pair	64 bits aligned to 64-bit boundary
__int40_t	Even/odd register pair	64 bits aligned to 64-bit boundary
unsigned __int40_t	Even/odd register pair	64 bits aligned to 64-bit boundary
long long	Even/odd register pair	64 bits aligned to 64-bit boundary
unsigned long long	Even/odd register pair	64 bits aligned to 64-bit boundary
double	Even/odd register pair	64 bits aligned to 64-bit boundary
long double	Even/odd register pair	64 bits aligned to 64-bit boundary
__x128_t (C6600 only) ⁽¹⁾	Register quad	128-bits aligned to 128-bit boundary
struct	Members are stored as their individual types require.	Multiple of 8 bits aligned to boundary of largest member type; members are stored and aligned as their individual types require.
array	Members are stored as their individual types require.	Members are stored as their individual types require. ⁽²⁾ All arrays inside a structure are aligned according to the type of each element in the array.
pointer to data member	Bits 0-31 of register	32 bits aligned to 32-bit boundary
pointer to member function	Components stored as their individual types require	64 bits aligned to 32-bit boundary





⁽¹⁾ For details on the __x128_t container type see [Section 7.5.6](#).

⁽²⁾ For C6400, C6400+, C6740, and C6600, aligned to a 64-bit boundary. For C6200, C6700, and C6700+, aligned to a 32-bit boundary for all types 32 bits and smaller, and to a 64-bit boundary for all types larger than 32 bits. For C6600, aligned to a 128-bit boundary.

7.2.1.1 char and short Data Types (signed and unsigned)

The char and unsigned char data types are stored in memory as a single byte and are loaded to and stored from bits 0-7 of a register (see [Figure 7-1](#)). Objects defined as short or unsigned short are stored in memory as two bytes at a halfword (2 byte) aligned address and they are loaded to and stored from bits 0-15 of a register (see [Figure 7-1](#)).



In big-endian mode, 2-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 8-15 of the register and moving the second byte of memory to bits 0-7. In little-endian mode, 2-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 0-7 of the register and moving the second byte of memory to bits 8-15.

Figure 7-1. Char and Short Data Storage Format
Signed 8-bit char

Unsigned 8-bit char

Signed 16-bit short

Unsigned 16-bit short


LEGEND: S = sign, I = signed integer, U = unsigned integer, MS = most significant, LS = least significant

7.2.1.2 enum, int, and long (EABI) Data Types (signed and unsigned)


The int, unsigned int, and enum data types are stored in memory as 32-bit objects (see Figure 7-2). Objects of these types are loaded to and stored from bits 0-31 of a register. In big-endian mode, 4-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 24-31 of the register, moving the second byte of memory to bits 16-23, moving the third byte to bits 8-15, and moving the fourth byte to bits 0-7. In little-endian mode, 4-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 0-7 of the register, moving the second byte to bits 8-15, moving the third byte to bits 16-23, and moving the fourth byte to bits 24-31.

Figure 7-2. 32-Bit Data Storage Format
Signed 32-bit integer, or enum char

Unsigned 32-bit integer


LEGEND: S = sign, U = unsigned integer, I = signed integer, MS = most significant, LS = least significant

7.2.1.3 float Data Type

The float data type is stored in memory as 32-bit objects (see Figure 7-3). Objects defined as float are loaded to and stored from bits 0-31 of a register. In big-endian mode, 4-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 24-31 of the register, moving the second byte of memory to bits 16-23, moving the third byte to bits 8-15, and moving the fourth byte to bits 0-7. In little-endian mode, 4-byte objects are loaded to registers by moving the first byte (that is, the lower address) of memory to bits 0-7 of the register, moving the second byte to bits 8-15, moving the third byte to bits 16-23, and moving the fourth byte to bits 24-31.


Figure 7-3. Single-Precision Floating-Point Char Data Storage Format


LEGEND: S = sign, M = mantissa, E = exponent, MS = most significant, LS = least significant

7.2.1.6 double and long double Data Types

Double and long double data types are stored in an odd/even pair of registers (see Figure 7-8) and can only exist in a register in one format: as a pair in the format of odd register:even register (for example, A1:A0). The odd memory word contains the sign bit, exponent, and the most significant part of the mantissa. The even memory word contains the least significant part of the mantissa. In little-endian mode, the lower address is loaded into the even register and the higher address is loaded into the odd register. In big-endian mode, the higher address is loaded into the even register and the lower address is loaded into the odd register. In little-endian mode, if code is loaded from location 0, then the byte at 0 is the lowest byte of the even register. In big-endian mode, if code is loaded from location 0, then the byte at 0 is the highest byte of the odd register.

Figure 7-8. Double-Precision Floating-Point Data Storage Format



LEGEND: S = sign, M = mantissa, E = exponent, MS = most significant, LS = least significant

7.2.1.7 Pointer to Data Member Types

Pointer to data member objects are stored in memory like an unsigned int (32 bit) integral type. Its value is the byte offset to the data member in the class, plus 1. The zero value is reserved to represent the NULL pointer.

7.2.1.8 Pointer to Member Function Types

Pointer to member function objects are stored as a structure with three members, and the layout is equivalent to:

```
struct {
    short int d;
    short int i;
    union {
        void (*f) ();
        int 0; }
};
```

The parameter d is the offset to be added to the beginning of the class object for this pointer. The parameter I is the index into the virtual function table, offset by 1. The index enables the NULL pointer to be represented. Its value is -1 if the function is nonvirtual. The parameter f is the pointer to the member function if it is nonvirtual, when I is 0. The 0 is the offset to the virtual function pointer within the class object.

7.2.1.9 Structures and Arrays

A nested structure is aligned to a boundary required by the largest type it contains. For example, if the largest type in a nested structure is of type short, then the nested structure is aligned to a 2-byte boundary. If the largest type in a nested structure is of type long, unsigned long, double, or long double, then the nested structure is aligned to an 8-byte boundary.

Structures always reserve memory in multiples of the size of the largest element type. For example, if a structure contains an int, unsigned int, or float, a multiple of 4 bytes of storage is reserved in memory. Members of structures are stored in the same manner as if they were individual objects. An array member in a struct is aligned to the natural boundary of its elements.

Top-level arrays are aligned on an 8-byte boundary for C6400, C6400+, C6740, and C6600, and either a 4-byte (for all element types of 32 bits or smaller) or an 8-byte boundary for C6200, C6700, or C6700+. Top-level arrays are aligned on a 16-byte boundary for C6600. Elements of arrays are stored in the same manner as if they were individual objects.

7.2.2 Bit Fields

Bit fields are handled differently in COFF ABI and EABI modes. [Section 7.2.2.1](#) details how bit fields are handled in all modes. [Section 7.2.2.2](#) details how bit fields differ in EABI mode.

7.2.2.1 Generic Bit Fields

Bit fields are the only objects that are packed within a byte. That is, two bit fields can be stored in the same byte. Bit fields can range in size from 1 to 32 bits for COFF ABI, and 1 to 64 bits in C or larger in C++ for EABI.

For big-endian mode, bit fields are packed into registers from most significant bit (MSB) to least significant bit (LSB) in the order in which they are defined. Bit fields are packed in memory from most significant byte (MSbyte) to least significant byte (LSbyte). For little-endian mode, bit fields are packed into registers from the LSB to the MSB in the order in which they are defined, and packed in memory from LSbyte to MSbyte.

[Figure 7-9](#) illustrates bit-field packing, using the following bit field definitions:

```
struct{
    int A:7
    int B:10
    int C:3
    int D:2
    int E:9
}x;
```

A0 represents the least significant bit of the field A; A1 represents the next least significant bit, etc. Again, storage of bit fields in memory is done with a byte-by-byte, rather than bit-by-bit, transfer.

Figure 7-9. Bit-Field Packing in Big-Endian and Little-Endian Formats

Big-endian register																																		
MS																															LS			
A	A	A	A	A	A	A	A	B	B	B	B	B	B	B	B	B	B	B	C	C	C	D	D	E	E	E	E	E	E	E	E	E	E	X
6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	2	1	0	1	0	8	7	6	5	4	3	2	1	0	X			
31																															0			
Big-endian memory																																		
Byte 0							Byte 1							Byte 2							Byte 3													
A	A	A	A	A	A	B	B	B	B	B	B	B	B	B	B	C	C	C	D	D	E	E	E	E	E	E	E	E	E	E	X			
6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	2	1	0	1	0	8	7	6	5	4	3	2	1	0	X			
Little-endian register																																		
MS																															LS			
X	E	E	E	E	E	E	E	E	E	D	D	C	C	C	B	B	B	B	B	B	B	B	B	B	A	A	A	A	A	A	A	A		
X	8	7	6	5	4	3	2	1	0	1	0	2	1	0	9	8	7	6	5	4	3	2	1	0	6	5	4	3	2	1	0			
31																															0			
Little-endian memory																																		
Byte 0							Byte 1							Byte 2							Byte 3													
B	A	A	A	A	A	A	B	B	B	B	B	B	B	B	E	E	D	D	C	C	C	B	X	E	E	E	E	E	E	E	E			
0	6	5	4	3	2	1	0	8	7	6	5	4	3	2	1	1	0	1	0	2	1	0	9	X	8	7	6	5	4	3	2			

LEGEND: X = not used, MS = most significant, LS = least significant

7.2.2.2 EABI Bit Field Differences

Bit fields are handled differently in TIABI mode versus EABI mode in these ways:

- In COFF ABI, bit fields of type long long are not allowed. In EABI, long long bit fields are supported.
- In COFF ABI, all bit fields are treated as signed or unsigned int type. In EABI, bit fields are treated as the declared type.
- In COFF ABI, the size and alignment a bit field contributes to the struct containing it depends on the number of bits in the bit field. In EABI, the size and alignment of the struct containing the bit field depends on the declared type of the bit field. For example, consider the struct:

```
struct st
{
    int a:4
};
```

In COFF ABI, this struct takes up 1 byte and is aligned at 1 byte. In EABI, this struct uses up 4 bytes and is aligned at 4 bytes.

- In COFF ABI, unnamed bit fields are zero-sized bit fields do not affect the struct or union alignment. In EABI, such fields affect the alignment of the struct or union. For example, consider the struct:

```
struct st
{
    char a:4;
    int :22;
};
```

In COFF ABI, this struct uses 4 bytes and is aligned at a 1-byte boundary. In EABI, this struct uses 4 bytes and is aligned at a 4-byte boundary.

- With EABI, bit fields declared volatile are accessed according to the bit field's declared type. A volatile bit field reference generates exactly one reference to its storage; multiple volatile bit field accesses are not merged.

7.2.3 Character String Constants

In C, a character string constant is used in one of the following ways:

- To initialize an array of characters. For example:

```
char s[] = "abc";
```

When a string is used as an initializer, it is simply treated as an initialized array; each character is a separate initializer. For more information about initialization, see [Section 7.8](#).

- In an expression. For example:

```
strcpy (s, "abc");
```

When a string is used in an expression, the string itself is defined in the `.const:string` section with the `.string` assembler directive, along with a unique label that points to the string; the terminating 0 byte is explicitly added by the compiler. For example, the following lines define the string `abc`, and the terminating 0 byte (the label `SL5` points to the string):

```
.sect ".const:string"
$$SL5: .string "abc",0
```

String labels have the form `$$SLn`, where `$$` is the compiler-generated symbol prefix and `n` is a number assigned by the compiler to make the label unique. The number begins at 0 and is increased by 1 for each string defined. All strings used in a source module are defined at the end of the compiled assembly language module.

The label `$$SLn` represents the address of the string constant. The compiler uses this label to reference the string expression.

Because strings are stored in the `.const` section (possibly in ROM) and shared, it is bad practice for a program to modify a string constant. The following code is an example of incorrect string use:

```
const char *a = "abc"
a[1] = 'x'; /* Incorrect! */
```

7.3 Register Conventions

Strict conventions associate specific registers with specific operations in the C/C++ environment. If you plan to interface an assembly language routine to a C/C++ program, you must understand and follow these register conventions.

The register conventions dictate how the compiler uses registers and how values are preserved across function calls. [Table 7-2](#) summarizes how the compiler uses the TMS320C6000 registers.

The registers in [Table 7-2](#) are available to the compiler for allocation to register variables and temporary expression results. If the compiler cannot allocate a register of a required type, spilling occurs. Spilling is the process of moving a register's contents to memory to free the register for another purpose.

Objects of type double, long, long long, or long double are allocated into an odd/even register pair and are always referenced as a register pair (for example, A1:A0). The odd register contains the sign bit, the exponent, and the most significant part of the mantissa. The even register contains the least significant part of the mantissa. The A4 register is used with A5 for passing the first argument if the first argument is a double, long, long long, or long double. The same is true for B4 and B5 for the second parameter, and so on. For more information about argument-passing registers and return registers, see [Section 7.4](#).

Table 7-2. Register Usage

Register	Function Preserved By	Special Uses	Register	Function Preserved By	Special Uses
A0	Parent	–	B0	Parent	–
A1	Parent	–	B1	Parent	–
A2	Parent	–	B2	Parent	–
A3	Parent	Structure register (pointer to a returned structure) ⁽¹⁾	B3	Parent	Return register (address to return to)
A4	Parent	Argument 1 or return value	B4	Parent	Argument 2
A5	Parent	Argument 1 or return value with A4 for doubles, longs and long longs	B5	Parent	Argument 2 with B4 for doubles, longs and long longs
A6	Parent	Argument 3	B6	Parent	Argument 4
A7	Parent	Argument 3 with A6 for doubles, longs, and long longs	B7	Parent	Argument 4 with B6 for doubles, longs, and long longs
A8	Parent	Argument 5	B8	Parent	Argument 6
A9	Parent	Argument 5 with A8 for doubles, longs, and long longs	B9	Parent	Argument 6 with B8 for doubles, longs, and long longs
A10	Child	Argument 7	B10	Child	Argument 8
A11	Child	Argument 7 with A10 for doubles, longs, and long longs	B11	Child	Argument 8 with B10 for doubles, longs, and long longs
A12	Child	Argument 9	B12	Child	Argument 10
A13	Child	Argument 9 with A12 for doubles, longs, and long longs	B13	Child	Argument 10 with B12 for doubles, longs, and long longs
A14	Child	–	B14	Child	Data page pointer (DP)
A15	Child	Frame pointer (FP)	B15	Child	Stack pointer (SP)
A16-A31	Parent	C6400, C6400+, C6700+, C6740, and C6600 only	B16-B31	Parent	C6400, C6400+, C6700+, C6740, and C6600 only
ILC	Child	C6400+, C6740, and C6600 only, loop buffer counter	NRP	Parent	
IRP	Parent		RILC	Child	C6400+, C6740, and C6600 only, loop buffer counter

⁽¹⁾ For EABI, structs of size 64 or less are passed by value in registers instead of by reference using a pointer in A3.

All other control registers are not saved or restored by the compiler.

The compiler assumes that control registers not listed in [Table 7-2](#) that can have an effect on compiled code have default values. For example, the compiler assumes all circular addressing-enabled registers are set for linear addressing (the AMR is used to enable circular addressing). Enabling circular addressing and then calling a C/C++ function without restoring the AMR to a default setting violates the calling convention. You must be certain that control registers which affect compiler-generated code have a default value when calling a C/C++ function from assembly.

Assembly language programmers must be aware that the linker assumes B15 contains the stack pointer. The linker needs to save and restore values on the stack in trampoline code that it generates. If you do not use B15 as the stack pointer in assembly code, you should use the linker option that disables trampolines, `--trampolines=off`. Otherwise, trampolines could corrupt memory and overwrite register values.

7.4 Function Structure and Calling Conventions

The C/C++ compiler imposes a strict set of rules on function calls. Except for special run-time support functions, any function that calls or is called by a C/C++ function must follow these rules. Failure to adhere to these rules can disrupt the C/C++ environment and cause a program to fail.

For details on the calling conventions in EABI mode, refer to *The C6000 Embedded Application Binary Interface Application Report* ([SPRAB89](#)).

7.4.1 How a Function Makes a Call

A function (parent function) performs the following tasks when it calls another function (child function).

1. Arguments passed to a function are placed in registers or on the stack.

A function (parent function) performs the following tasks when it calls another function (child function):

If arguments are passed to a function, up to the first ten arguments are placed in registers A4, B4, A6, B6, A8, B8, A10, B10, A12, and B12. If longs, long longs, doubles, or long doubles are passed, they are placed in register pairs A5:A4, B5:B4, A7:A6, and so on.

However, for C6600, if one or more `__x128_t` arguments are passed, the next `__x128_t` argument is passed in the first available quad, where the list of available quads has the ordering: A7:A6:A5:A4, B7:B6:B5:B4, A11:A10:A9:A8, B11:B10:B9:B8. If there are no more available quads, the `__x128_t` goes onto the stack. A subsequent 32-bit, 40-bit, or 64-bit argument can take the first available register or register pair even if an earlier `__x128_t` argument has been put on the stack.

Any remaining arguments are placed on the stack (that is, the stack pointer points to the next free location; `SP + offset` points to the eleventh argument, and so on, assuming for C6600 an `__x128_t` is not passed.) Arguments placed on the stack must be aligned to a value appropriate for their size. An argument that is not declared in a prototype and whose size is less than the size of `int` is passed as an `int`. An argument that is a float is passed as double if it has no prototype declared.

A structure argument is passed as the address of the structure. It is up to the called function to make a local copy.

For a function declared with an ellipsis indicating that it is called with varying numbers of arguments, the convention is slightly modified. The last explicitly declared argument is passed on the stack, so that its stack address can act as a reference for accessing the undeclared arguments.

[Figure 7-10](#) shows the register argument conventions.

2. The calling function must save registers A0 to A9 and B0 to B9 (and A16 to A31 and B16 to B31 for C6400, C6400+, and C6700+), if their values are needed after the call, by pushing the values onto the stack.
3. The caller (parent) calls the function (child).
4. Upon returning, the caller reclaims any stack space needed for arguments by adding to the stack pointer. This step is needed only in assembly programs that were not compiled from C/C++ code. This is because the C/C++ compiler allocates the stack space needed for all calls at the beginning of the function and deallocates the space at the end of the function.

Figure 7-10. Register Argument Conventions

<code>int func1(int a,</code>	<code>int b,</code>	<code>int c);</code>							
A4	A4	B4	A6						
<code>int func2(int a,</code>	<code>float b,</code>	<code>int c)</code>	<code>struct A d,</code>	<code>float e,</code>	<code>int f,</code>	<code>int g);</code>			
A4	A4	B4	A6	B6	A8	B8	A10		
<code>int func3(int a,</code>	<code>double b,</code>	<code>float c)</code>	<code>long double d);</code>						
A4	A4	B5:B4	A6	B7:B6					
/*NOTE: The following function has a variable number of arguments. */									
<code>int vararg(int a,</code>	<code>int b,</code>	<code>int c,</code>	<code>int d);</code>						
A4	A4	B4	A6	stack					
<code>struct A func4(</code>	<code>int y);</code>								
A3	A4								
<code>__x128_t func5(</code>	<code>__x128_t a);</code>								
A7:A6:A5:A4	A7:A6:A5:A4								
<code>void func6(int a,</code>	<code>int b,</code>	<code>__x128_t c);</code>							
A4	B4	A11:A10:A9:A8							
<code>void func7(int a,</code>	<code>int b,</code>	<code>__x128_t c,</code>	<code>int d,</code>	<code>int e,</code>	<code>int f,</code>	<code>__x128_t g,</code>	<code>int h);</code>		
A4	B4	A11:A10:A9:A8	A6	B6	B8	stack	B10		

7.4.2 How a Called Function Responds

A called function (child function) must perform the following tasks:

1. The called function (child) allocates enough space on the stack for any local variables, temporary storage areas, and arguments to functions that this function might call. This allocation occurs once at the beginning of the function and may include the allocation of the frame pointer (FP).
The frame pointer is used to read arguments from the stack and to handle register spilling instructions. If any arguments are placed on the stack or if the frame size exceeds 128K bytes, the frame pointer (A15) is allocated in the following manner:
 - (a) The old A15 is saved on the stack.
 - (b) The new frame pointer is set to the current SP (B15).
 - (c) The frame is allocated by decrementing SP by a constant.
 - (d) Neither A15 (FP) nor B15 (SP) is decremented anywhere else within this function.
 If the above conditions are not met, the frame pointer (A15) is not allocated. In this situation, the frame is allocated by subtracting a constant from register B15 (SP). Register B15 (SP) is not decremented anywhere else within this function.
2. If the called function calls any other functions, the return address must be saved on the stack. Otherwise, it is left in the return register (B3) and is overwritten by the next function call.
3. If the called function modifies any registers numbered A10 to A15 or B10 to B15, it must save them, either in other registers or on the stack. The called function can modify any other registers without saving them.
4. If the called function expects a structure argument, it receives a pointer to the structure instead. If writes are made to the structure from within the called function, space for a local copy of the structure must be allocated on the stack and the local structure must be copied from the passed pointer to the structure. If no writes are made to the structure, it can be referenced in the called function indirectly through the pointer argument.

You must be careful to declare functions properly that accept structure arguments, both at the point where they are called (so that the structure argument is passed as an address) and at the point where they are declared (so the function knows to copy the structure to a local copy).

5. The called function executes the code for the function.
6. If the called function returns any integer, pointer, or float type, the return value is placed in the A4 register. If the function returns a double, long double, long, or long long type, the value is placed in the A5:A4 register pair. For C6600 if the function returns a `__x128_t`, the value is placed in A7:A6:A5:A4. If the function returns a structure, the caller allocates space for the structure and passes the address of the return space to the called function in A3. To return a structure, the called function copies the structure to the memory block pointed to by the extra argument.

In this way, the caller can be smart about telling the called function where to return the structure. For example, in the statement `s = f(x)`, where `s` is a structure and `f` is a function that returns a structure, the caller can actually make the call as `f(&s, x)`. The function `f` then copies the return structure directly into `s`, performing the assignment automatically.

If the caller does not use the return structure value, an address value of 0 can be passed as the first argument. This directs the called function not to copy the return structure.

You must be careful to declare functions properly that return structures, both at the point where they are called (so that the extra argument is passed) and at the point where they are declared (so the function knows to copy the result).

7. Any register numbered A10 to A15 or B10 to B15 that was saved in is restored.
8. If A15 was used as a frame pointer (FP), the old value of A15 is restored from the stack. The space allocated for the function in is reclaimed at the end of the function by adding a constant to register B15 (SP).
9. The function returns by jumping to the value of the return register (B3) or the saved value of the return register.

7.4.3 Accessing Arguments and Local Variables

A function accesses its stack arguments and local nonregister variables indirectly through register A15 (FP) or through register B15 (SP), one of which points to the top of the stack. Since the stack grows toward smaller addresses, the local and argument data for a function are accessed with a positive offset from FP or SP. Local variables, temporary storage, and the area reserved for stack arguments to functions called by this function are accessed with offsets smaller than the constant subtracted from FP or SP at the beginning of the function.

Stack arguments passed to this function are accessed with offsets greater than or equal to the constant subtracted from register FP or SP at the beginning of the function. The compiler attempts to keep register arguments in their original registers if optimization is used or if they are defined with the register keyword. Otherwise, the arguments are copied to the stack to free those registers for further allocation.

For information on whether FP or SP is used to access local variables, temporary storage, and stack arguments, see [Section 7.4.2](#). For more information on the C/C++ System stack, see [Section 7.1.2](#).

7.5 Interfacing C and C++ With Assembly Language

The following are ways to use assembly language with C/C++ code:

- Use separate modules of assembled code and link them with compiled C/C++ modules (see [Section 7.5.1](#)).
- Use assembly language variables and constants in C/C++ source (see [Section 7.5.2](#)).
- Use inline assembly language embedded directly in the C/C++ source (see [Section 7.5.4](#)).
- Use intrinsics in C/C++ source to directly call an assembly language statement (see [Section 7.5.5](#)).

7.5.1 Using Assembly Language Modules With C/C++ Code

Interfacing C/C++ with assembly language functions is straightforward if you follow the calling conventions defined in [Section 7.4](#), and the register conventions defined in [Section 7.3](#). C/C++ code can access variables and call functions defined in assembly language, and assembly code can access C/C++ variables and call C/C++ functions.

Follow these guidelines to interface assembly language and C:

- All functions, whether they are written in C/C++ or assembly language, must follow the register conventions outlined in [Section 7.3](#).
- You must preserve registers A10 to A15, B3, and B10 to B15, and you may need to preserve A3. If you use the stack normally, you do not need to explicitly preserve the stack. In other words, you are free to use the stack inside a function as long as you pop everything you pushed before your function exits. You can use all other registers freely without preserving their contents.
- A10 to A15 and B10 to B15 need to be restored before a function returns, even if any of A10 to A13 and B10 to B13 are being used for passing arguments.
- Interrupt routines must save *all* the registers they use. For more information, see [Section 7.6](#).
- When you call a C/C++ function from assembly language, load the designated registers with arguments and push the remaining arguments onto the stack as described in [Section 7.4.1](#).
Remember that only A10 to A15 and B10 to B15 are preserved by the C/C++ compiler. C/C++ functions can alter any other registers, save any other registers whose contents need to be preserved by pushing them onto the stack before the function is called, and restore them after the function returns.
- Functions must return values correctly according to their C/C++ declarations. Integers and 32-bit floating-point (float) values are returned in A4. Doubles, long doubles, longs, and long longs are returned in A5:A4. For C6600 `__x128_t` values are returned in A7:A6:A5:A4. Structures are returned by copying them to the address in A3.
- No assembly module should use the `.cinit` section for any purpose other than autoinitialization of global variables. The C/C++ startup routine assumes that the `.cinit` section consists *entirely* of initialization tables. Disrupting the tables by putting other information in `.cinit` can cause unpredictable results.
- The compiler assigns linknames to all external objects. Thus, when you are writing assembly language code, you must use the same linknames as those assigned by the compiler. See [Section 6.12](#) for more information.
- Any object or function declared in assembly language that is accessed or called from C/C++ must be declared with the `.def` or `.global` directive in the assembly language modifier. This declares the symbol as external and allows the linker to resolve references to it.
Likewise, to access a C/C++ function or object from assembly language, declare the C/C++ object with the `.ref` or `.global` directive in the assembly language module. This creates an undeclared external reference that the linker resolves.
- The SGIE bit of the TSR control register may need to be saved. Please see [Section 7.6.1](#) for more information.
- The compiler assumes that control registers not listed in [Table 7-2](#) that can have an effect on compiled code have default values. For example, the compiler assumes all circular-addressing-enabled registers are set for linear addressing (the AMR is used to enable circular addressing). Enabling circular addressing and then calling a C/C++ function without restoring the AMR to a default setting violates the calling convention. Also, enabling circular addressing and having interrupts enabled violates the calling

convention. You must be certain that control registers that affect compiler-generated code have a default value when calling a C/C++ function from assembly.

- Assembly language programmers must be aware that the linker assumes B15 contains the stack pointer. The linker needs to save and restore values on the stack in trampoline code that it generates. If you do not use B15 as the stack pointer in your assembly code, you should use the linker option that disables trampolines, `--trampolines=off`. Otherwise, trampolines could corrupt memory and overwrite register values.
- Assembly code that utilizes B14 and/or B15 for localized purposes other than the data-page pointer and stack pointer may violate the calling convention. The assembly programmer needs to protect these areas of non-standard use of B14 and B15 by turning off interrupts around this code. Because interrupt handling routines need the stack (and thus assume the stack pointer is in B15) interrupts need to be turned off around this code. Furthermore, because interrupt service routines may access global data and may call other functions which access global data, this special treatment also applies to B14. After the data-page pointer and stack pointer have been restored, interrupts may be turned back on.

Example 7-1 illustrates a C++ function called `main`, which calls an assembly language function called `asmfunc`, **Example 7-2**. The `asmfunc` function takes its single argument, adds it to the C++ global variable called `gvar`, and returns the result.

Example 7-1. Calling an Assembly Language Function From a C/C++ Program

```
extern "C" {
extern int asmfunc(int a); /* declare external asm function */
int gvar = 0;             /* define global variable          */
}

void main()
{
    int I = 5;

    I = asmfunc(I);      /* call function normally      */
}
```

Example 7-2. Assembly Language Program Called by [Example 7-1](#)

```
.global _asmfunc
.global _gvar
_asmfunc:
LDW    *+b14(_gvar),A3
NOP    4
ADD    a3,a4,a3
STW    a3,*b14(_gvar)
MV     a3,a4
B      b3
NOP    5
```

In the C++ program in **Example 7-1**, the `extern` declaration of `asmfunc` is optional because the return type is `int`. Like C/C++ functions, you need to declare assembly functions only if they return noninteger values or pass noninteger parameters.

NOTE: SP Semantics

The stack pointer must always be 8-byte aligned. For C6600 the stack pointer must always be 16-byte aligned. This is automatically performed by the C compiler and system initialization code in the run-time-support libraries. Any hand assembly code that has interrupts enabled or calls a function defined in C or linear assembly source should also reserve a multiple of 8 bytes (multiple of 16 bytes for C6600) on the stack.

NOTE: Stack Allocation

Even though the compiler guarantees a doubleword alignment of the stack and the stack pointer (SP) points to the next free location in the stack space, there is only enough guaranteed room to store one 32-bit word at that location. The called function must allocate space to store the doubleword.

NOTE: Alignment of __x128_t Data Objects (C6600 Only)

The C6600 introduces the 128-bit container type `__x128_t`. Global data objects of this type are aligned to a 16-byte boundary (128 bits). Local `__x128_t` variables are allocated on the stack, but are not necessarily aligned on a 16-byte boundary since its actual alignment will depend on the alignment of the stack and the SP-relative offset of the local `__x128_t` type object. For the C6400/C6400+/C6740/C6600 ISAs, the compiler will always align the stack to an 8-byte boundary.

7.5.2 Accessing Assembly Language Variables From C/C++

It is sometimes useful for a C/C++ program to access variables or constants defined in assembly language. There are several methods that you can use to accomplish this, depending on where and how the item is defined: a variable defined in the `.bss` section, a variable not defined in the `.bss` section, or a constant.

7.5.2.1 Accessing Assembly Language Global Variables

Accessing uninitialized variables from the `.bss` section or a section named with `.usect` is straightforward:

1. Use the `.bss` or `.usect` directive to define the variable.
2. When you use `.usect`, the variable is defined in a section other than `.bss` and therefore must be declared far in C.
3. Use the `.def` or `.global` directive to make the definition external.
4. Use the appropriate linkname in assembly language.
5. In C/C++, declare the variable as *extern* and access it normally.

[Example 7-4](#) and [Example 7-3](#) show how you can access a variable defined in `.bss`.

Example 7-3. Assembly Language Variable Program

```
* Note the use of underscores in the following lines

        .bss      _var1,4,4      ; Define the variable
        .global  var1           ; Declare it as external

_var2   .usect   "mysect",4,4   ; Define the variable
        .global  _var2          ; Declare it as external
```

Example 7-4. C Program to Access Assembly Language From [Example 7-3](#)

```
extern int var1;      /* External variable */
extern far int var2;  /* External variable */
var1 = 1;            /* Use the variable */
var2 = 1;            /* Use the variable */
```

7.5.2.2 Accessing Assembly Language Constants

You can define global constants in assembly language by using the `.set`, `.def`, and `.global` directives, or you can define them in a linker command file using a linker assignment statement. These constants are accessible from C/C++ only with the use of special operators.

For normal variables defined in C/C++ or assembly language, the symbol table contains the *address of the value* of the variable. For assembler constants, however, the symbol table contains the *value* of the constant. The compiler cannot tell which items in the symbol table are values and which are addresses.

If you try to access an assembler (or linker) constant by name, the compiler attempts to fetch a value from the address represented in the symbol table. To prevent this unwanted fetch, you must use the `&` (address of) operator to get the value. In other words, if `x` is an assembly language constant, its value in C/C++ is `&x`.

You can use casts and `#defines` to ease the use of these symbols in your program, as in [Example 7-5](#) and [Example 7-6](#).

Example 7-5. Accessing an Assembly Language Constant From C

```
extern int table_size;          /*external ref */
#define TABLE_SIZE ((int) (&table_size))
.                               /* use cast to hide address-of */
.
.
for (I=0; i<TABLE_SIZE; ++I) /* use like normal symbol */
```

Example 7-6. Assembly Language Program for [Example 7-5](#)

```
_table_size .set 10000 ; define the constant
.global _table_size ; make it global
```

Because you are referencing only the symbol's value as stored in the symbol table, the symbol's declared type is unimportant. In [Example 7-5](#), `int` is used. You can reference linker-defined symbols in a similar manner.

7.5.3 Sharing C/C++ Header Files With Assembly Source

You can use the `.cdecls` assembler directive to share C headers containing declarations and prototypes between C and assembly code. Any legal C/C++ can be used in a `.cdecls` block and the C/C++ declarations will cause suitable assembly to be generated automatically, allowing you to reference the C/C++ constructs in assembly code. For more information, see the C/C++ header files chapter in the *TMS320C6000 Assembly Language Tools User's Guide*.

7.5.4 Using Inline Assembly Language

Within a C/C++ program, you can use the `asm` statement to insert a single line of assembly language into the assembly language file created by the compiler. A series of `asm` statements places sequential lines of assembly language into the compiler output with no intervening code. For more information, see [Section 6.8](#).

The `asm` statement is useful for inserting comments in the compiler output. Simply start the assembly code string with a semicolon (;) as shown below:

```
asm(" ;*** this is an assembly language comment");
```

NOTE: Using the `asm` Statement

Keep the following in mind when using the `asm` statement:

- Be extremely careful not to disrupt the C/C++ environment. The compiler does not check or analyze the inserted instructions.
 - Avoid inserting jumps or labels into C/C++ code because they can produce unpredictable results by confusing the register-tracking algorithms that the code generator uses.
 - Do not change the value of a C/C++ variable when using an `asm` statement. This is because the compiler does not verify such statements. They are inserted as is into the assembly code, and potentially can cause problems if you are not sure of their effect.
 - Do not use the `asm` statement to insert assembler directives that change the assembly environment.
 - Avoid creating assembly macros in C code and compiling with the `--symdebug:dwarf` (or `-g`) option. The C environment's debug information and the assembly macro expansion are not compatible.
-

7.5.5 Using Intrinsic to Access Assembly Language Statements

The C6000 compiler recognizes a number of intrinsic operators. Intrinsic allow you to express the meaning of certain assembly statements that would otherwise be cumbersome or inexpressible in C/C++. Intrinsic are used like functions; you can use C/C++ variables with these intrinsic, just as you would with any normal function.

The intrinsic are specified with a leading underscore, and are accessed by calling them as you do a function. For example:

```
int x1, x2, y;
y = _sadd(x1, x2);
```

Intrinsic Instructions in C Versus Assembly Language

NOTE: In some instances, an intrinsic's exact corresponding assembly language instruction may not be used by the compiler. When this is the case, the meaning of the program does not change.

[Table 7-4](#) provides a summary of the C6000 intrinsic clarifying which devices support which intrinsic.

Table 7-3. C6000 C/C++ Intrinsic Support by Device

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_abs	Yes	Yes	Yes	Yes	Yes	Yes
_abs2		Yes	Yes	Yes		Yes
_add2	Yes	Yes	Yes	Yes	Yes	Yes
_add4		Yes	Yes	Yes		Yes
_addsub			Yes	Yes		Yes
_addsub2			Yes	Yes		Yes
_amem2	Yes	Yes	Yes	Yes	Yes	Yes
_amem2_const	Yes	Yes	Yes	Yes	Yes	Yes
_amem4	Yes	Yes	Yes	Yes	Yes	Yes
_amem4_const	Yes	Yes	Yes	Yes	Yes	Yes
_amem8		Yes	Yes	Yes		Yes
_amem8_const		Yes	Yes	Yes		Yes
_amem8_f2		Yes	Yes	Yes		Yes
_amem8_f2_const		Yes	Yes	Yes		Yes
_amemd8	Yes	Yes	Yes	Yes	Yes	Yes
_amemd8_const	Yes	Yes	Yes	Yes	Yes	Yes
_avg2		Yes	Yes	Yes		Yes
_avgu4		Yes	Yes	Yes		Yes
_bitc4		Yes	Yes	Yes		Yes
_bitr		Yes	Yes	Yes		Yes
_ccmatmpy				Yes		
_ccmatmpyr1				Yes		
_ccmpy32r1				Yes		
_clr	Yes	Yes	Yes	Yes	Yes	Yes
_clrr	Yes	Yes	Yes	Yes	Yes	Yes
_cmatmpy				Yes		
_cmatmpyr1				Yes		
_cmpeq2		Yes	Yes	Yes		Yes
_cmpeg4		Yes	Yes	Yes		Yes
_cmpgt2		Yes	Yes	Yes		Yes
_cmpgtu4		Yes	Yes	Yes		Yes
_cmplt2		Yes	Yes	Yes		Yes
_cmpltu4		Yes	Yes	Yes		Yes
_cmpy			Yes	Yes		Yes
_cmpy32r1				Yes		
_cmpyr			Yes	Yes		Yes
_cmpyr1			Yes	Yes		Yes
_cmpyisp				Yes		
_complex_conjugate_mpyisp				Yes		
_complex_mpyisp				Yes		
_crot270				Yes		
_crot90				Yes		
_dadd				Yes		
_dadd2				Yes		
_daddsp				Yes		
_dadd_c				Yes		
_dapys2				Yes		

Table 7-3. C6000 C/C++ Intrinsic Support by Device (continued)

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_davg2				Yes		
_davgnr2				Yes		
_davgnr4				Yes		
_davgu4				Yes		
_dccmpyr1				Yes		
_dcmpeq2				Yes		
_dcmpeq4				Yes		
_dcmpgt2				Yes		
_dcmpgtu4				Yes		
_dccmpy				Yes		
_dcmpy				Yes		
_dcmpyr1				Yes		
_dcrot90				Yes		
_dcrot270				Yes		
_ddotp4			Yes	Yes		Yes
_ddotp4h				Yes		
_ddotph2			Yes	Yes		Yes
_ddotph2r			Yes	Yes		Yes
_ddotpl2			Yes	Yes		Yes
_ddotpl2r			Yes	Yes		Yes
_ddotpsu4h				Yes		
_deal		Yes	Yes	Yes		Yes
_dinthsp				Yes		
_dinthspu				Yes		
_dintsp				Yes		
_dintspu				Yes		
_dmax2				Yes		
_dmaxu4				Yes		
_dmin2				Yes		
_dminu4				Yes		
_dmpy2				Yes		
_dmpysp				Yes		
_dmpysu4				Yes		
_dmpyu2				Yes		
_dmpyu4				Yes		
_dmv			Yes	Yes		Yes
_dmvd				Yes		
_dotp2		Yes	Yes	Yes		Yes
_dotp4h				Yes		
_dotp4hll				Yes		
_dotpn2		Yes	Yes	Yes		Yes
_dotpnrsu2		Yes	Yes	Yes		Yes
_dotpnrus2		Yes	Yes	Yes		Yes
_dotprsu2		Yes	Yes	Yes		Yes
_dotpsu4		Yes	Yes	Yes		Yes
_dotpus4		Yes	Yes	Yes		Yes
_dotpsu4h				Yes		

Table 7-3. C6000 C/C++ Intrinsic Support by Device (continued)

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_dotpsu4hll				Yes		
_dotpu4		Yes	Yes	Yes		Yes
_dpack2			Yes	Yes		Yes
_dpackh2				Yes		
_dpackh4				Yes		
_dpacklh2				Yes		
_dpacklh4				Yes		
_dpackl2				Yes		
_dpackl4				Yes		
_dpackx2			Yes	Yes		Yes
_dpint				Yes	Yes	Yes
_dsadd				Yes		
_dsadd2				Yes		
_dshl				Yes		
_dshl2				Yes		
_dshr				Yes		
_dshr2				Yes		
_dshru				Yes		
_dshru2				Yes		
_dsmpy2				Yes		
_dspacku4				Yes		
_dspint				Yes		
_dspinth				Yes		
_dssub				Yes		
_dssub2				Yes		
_dsub				Yes		
_dsub2				Yes		
_dsubsp				Yes		
_dtol	Yes	Yes	Yes	Yes	Yes	Yes
_dtoll	Yes	Yes	Yes	Yes	Yes	Yes
_d xpnd2				Yes		
_d xpnd4				Yes		
_ext	Yes	Yes	Yes	Yes	Yes	Yes
_extr	Yes	Yes	Yes	Yes	Yes	Yes
_extu	Yes	Yes	Yes	Yes	Yes	Yes
_extur	Yes	Yes	Yes	Yes	Yes	Yes
_f2tol				Yes	Yes	Yes
_f2toll				Yes	Yes	Yes
_fabs				Yes	Yes	Yes
_fabsf				Yes	Yes	Yes
_fdmvd_f2				Yes		
_fmdv_f2			Yes	Yes		Yes
_ftoi	Yes	Yes	Yes	Yes	Yes	Yes
_gmpy			Yes	Yes		Yes
_gmpy4		Yes	Yes	Yes		Yes
_hi	Yes	Yes	Yes	Yes	Yes	Yes
_hill	Yes	Yes	Yes	Yes	Yes	Yes

Table 7-3. C6000 C/C++ Intrinsic Support by Device (continued)

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_itod	Yes	Yes	Yes	Yes	Yes	Yes
_itof	Yes	Yes	Yes	Yes	Yes	Yes
_itoll	Yes	Yes	Yes	Yes	Yes	Yes
_labs	Yes	Yes	Yes	Yes	Yes	Yes
_land				Yes		
_landn				Yes		
_ldotp2		Yes	Yes	Yes		Yes
_lmbd	Yes	Yes	Yes	Yes	Yes	Yes
_lnorm	Yes	Yes	Yes	Yes	Yes	Yes
_lo	Yes	Yes	Yes	Yes	Yes	Yes
_loll	Yes	Yes	Yes	Yes	Yes	Yes
_lor				Yes		
_lsadd	Yes	Yes	Yes	Yes	Yes	Yes
_lssub	Yes	Yes	Yes	Yes	Yes	Yes
_ltod	Yes	Yes	Yes	Yes	Yes	Yes
_lltod	Yes	Yes	Yes	Yes	Yes	Yes
_lltof2				Yes	Yes	Yes
_ltof2				Yes	Yes	Yes
_max2		Yes	Yes	Yes		Yes
_maxu4		Yes	Yes	Yes		Yes
_mfence				Yes		
_min2		Yes	Yes	Yes		Yes
_minu4		Yes	Yes	Yes		Yes
_mem2		Yes	Yes	Yes		Yes
_mem2_const		Yes	Yes	Yes		Yes
_mem4		Yes	Yes	Yes		Yes
_mem4_const		Yes	Yes	Yes		Yes
_mem8		Yes	Yes	Yes		Yes
_mem8_const		Yes	Yes	Yes		Yes
_mem8_f2				Yes	Yes	Yes
_mem8_f2_const				Yes	Yes	Yes
_memd8		Yes	Yes	Yes		Yes
_memd8_const		Yes	Yes	Yes		Yes
_mpy	Yes	Yes	Yes	Yes	Yes	Yes
_mpy2ir			Yes	Yes		Yes
_mpy2ll		Yes	Yes	Yes		Yes
_mpy32			Yes	Yes		Yes
_mpy32ll			Yes	Yes		Yes
_mpy32su			Yes	Yes		Yes
_mpy32u			Yes	Yes		Yes
_mpy32us			Yes	Yes		Yes
_mpyh	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhill		Yes	Yes	Yes		Yes
_mpyihll		Yes	Yes	Yes		Yes
_mpyilll		Yes	Yes	Yes		Yes
_mpyhir		Yes	Yes	Yes		Yes
_mpyihr		Yes	Yes	Yes		Yes

Table 7-3. C6000 C/C++ Intrinsic Support by Device (continued)

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_mpylr		Yes	Yes	Yes		Yes
_mpyhl	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhlu	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhslu	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhsu	Yes	Yes	Yes	Yes	Yes	Yes
_myphu	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhuls	Yes	Yes	Yes	Yes	Yes	Yes
_mpyhus	Yes	Yes	Yes	Yes	Yes	Yes
_mpyidl				Yes	Yes	Yes
_mpylh	Yes	Yes	Yes	Yes	Yes	Yes
_mpylhu	Yes	Yes	Yes	Yes	Yes	Yes
_mpylill		Yes	Yes	Yes		Yes
_mpylir		Yes	Yes	Yes		Yes
_mpylshu	Yes	Yes	Yes	Yes	Yes	Yes
_mpyluhs	Yes	Yes	Yes	Yes	Yes	Yes
_mpysp2dp				Yes	Yes	Yes
_mpyspdp				Yes	Yes	Yes
_mpysu	Yes	Yes	Yes	Yes	Yes	Yes
_mpysu4ll		Yes	Yes	Yes		Yes
_mpyus4ll		Yes	Yes	Yes		Yes
_mpyu	Yes	Yes	Yes	Yes	Yes	Yes
_mpyu2				Yes		
_mpyu4ll		Yes	Yes	Yes		Yes
_mpyus	Yes	Yes	Yes	Yes	Yes	Yes
_mvd		Yes	Yes	Yes		Yes
_nassert	Yes	Yes	Yes	Yes	Yes	Yes
_norm	Yes	Yes	Yes	Yes	Yes	Yes
_pack2		Yes	Yes	Yes		Yes
_packh2		Yes	Yes	Yes		Yes
_packh4		Yes	Yes	Yes		Yes
_packhl2		Yes	Yes	Yes		Yes
_packl4		Yes	Yes	Yes		Yes
_packlh2		Yes	Yes	Yes		Yes
_qmpy32				Yes		
_qmpysp				Yes		
_qsmphy32r1				Yes		
_rcpdp				Yes	Yes	Yes
_rcpsp				Yes	Yes	Yes
_rsqrdp				Yes	Yes	Yes
_rsqrsp				Yes	Yes	Yes
_rotl		Yes	Yes	Yes		Yes
_rpack2			Yes	Yes		Yes
_sadd	Yes	Yes	Yes	Yes	Yes	Yes
_sadd2		Yes	Yes	Yes		Yes
_saddsub			Yes	Yes		Yes
_saddsub2			Yes	Yes		Yes
_saddu4		Yes	Yes	Yes		Yes

Table 7-3. C6000 C/C++ Intrinsic Support by Device (continued)

Intrinsic	C6200	C6400	C6400+	C6600	C6700/C6700 +	C6740
_saddus2		Yes	Yes	Yes		Yes
_saddsu2		Yes	Yes	Yes		Yes
_sat	Yes	Yes	Yes	Yes	Yes	Yes
_set	Yes	Yes	Yes	Yes	Yes	Yes
_setr	Yes	Yes	Yes	Yes	Yes	Yes
_shfl		Yes	Yes	Yes		Yes
_shfl3			Yes	Yes		Yes
_shl2				Yes		
_shlmb		Yes	Yes	Yes		Yes
_shr2		Yes	Yes	Yes		Yes
_shrmb		Yes	Yes	Yes		Yes
_shru2		Yes	Yes	Yes		Yes
_smpy	Yes	Yes	Yes	Yes	Yes	Yes
_smpy2ll		Yes	Yes	Yes		Yes
_smpy32			Yes	Yes		Yes
_smpyh	Yes	Yes	Yes	Yes	Yes	Yes
_smpyhl	Yes	Yes	Yes	Yes	Yes	Yes
_smpylh	Yes	Yes	Yes	Yes	Yes	Yes
_spack2		Yes	Yes	Yes		Yes
_spacku4		Yes	Yes	Yes		Yes
_spint				Yes	Yes	Yes
_sshl	Yes	Yes	Yes	Yes	Yes	Yes
_sshvl		Yes	Yes	Yes		Yes
_sshvr		Yes	Yes	Yes		Yes
_ssub	Yes	Yes	Yes	Yes	Yes	Yes
_ssub2			Yes	Yes		Yes
_sub2	Yes	Yes	Yes	Yes	Yes	Yes
_sub4		Yes	Yes	Yes		Yes
_subabs4		Yes	Yes	Yes		Yes
_subc	Yes	Yes	Yes	Yes	Yes	Yes
_swap4		Yes	Yes	Yes		Yes
_swap2		Yes	Yes	Yes		Yes
_unpkbu4				Yes		
_unpkh2				Yes		
_unpkhu2				Yes		
_unpkhu4		Yes	Yes	Yes		Yes
_unpklu4		Yes	Yes	Yes		Yes
_xorll_c				Yes		
_xormpy			Yes	Yes		Yes
_xpnd2		Yes	Yes	Yes		Yes
_xpnd4		Yes	Yes	Yes		Yes

The intrinsics listed in [Table 7-4](#) can be used on all C6000 devices. They correspond to the indicated C6000 assembly language instruction(s). See the *TMS320C6000 CPU and Instruction Set Reference Guide* for more information.

See [Table 7-5](#) for the listing of C6400-specific intrinsics, which are also compatible with C6400+, C6740, and C6600. See [Table 7-6](#) for the listing of C6400+-specific intrinsics, which are also compatible with C6740 and C6600 devices. See [Table 7-7](#) for the listing of C6700-specific intrinsics. See [Table 7-8](#) for a listing of C6600-specific intrinsics.

Table 7-4. TMS320C6000 C/C++ Compiler Intrinsics

C/C++ Compiler Intrinsic	Assembly Instruction	Description
int _abs (int <i>src</i>); int _labs (__int40_t <i>src</i>);	ABS	Returns the saturated absolute value of <i>src</i>
int _add2 (int <i>src1</i> , int <i>src2</i>);	ADD2	Adds the upper and lower halves of <i>src1</i> to the upper and lower halves of <i>src2</i> and returns the result. Any overflow from the lower half add does not affect the upper half add.
ushort & _amem2 (void * <i>ptr</i>);	LDHU STHU	Allows aligned loads and stores of 2 bytes to memory. The pointer must be aligned to a two-byte boundary. ⁽¹⁾
const ushort & _amem2_const (const void * <i>ptr</i>);	LDHU	Allows aligned loads of 2 bytes from memory. The pointer must be aligned to a two-byte boundary. ⁽¹⁾
unsigned & _amem4 (void * <i>ptr</i>);	LDW STW	Allows aligned loads and stores of 4 bytes to memory. The pointer must be aligned to a four-byte boundary. ⁽¹⁾
const unsigned & _amem4_const (const void * <i>ptr</i>);	LDW	Allows aligned loads of 4 bytes from memory. The pointer must be aligned to a four-byte boundary. ⁽¹⁾
double & _amemd8 (void * <i>ptr</i>);	LDW/LDW STW/STW	Allows aligned loads and stores of 8 bytes to memory. The pointer must be aligned to an eight-byte boundary. ⁽¹⁾⁽²⁾ For C6400 _amemd8 corresponds to different assembly instructions than when used with other C6000 devices; see Table 7-5 for specifics.
const double & _amemd8_const (const void * <i>ptr</i>);	LDDW	Allows aligned loads of 8 bytes from memory. The pointer must be aligned to an eight-byte boundary. ⁽¹⁾⁽²⁾
unsigned _clr (unsigned <i>src2</i> , unsigned <i>csta</i> , unsigned <i>cstb</i>);	CLR	Clears the specified field in <i>src2</i> . The beginning and ending bits of the field to be cleared are specified by <i>csta</i> and <i>cstb</i> , respectively.
unsigned _clrr (unsigned <i>src2</i> , int <i>src1</i>);	CLR	Clears the specified field in <i>src2</i> . The beginning and ending bits of the field to be cleared are specified by the lower 10 bits of <i>src1</i> .
_int40_t dtol (double <i>src</i>);		Reinterprets double register pair <i>src</i> as an _int40_t (stored as a register pair).
long long _dtoll (double <i>src</i>);		Reinterprets double register pair <i>src</i> as a long long register pair.
int _ext (int <i>src2</i> , unsigned <i>csta</i> , unsigned <i>cstb</i>);	EXT	Extracts the specified field in <i>src2</i> , sign-extended to 32 bits. The extract is performed by a shift left followed by a signed shift right; <i>csta</i> and <i>cstb</i> are the shift left and shift right amounts, respectively.
int _extr (int <i>src2</i> , int <i>src1</i>);	EXT	Extracts the specified field in <i>src2</i> , sign-extended to 32 bits. The extract is performed by a shift left followed by a signed shift right; the shift left and shift right amounts are specified by the lower 10 bits of <i>src1</i> .
unsigned _extu (unsigned <i>src2</i> , unsigned <i>csta</i> , unsigned <i>cstb</i>);	EXTU	Extracts the specified field in <i>src2</i> , zero-extended to 32 bits. The extract is performed by a shift left followed by an unsigned shift right; <i>csta</i> and <i>cstb</i> are the shift left and shift right amounts, respectively.
unsigned _extur (unsigned <i>src2</i> , int <i>src1</i>);	EXTU	Extracts the specified field in <i>src2</i> , zero-extended to 32 bits. The extract is performed by a shift left followed by an unsigned shift right; the shift left and shift right amounts are specified by the lower 10 bits of <i>src1</i> .
unsigned _ftoi (float <i>src</i>);		Reinterprets the bits in the float as an unsigned. For example: _ftoi (1.0) == 1065353216U
unsigned _hi (double <i>src</i>);		Returns the high (odd) register of a double register pair
unsigned _hill (long long <i>src</i>);		Returns the high (odd) register of a long long register pair
double _itod (unsigned <i>src2</i> , unsigned <i>src1</i>);		Builds a new double register pair by reinterpreting two unsigned values, where <i>src2</i> is the high (odd) register and <i>src1</i> is the low (even) register
float _itof (unsigned <i>src</i>);		Reinterprets the bits in the unsigned as a float. For example: _itof (0x3f800000) = 1.0
long long _itoll (unsigned <i>src2</i> , unsigned <i>src1</i>);		Builds a new long long register pair by reinterpreting two unsigned values, where <i>src2</i> is the high (odd) register and <i>src1</i> is the low (even) register
unsigned _lmbd (unsigned <i>src1</i> , unsigned <i>src2</i>);	LMBD	Searches for a leftmost 1 or 0 of <i>src2</i> determined by the LSB of <i>src1</i> . Returns the number of bits up to the bit change.
unsigned _lo (double <i>src</i>);		Returns the low (even) register of a double register pair

⁽¹⁾ See the *TMS320C6000 Programmer's Guide* for more information.

⁽²⁾ See [Section 7.5.9](#) for details on manipulating 8-byte data quantities.

Table 7-4. TMS320C6000 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
unsigned _loll (long long <i>src</i>);		Returns the low (even) register of a long long register pair
double _ltod (__int40_t <i>src</i>);		Reinterprets an __int40_t register pair <i>src</i> as a double register pair.
double _lltod (long long <i>src</i>);		Reinterprets long long register pair <i>src</i> as a double register pair.
int _mpy (int <i>src1</i> , int <i>src2</i>); int _mpyus (unsigned <i>src1</i> , int <i>src2</i>); int _mpysu (int <i>src1</i> , unsigned <i>src2</i>); unsigned _mpyu (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPY MPYUS MPYSU MPYU	Multiplies the 16 LSBs of <i>src1</i> by the 16 LSBs of <i>src2</i> and returns the result. Values can be signed or unsigned.
int _mpyh (int <i>src1</i> , int <i>src2</i>); int _mpyhys (unsigned <i>src1</i> , int <i>src2</i>); int _mpyhysu (int <i>src1</i> , unsigned <i>src2</i>); unsigned _mpyhu (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPYH MPYHUS MPYHSU MPYHU	Multiplies the 16 MSBs of <i>src1</i> by the 16 MSBs of <i>src2</i> and returns the result. Values can be signed or unsigned.
int _mpyhl (int <i>src1</i> , int <i>src2</i>); int _mpyhuls (unsigned <i>src1</i> , int <i>src2</i>); int _mpyhslu (int <i>src1</i> , unsigned <i>src2</i>); unsigned _mpyhlu (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPYHL MPYHULS MPYHSLU MPYHLU	Multiplies the 16 MSBs of <i>src1</i> by the 16 LSBs of <i>src2</i> and returns the result. Values can be signed or unsigned.
int _mpylh (int <i>src1</i> , int <i>src2</i>); int _mpyluhs (unsigned <i>src1</i> , int <i>src2</i>); int _mpylshu (int <i>src1</i> , unsigned <i>src2</i>); unsigned _mpylhu (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPYLH MPYLUHS MPYLSHU MPYLHU	Multiplies the 16 LSBs of <i>src1</i> by the 16 MSBs of <i>src2</i> and returns the result. Values can be signed or unsigned.
void _nassert (int <i>src</i>);		Generates no code. Tells the optimizer that the expression declared with the assert function is true; this gives a hint to the optimizer as to what optimizations might be valid.
unsigned _norm (int <i>src</i>); unsigned _lnorm (__int40_t <i>src</i>);	NORM	Returns the number of bits up to the first nonredundant sign bit of <i>src</i>
int _sadd (int <i>src1</i> , int <i>src2</i>); long _lsadd (int <i>src1</i> , __int40_t <i>src2</i>);	SADD	Adds <i>src1</i> to <i>src2</i> and saturates the result. Returns the result.
int _sat (__int40_t <i>src2</i>);	SAT	Converts a 40-bit long to a 32-bit signed int and saturates if necessary.
unsigned _set (unsigned <i>src2</i> , unsigned <i>csta</i> , unsigned <i>cstb</i>);	SET	Sets the specified field in <i>src2</i> to all 1s and returns the <i>src2</i> value. The beginning and ending bits of the field to be set are specified by <i>csta</i> and <i>cstb</i> , respectively.
unsigned _setr (unit <i>src2</i> , int <i>src1</i>);	SET	Sets the specified field in <i>src2</i> to all 1s and returns the <i>src2</i> value. The beginning and ending bits of the field to be set are specified by the lower ten bits of <i>src1</i> .
int _smpy (int <i>src1</i> , int <i>src2</i>); int _smpyh (int <i>src1</i> , int <i>src2</i>); int _smpyhl (int <i>src1</i> , int <i>src2</i>); int _smpylh (int <i>src1</i> , int <i>src2</i>);	SMPY SMPYH SMPYHL SMPYLH	Multiplies <i>src1</i> by <i>src2</i> , left shifts the result by 1, and returns the result. If the result is 0x80000000, saturates the result to 0x7FFFFFFF
int _sshl (int <i>src2</i> , unsigned <i>src1</i>);	SSHL	Shifts <i>src2</i> left by the contents of <i>src1</i> , saturates the result to 32 bits, and returns the result
int _ssub (int <i>src1</i> , int <i>src2</i>); __int40_t _lssub (int <i>src1</i> , __int40_t <i>src2</i>);	SSUB	Subtracts <i>src2</i> from <i>src1</i> , saturates the result, and returns the result.
unsigned _subc (unsigned <i>src1</i> , unsigned <i>src2</i>);	SUBC	Conditional subtract divide step
int _sub2 (int <i>src1</i> , int <i>src2</i>);	SUB2	Subtracts the upper and lower halves of <i>src2</i> from the upper and lower halves of <i>src1</i> , and returns the result. Borrowing in the lower half subtract does not affect the upper half subtract.

The intrinsics listed in [Table 7-5](#) can be used for C6400, C6400+, C6740, and C6600 devices. The intrinsics shown correspond to the indicated C6000 assembly language instruction(s). See the *TMS320C6000 CPU and Instruction Set Reference Guide* for more information.

See [Table 7-4](#) for the listing of generic C6000 intrinsics. See [Table 7-6](#) for the listing of C6400+, C6740-, and C6600-specific intrinsics. See [Table 7-7](#) for the listing of C6700-specific intrinsics. See [Table 7-8](#) for the listing of C6600-specific intrinsics.

Table 7-5. TMS320C6400, C6400+, C6740, and C6600 C/C++ Compiler Intrinsics

C/C++ Compiler Intrinsic	Assembly Instruction	Description
int _abs2 (int <i>src</i>);	ABS2	Calculates the absolute value for each 16-bit value
int _add4 (int <i>src1</i> , int <i>src2</i>);	ADD4	Performs 2s-complement addition to pairs of packed 8-bit numbers
long long & _amem8 (void * <i>ptr</i>);	LDDW STDW	Allows aligned loads and stores of 8 bytes to memory. The pointer must be aligned to an eight-byte boundary.
const long long & _amem8_const (const void * <i>ptr</i>);	LDDW	Allows aligned loads of 8 bytes from memory. The pointer must be aligned to an eight-byte boundary. ⁽¹⁾
_float2_t & _amem8_f2 (void * <i>ptr</i>);	LDDW STDW	Allows aligned loads and stores of 8 bytes to memory. The pointer must be aligned to an eight-byte boundary. You must include <code>c6x.h</code> . ⁽¹⁾⁽²⁾
const _float2_t & _amem8_f2_const (void * <i>ptr</i>);	LDDW	Allows aligned loads of 8 bytes from memory. The pointer must be aligned to an eight-byte boundary. You must include <code>c6x.h</code> . ⁽¹⁾⁽²⁾
double & _amemd8 (void * <i>ptr</i>);	LDDW STDW	Allows aligned loads and stores of 8 bytes to memory. The pointer must be aligned to an eight-byte boundary. ⁽¹⁾⁽²⁾ For C6400 <code>_amemd8</code> corresponds to different assembly instructions than when used with other C6000 devices; see Table 7-4 .
const double & _amemd8_const (const void * <i>ptr</i>);	LDDW	Allows aligned loads of 8 bytes from memory. The pointer must be aligned to an eight-byte boundary. ⁽¹⁾⁽²⁾
int _avg2 (int <i>src1</i> , int <i>src2</i>);	AVG2	Calculates the average for each pair of signed 16-bit values
unsigned _avgu4 (unsigned, unsigned);	AVGU4	Calculates the average for each pair of signed 8-bit values
unsigned _bitc4 (unsigned <i>src</i>);	BITC4	For each of the 8-bit quantities in <i>src</i> , the number of 1 bits is written to the corresponding position in the return value
unsigned _bitr (unsigned <i>src</i>);	BITR	Reverses the order of the bits
int _cmpeq2 (int <i>src1</i> , int <i>src2</i>);	CMPEQ2	Performs equality comparisons on each pair of 16-bit values. Equality results are packed into the two least-significant bits of the return value.
int _cmpeq4 (int <i>src1</i> , int <i>src2</i>);	CMPEQ4	Performs equality comparisons on each pair of 8-bit values. Equality results are packed into the four least-significant bits of the return value.
int _cmpgt2 (int <i>src1</i> , int <i>src2</i>);	CMPGT2	Compares each pair of signed 16-bit values. Results are packed into the two least-significant bits of the return value.
unsigned _cmpgtu4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	CMPGTU4	Compares each pair of 8-bit values. Results are packed into the four least-significant bits of the return value.
int _cmplt2 (int <i>src1</i> , int <i>src2</i>);	CMPLT2	Swaps operands and calls <code>_cmpgt2</code> .
unsigned _cmpltu4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	CMPLTU4	Swaps operands and calls <code>_cmpgtu4</code> .
unsigned _deal (unsigned <i>src</i>);	DEAL	The odd and even bits of <i>src</i> are extracted into two separate 16-bit values.
int _dotp2 (int <i>src1</i> , int <i>src2</i>); _int40_t _ldotp2 (int <i>src1</i> , int <i>src2</i>);	DOTP2 DOTP2	The product of the signed lower 16-bit values of <i>src1</i> and <i>src2</i> is added to the product of the signed upper 16-bit values of <i>src1</i> and <i>src2</i> . The <code>_lo</code> and <code>_hi</code> intrinsics are needed to access each half of the 64-bit integer result.
int _dotpn2 (int <i>src1</i> , int <i>src2</i>);	DOTPN2	The product of the signed lower 16-bit values of <i>src1</i> and <i>src2</i> is subtracted from the product of the signed upper 16-bit values of <i>src1</i> and <i>src2</i> .
int _dotpnrsu2 (int <i>src1</i> , unsigned <i>src2</i>);	DOTPNRSU2	The product of the lower 16-bit values of <i>src1</i> and <i>src2</i> is subtracted from the product of the upper 16-bit values of <i>src1</i> and <i>src2</i> . The values in <i>src1</i> are treated as signed packed quantities; the values in <i>src2</i> are treated as unsigned packed quantities. 2^{15} is added and the result is sign shifted right by 16.
int _dotpnrus2 (unsigned <i>src1</i> , int <i>src2</i>);	DOTPNRUS2	Swaps the operands and calls <code>_dotpnrsu2</code> .
int _dotprs2 (int <i>src1</i> , unsigned <i>src2</i>);	DOTPRS2	The product of the lower 16-bit values of <i>src1</i> and <i>src2</i> is added to the product of the upper 16-bit values of <i>src1</i> and <i>src2</i> . The values in <i>src1</i> are treated as signed packed quantities; the values in <i>src2</i> are treated as unsigned packed quantities. 2^{15} is added and the result is sign shifted by 16.

⁽¹⁾ See [Section 7.5.9](#) for details on manipulating 8-byte data quantities.

⁽²⁾ See the *TMS320C6000 Programmer's Guide* for more information.

Table 7-5. TMS320C6400, C6400+, C6740, and C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
int _dotpsu4 (int <i>src1</i> , unsigned <i>src2</i>); int _dotpus4 (unsigned <i>src1</i> , int <i>src2</i>); unsigned _dotpu4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	DOTPSU4 DOTPUS4 DOTPU4	For each pair of 8-bit values in <i>src1</i> and <i>src2</i> , the 8-bit value from <i>src1</i> is multiplied with the 8-bit value from <i>src2</i> . The four products are summed together.
int _gmpy4 (int <i>src1</i> , int <i>src2</i>);	GMPY4	Performs the Galois Field multiply on four values in <i>src1</i> with four parallel values in <i>src2</i> . The four products are packed into the return value.
int _max2 (int <i>src1</i> , int <i>src2</i>); int _min2 (int <i>src1</i> , int <i>src2</i>); unsigned _maxu4 (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _minu4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	MAX2 MIN2 MAX4 MINU4	Places the larger/smaller of each pair of values in the corresponding position in the return value. Values can be 16-bit signed or 8-bit unsigned.
ushort & _mem2 (void * <i>ptr</i>);	LDB/LDB STB/STB	Allows unaligned loads and stores of 2 bytes to memory ⁽²⁾
const ushort & _mem2_const (const void * <i>ptr</i>);	LDB/LDB	Allows unaligned loads of 2 bytes to memory ⁽²⁾
unsigned & _mem4 (void * <i>ptr</i>);	LDNW STNW	Allows unaligned loads and stores of 4 bytes to memory ⁽²⁾
const unsigned & _mem4_const (const void * <i>ptr</i>);	LDNW	Allows unaligned loads of 4 bytes from memory ⁽²⁾
long long & _mem8 (void * <i>ptr</i>);	LDNDW STNDW	Allows unaligned loads and stores of 8 bytes to memory ⁽²⁾
const long long & _mem8_const (const void * <i>ptr</i>);	LDNDW	Allows unaligned loads of 8 bytes from memory ⁽³⁾
double & _memd8 (void * <i>ptr</i>);	LDNDW STNDW	Allows unaligned loads and stores of 8 bytes to memory ⁽⁴⁾⁽³⁾
const double & _memd8_const (const void * <i>ptr</i>);	LDNDW	Allows unaligned loads of 8 bytes from memory ⁽⁴⁾⁽³⁾
long long _mpy2ll (int <i>src1</i> , int <i>src2</i>);	MPY2	Returns the products of the lower and higher 16-bit values in <i>src1</i> and <i>src2</i>
long long _mpyhill (int <i>src1</i> , int <i>src2</i>); long long _mpylill (int <i>src1</i> , int <i>src2</i>);	MPYHI MPYLI	Produces a 16 by 32 multiply. The result is placed into the lower 48 bits of the return type. Can use the upper or lower 16 bits of <i>src1</i> .
long long _mpyihll (int <i>src1</i> , int <i>src2</i>); long long _mpyilll (int <i>src1</i> , int <i>src2</i>);	MPYIH MPYIL	Swaps operands and calls _mpyhill . Swaps operands and calls _mpylill .
int _mpyhir (int <i>src1</i> , int <i>src2</i>); int _mpylir (int <i>src1</i> , int <i>src2</i>);	MPYHIR MPYLIR	Produces a signed 16 by 32 multiply. The result is shifted right by 15 bits. Can use the upper or lower 16 bits of <i>src1</i> .
int _mpyihr (int <i>src1</i> , int <i>src2</i>); int _mpyilr (int <i>src1</i> , int <i>src2</i>);	MPYIHR MPYILR	Swaps operands and calls _mpyhir . Swaps operands and calls _mpylir .
long long _mpysu4ll (int <i>src1</i> , unsigned <i>src2</i>); long long _mpyus4ll (unsigned <i>src1</i> , int <i>src2</i>); long long _mpyu4ll (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPYSU4 MPYUS4 MPYU4	For each 8-bit quantity in <i>src1</i> and <i>src2</i> , performs an 8-bit by 8-bit multiply. The four 16-bit results are packed into a 64-bit result. The results can be signed or unsigned.
int _mvd (int <i>src2</i>);	MVD	Moves the data from <i>src2</i> to the return value over four cycles using the multiplier pipeline
unsigned _pack2 (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _packh2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	PACK2 PACKH2	The lower/upper halfwords of <i>src1</i> and <i>src2</i> are placed in the return value.
unsigned _packh4 (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _packl4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	PACKH4 PACKL4	Packs alternate bytes into return value. Can pack high or low bytes.
unsigned _packhl2 (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _packlh2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	PACKHL2 PACKLH2	The upper/lower halfword of <i>src1</i> is placed in the upper halfword the return value. The lower/upper halfword of <i>src2</i> is placed in the lower halfword the return value.
unsigned _rotl (unsigned <i>src1</i> , unsigned <i>src2</i>);	ROTL	Rotates <i>src2</i> to the left by the amount in <i>src1</i>
int _sadd2 (int <i>src1</i> , int <i>src2</i>); int _saddus2 (unsigned <i>src1</i> , int <i>src2</i>); int _saddsu2 (int <i>src1</i> , unsigned <i>src2</i>);	SADD2 SADDUS2 SADDSU2	Performs saturated addition between pairs of 16-bit values in <i>src1</i> and <i>src2</i> . Values for <i>src1</i> can be signed or unsigned.
unsigned _saddu4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SADDU4	Performs saturated addition between pairs of 8-bit unsigned values in <i>src1</i> and <i>src2</i> .
unsigned _shfl (unsigned <i>src2</i>);	SHFL	The lower 16 bits of <i>src2</i> are placed in the even bit positions, and the upper 16 bits of <i>src</i> are placed in the odd bit positions.
unsigned _shlmb (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _shrmb (unsigned <i>src1</i> , unsigned <i>src2</i>);	SHLMB SHRMB	Shifts <i>src2</i> left/right by one byte, and the most/least significant byte of <i>src1</i> is merged into the least/most significant byte position.

⁽³⁾ See the *TMS320C6000 Programmer's Guide* for more information.⁽⁴⁾ See [Section 7.5.9](#) for details on manipulating 8-byte data quantities.

Table 7-5. TMS320C6400, C6400+, C6740, and C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
int _shr2 (int <i>src1</i> , unsigned <i>src2</i>); unsigned shru2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SHR2 SHRU2	For each 16-bit quantity in <i>src2</i> , the quantity is arithmetically or logically shifted right by <i>src1</i> number of bits. <i>src2</i> can contain signed or unsigned values
long long _smpy2ll (int <i>src1</i> , int <i>src2</i>);	SMPY2	Performs 16-bit multiplication between pairs of signed packed 16-bit values, with an additional 1 bit left-shift and saturate into a 64-bit result.
int _spack2 (int <i>src1</i> , int <i>src2</i>);	SPACK2	Two signed 32-bit values are saturated to 16-bit values and packed into the return value
unsigned _spacku4 (int <i>src1</i> , int <i>src2</i>);	SPACKU4	Four signed 16-bit values are saturated to 8-bit values and packed into the return value
int _sshvl (int <i>src2</i> , int <i>src1</i>); int _sshvr (int <i>src2</i> , int <i>src1</i>);	SSHVL SSHVR	Shifts <i>src2</i> to the left/right <i>src1</i> bits. Saturates the result if the shifted value is greater than MAX_INT or less than MIN_INT.
int _sub4 (int <i>src1</i> , int <i>src2</i>);	SUB4	Performs 2s-complement subtraction between pairs of packed 8-bit values
int _subabs4 (int <i>src1</i> , int <i>src2</i>);	SUBABS4	Calculates the absolute value of the differences for each pair of packed 8-bit values
unsigned _swap4 (unsigned <i>src</i>);	SWAP4	Exchanges pairs of bytes (an endian swap) within each 16-bit value
unsigned _swap2 (unsigned <i>src</i>);	SWAP2	Calls <code>_packh2</code> .
unsigned _unpkhu4 (unsigned <i>src</i>);	UNPKHU4	Unpacks the two high unsigned 8-bit values into unsigned packed 16-bit values
unsigned _unpklu4 (unsigned <i>src</i>);	UNPKLU4	Unpacks the two low unsigned 8-bit values into unsigned packed 16-bit values
unsigned _xpnd2 (unsigned <i>src</i>);	XPND2	Bits 1 and 0 of <i>src</i> are replicated to the upper and lower halfwords of the result, respectively.
unsigned _xpnd4 (unsigned <i>src</i>);	XPND4	Bits 3 and 0 of <i>src</i> are replicated to bytes 3 through 0 of the result.

The intrinsics listed in [Table 7-6](#) are included only for C6400+, C6740, and C6600 devices. The intrinsics shown correspond to the indicated C6000 assembly language instruction(s). See the *TMS320C6000 CPU and Instruction Set Reference Guide* for more information.

See [Table 7-4](#) for the listing of generic C6000 intrinsics. See [Table 7-5](#) for the general listing of intrinsics for C6400 devices, which includes C6400, C6400+, C6740 and C6600. See [Table 7-7](#) for the listing of C6700-specific intrinsics. See [Table 7-8](#) for a listing of additional intrinsics only for C6600.

Table 7-6. TMS320C6400+, C6740, and C6600 C/C++ Compiler Intrinsics

C/C++ Compiler Intrinsic	Assembly Instruction	Description
long long _addsub (int <i>src1</i> , int <i>src2</i>);	ADDSUB	Performs an addition and subtraction in parallel.
long long _addsub2 (int <i>src1</i> , int <i>src2</i>);	ADDSUB2	Performs an ADD2 and SUB2 in parallel.
long long _cmpy (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _cmpyr (unsigned <i>src1</i> , unsigned <i>src2</i>); unsigned _cmpyr1 (unsigned <i>src1</i> , unsigned <i>src2</i>);	CMPY CMPYR CMPYR1	Performs various complex multiply operations.
long long _ddotp4 (unsigned <i>src1</i> , unsigned <i>src2</i>);	DDOTP4	Performs two DOTP2 operations simultaneously.
long long _ddotph2 (long long <i>src1</i> , unsigned <i>src2</i>); long long _ddotpl2 (long long <i>src1</i> , unsigned <i>src2</i>); unsigned _ddotph2r (long long <i>src1</i> , unsigned <i>src2</i>); unsigned _ddotpl2r (long long <i>src1</i> , unsigned <i>src2</i>);	DDOTPH2 DDOTPL2 DDOTPH2R DDOTPL2R	Performs various dual dot-product operations between two pairs of signed, packed 16-bit values.
long long _dmv (int <i>src1</i> , int <i>src2</i>);	DMV	Places <i>src1</i> in the 32 LSBs of the long long and <i>src2</i> in the 32 MSBs of the long long. See also <code>_itoll()</code> .
long long _dpack2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	DPACK2	PACK2 and PACKH2 operations performed in parallel.
long long _dpackx2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	DPACKX2	PACKLH2 and PACKX2 operations performed in parallel.
<code>__float2_t</code> _fmdv_f2 (float <i>src1</i> , float <i>src2</i>)	DMV	Places <i>src1</i> in the 32 LSBs of the <code>__float2_t</code> and <i>src2</i> in the 32 MSBs of the <code>__float2_t</code> . See also <code>_itoll()</code> .

Table 7-6. TMS320C6400+, C6740, and C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
unsigned _gmpy (unsigned <i>src1</i> , unsigned <i>src2</i>);	GMPY	Performs the Galois Field multiply.
long long _mpy2ir (int <i>src1</i> , int <i>src2</i>);	MPY2IR	Performs two 16 by 32 multiplies. Both results are shifted right by 15 bits to produce a rounded result.
int _mpy32 (int <i>src1</i> , int <i>src2</i>);	MPY32	Returns the 32 LSBs of a 32 by 32 multiply.
long long _mpy32ll (int <i>src1</i> , int <i>src2</i>); long long _mpy32su (int <i>src1</i> , int <i>src2</i>); long long _mpy32us (unsigned <i>src1</i> , int <i>src2</i>); long long _mpy32u (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPY32 MPY32SU MPY32US MPY32U	Returns all 64 bits of a 32 by 32 multiply. Values can be signed or unsigned.
int _rpack2 (int <i>src1</i> , int <i>src2</i>);	RPACK2	Shifts <i>src1</i> and <i>src2</i> left by 1 with saturation. The 16 MSBs of the shifted <i>src1</i> is placed in the 16 MSBs of the long long. The 16 MSBs of the shifted <i>src2</i> is placed in the 16 LSBs of the long long.
long long _saddsub (unsigned <i>src1</i> , unsigned <i>src2</i>);	SADDSUB	Performs a saturated addition and a saturated subtraction in parallel.
long long _saddsub2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SADDSUB2	Performs a SADD2 and a SSUB2 in parallel.
long long _shfl3 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SHFL3	Takes two 16-bit values from <i>src1</i> and 16 LSBs from <i>src2</i> to perform a 3-way interleave, creating a 48-bit result.
int _smpy32 (int <i>src1</i> , int <i>src2</i>);	SMPY32	Returns the 32 MSBs of a 32 by 32 multiply shifted left by 1.
int _ssub2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SSUB2	Subtracts the upper and lower halves of <i>src2</i> from the upper and lower halves of <i>src1</i> and saturates each result.
unsigned _xormpy (unsigned <i>src1</i> , unsigned <i>src2</i>);	XORMPY	Performs a Galois Field multiply

The intrinsics listed in [Table 7-7](#) can be used for C6700, C6700+, C6740, and C6600 devices. The intrinsics shown correspond to the indicated C6000 assembly language instruction(s). See the *TMS320C6000 CPU and Instruction Set Reference Guide* for more information.

See [Table 7-4](#) for the listing of generic C6000 intrinsics. See [Table 7-5](#) for the listing of C6400-specific intrinsics, which are also compatible with C6400+, C6740 and C6600. See [Table 7-6](#) for the listing of C6400+-specific intrinsics, which are also compatible with C6740 and C6600 devices. See [Table 7-8](#) for the listing of C6600-specific intrinsics.

Table 7-7. TMS320C6700, C6700+, C6740, and C6600 C/C++ Compiler Intrinsics

C/C++ Compiler Intrinsic	Assembly Instruction	Description
int _dpint (double <i>src</i>);	DPINT	Converts 64-bit double to 32-bit signed integer, using the rounding mode set by the CSR register
__int40_t _f2tol (__float2_t <i>src</i>);		Reinterprets a __float2_t register pair <i>src</i> as an __int40_t (stored as a register pair). You must include <i>c6x.h</i> .
__float2_t _f2toll (__float2_t <i>src</i>);		Reinterprets a __float2_t register pair as a long long register pair. You must include <i>c6x.h</i> .
double _fabs (double <i>src</i>); float _fabsf (float <i>src</i>);	ABSDP ABSSP	Returns absolute value of <i>src</i>
__float2_t _ltof2 (long long <i>src</i>);		Reinterprets a long long register pair as a __float2_t register pair. You must include <i>c6x.h</i> .
__float2_t _ltof2 (__int40_t <i>src</i>);		Reinterprets an __int40_t register pair as a __float2_t register pair. You must include <i>c6x.h</i> .
__float2_t & _mem8_f2 (void * <i>ptr</i>);	LDNDW STNDW	Allows unaligned loads and stores of 8 bytes to memory ⁽¹⁾
const __float2_t & _mem8_f2_const (void * <i>ptr</i>);	LDNDW STNDW	Allows unaligned loads of 8 bytes from memory ⁽¹⁾
long long _mpyidll (int <i>src1</i> , int <i>src2</i>);	MPYID	Produces a signed integer multiply. The result is placed in a register pair.

⁽¹⁾ See the *TMS320C6000 Programmer's Guide* for more information.

Table 7-7. TMS320C6700, C6700+, C6740, and C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
double _mpysp2dp (float <i>src1</i> , float <i>src2</i>);	MPYSP2DP	(C6700+, C6740, and C6600 only) Produces a double-precision floating-point multiply. The result is placed in a register pair.
double _mpyspdp (float <i>src1</i> , double <i>src2</i>);	MPYSPDP	(C6700+, C6740, and C6600 only) Produces a double-precision floating-point multiply. The result is placed in a register pair.
double _rcpdp (double <i>src</i>);	RCPDP	Computes the approximate 64-bit double reciprocal
float _rcpfp (float <i>src</i>);	RCPFP	Computes the approximate 32-bit float reciprocal
double _rsqrdp (double <i>src</i>);	RSQRDP	Computes the approximate 64-bit double square root reciprocal
float _rsqrfp (float <i>src</i>);	RSQRFP	Computes the approximate 32-bit float square root reciprocal
int _spint (float);	SPINT	Converts 32-bit float to 32-bit signed integer, using the rounding mode set by the CSR register

The intrinsics listed in [Table 7-8](#) are included only for C6600 devices. These intrinsics are in addition to those listed in [Table 7-5](#) and [Table 7-6](#). The intrinsics shown correspond to the indicated C6000 assembly language instruction(s). See the *TMS320C6000 CPU and Instruction Set Reference Guide* for more information.

See [Table 7-4](#) for the listing of generic C6000 intrinsics. See [Table 7-5](#) for the listing of C6400-specific intrinsics, which are also compatible with C6400+, C6740 and C6600. See [Table 7-6](#) for the listing of C6400+-specific intrinsics, which are also compatible with C6740 and C6600 devices. See [Table 7-7](#) for the listing of C6700-specific intrinsics.

Table 7-8. TMS320C6600 C/C++ Compiler Intrinsics

C/C++ Compiler Intrinsic	Assembly Instruction	Description
	ADDDP	No intrinsic. Use native C: $a + b$ where a and b are doubles
	ADDSP	No intrinsic. Use native C: $a + b$ where a and b are floats.
	AND	No intrinsic: Use native C: " $a \& b$ " where a and b are long longs
	ANDN	No intrinsic: Use native C: " $a \& \sim b$ " where a and b are long longs
	MPYSP	No intrinsic. Use native C: $a * b$ where a and b are doubles
	OR	No intrinsic: Use native C: " $a b$ " where a and b are long longs
	SUBDP	No intrinsic. Use native C: $a - b$ where a and b are doubles
	SUBSP	No intrinsic. Use native C: $a - b$ where a and b are floats
	XOR	No intrinsic: Use native C: " $a \wedge b$ " where a and b are long longs. See also <code>_xorll_c()</code> .
<code>__x128_t</code> _ccmatmpy (long long <i>src1</i> , <code>__x128_t</code> <i>src2</i>);	CMATMPY	Multiply the conjugate of 1x2 complex vector by a 2x2 complex matrix, producing two 64-bit results.
long long _ccmatmpyr1 (long long <i>src1</i> , <code>__x128_t</code> <i>src2</i>);	CCMATCMPYR1	Multiply the complex conjugate of a 1x2 complex vector by a 2x2 complex matrix, producing two 32-bit complex results.
long long _ccmpy32r1 (long long <i>src1</i> , long long <i>src2</i>);	CCMPY32R1	32-bit complex conjugate multiply of Q31 numbers with rounding
<code>__x128_t</code> _cmatmpy (long long <i>src1</i> , <code>__x128_t</code> <i>src2</i>);	CMATMPY	Multiply a 1x2 vector by a 2x2 complex matrix, producing two 64-bit complex results.
long long _cmatmpyr1 (long long <i>src1</i> , <code>__x128_t</code> <i>src2</i>);	CMATMPYR1	Multiply a 1x2 complex vector by a 2x2 complex matrix, producing two 32-bit complex results.
long long _cmpy32r1 (long long <i>src1</i> , long long <i>src2</i>);	CMPY32R1	32-bit complex multiply of Q31 numbers with rounding
<code>__x128_t</code> _cmpysp (<code>__float2_t</code> <i>src1</i> , <code>__float2_t</code> <i>src2</i>);	CMPYSP	Perform the multiply operations for a complex multiply of two complex numbers (See also <code>_complex_mpyisp</code> and <code>_complex_conjugate_mpyisp</code> .)

Table 7-8. TMS320C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
double _complex_conjugate_mpysp (double <i>src1</i> , double <i>src2</i>);	CMPYSP DSUBSP	Performs a complex conjugate multiply by performing a CMPYSP and DSUBSP
double _complex_mpysp (double <i>src1</i> , double <i>src2</i>);	CMPYSP DADDSP	Performs a complex multiply by performing a CMPYSP and DADDSP
int _crot90 (int <i>src</i>);	CROT90	Rotate complex number by 90 degrees
int _crot270 (int <i>src</i>);	CROT270	Rotate complex number by 270 degrees
long long _dadd (long long <i>src1</i> , long long <i>src2</i>);	DADD	Two-way SIMD addition of signed 32-bit values producing two signed 32-bit results.
long long _dadd2 (long long <i>src1</i> , long long <i>src2</i>);	DADD2	Four-way SIMD addition of signed 16-bit values producing four signed 32-bit results. (Two-way _add2)
_float2_t_daddsp (_float2_t <i>src1</i> , _float2_t <i>src2</i>);	DADDSP	Two-way SIMD addition of 32-bit single precision numbers
long long _dadd_c (scst5 immediate <i>src1</i> , long long <i>src2</i>);	DADD	Addition of two signed 32-bit values by a single constant in <i>src2</i> (-16 to 15) producing two signed 32-bit results.
long long _dapys2 (long long <i>src1</i> , long long <i>src2</i>);	DAPYS2	Use the sign bit of <i>src1</i> to determine whether to multiply the four 16-bit values in <i>src2</i> by 1 or -1. Yields four signed 16-bit results. (If <i>src1</i> and <i>src2</i> are the same register pair, it is equivalent to a two-way _abs2).
long long _davg2 (long long <i>src1</i> , long long <i>src2</i>);	DAVG2	Four-way SIMD average of signed 16-bit values, with rounding. (Two-way _avg2)
long long _davgnr2 (long long <i>src1</i> , long long <i>src2</i>);	DAVGNR2	Four-way SIMD average of signed 16-bit values, without rounding
long long _davgnr4 (long long <i>src1</i> , long long <i>src2</i>);	DAVGNRU4	Eight-way SIMD average of unsigned 8-bit values, without rounding
long long _davg4 (long long <i>src1</i> , long long <i>src2</i>);	DAVGU4	Eight-way SIMD average of unsigned 8-bit values, with rounding. (Two-way _avg4)
long long _dccmpyr1 (long long <i>src1</i> , long long <i>src2</i>);	DCCMPYR1	Two-way SIMD complex multiply with rounding (_cmpyr1) with complex conjugate of <i>src2</i>
unsigned _dcmpeq2 (long long <i>src1</i> , long long <i>src2</i>);	DCMPEQ2	Four-way SIMD comparison of signed 16-bit values. Results are packed into the four least-significant bits of the return value. (Two-way _cmpeq2)
unsigned _dcmpeq4 (long long <i>src1</i> , long long <i>src2</i>);	DCMPEQ4	Eight-way SIMD comparison of unsigned 8-bit values. Results are packed into the eight least-significant bits of the return value. (Two-way _cmpeq4)
unsigned _dcmpgt2 (long long <i>src1</i> , long long <i>src2</i>);	DCMPGT2	Four-way SIMD comparison of signed 16-bit values. Results are packed into the four least-significant bits of the return value. (Two-way _cmpgt2)
unsigned _dcmpgt4 (long long <i>src1</i> , long long <i>src2</i>);	DCMPGTU4	Eight-way SIMD comparison of unsigned 8-bit values. Results are packed into the eight least-significant bits of the return value. (Two-way _cmpgt4)
_x128_t_dccmpy (long long <i>src1</i> , long long <i>src2</i>);	DCCMPY	Two complex multiply operations on two sets of packed complex numbers, with complex conjugate of <i>src2</i> .
_x128_t_dcmpy (long long <i>src1</i> , long long <i>src2</i>);	DCMPY	Performs two complex multiply operations on two sets of packed complex numbers. (Two-way SIMD _cmpy).
long long _dccmpyr1 (long long <i>src1</i> , long long <i>src2</i>);	DCMPYR1	Two-way SIMD complex multiply with rounding (_cmpyr1)
long long _dcrot90 (long long <i>src</i>);	DCROT90	Two-way SIMD version of _crot90
long long _dcrot270 (long long <i>src</i>);	DCROT270	Two-way SIMD version of _crot270
long long _ddotp4h (_x128_t <i>src1</i> , _x128_t <i>src2</i>);	DDOTP4H	Performs two dot-products between four sets of packed 16-bit values. (Two-way _dotp4h)
long long _ddotpsu4h (_x128_t <i>src1</i> , _x128_t <i>src2</i>);	DDOTPSU4H	Performs two dot-products between four sets of packed 16-bit values. (Two-way _dotpsu4h)
_float2_t_dinthsp (int <i>src</i>);	DINTHSP	Converts two packed signed 16-bit values into two single-precision floating point values
_float2_t_dinthspu (unsigned <i>src</i>);	DINTHSPU	Converts two packed unsigned 16-bit values into two single-precision float point values

Table 7-8. TMS320C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
<code>__float2_t __dintsp(long long src);</code>	DINTSP	Converts two 32-bit signed integers to two single-precision float point values.
<code>__float2_t __dintspu(long long src);</code>	DINTSPU	Converts two 32-bit unsigned integers to two single-precision float point values.
<code>long long __dmax2 (long long src1, long long src2);</code>	DMAX2	Four-way SIMD maximum of 16-bit signed values producing four signed 16-bit results. (Two-way <code>__max2</code>)
<code>long long __dmaxu4 (long long src1, long long src2);</code>	DMAXU4	8-way SIMD maximum of unsigned 8-bit values producing eight unsigned 8-bit results. (Two-way <code>__maxu4</code>)
<code>long long __dmin2 (long long src1, long long src2);</code>	DMIN2	Four-way SIMD minimum of signed 16-bit values producing four signed 16-bit results. (Two-way <code>__min2</code>)
<code>long long __dminu4 (long long src1, long long src2);</code>	DMINU4	8-way SIMD minimum of unsigned 8-bit values producing eight unsigned 8-bit results. (Two-way <code>__minu4</code>)
<code>__x128_t __dmpy2 (long long src1, long long src2);</code>	DMPY2	Four-way SIMD multiply of signed 16-bit values producing four signed 32-bit results. (Two-way <code>__mpy2</code>)
<code>__float2_t __dmpysp (__float2_t src1, __float2_t src2);</code>	DMPYSP	Two-way single precision floating point multiply producing two single-precision results
<code>__x128_t __dmpysu4 (long long src1, long long src2);</code>	DMPYSU4	Eight-way SIMD multiply of signed 8-bit values by unsigned 8-bit values producing eight signed 16-bit results. (Two-way <code>__mpysu4</code>)
<code>__x128_t __dmpyu2 (long long src1, long long src2);</code>	DMPYU2	Four-way SIMD multiply of unsigned 16-bit values producing four unsigned 32-bit results. (Two-way <code>__mpyu2</code>)
<code>__x128_t __dmpyu4 (long long src1, long long src2);</code>	DMPYU4	Eight-way SIMD multiply of signed 8-bit values producing eight signed 16-bit results. (Two-way <code>__mpyu4</code>)
<code>long long __dmvd (long long src1, unsigned src2);</code>	DMVD	Places src1 in the low register of the long long and src2 in the high register of the long long. Takes four cycles. See also <code>_dmv()</code> , <code>_fdmv_f2</code> , and <code>_itoll()</code> .
<code>int __dotp4h (long long src1, long long src2);</code>	DDOTP4H	Multiply two sets of four signed 16-bit values and return the 32-bit sum.
<code>long long __dotp4hl (long long src1, long long src2);</code>	DOTP4H	Multiply two sets of four signed 16-bit values and return the 64-bit sum.
<code>int __dotpsu4h (long long src1, long long src2);</code>	DOTPSU4H	Multiply four signed 16-bit values by four unsigned 16-bit values and return the 32-bit sum.
<code>long long __dotpsu4hl (long long src1, long long src2);</code>	DOTPSU4H	Multiply four signed 16-bit values by four unsigned 16-bit values and return the 64-bit sum.
<code>long long __dpackh2 (long long src1, long long src2);</code>	DPACKH2	Two-way <code>__packh2</code>
<code>long long __dpackh4 (long long src1, long long src2);</code>	DPACKH4	Two-way <code>__packh4</code>
<code>long long __dpacklh2 (long long src1, long long src2);</code>	DPACKLH2	Two-way <code>__packlh2</code>
<code>long long __dpacklh4 (unsigned src1, unsigned src2);</code>	DPACKLH4	Performs a <code>__packl4</code> and a <code>__packh4</code> . The output of the <code>__packl4</code> is in the low register of the result and the output of the <code>__packh4</code> is in the high register of the result.
<code>long long __dpackl2 (long long src1, long long src2);</code>	DPACKL2	Two-way <code>__packl2</code>
<code>long long __dpackl4 (long long src1, long long src2);</code>	DPACKL4	Two-way <code>__packl4</code>
<code>long long __dsadd (long long src1, long long src2);</code>	DSADD	Two-way SIMD saturated addition of signed 32-bit values producing two signed 32-bit results. (Two-way <code>__sadd</code>)
<code>long long __dsadd2 (long long src1, long long src2);</code>	DSADD2	Four-way SIMD saturated addition of signed 16-bit values producing four signed 16-bit results. (Two-way <code>__sadd2</code>)
<code>long long __dshl (long long src1, unsigned src2);</code>	DSHL	Shift-left of two signed 32-bit values by a single value in the src2 argument.
<code>long long __dshl2 (long long src1, unsigned src2);</code>	DSHL2	Shift-left of four signed 16-bit values by a single value in the src2 argument. (Two-way <code>__shl2</code>)
<code>long long __dshr (long long src1, unsigned src2);</code>	DSHR	Shift-right of two signed 32-bit values by a single value in the src2 argument.
<code>long long __dshr2 (long long src1, unsigned src2);</code>	DSHR2	Shift-right of four signed 16-bit values by a single value in the src2 argument. (Two-way <code>__shr2</code>)

Table 7-8. TMS320C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
long long _dshru (long long <i>src1</i> , unsigned <i>src2</i>);	DSHRU	Shift-right of two unsigned 32-bit values by a single value in the <i>src2</i> argument.
long long _dshru2 (long long <i>src1</i> , unsigned <i>src2</i>);	DSHRU2	Shift-right of four unsigned 16-bit values by a single value in the <i>src2</i> argument. (Two-way _shru2)
_x128_t_dsmPY2 (long long <i>src1</i> , long long <i>src2</i>);	DSMPY2	Four-way SIMD multiply of signed 16-bit values with 1-bit left-shift and saturate producing four signed 32-bit results. (Two-way _smPY2)
long long _dspackU4 (long long <i>src1</i> , long long <i>src2</i>);	DSPACKU4	Two-way _spackU4
long long _dspint (_float2_t <i>src</i>);	DSPINT	Converts two packed single-precision floating point values to two signed 32-bit values
unsigned _dspinh (_float2_t <i>src</i>);	DSPINTH	Converts two packed single-precision floating point values to two packed signed 16-bit values
long long _dssub (long long <i>src1</i> , long long <i>src2</i>);	DSSUB	Two-way SIMD saturated subtraction of 32-bit signed values producing two signed 32-bit results.
long long _dssub2 (long long <i>src1</i> , long long <i>src2</i>);	DSSUB2	Four-way SIMD saturated subtraction of signed 16-bit values producing four signed 16-bit results. (Two-way _ssub2)
long long _dsUB (long long <i>src1</i> , long long <i>src2</i>);	DSUB	Two-way SIMD subtraction of 32-bit signed values producing two signed 32-bit results.
long long _dsUB2 (long long <i>src1</i> , long long <i>src2</i>);	DSUB2	Four-way SIMD subtraction of signed 16-bit values producing four signed 16-bit results. (Two-way _sub2)
_float2_t_dsubSP (_float2_t <i>src1</i> , _float2_t <i>src2</i>);	DSUBSP	Two-way SIMD subtraction of 32-bit single precision numbers
long long _dxpnd2 (unsigned <i>src</i>);	DXPND2	Expand four lower bits to four 16-bit fields.
long long _dxpnd4 (unsigned <i>src</i>);	DXPND4	Expand eight lower bits to eight 8-bit fields.
_float2_t_fmvd_f2 (float <i>src1</i> , float <i>src2</i>);	DMVD	Places <i>src1</i> in the low register of the _float2_t and <i>src2</i> in the high register of the _float2_t . Takes four cycles. See also _dmv() , _dmvd() , and _itoll() . You must include c6x.h .
int _land (int <i>src1</i> , int <i>src2</i>);	LAND	Logical AND of <i>src1</i> and <i>src2</i>
int _landn (int <i>src1</i> , int <i>src2</i>);	LANDN	Logical AND of <i>src1</i> and NOT of <i>src2</i> ; i.e. <i>src1</i> AND \sim <i>src2</i>
int _lor (int <i>src1</i> , int <i>src2</i>);	LOR	Logical OR of <i>src1</i> and <i>src2</i>
void _mfence ();	MFENCE	Stall CPU while memory system is busy
double _mpysp2dp (float <i>src1</i> , float <i>src2</i>);	MPYSP2DP	(C6600 and C6700+ only) Produces a double-precision floating-point multiply. The result is placed in a register pair.
double _mpyspdp (float <i>src1</i> , double <i>src2</i>);	MPYSPDP	(C6600 and C6700+ only) Produces a double-precision floating-point multiply. The result is placed in a register pair.
long long _mpyu2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	MPYU2	Two-way SIMD multiply of unsigned 16-bit values producing two unsigned 32-bit results.
_x128_t_qmpy32 (_x128_t <i>src1</i> , _x128_t <i>src2</i>);	QMPY32	Four-way SIMD multiply of signed 32-bit values producing four 32-bit results. (Four-way _mpy32)
_x128_t_qmpysp (_x128_t <i>src1</i> , _x128_t <i>src2</i>);	QMPYSP	Four-way SIMD 32-bit single precision multiply producing four 32-bit single precision results
_x128_t_qsmpy32r1 (_x128_t <i>src1</i> , _x128_t <i>src2</i>);	QSMPY32R1	4-way SIMD fractional 32-bit by 32-bit multiply where each result value is shifted right by 31 bits and rounded. This normalizes the result to lie within -1 and 1 in a Q31 fractional number system.
unsigned _shl2 (unsigned <i>src1</i> , unsigned <i>src2</i>);	SHL2	Shift-left of two signed 16-bit values by a single value in the <i>src2</i> argument.
long long _unpkbu4 (unsigned <i>src</i>);	UNPKBU4	Unpack four unsigned 8-bit values into four unsigned 16-bit values. (See also _unpklu4 and _unpkhu4)
long long _unpkh2 (unsigned <i>src</i>);	UNPKH2	Unpack two signed 16-bit values to two signed 32-bit values

Table 7-8. TMS320C6600 C/C++ Compiler Intrinsics (continued)

C/C++ Compiler Intrinsic	Assembly Instruction	Description
long long _unpkhu2 (unsigned <i>src</i>);	UNPKHU2	Unpack two unsigned 16-bit values to two unsigned 32-bit values
long long _xorll_c (scst5 immediate <i>src1</i> , long long <i>src2</i>);	XOR	XOR <i>src1</i> with the upper and lower 32-bit portions of <i>src2</i> (SIMD XOR by constant)

For details on the `__x128_t` container type see [Section 7.5.6](#).

7.5.6 The `__x128_t` Container Type

The `__x128_t` container type is for storing 128-bits of data and its use is necessary when performing certain SIMD operations on C6600. This type can only be used when compiling for C6600. Also, note the leading double underscore. When using the `__x128_t` container type, you must include `c6x.h`.

This type can be used to define objects that can be used with certain C6600 intrinsics. (See [Table 7-8](#).) The object can be filled and manipulated using various intrinsics. The type is not a full-fledged built-in type (like `long long`), and so various native C operations are not allowed. Think of this type as a struct with private members and special manipulation functions.

When the compiler puts a `__x128_t` object in the register file, the `__x128_t` object takes four registers (a register *quad*).

Objects of type `__x128_t` are aligned to a 128-bit boundary in memory.

The following operations are supported:

- Declare a `__x128_t` global object (for example: `__x128_t a`). By default, it will be put in the `.far` section.
- Declare a `__x128_t` local object (for example: `__x128_t a`). It will be put on the stack.
- Declare a `__x128_t` global/local pointer (for example: `__x128_t *a`).
- Declare an array of `__x128_t` objects (for example: `__x128_t a[10]`).
- Declare a `__x128_t` type as a member of a struct, class, or union.
- Assign a `__x128_t` object to another `__x128_t` object.
- Pass a `__x128_t` object to a function (including variadic argument functions). (Pass by value.)
- Return a `__x128_t` object from a function.
- Use 128-bit manipulation intrinsics to set and extract contents (see [Table 7-9](#)).

The following operations are not supported:

- Native-type operations on `__x128_t` objects, such as `+`, `-`, `*`, etc.
- Cast an object to a `__x128_t` type.
- Access the elements of a `__x128_t` using array or struct notation.
- Pass a `__x128_t` object to I/O functions like `printf`. Instead, extract the values from the `__x128_t` object by using appropriate intrinsics.

Example 7-7. The `__x128_t` Container Type

```
#include <c6x.h>
#include <stdio.h>
__x128_t mpy_four_way_example(__x128_t s, int a, int b, int c, int d)
{
    __x128_t t = _ito128(a, b, c, d); // Pack values into a __x128_t
    __x128_t results = _qmpy32(s, t); // Perform a four-way SIMD multiply

    int lowest32 = _get32_128(results, 0); // Extract lowest reg of __x128_t
    int highest32 = _get32_128(results, 3); // Extract highest reg of __x128_t

    printf("lowest = %d\n", lowest32);
    printf("highest = %d\n", highest32);

    return results;
}
```

Include `c6x.h` With Type `__x128_t` or `__float2_t`

NOTE: When using the `__x128_t` container type, or `__float2_t` typedef, or any intrinsics involving `__float2_t`, you must include `c6x.h`.

Table 7-9. Vector-in-Scalar Support C/C++ Compiler v7.2 Intrinsics

C/C++ Compiler Intrinsic	Description
Creation	
<code>__x128_t _ito128 (unsigned src1, unsigned src2, unsigned src3, unsigned src4);</code>	Creates <code>__x128_t</code> from (u)int (reg+3, reg+2, reg+1, reg+0)
<code>__x128_t _fto128 (float src1, float src2, float src3, float src4);</code>	Creates <code>__x128_t</code> from float (reg+3, reg+2, reg+1, reg+0)
<code>__x128_t _llto128 (long long src1, long long src2);</code>	Creates <code>__x128_t</code> from two long longs
<code>__x128_t _dto128 (double src1, double src2);</code>	Creates <code>__x128_t</code> from two doubles
<code>__x128_t _f2to128(__float2_t src1, __float2_t src2);</code>	Creates <code>__x128_t</code> from two <code>__float2_t</code> objects
<code>__x128_t _dup32_128 (int src);</code>	Creates <code>__x128_t</code> from duplicating src1
<code>__float2_t _ftof2(float src1, float src2);</code>	Creates <code>__float2_t</code> from two floats
Extraction	
<code>float _hif (double src);</code>	Extracts upper float from double
<code>float _lof (double src);</code>	Extracts lower float from double
<code>float _hif2(__float2_t src);</code>	Extracts upper float from <code>__float2_t</code>
<code>float _lof2(__float2_t src);</code>	Extracts lower float from <code>__float2_t</code>
<code>long long _hi128 (__x128_t src);</code>	Extracts upper two registers of register quad
<code>double _hid128 (__x128_t src);</code>	Extracts upper two registers of register quad
<code>__float2_t _hif2_128(__x128_t src);</code>	Extracts upper two registers of register quad
<code>long long _lo128 (__x128_t src);</code>	Extracts lower two registers of register quad
<code>double _lod128 (__x128_t src);</code>	Extracts lower two registers of register quad
<code>__float2_t _lof2_128(__x128_t src);</code>	Extracts lower two registers of register quad
<code>unsigned _get32_128 (__x128_t src, 0);</code>	Extracts first register of register quad (base reg + 0)
<code>unsigned _get32_128 (__x128_t src, 1);</code>	Extracts second register of register quad (base reg + 1)
<code>unsigned _get32_128 (__x128_t src, 2);</code>	Extracts third register of register quad (base reg + 2)
<code>unsigned _get32_128 (__x128_t src, 3);</code>	Extracts fourth register of register quad (base reg + 3)
<code>float _get32f_128 (__x128_t src, 0);</code>	Extracts first register of register quad (base reg + 0)
<code>float _get32f_128 (__x128_t src, 1);</code>	Extracts second register of register quad (base reg + 1)
<code>float _get32f_128 (__x128_t src, 2);</code>	Extracts third register of register quad (base reg + 2)
<code>float _get32f_128 (__x128_t src, 3);</code>	Extracts fourth register of register quad (base reg + 3)

7.5.7 The `__float2_t` Container Type

The `__float2_t` container type should be used (instead of double) to store two floats. There are manipulation intrinsics to create and manipulate objects with the `__float2_t` type (see [Table 7-9](#)). The run-time-support file, `c6x.h`, must be included when using `__float2_t` or when using any of the `__float2_t` manipulation intrinsics.

Recommendations for using the `__float2_t` type:

- Use `__float2_t` to store two floats. Do not use double.
- Use long long to store 64-bit packed integer data. Do not use double or `__float2_t` for packed integer data.

7.5.8 Using Intrinsic for Interrupt Control and Atomic Sections

The C/C++ compiler supports three intrinsics for enabling, disabling, and restoring interrupts. The syntaxes are:

```
unsigned int    _disable_interrupts ( );
unsigned int    _enable_interrupts ( );
void           _restore_interrupts (unsigned int);
```

The `_disable_interrupts()` and `_enable_interrupts()` intrinsics both return an unsigned int that can be subsequently passed to `_restore_interrupts()` to restore the previous interrupt state. These intrinsics provide a barrier to optimization and are therefore appropriate for implementing a critical (or atomic) section. For example,

```
unsigned int restore_value;

restore_value = _disable_interrupts();
if (sem) sem--;
_restore_interrupts(restore_value);
```

The example code disables interrupts so that the value of `sem` read for the conditional clause does not change before the modification of `sem` in the then clause. The intrinsics are barriers to optimization, so the memory reads and writes of `sem` do not cross the `_disable_interrupts` or `_restore_interrupts` locations.

Overwrites CSR

NOTE: The `_restore_interrupts()` intrinsic overwrites the CSR control register with the value in the argument. Any CSR bits changed since the `_disable_interrupts()` intrinsic or `_enable_interrupts()` intrinsic will be lost.

On C6400+, C6740, and C6600, the `_restore_interrupts()` intrinsic does not use the RINT instruction.

7.5.9 Using Unaligned Data and 64-Bit Values

The C6400, C6400+, C6740, and C6600 families have support for unaligned loads and stores of 64-bit and 32-bit values via the use of the `_mem8`, `_memd8`, and `_mem4` intrinsics. The `_lo` and `_hi` intrinsics are useful for extracting the two 32-bit portions from a 64-bit double. The `_loll` and `_hill` intrinsics are useful for extracting the two 32-bit portions from a 64-bit long long.

For the C6400+, C6740, and C6600 intrinsics that use 64-bit types, the equivalent C type is long long. Do not use the C type double or the compiler performs a call to a run-time-support math function to do the floating-point conversion. Here are ways to access 64-bit and 32-bit values:

- To get the upper 32 bits of a long long in C code, use `>> 32` or the `_hill()` intrinsic.
- To get the lower 32 bits of a long long in C code, use a cast to int or unsigned, or use the `_loll` intrinsic.
- To get the upper 32 bits of a double (interpreted as an int), use `_hi()`.
- To get the lower 32 bits of a double (interpreted as an int), use `_lo()`.
- To get the upper 32 bits of a `__float2_t`, use `_hif2()`.
- To get the lower 32 bits of a `__float2_t`, use `_lof2()`.
- To create a long long value, use the `_itoll(int high32bits, int low32bits)` intrinsic.
- To create a `__float2_t` value, use the `_ftof2(float high32bits, float low32bits)` intrinsic.

[Example 7-8](#) shows the usage of the `_mem8` intrinsic.

Example 7-8. Using the `_mem8` Intrinsic

```
void alt_load_longlong_unaligned(void *a, int *high, int *low)
{
    long long p = _mem8(a);
    *high = p >> 32;
    *low = (unsigned int) p;
}
```

7.5.10 Using `MUST_ITERATE` and `_nassert` to Enable SIMD and Expand Compiler Knowledge of Loops

Through the use of `MUST_ITERATE` and `_nassert`, you can guarantee that a loop executes a certain number of times.

This example tells the compiler that the loop is guaranteed to run exactly 10 times:

```
#pragma MUST_ITERATE(10,10);
for (I = 0; I < trip_count; I++) { ...
```

`MUST_ITERATE` can also be used to specify a range for the trip count as well as a factor of the trip count. For example:

```
#pragma MUST_ITERATE(8,48,8);
for (I = 0; I < trip; I++) { ...
```

This example tells the compiler that the loop executes between 8 and 48 times and that the trip variable is a multiple of 8 (8, 16, 24, 32, 40, 48). The compiler can now use all this information to generate the best loop possible by unrolling better even when the `--interrupt_thresholdn` option is used to specify that interrupts do occur every n cycles.

The *TMS320C6000 Programmer's Guide* states that one of the ways to refine C/C++ code is to use word accesses to operate on 16-bit data stored in the high and low parts of a 32-bit register. Examples using casts to int pointers are shown with the use of intrinsics to use certain instructions like `_mpyh`. This can be automated by using the `_nassert()`; intrinsic to specify that 16-bit short arrays are aligned on a 32-bit (word) boundary.

The following examples generate the same assembly code:

- **Example 1**

```
int dot_product(short *x, short *y, short z)
{
    int *w_x = (int *)x;
    int *w_y = (int *)y;
    int sum1 = 0, sum2 = 0, I;
    for (I = 0; I < z/2; I++)
    {
        sum1 += _mpy(w_x[i], w_y[i]);
        sum2 += _mpyh(w_x[i], w_y[i]);
    }
    return (sum1 + sum2);
}
```

- **Example 2**

```
int dot_product(short *x, short *y, short z)
{
    int sum = 0, I;

    _nassert (((int)(x) & 0x3) == 0);
    _nassert (((int)(y) & 0x3) == 0);
    #pragma MUST_ITERATE(20, , 4);
    for (I = 0; I < z; I++) sum += x[i] * y[i];
    return sum;
}
```

C++ Syntax for `_nassert`

NOTE: In C++ code, `_nassert` is part of the standard namespace. Thus, the correct syntax is `std::_nassert()`.

7.5.11 Methods to Align Data

In the following code, the `_nassert` tells the compiler, for every invocation of `f()`, that `ptr` is aligned to an 8-byte boundary. Such an assertion often leads to the compiler producing code which operates on multiple data values with a single instruction, also known as SIMD (single instruction multiple data) optimization.

```
void f(short *ptr)
{
    _nassert((int) ptr % 8 == 0)

    ; a loop operating on data accessed by ptr
}
```

The following subsections describe methods you can use to ensure the data referenced by `ptr` is aligned. You have to employ one of these methods at every place in your code where `f()` is called.

7.5.11.1 Base Address of an Array

An argument such as `ptr` is most commonly passed the base address of an array, for example:

```
short buffer[100];
...
f(buffer);
```

When compiling for C6600 devices, such an array is automatically aligned to a 16-byte boundary. When compiling for C6400, C6400+, C6740, and C6600 devices, such an array is automatically aligned to an 8-byte boundary. When compiling for C6200 or C6700, such an array is automatically aligned to 4-byte boundary, or, if the base type requires it, an 8-byte boundary. This is true whether the array is global, static, or local. This automatic alignment is all that is required to achieve SIMD optimization on those respective devices. You still need to include the `_nassert` because, in the general case, the compiler cannot guarantee that `ptr` holds the address of a properly aligned array.

If you always pass the base address of an array to pointers like `ptr`, then you can use the following macro to reflect that fact.

```
#if defined(_TMS320C6600)
    #define ALIGNED_ARRAY(ptr) _nassert((int) ptr % 16 == 0)
#elif defined(_TMS320C6400)
    #define ALIGNED_ARRAY(ptr) _nassert((int) ptr % 8 == 0)
#elif defined(_TMS320C6200) || defined(_TMS320C6700)
    #define ALIGNED_ARRAY(ptr) _nassert((int) ptr % 4 == 0)
#else
    #define ALIGNED_ARRAY(ptr) /* empty */
#endif

void f(short *ptr)
{
    ALIGNED_ARRAY(ptr);
    ; a loop operating on data accessed by ptr
}
```

The macro works regardless of which C6000 device you build for, or if you port the code to another target.

7.5.11.2 Offset from the Base of an Array

A more rare case is to pass the address of an offset from an array, for example:

```
f(&buffer[3]);
```

This code passes an unaligned address to `ptr`, thus violating the presumption coded in the `_nassert()`. There is no direct remedy for this case. Avoid this practice whenever possible.

7.5.11.3 Dynamic Memory Allocation

Ordinary dynamic memory allocation guarantees that the allocated memory is properly aligned for any scalar object of a native type (for instance, it is correctly aligned for a long double or long long int), but does not guarantee any larger alignment. For example:

```
buffer = calloc(100, sizeof(short))
```

To get a stricter alignment, use the function `memalign` with the desired alignment. To get an alignment of 256 bytes for example:

```
buffer = memalign(256, 100 * sizeof(short));
```

If you are using BIOS memory allocation routines, be sure to pass the alignment factor as the last argument using the syntax that follows:

```
buffer = MEM_alloc( segid , 100 * sizeof(short), 256);
```

See the *TMS320C6000 DSP/BIOS Help* for more information about BIOS memory allocation routines and the *segid* parameter in particular.

7.5.11.4 Member of a Structure or Class

Arrays which are members of a structure or a class are aligned only as the base type of the array requires. The automatic alignment described in [Section 7.5.11.1](#) does not occur.

Example 7-9. An Array in a Structure

```
struct s
{
    ...
    short buf1[50];
    ...
} g;

...

f(g.buf1);
```

Example 7-10. An Array in a Class

```
class c
{
    public :
        short buf1[50];
        void mfunc(void);
        ...
};

void c::mfunc()
{
    f(buf1);
    ...
}
```

To align an array in a structure, place it inside a union with a dummy object that has the desired alignment. If you want 8 byte alignment, use a "long long" dummy field. For example:

```
struct s
{
    union u
    {
        long long dummy; /* 8-byte alignment */
        short buffer[50]; /* also 8-byte alignment */
    } u;
    ...
};
```

If you want to declare several arrays contiguously, and maintain a given alignment, you can do so by keeping the array size, measured in bytes, an even multiple of the desired alignment. For example:

```
struct s
{
    long long dummy; /* 8-byte alignment */
    short buffer[50]; /* also 8-byte alignment */
    short buf2[50]; /* 4-byte alignment */
    ...
};
```

Because the size of buf1 is 50 * 2-bytes per short = 100 bytes, and 100 is an even multiple of 4, not 8, buf2 is only aligned on a 4-byte boundary. Padding buf1 out to 52 elements makes buf2 8-byte aligned.

Within a structure or class, there is no way to enforce an array alignment greater than 8. For the purposes of SIMD optimization, this is not necessary.

Alignment With Program-Level Optimization

NOTE: In most cases program-level optimization (see [Section 3.7](#)) entails compiling all of your source files with a single invocation of the compiler, while using the `-pm -o3` options. This allows the compiler to see all of your source code at once, thus enabling optimizations that are rarely applied otherwise. Among these optimizations is seeing that, for instance, all of the calls to the function `f()` are passing the base address of an array to `ptr`, and thus `ptr` is always correctly aligned for SIMD optimization. In such a case, the `_nassert()` is not required. The compiler automatically determines that `ptr` must be aligned, and produces the optimized SIMD instructions.

7.5.12 SAT Bit Side Effects

The saturated intrinsic operations define the SAT bit if saturation occurs. The SAT bit can be set and cleared from C/C++ code by accessing the control status register (CSR). The compiler uses the following steps for generating code that accesses the SAT bit:

1. The SAT bit becomes undefined by a function call or a function return. This means that the SAT bit in the CSR is valid and can be read in C/C++ code until a function call or until a function returns.
2. If the code in a function accesses the CSR, then the compiler assumes that the SAT bit is live across the function, which means:
 - The SAT bit is maintained by the code that disables interrupts around software pipelined loops.
 - Saturated instructions cannot be speculatively executed.
3. If an interrupt service routine modifies the SAT bit, then the routine should be written to save and restore the CSR.

7.5.13 IRP and AMR Conventions

There are certain assumptions that the compiler makes about the IRP and AMR control registers. The assumptions should be enforced in all programs and are as follows:

1. The AMR must be set to 0 upon calling or returning from a function. A function does not have to save and restore the AMR, but must ensure that the AMR is 0 before returning.
2. The AMR must be set to 0 when interrupts are enabled, or the `SAVE_AMR` and `STORE_AMR` macros should be used in all interrupts (see [Section 7.6.3](#)).
3. The IRP can be safely modified only when interrupts are disabled.
4. The IRP's value must be saved and restored if you use the IRP as a temporary register.

7.5.14 Floating Point Control Register Side Effects

When performing floating point operations on a floating-point architecture, status bits in certain control registers may be set. In particular, status bits may be set in the FADCR, FAUCR, and FMCR registers, hereafter referred to as the "floating point control registers". These bits can be set and cleared from C/C++ code by writing to or reading from these registers, as shown in example 6-1.

In compiler versions released after July of 2009, the compiler uses the following steps for generating code that accesses any of the floating point control registers.

1. The floating point control registers become undefined by a function call or a function return. This means the data in the floating point control registers is valid and can be read in C/C++ code until a function call or a function returns.
2. If the code in a function accesses any of the floating point control registers, the compiler assumes that those registers are live across the function, which means that floating point instructions that may set bits in those floating point control registers cannot be speculatively executed.
3. If an interrupt service routine modifies any of the bits in a floating point control register, the interrupt service routine should be written to save and restore that floating point control register.

7.6 Interrupt Handling

As long as you follow the guidelines in this section, you can interrupt and return to C/C++ code without disrupting the C/C++ environment. When the C/C++ environment is initialized, the startup routine does not enable or disable interrupts. If the system is initialized by way of a hardware reset, interrupts are disabled. If your system uses interrupts, you must handle any required enabling or masking of interrupts. Such operations have no effect on the C/C++ environment and are easily incorporated with `asm` statements or calling an assembly language function.

7.6.1 Saving the SGIE Bit

When compiling for C6400+, C6740, and C6600, the compiler may use the C6400+-, C6740-, and C6600-specific instructions `DINT` and `RINT` to disable and restore interrupts around software-pipelined loops. These instructions utilize the CSR control register as well as the SGIE bit in the TSR control register. Therefore, the SGIE bit is considered to be save-on-call. If you have assembly code that calls compiler-generated code, the SGIE bit should be saved (e.g. to the stack) if it is needed later. The SGIE bit should then be restored upon return from compiler generated code.

7.6.2 Saving Registers During Interrupts

When C/C++ code is interrupted, the interrupt routine must preserve the contents of all machine registers that are used by the routine or by any functions called by the routine. The compiler handles register preservation if the interrupt service routine is written in C/C++ and declared with the `interrupt` keyword. For C6400+, C6740, and C6600, the compiler will save and restore the ILC and RILC control registers if needed.

7.6.3 Using C/C++ Interrupt Routines

A C/C++ interrupt routine is like any other C/C++ function in that it can have local variables and register variables; however, it should be declared with no arguments and should return void. C/C++ interrupt routines can allocate up to 32K on the stack for local variables. For example:

```
interrupt void example (void)
{
    ...
}
```

If a C/C++ interrupt routine does not call any other functions, only those registers that the interrupt handler attempts to define are saved and restored. However, if a C/C++ interrupt routine *does* call other functions, these functions can modify unknown registers that the interrupt handler does not use. For this reason, the routine saves all usable registers if any other functions are called. Interrupts branch to the interrupt return pointer (IRP). Do not call interrupt handling functions directly.

Interrupts can be handled *directly* with C/C++ functions by using the interrupt pragma or the interrupt keyword. For more information, see [Section 6.9.18](#) and [Section 6.5.3](#), respectively.

You are responsible for handling the AMR control register and the SAT bit in the CSR correctly inside an interrupt. By default, the compiler does not do anything extra to save/restore the AMR and the SAT bit. Macros for handling the SAT bit and the AMR register are included in the c6x.h header file.

For example, you are using circular addressing in some hand assembly code (that is, the AMR does not equal 0). This hand assembly code can be interrupted into a C code interrupt service routine. The C code interrupt service routine assumes that the AMR is set to 0. You need to define a local unsigned int temporary variable and call the SAVE_AMR and RESTORE_AMR macros at the beginning and end of your C interrupt service routine to correctly save/restore the AMR inside the C interrupt service routine.

Example 7-11. AMR and SAT Handling

```
#include <c6x.h>

interrupt void interrupt_func()
{
    unsigned int temp_amr;
    /* define other local variables used inside interrupt */

    /* save the AMR to a temp location and set it to 0 */
    SAVE_AMR(temp_amr);

    /* code and function calls for interrupt service routine */
    ...

    /* restore the AMR for you hand assembly code before exiting */
    RESTORE_AMR(temp_amr);
}
```

If you need to save/restore the SAT bit (i.e. you were performing saturated arithmetic when interrupted into the C interrupt service routine which may also perform some saturated arithmetic) in your C interrupt service routine, it can be done in a similar way as the above example using the SAVE_SAT and RESTORE_SAT macros.

For C6400+, C6740, and C6600, the compiler saves and restores the ILC and RILC control registers if needed.

For floating point architectures, you are responsible for handling the floating point control registers FADCR, FAUCR and FMCR. If you are reading bits out of the floating pointer control registers, and if the interrupt service routine (or any called function) performs floating point operations, then the relevant floating point control registers should be saved and restored. No macros are provided for these registers, as simple assignment to and from an unsigned int temporary will suffice.

7.6.4 Using Assembly Language Interrupt Routines

You can handle interrupts with assembly language code as long as you follow the same register conventions the compiler does. Like all assembly functions, interrupt routines can use the stack, access global C/C++ variables, and call C/C++ functions normally. When calling C/C++ functions, be sure that any registers listed in [Table 7-2](#) are saved, because the C/C++ function can modify them.

7.7 Run-Time-Support Arithmetic Routines

The run-time-support library contains a number of assembly language functions that provide arithmetic routines for C/C++ math operations that the C6000 instruction set does not provide, such as integer division, integer remainder, and floating-point operations.

These routines follow the standard C/C++ calling sequence. The compiler automatically adds these routines when appropriate; they are not intended to be called directly by your programs.

The source code for these functions is in the source library rts.src. The source code has comments that describe the operation of the functions. You can extract, inspect, and modify any of the math functions. Be sure, however, that you follow the calling conventions and register-saving rules outlined in this chapter.

[Table 7-10](#) summarizes the run-time-support functions used for arithmetic.

Table 7-10. Summary of Run-Time-Support Arithmetic Functions

Type	Function in COFF ABI	Function in EABI	Description
float	_cvtidf (double)	__c6xabi_cvtidf (double)	Convert double to float
int	_fixdi (double)	__c6xabi_fixdi (double)	Convert double to signed integer
long	_fixdi (double)	__c6xabi_fixdi (double)	Convert double to long
long long	_fixdlli (double)	__c6xabi_fixdlli (double)	Convert double to long long
uint	_fixdlli (double)	__c6xabi_fixdlli (double)	Convert double to unsigned integer
ulong	_fixdul (double)	__c6xabi_fixdul (double)	Convert double to unsigned long
ulong long	_fixdull (double)	__c6xabi_fixdull (double)	Convert double to unsigned long long
double	_cvtfd (float)	__c6xabi_cvtfd (float)	Convert float to double
int	_fixfi (float)	__c6xabi_fixfi (float)	Convert float to signed integer
long	_fixfli (float)	__c6xabi_fixfli (float)	Convert float to long
long long	_fixfli (float)	__c6xabi_fixfli (float)	Convert float to long long
uint	_fixfu (float)	__c6xabi_fixfu (float)	Convert float to unsigned integer
ulong	_fixful (float)	__c6xabi_fixful (float)	Convert float to unsigned long
ulong long	_fixfull (float)	__c6xabi_fixfull (float)	Convert float to unsigned long long
double	_fltld (int)	__c6xabi_fltld (int)	Convert signed integer to double
float	_fltif (int)	__c6xabi_fltif (int)	Convert signed integer to float
double	_fltud (uint)	__c6xabi_fltud (uint)	Convert unsigned integer to double
float	_fltuf (uint)	__c6xabi_fltuf (uint)	Convert unsigned integer to float
double	_ftlld (long)	__c6xabi_ftlld (long)	Convert signed long to double
float	_ftlif (long)	__c6xabi_ftlif (long)	Convert signed long to float
double	_ftuld (ulong)	__c6xabi_ftuld (ulong)	Convert unsigned long to double
float	_ftulf (ulong)	__c6xabi_ftulf (ulong)	Convert unsigned long to float
double	_ftllld (long long)	__c6xabi_ftllld (long long)	Convert signed long long to double
float	_ftllif (long long)	__c6xabi_ftllif (long long)	Convert signed long long to float
double	_ftulld (ulong long)	__c6xabi_ftulld (ulong long)	Convert unsigned long long to double
float	_ftullf (ulong long)	__c6xabi_ftullf (ulong long)	Convert unsigned long long to float
double	_absd (double)	__c6xabi_absd (double)	Double absolute value
float	_absf (float)	__c6xabi_absf (float)	Float absolute value
long	_labs (long)	__c6xabi_labs (long)	Long absolute value
long long	_llabs (long long)	__c6xabi_llabs (long long)	Long long absolute value
double	_negd (double)	__c6xabi_negd (double)	Double negate value

Table 7-10. Summary of Run-Time-Support Arithmetic Functions (continued)

Type	Function in COFF ABI	Function in EABI	Description
float	_negf (float)	__c6xabi_negf (float)	Float negate value
long long	_negll (long)	__c6xabi_negll (long)	Long long negate value
long long	_llshl (long long)	__c6xabi_llshl (long long)	Long long shift left
long long	_llshr (long long)	__c6xabi_llshr (long long)	Long long shift right
ulong long	_llshru (ulong long)	__c6xabi_llshru (ulong long)	Unsigned long long shift right
double	_addd (double, double)	__c6xabi_addd (double, double)	Double addition
double	_cmpd (double, double)	__c6xabi_cmpd (double, double)	Double comparison
double	_divd (double, double)	__c6xabi_divd (double, double)	Double division
double	_mpyd (double, double)	__c6xabi_mpyd (double, double)	Double multiplication
double	_subd (double, double)	__c6xabi_subd (double, double)	Double subtraction
float	_addf (float, float)	__c6xabi_addf (float, float)	Float addition
float	_cmpf (float, float)	__c6xabi_cmpf (float, float)	Float comparison
float	_divf (float, float)	__c6xabi_divf (float, float)	Float division
float	_mpyf (float, float)	__c6xabi_mpyf (float, float)	Float multiplication
float	_subf (float, float)	__c6xabi_subf (float, float)	Float subtraction
int	_divi (int, int)	__c6xabi_divi (int, int)	Signed integer division
int	_remi (int, int)	__c6xabi_remi (int, int)	Signed integer remainder
uint	_divu (uint, uint)	__c6xabi_divu (uint, uint)	Unsigned integer division
uint	_remu (uint, uint)	__c6xabi_remu (uint, uint)	Unsigned integer remainder
long	_divli (long, long)	__c6xabi_divli (long, long)	Signed long division
long	_reml (long, long)	__c6xabi_reml (long, long)	Signed long remainder
ulong	_divul (ulong, ulong)	__c6xabi_divul (ulong, ulong)	Unsigned long division
ulong	_remul (ulong, ulong)	__c6xabi_remul (ulong, ulong)	Unsigned long remainder
long long	_divlli (long long, long long)	__c6xabi_divlli (long long, long long)	Signed long long division
long long	_remlli (long long, long long)	__c6xabi_remlli (long long, long long)	Signed long long remainder
ulong long	_mpyll (ulong long, ulong long)	__c6xabi_mpyll (ulong long, ulong long)	Unsigned long long multiplication
ulong long	_divull (ulong long, ulong long)	__c6xabi_divull (ulong long, ulong long)	Unsigned long long division
ulong long	_remull (ulong long, ulong long)	__c6xabi_remull (ulong long, ulong long)	Unsigned long long remainder

7.8 System Initialization

Before you can run a C/C++ program, you must create the C/C++ run-time environment. The C/C++ boot routine performs this task using a function called `c_int00` (or `_c_int00`). The run-time-support source library, `rts.src`, contains the source to this routine in a module named `boot.c` (or `boot.asm`).

To begin running the system, the `c_int00` function can be branched to or called, but it is usually vectored to by reset hardware. You must link the `c_int00` function with the other object files. This occurs automatically when you use the `--rom_model` or `--ram_model` link option and include a standard run-time-support library as one of the linker input files.

When C/C++ programs are linked, the linker sets the entry point value in the executable output file to the symbol `c_int00`. This does not, however, set the hardware to automatically vector to `c_int00` at reset (see the *TMS320C2x DSP CPU and Instruction Set Reference Guide*, the *TMS320C64x/C64x+ DSP CPU and Instruction Set Reference Guide*, the *TMS320C67x/C67x+ DSP CPU and Instruction Set Reference Guide*, or the *TMS320C66x+ DSP CPU and Instruction Set Reference Guide*).

The `c_int00` function performs the following tasks to initialize the environment:

1. Defines a section called `.stack` for the system stack and sets up the initial stack pointers
2. Performs C autoinitialization of global/static variables. For more information, see [Section 7.8.1](#) for COFF ABI mode and [Section 7.8.2](#) for EABI mode.
3. Initializes global variables by copying the data from the initialization tables to the storage allocated for the variables in the `.bss` section. If you are initializing variables at load time (`--ram_model` option), a loader performs this step before the program runs (it is not performed by the boot routine). For more information, see [Section 7.8.1](#).
4. Calls C++ initialization routines for file scope construction from the global constructor table. For more information, see [Section 7.8.2.6](#) for EABI mode and [Section 7.8.1.3](#) for COFF ABI mode.
5. Calls the function `main` to run the C/C++ program

You can replace or modify the boot routine to meet your system requirements. However, the boot routine *must* perform the operations listed above to correctly initialize the C/C++ environment.

7.8.1 COFF ABI Automatic Initialization of Variables

Some global variables must have initial values assigned to them before a C/C++ program starts running. The process of retrieving these variables' data and initializing the variables with the data is called autoinitialization.

The COFF ABI compiler builds tables in a special section called `.cinit` that contains data for initializing global and static variables. Each compiled module contains these initialization tables. The linker combines them into a single table (a single `.cinit` section). The boot routine or a loader uses this table to initialize all the system variables.

Initializing Variables

NOTE: In ANSI/ISO C, global and static variables that are not explicitly initialized must be set to 0 before program execution. The COFF ABI C/C++ compiler does not perform any preinitialization of uninitialized variables. Explicitly initialize any variable that must have an initial value of 0.

Global variables are either autoinitialized at run time or at load time; see [Section 7.8.1.1](#) and [Section 7.8.1.2](#). Also see [Section 6.13](#). In EABI mode, the compiler automatically zero initializes the uninitialized variables. See [Section 7.8.2](#) for details.


7.8.1.1 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 `--rom_model` option.

Using this method, the `.cinit` section is loaded into memory along with all the other initialized sections, and global variables are initialized at run time. 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/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 ROM and copied to RAM each time the program starts.

Figure 7-11 illustrates autoinitialization at run time. Use this method in any system where your application runs from code burned into ROM.

Figure 7-11. Autoinitialization at Run Time



7.8.1.2 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 `--ram_model` option.

When you use the `--ram_model` link 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 (which is not part of the compiler package) 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 7-12 illustrates the initialization of variables at load time.

Figure 7-12. Initialization at Load Time


Regardless of the use of the `--rom_model` or `--ram_model` options, the `.pinit` section is always loaded and processed at run time.

7.8.1.3 Global Constructors for COFF

All global C++ variables that have constructors must have their constructor called before `main`. The compiler builds a table in a section called `.pinit` of global constructor addresses that must be called, in order, before `main`. The linker combines the `.pinit` section from each input file to form a single table in the `.pinit` section. The boot routine uses this table to execute the constructors. (See [Section 7.8.2.6](#) for global constructor details for EABI.)

7.8.2 EABI Automatic Initialization of Variables

Any global variables declared as preinitialized must have initial values assigned to them before a C/C++ program starts running. The process of retrieving these variables' data and initializing the variables with the data is called autoinitialization.

7.8.2.1 Zero Initializing Variables

In ANSI C, global and static variables that are not explicitly initialized, must be set to 0 before program execution. The C/C++ EABI compiler supports preinitialization of uninitialized variables by default. This can be turned off by specifying the linker option `--zero_init=off`. COFF ABI does not support zero initialization.

7.8.2.2 EABI Direct Initialization

The EABI compiler uses direct initialization to initialize global variables. For example, consider the following C code:

```
int i = 23;
int a[5] = { 1, 2, 3, 4, 5 };
```

The compiler allocates the variables `'i'` and `'a[]'` to `.data` section and the initial values are placed directly.

```
.global i
.data
.align 4
i:
    .field      23,32          ; i @ 0

.global a
.data
.align 4
a:
    .field      1,32          ; a[0] @ 0
    .field      2,32          ; a[1] @ 32
    .field      3,32          ; a[2] @ 64
```

```
.field      4,32      ; a[3] @ 96
.field      5,32      ; a[4] @ 128
```

Each compiled module that defines static or global variables contains these .data sections. The linker treats the .data section like any other initialized section and creates an output section. In the load-time initialization model, the sections are loaded into memory and used by the program. See [Section 7.8.2.5](#).

In the run-time initialization model, the linker uses the data in these sections to create initialization data and an additional initialization table. The boot routine processes the initialization table to copy data from load addresses to run addresses. See [Section 7.8.2.3](#).

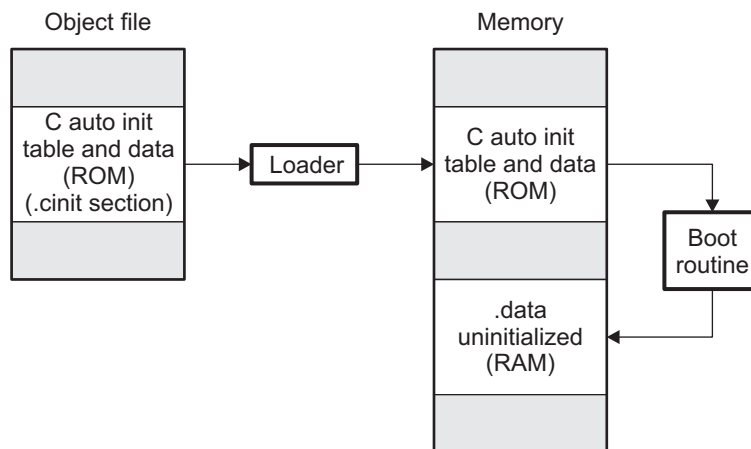
7.8.2.3 Autoinitialization of Variables at Run Time in EABI Mode

Autoinitializing variables at run time is the default method of autoinitialization. To use this method, invoke the linker with the --rom_model option.

Using this method, the linker creates an initialization table and initialization data from the direct initialized sections in the compiled module. The table and data are used by the C/C++ boot routine to initialize variables in RAM using the table and data in ROM.

[Figure 7-13](#) illustrates autoinitialization at run time in EABI Mode. Use this method in any system where your application runs from code burned into ROM.

Figure 7-13. Autoinitialization at Run Time in EABI Mode



7.8.2.4 Autoinitialization Tables

In EABI mode, the compiled object files do not have initialization tables. The variables are initialized directly. The linker, when the --rom_model option is specified, creates C auto initialization table and the initialization data. The linker creates both the table and the initialization data in an output section named .cinit.

Migration from COFF to ELF Initialization

NOTE: The name .cinit is used primarily to simplify migration from COFF to ELF format and the .cinit section created by the linker has nothing in common (except the name) with the COFF cinit records.

The autoinitialization table has the following format:

`__TI_CINIT_Base`:

32-bit load address	32-bit run address
⋮	⋮
32-bit load address	32-bit run address

`__TI_CINIT_Limit`:

The linker defined symbols `__TI_CINIT_Base` and `__TI_CINIT_Limit` point to the start and end of the table, respectively. Each entry in this table corresponds to one output section that needs to be initialized. The initialization data for each output section could be encoded using different encoding.

The load address in the C auto initialization record points to initialization data with the following format:

8-bit index	Encoded data
-------------	--------------

The first 8-bits of the initialization data is the handler index. It indexes into a handler table to get the address of a handler function that knows how to decode the following data.

The handler table is a list of 32-bit function pointers.

`__TI_Handler_Table_Base`:

32-bit handler 1 address
⋮
32-bit handler n address

`__TI_Handler_Table_Limit`:

The *encoded data* that follows the 8-bit index can be in one of the following format types. For clarity the 8-bit index is also depicted for each format.

7.8.2.4.1 Length Followed by Data Format

8-bit index	24-bit padding	32-bit length (N)	N byte initialization data (not compressed)
-------------	----------------	-------------------	---

The compiler uses 24-bit padding to align the length field to a 32-bit boundary. The 32-bit length field encodes the length of the initialization data in bytes (N). N byte initialization data is not compressed and is copied to the run address as is.

The run-time support library has a function `__TI_zero_init()` to process this type of initialization data. The first argument to this function is the address pointing to the byte after the 8-bit index. The second argument is the run address from the C auto initialization record.

7.8.2.4.2 Zero Initialization Format

8-bit index	24-bit padding	32-bit length (N)
-------------	----------------	-------------------

The compiler uses 24-bit padding to align the length field to a 32-bit boundary. The 32-bit length field encodes the number of bytes to be zero initialized.

The run-time support library has a function `__TI_zero_init()` to process the zero initialization. The first argument to this function is the address pointing to the byte after the 8-bit index. The second argument is the run address from the C auto initialization record.

7.8.2.4.3 Run Length Encoded (RLE) Format

8-bit index	Initialization data compressed using run length encoding
-------------	--

The data following the 8-bit index is compressed using Run Length Encoded (RLE) format. uses a simple run length encoding that can be decompressed using the following algorithm:

1. Read the first byte, Delimiter (D).
2. Read the next byte (B).
3. If B != D, copy B to the output buffer and go to step 2.
4. Read the next byte (L).
 - (a) If L == 0, then length is either a 16-bit, a 24-bit value, or we've reached the end of the data, read next byte (L).
 - (i) If L == 0, length is a 24-bit value or the end of the data is reached, read next byte (L).
 - (i) If L == 0, the end of the data is reached, go to step 7.
 - (ii) Else L <= 16, read next two bytes into lower 16 bits of L to complete 24-bit value for L.
 - (ii) Else L <= 8, read next byte into lower 8 bits of L to complete 16-bit value for L.
 - (b) Else if L > 0 and L < 4, copy D to the output buffer L times. Go to step 2.
 - (c) Else, length is 8-bit value (L).
5. Read the next byte (C); C is the repeat character.
6. Write C to the output buffer L times; go to step 2.
7. End of processing.

The run-time support library has a routine `__TI_decompress_rle24()` to decompress data compressed using RLE. The first argument to this function is the address pointing to the byte after the 8-bit index. The second argument is the run address from the C auto initialization record.

RLE Decompression Routine

NOTE: The previous decompression routine, `__TI_decompress_rle()`, is included in the run-time-support library for decompressing RLE encodings that are generated by older versions of the linker.

7.8.2.4.4 Lempel-Ziv-Storer-Szymanski Compression (LZSS) Format

8-bit index	Initialization data compressed using LZSS
-------------	---

The data following the 8-bit index is compressed using LZSS compression. The run-time support library has the routine `__TI_decompress_lzss()` to decompress the data compressed using LZSS. The first argument to this function is the address pointing to the byte after the 8-bit index. The second argument is the run address from the C auto initialization record.

7.8.2.4.5 Sample C Code to Process the C Autoinitialization Table

The run-time support boot routine has code to process the C autoinitialization table. The following C code illustrates how the autoinitialization table can be processed on the target.

Example 7-12. Processing the C Autoinitialization Table

```
typedef void (*handler_fptr)(const unsigned char *in,
unsigned char *out);

#define HANDLER_TABLE __TI_Handler_Table_Base
#pragma WEAK(HANDLER_TABLE)
extern unsigned int HANDLER_TABLE;
```

Example 7-12. Processing the C Autoinitialization Table (continued)

```

extern unsigned char *__TI_CINIT_Base;
extern unsigned char *__TI_CINIT_Limit;

void auto_initialize()
{
    unsigned char **table_ptr;
    unsigned char **table_limit;

    /*-----*/
    /* Check if Handler table has entries.                */
    /*-----*/
    if (&__TI_Handler_Table_Base >= &__TI_Handler_Table_Limit)
        return;

    /*-----*/
    /* Get the Start and End of the CINIT Table.          */
    /*-----*/
    table_ptr = (unsigned char **)&__TI_CINIT_Base;
    table_limit = (unsigned char **)&__TI_CINIT_Limit;
    while (table_ptr < table_limit)
    {
        /*-----*/
        /* 1. Get the Load and Run address.                */
        /* 2. Read the 8-bit index from the load address.   */
        /* 3. Get the handler function pointer using the  */
        /*    index from handler table.                   */
        /*-----*/
        unsigned char *load_addr = *table_ptr++;
        unsigned char *run_addr  = *table_ptr++;
        unsigned char handler_idx = *load_addr++;
        handler_fptr handler =
            (handler_fptr)(&HANDLER_TABLE)[handler_idx];

        /*-----*/
        /* 4. Call the handler and pass the pointer to the */
        /*    load data after index and the run address.   */
        /*-----*/
        (*handler)((const unsigned char *)load_addr, run_addr);
    }
}

```

7.8.2.5 Initialization of Variables at Load Time in EABI Mode

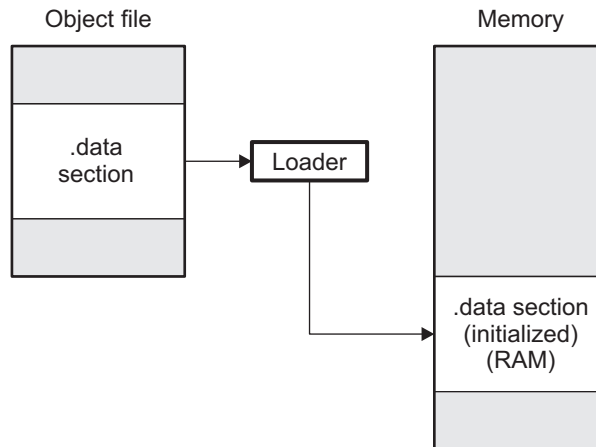
Initialization of variables at load time enhances performance by reducing boot time and by saving the memory used by the initialization tables. To use this method, invoke the linker with the `--ram_model` option.

When you use the `--ram_model` link option, the linker does not generate C autoinitialization tables and data. The direct initialized sections (`.data`) in the compiled object files are combined according to the linker command file to generate initialized output sections. The loader loads the initialized output sections into memory. After the load, the variables are assigned their initial values.

Since the linker does not generate the C autoinitialization tables, no boot time initialization is performed.

[Figure 7-14](#) illustrates the initialization of variables at load time in EABI mode.

Figure 7-14. Initialization at Load Time in EABI Mode



7.8.2.6 Global Constructors

All global C++ variables that have constructors must have their constructor called before main. The compiler builds a table of global constructor addresses that must be called, in order, before main in a section called `.init_array`. The linker combines the `.init_array` section from each input file to form a single table in the `.init_array` section. The boot routine uses this table to execute the constructors. The linker defines two symbols to identify the combined `.init_array` table as shown below. This table is not null terminated by the linker.

Figure 7-15. Constructor Table for EABI Mode

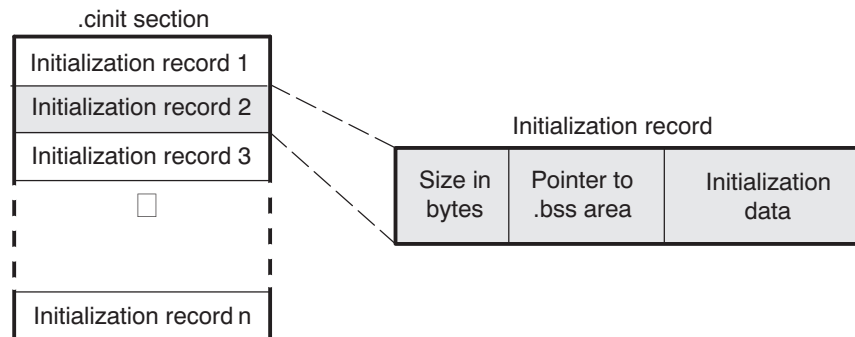
`__TI_INITARRAY_Base:`

Address of constructor 1
Address of constructor 2
⋮
Address of constructor n

`__TI_INITARRAY_Limit:`

7.8.3 Initialization Tables

The tables in the `.cinit` section consist of variable-size initialization records. Each variable that must be autoinitialized has a record in the `.cinit` section. Figure 7-16 shows the format of the `.cinit` section and the initialization records.

Figure 7-16. Format of Initialization Records in the .cinit Section


The fields of an initialization record contain the following information:

- The first field of an initialization record contains the size (in bytes) of the initialization data.
- The second field contains the starting address of the area within the .bss section where the initialization data must be copied.
- The third field contains the data that is copied into the .bss section to initialize the variable.

Each variable that must be autoinitialized has an initialization record in the .cinit section.

[Example 7-13](#) shows initialized global variables defined in C. [Example 7-14](#) shows the corresponding initialization table. The section .cinit:c is a subsection in the .cinit section that contains all scalar data. The subsection is handled as one record during initialization, which minimizes the overall size of the .cinit section.

Example 7-13. Initialized Variables Defined in C

```
int x;
short i = 23;
int *p =
int a[5] = {1,2,3,4,5};
```

Example 7-14. Initialized Information for Variables Defined in [Example 7-13](#)

```

.global _x
.bss    _x,4,4

.sect   ".cinit:c"
.align  8
.field   (CIR - $) - 8, 32
.field   _I+0,32
.field   23,16                ; _I @ 0

.sect   ".text"
.global _I
_I:     .usect  ".bss:c",2,2

.sect   ".cinit:c"
.align  4
.field   _x,32                ; _p @ 0

.sect   ".text"
.global _p
_p:     .usect  ".bss:c",4,4

.sect   ".cinit"
.align  8
.field   IR_1,32
.field   _a+0,32
.field   1,32                ; _a[0] @ 0
.field   2,32                ; _a[1] @ 32
.field   3,32                ; _a[2] @ 64
.field   4,32                ; _a[3] @ 96
.field   5,32                ; _a[4] @ 128
IR_1:   .set    20


.sect   ".text"
.global _a
.bss    _a,20,4
;*****
;* MARK THE END OF THE SCALAR INIT RECORD IN CINIT:C      *
;*****
CIR:    .sect   ".cinit:c"

```

The .cinit section must contain only initialization tables in this format. When interfacing assembly language modules, do not use the .cinit section for any other purpose.

The table in the .pinit section simply consists of a list of addresses of constructors to be called (see [Figure 7-17](#)). The constructors appear in the table after the .cinit initialization.

Figure 7-17. Format of Initialization Records in the .pinit Section



When you use the `--rom_model` or `--ram_model` option, the linker combines the `.cinit` sections from all the C/C++ modules and appends a null word to the end of the composite `.cinit` section. This terminating record appears as a record with a size field of 0 and marks the end of the initialization tables.

Likewise, the `--rom_model` or `--ram_model` link option causes the linker to combine all of the `.pinit` sections from all C/C++ modules and append a null word to the end of the composite `.pinit` section. The boot routine knows the end of the global constructor table when it encounters a null constructor address.

The const-qualified variables are initialized differently; see [Section 6.5.1](#).

7.9 Support for Multi-Threaded Applications

The compiler supports various features for multi-threaded applications. These features assume an underlying runtime operating system that provides thread management services.

7.9.1 Compiling with OpenMP

The compiler implements support for the OpenMP 3.0 API. To enable support for OpenMP use the `--openmp` or `--omp` compiler option.

OpenMP is the industry standard for shared memory parallel programming. It provides portable high-level programming constructs that enable users to easily expose a program's task and loop level parallelism in an incremental fashion. With OpenMP, users specify the parallelization strategy for a program at a high level by annotating the program code with compiler directives that specify how a region of code is executed by a team of threads. The compiler works out the detailed mapping of the computation to the machine. The OpenMP programming API enables the programmer to perform the following:

- Create and manage threads
- Assign and distribute work (tasks) to threads
- Specify which data is shared among threads and which data is private
- Coordinate thread access to shared data

OpenMP is a thread-based programming language. The master thread executes the sequential parts of a program. When the master thread encounters a parallel region, it forks a team of worker threads that along with the master thread execute in parallel.

The OpenMP API is made up of directives (`#pragmas`), function calls, and environment variables. The compiler translates the OpenMP API into multi-threaded code with calls to a custom runtime library that implements support for thread management, shared memory and synchronization. For further details on the OpenMP API (including the API specification), please refer to <http://www.openmp.org>.

The OpenMP runtime for SYS/BIOS (OMP) library implements the OpenMP solution stack. Currently, OpenMP is supported on TI DSPs only for SYS/BIOS operating system. All OpenMP programs must be linked with the OMP run-time library found in the BIOS-MCSDK 2.1.

7.9.2 Multi-Threading Runtime Support

7.9.2.1 Runtime Thread Safety

Thread safety involves creation, initialization, maintenance, and destruction of thread-private data. It also requires that accesses to data that is shared between threads must be protected. That is, only one thread should be allowed to access a piece of shared data at a given time. An additional issue that needs to be addressed on C6000 multi-core devices that have a private data cache (L1D cache on C6600 devices, for example) is that copies of shared data that exist among private data caches and shared memory must be kept coherent. This means that if a thread reads a piece of shared data into the private data cache on a processor, it must invalidate any local copies of that data that exist in the local data cache before accessing or modifying the data. This ensures that the currently executing thread will only access the latest available copy of the shared data.

7.9.2.1.1 Thread Creation, Initialization, and Termination

Thread libraries are responsible for allocating a thread-local area of memory when a thread is created, then initializing any thread-private data objects that reside in that thread-local area, and when the thread is terminated, the thread library must free the thread-local area that was allocated for thread-private data.

7.9.2.1.2 Thread Local Storage (TLS)

The compiler now supports the `__thread` qualifier to identify a data object that is to be given thread-private storage. To access a variable that has been identified as thread-local, the compiler will rely on support of the runtime operating system's thread library to find the run-time location of a thread-local variable. Specifically, the thread library must provide an implementation of the `__c6xabi_get_tp()` function. The runtime operating system's thread library will provide the address of a TLS block that has been allocated on behalf of the currently executed thread and then the compiler can access data in the TLS block with knowledge of where a given thread-local variable exists within the TLS block. Further information about TLS data objects can be found in the "Thread Local Storage Specification for C6000 EABI" document (C6000_TLS.docx in the 'doc' sub-directory where your C6000 CGT is installed).

7.9.2.1.3 Accessing Shared Data

Accesses to shared data objects must be protected within a critical region which prevents a second thread from entering the critical region of code while the first thread is accessing a shared data object. We also need to be concerned about data coherency between copies of shared data that may exist in both shared memory and in private data caches as mentioned above in [Section 7.9.2.1](#).

Using Run-Time-Support Functions and Building Libraries

Some of the features of C/C++ (such as I/O, dynamic memory allocation, string operations, and trigonometric functions) are provided as an ANSI/ISO C/C++ standard library, rather than as part of the compiler itself. The TI implementation of this library is the run-time-support library (RTS). The C/C++ compiler implements the complete ISO standard library except for those facilities that handle signal and locale issues (properties that depend on local language, nationality, or culture). Using the ANSI/ISO standard library ensures a consistent set of functions that provide for greater portability.

In addition to the ANSI/ISO-specified functions, the run-time-support library includes routines that give you processor-specific commands and direct C language I/O requests. These are detailed in [Section 8.1](#) and [Section 8.2](#).

A library-build utility is provided with the code generation tools that lets you create customized run-time-support libraries. This process is described in [Section 8.5](#).

Topic	Page
8.1 C and C++ Run-Time Support Libraries	248
8.2 The C I/O Functions	250
8.3 Handling Reentrancy (_register_lock() and _register_unlock() Functions)	262
8.4 C6700 FastMath Library	263
8.5 Library-Build Process	264

8.1 C and C++ Run-Time Support Libraries

TMS320C6000 compiler releases include pre-built run-time libraries that provide all the standard capabilities. Separate libraries are provided for each target CPU version, big and little endian support, each ABI, and C++ exception support. See [Section 8.5](#) for information on the library-naming conventions.

The run-time-support library contains the following:

- ANSI/ISO C/C++ standard library
- C I/O library
- Low-level support functions that provide I/O to the host operating system
- Fundamental arithmetic routines
- System startup routine, `_c_int00`
- Functions and macros that allow C/C++ to access specific instructions

The run-time-support libraries do not contain functions involving signals and locale issues.

The C++ library supports wide chars, in that template functions and classes that are defined for `char` are also available for wide char. For example, wide char stream classes `wios`, `wiostream`, `wstreambuf` and so on (corresponding to `char` classes `ios`, `iostream`, `streambuf`) are implemented. However, there is no low-level file I/O for wide chars. Also, the C library interface to wide char support (through the C++ headers `<wchar>` and `<cwctype>`) is limited as described in [Section 6.1](#).

The C++ library included with the compiler is licensed from [Dinkumware, Ltd.](#) The Dinkumware C++ library is a fully conforming, industry-leading implementation of the standard C++ library.

TI does not provide documentation that covers the functionality of the C++ library. TI suggests referring to one of the following sources:

- *The Standard C++ Library: A Tutorial and Reference*, Nicolai M. Josuttis, Addison-Wesley, ISBN 0-201-37926-0
- *The C++ Programming Language* (Third or Special Editions), Bjarne Stroustrup, Addison-Wesley, ISBN 0-201-88954-4 or 0-201-70073-5
- Dinkumware's online reference at <http://dinkumware.com/manuals>

8.1.1 Linking Code With the Object Library

When you link your program, you must specify the object library as one of the linker input files so that references to the I/O and run-time-support functions can be resolved. You can either specify the library or allow the compiler to select one for you. See [Section 5.3.1](#) for further information.

When a library is linked, the linker includes only those library members required to resolve undefined references. For more information about linking, see the *TMS320C6000 Assembly Language Tools User's Guide*.

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.

8.1.2 Header Files

You must use the header files provided with the compiler run-time support when using functions from C/C++ standard library. Set the C6X_C_DIR environment variable to the include directory where the tools are installed.

8.1.3 Modifying a Library Function

You can inspect or modify library functions by unzipping the source file (rtssrc.zip), changing the specific function file, and rebuilding the library. When extracted (with any standard unzip tool on windows, linux, or unix), this zip file recreates the run-time source tree for the run-time library.

The source for the libraries is included in the rtssrc.zip file. See [Section 8.5](#) for details on rebuilding.

You can also build a new library this way, rather than rebuilding into rts6200.lib. See [Section 8.5](#).

8.1.4 Minimal Support for Internationalization

The library now includes the header files <locale.h>, <wchar.h>, and <wctype.h>, which provide APIs to support non-ASCII character sets and conventions. Our implementation of these APIs is limited in the following ways:

- The library has minimal support for wide and multi-byte characters. The type `wchar_t` is implemented as `int`. The wide character set is equivalent to the set of values of type `char`. The library includes the header files <wchar.h> and <wctype.h> but does not include all the functions specified in the standard. So-called multi-byte characters are limited to single characters. There are no shift states. The mapping between multi-byte characters and wide characters is simple equivalence; that is, each wide character maps to and from exactly a single multi-byte character having the same value.
- The C library includes the header file <locale.h> but with a minimal implementation. The only supported locale is the C locale. That is, library behavior that is specified to vary by locale is hard-coded to the behavior of the C locale, and attempting to install a different locale via a call to `setlocale()` will return `NULL`.

8.1.5 Allowable Number of Open Files

In the <stdio.h> header file, the value for the macro `FOPEN_MAX` has been changed from 12 to the value of the macro `_NFILE`, which is set to 10. The impact is that you can only have 10 files simultaneously open at one time (including the pre-defined streams - `stdin`, `stdout`, `stderr`).

The C standard requires that the minimum value for the `FOPEN_MAX` macro is 8. The macro determines the maximum number of files that can be opened at one time. The macro is defined in the `stdio.h` header file and can be modified by changing the value of the `_NFILE` macro.

8.1.6 Library Naming Conventions

By default, the linker uses automatic library selection to select the correct run-time-support library (see [Section 5.3.1.1](#)) for your application. If you select the library manually, you must select the matching library according to the following naming scheme:

`rtstrg[endian][abi][eh].lib`

<i>trg</i>	The device family of the C6000 architecture that the library was built for. This can be one of the following: 6200, 6400, 64plus, 6600, 6700, 6740, 67plus.
<i>endian</i>	Indicates endianness: (blank) Little-endian library e Big-endian library
<i>abi</i>	Indicates the application binary interface (ABI) used: (blank) COFF ABI _elf EABI
<i>eh</i>	Indicates whether the library has exception handling support (blank) exception handling not supported _eh exception handling support

For information on the C6700 FastMath source library, `fastmathc67x.src`, see [Section 8.4](#).

8.2 The C I/O Functions

The C I/O functions make it possible to access the host's operating system to perform I/O. The capability to perform I/O on the host gives you more options when debugging and testing code.

The I/O functions are logically divided into layers: high level, low level, and device-driver level.

With properly written device drivers, the C-standard high-level I/O functions can be used to perform I/O on custom user-defined devices. This provides an easy way to use the sophisticated buffering of the high-level I/O functions on an arbitrary device.

Debugger Required for Default HOST

NOTE: For the default HOST device to work, there must be a debugger to handle the C I/O requests; the default HOST device cannot work by itself in an embedded system. To work in an embedded system, you will need to provide an appropriate driver for your system.

NOTE: C I/O Mysteriously Fails

If there is not enough space on the heap for a C I/O buffer, operations on the file will silently fail. If a call to `printf()` mysteriously fails, this may be the reason. The heap needs to be at least large enough to allocate a block of size `BUFSIZ` (defined in `stdio.h`) for every file on which I/O is performed, including `stdout`, `stdin`, and `stderr`, plus allocations performed by the user's code, plus allocation bookkeeping overhead. Alternately, declare a char array of size `BUFSIZ` and pass it to `setvbuf` to avoid dynamic allocation. To set the heap size, use the `--heap_size` option when linking (refer to the *Linker Description* chapter in the *TMS320C6000 Assembly Language Tools User's Guide*).

NOTE: Open Mysteriously Fails

The run-time support limits the total number of open files to a small number relative to general-purpose processors. If you attempt to open more files than the maximum, you may find that the open will mysteriously fail. You can increase the number of open files by extracting the source code from `rts.src` and editing the constants controlling the size of some of the C I/O data structures. The macro `_NFILE` controls how many FILE (fopen) objects can be open at one time (stdin, stdout, and stderr count against this total). (See also `FOPEN_MAX`.) The macro `_NSTREAM` controls how many low-level file descriptors can be open at one time (the low-level files underlying stdin, stdout, and stderr count against this total). The macro `_NDEVICE` controls how many device drivers are installed at one time (the HOST device counts against this total).

8.2.1 High-Level I/O Functions

The high-level functions are the standard C library of stream I/O routines (printf, scanf, fopen, getchar, and so on). These functions call one or more low-level I/O functions to carry out the high-level I/O request. The high-level I/O routines operate on FILE pointers, also called *streams*.

Portable applications should use only the high-level I/O functions.

To use the high-level I/O functions, include the header file `stdio.h`, or `cstdio` for C++ code, for each module that references a C I/O function.

For example, given the following C program in a file named `main.c`:

```
#include <stdio.h>

void main()
{
    FILE *fid;

    fid = fopen("myfile","w");
    fprintf(fid,"Hello, world\n");
    fclose(fid);

    printf("Hello again, world\n");
}
```

Issuing the following compiler command compiles, links, and creates the file `main.out` from the run-time-support library:

```
cl6x main.c -z --heap_size=1000 --output_file=main.out
```

Executing `main.out` results in

```
Hello, world
```

being output to a file and

```
Hello again, world
```

being output to your host's stdout window.

8.2.2 Overview of Low-Level I/O Implementation

The low-level functions are comprised of seven basic I/O functions: open, read, write, close, lseek, rename, and unlink. These low-level routines provide the interface between the high-level functions and the device-level drivers that actually perform the I/O command on the specified device.

The low-level functions are designed to be appropriate for all I/O methods, even those which are not actually disk files. Abstractly, all I/O channels can be treated as files, although some operations (such as lseek) may not be appropriate. See [Section 8.2.3](#) for more details.

The low-level functions are inspired by, but not identical to, the POSIX functions of the same names.

The low-level functions operate on file descriptors. A file descriptor is an integer returned by open, representing an opened file. Multiple file descriptors may be associated with a file; each has its own independent file position indicator.

open	<i>Open File for I/O</i>
Syntax	<pre>#include <file.h> int open (const char * path , unsigned flags , int file_descriptor);</pre>
Description	<p>The open function opens the file specified by <i>path</i> and prepares it for I/O.</p> <ul style="list-style-type: none"> • The <i>path</i> is the filename of the file to be opened, including an optional directory path and an optional device specifier (see Section 8.2.5). • The <i>flags</i> are attributes that specify how the file is manipulated. The flags are specified using the following symbols: <pre>O_RDONLY (0x0000) /* open for reading */ O_WRONLY (0x0001) /* open for writing */ O_RDWR (0x0002) /* open for read & write */ O_APPEND (0x0008) /* append on each write */ O_CREAT (0x0200) /* open with file create */ O_TRUNC (0x0400) /* open with truncation */ O_BINARY (0x8000) /* open in binary mode */</pre> <p>Low-level I/O routines allow or disallow some operations depending on the flags used when the file was opened. Some flags may not be meaningful for some devices, depending on how the device implements files.</p> • The <i>file_descriptor</i> is assigned by open to an opened file. <p>The next available file descriptor is assigned to each new file opened.</p>
Return Value	<p>The function returns one of the following values:</p> <pre>non-negative file descriptor if successful -1 on failure</pre>

close ***Close File for I/O***

Syntax `#include <file.h>`
int close (int file_descriptor);

Description The close function closes the file associated with *file_descriptor*.
The *file_descriptor* is the number assigned by open to an opened file.

Return Value The return value is one of the following:
 0 if successful
 -1 on failure

read ***Read Characters from a File***

Syntax `#include <file.h>`
int read (int file_descriptor , char * buffer , unsigned count);

Description The read function reads *count* characters into the *buffer* from the file associated with *file_descriptor*.

- The *file_descriptor* is the number assigned by open to an opened file.
- The *buffer* is where the read characters are placed.
- The *count* is the number of characters to read from the file.

Return Value The function returns one of the following values:
 0 if EOF was encountered before any characters were read
 # number of characters read (may be less than *count*)
 -1 on failure

write ***Write Characters to a File***

Syntax `#include <file.h>`
int write (int file_descriptor , const char * buffer , unsigned count);

Description The write function writes the number of characters specified by *count* from the *buffer* to the file associated with *file_descriptor*.

- The *file_descriptor* is the number assigned by open to an opened file.
- The *buffer* is where the characters to be written are located.
- The *count* is the number of characters to write to the file.

Return Value The function returns one of the following values:
 # number of characters written if successful (may be less than *count*)
 -1 on failure

lseek	<i>Set File Position Indicator</i>
Syntax for C	<pre>#include <file.h> off_t lseek (int file_descriptor , off_t offset , int origin);</pre>
Description	<p>The lseek function sets the file position indicator for the given file to a location relative to the specified origin. The file position indicator measures the position in characters from the beginning of the file.</p> <ul style="list-style-type: none"> • The <i>file_descriptor</i> is the number assigned by open to an opened file. • The <i>offset</i> indicates the relative offset from the <i>origin</i> in characters. • The <i>origin</i> is used to indicate which of the base locations the <i>offset</i> is measured from. The <i>origin</i> must be one of the following macros: <ul style="list-style-type: none"> SEEK_SET (0x0000) Beginning of file SEEK_CUR (0x0001) Current value of the file position indicator SEEK_END (0x0002) End of file
Return Value	<p>The return value is one of the following:</p> <pre># new value of the file position indicator if successful (off_t)-1 on failure</pre>
unlink	<i>Delete File</i>
Syntax	<pre>#include <file.h> int unlink (const char * path);</pre>
Description	<p>The unlink function deletes the file specified by <i>path</i>. Depending on the device, a deleted file may still remain until all file descriptors which have been opened for that file have been closed. See Section 8.2.3.</p> <p>The <i>path</i> is the filename of the file, including path information and optional device prefix. (See Section 8.2.5.)</p>
Return Value	<p>The function returns one of the following values:</p> <pre>0 if successful -1 on failure</pre>

rename	<i>Rename File</i>				
Syntax for C	<pre>#include {<stdio.h> <file.h>} int rename (const char * old_name , const char * new_name);</pre>				
Syntax for C++	<pre>#include {<cstdio> <file.h>} int std::rename (const char * old_name , const char * new_name);</pre>				
Description	<p>The rename function changes the name of a file.</p> <ul style="list-style-type: none"> • The <i>old_name</i> is the current name of the file. • The <i>new_name</i> is the new name for the file. <hr/> <p>NOTE: The optional device specified in the new name must match the device of the old name. If they do not match, a file copy would be required to perform the rename, and rename is not capable of this action.</p> <hr/>				
Return Value	<p>The function returns one of the following values:</p> <table border="0"> <tr> <td>0</td> <td>if successful</td> </tr> <tr> <td>-1</td> <td>on failure</td> </tr> </table> <hr/> <p>NOTE: Although rename is a low-level function, it is defined by the C standard and can be used by portable applications.</p> <hr/>	0	if successful	-1	on failure
0	if successful				
-1	on failure				

8.2.3 Device-Driver Level I/O Functions

At the next level are the device-level drivers. They map directly to the low-level I/O functions. The default device driver is the HOST device driver, which uses the debugger to perform file operations. The HOST device driver is automatically used for the default C streams stdin, stdout, and stderr.

The HOST device driver shares a special protocol with the debugger running on a host system so that the host can perform the C I/O requested by the program. Instructions for C I/O operations that the program wants to perform are encoded in a special buffer named `_CIOBUF_` in the `.cio` section. The debugger halts the program at a special breakpoint (C\$\$IO\$\$), reads and decodes the target memory, and performs the requested operation. The result is encoded into `_CIOBUF_`, the program is resumed, and the target decodes the result.

The HOST device is implemented with seven functions, `HOSTopen`, `HOSTclose`, `HOSTread`, `HOSTwrite`, `HOSTlseek`, `HOSTunlink`, and `HOSTrename`, which perform the encoding. Each function is called from the low-level I/O function with a similar name.

A device driver is composed of seven required functions. Not all function need to be meaningful for all devices, but all seven must be defined. Here we show the names of all seven functions as starting with `DEV`, but you may chose any name except for `HOST`.

DEV_open
Open File for I/O
Syntax

```
int DEV_open (const char * path , unsigned flags , int llv_fd);
```

Description

This function finds a file matching *path* and opens it for I/O as requested by *flags*.

- The *path* is the filename of the file to be opened. If the name of a file passed to open has a device prefix, the device prefix will be stripped by open, so DEV_open will not see it. (See [Section 8.2.5](#) for details on the device prefix.)
- The *flags* are attributes that specify how the file is manipulated. The flags are specified using the following symbols:

```
O_RDONLY   (0x0000)  /* open for reading */
O_WRONLY   (0x0001)  /* open for writing */
O_RDWR     (0x0002)  /* open for read & write */
O_APPEND   (0x0008)  /* append on each write */
O_CREAT     (0x0200)  /* open with file create */
O_TRUNC    (0x0400)  /* open with truncation */
O_BINARY   (0x8000)  /* open in binary mode */
```

See POSIX for further explanation of the flags.

- The *llv_fd* is treated as a suggested low-level file descriptor. This is a historical artifact; newly-defined device drivers should ignore this argument. This differs from the low-level I/O open function.

This function must arrange for information to be saved for each file descriptor, typically including a file position indicator and any significant flags. For the HOST version, all the bookkeeping is handled by the debugger running on the host machine. If the device uses an internal buffer, the buffer can be created when a file is opened, or the buffer can be created during a read or write.

Return Value

This function must return -1 to indicate an error if for some reason the file could not be opened; such as the file does not exist, could not be created, or there are too many files open. The value of *errno* may optionally be set to indicate the exact error (the HOST device does not set *errno*). Some devices might have special failure conditions; for instance, if a device is read-only, a file cannot be opened O_WRONLY.

On success, this function must return a non-negative file descriptor unique among all open files handled by the specific device. It need not be unique across devices. Only the low-level I/O functions will see this device file descriptor; the low-level function open will assign its own unique file descriptor.

DEV_close	<i>Close File for I/O</i>
<hr/>	
Syntax	int DEV_close (int <i>dev_fd</i>);
Description	<p>This function closes a valid open file descriptor.</p> <p>On some devices, DEV_close may need to be responsible for checking if this is the last file descriptor pointing to a file that was unlinked. If so, it is responsible for ensuring that the file is actually removed from the device and the resources reclaimed, if appropriate.</p>
Return Value	<p>This function should return -1 to indicate an error if the file descriptor is invalid in some way, such as being out of range or already closed, but this is not required. The user should not call close() with an invalid file descriptor.</p>
DEV_read	<i>Read Characters from a File</i>
<hr/>	
Syntax	int DEV_read (int <i>dev_fd</i> , char * <i>bu</i> , unsigned <i>count</i>);
Description	<p>The read function reads <i>count</i> bytes from the input file associated with <i>dev_fd</i>.</p> <ul style="list-style-type: none"> • The <i>dev_fd</i> is the number assigned by open to an opened file. • The <i>buf</i> is where the read characters are placed. • The <i>count</i> is the number of characters to read from the file.
Return Value	<p>This function must return -1 to indicate an error if for some reason no bytes could be read from the file. This could be because of an attempt to read from a O_WRONLY file, or for device-specific reasons.</p> <p>If count is 0, no bytes are read and this function returns 0.</p> <p>This function returns the number of bytes read, from 0 to count. 0 indicates that EOF was reached before any bytes were read. It is not an error to read less than count bytes; this is common if there are not enough bytes left in the file or the request was larger than an internal device buffer size.</p>
DEV_write	<i>Write Characters to a File</i>
<hr/>	
Syntax	int DEV_write (int <i>dev_fd</i> , const char * <i>buf</i> , unsigned <i>count</i>);
Description	<p>This function writes <i>count</i> bytes to the output file.</p> <ul style="list-style-type: none"> • The <i>dev_fd</i> is the number assigned by open to an opened file. • The <i>buffer</i> is where the write characters are placed. • The <i>count</i> is the number of characters to write to the file.
Return Value	<p>This function must return -1 to indicate an error if for some reason no bytes could be written to the file. This could be because of an attempt to read from a O_RDONLY file, or for device-specific reasons.</p>

DEV_lseek ***Set File Position Indicator***

Syntax **off_t lseek (int dev_fd , off_t offset , int origin);**

Description This function sets the file's position indicator for this file descriptor as [lseek](#).

If lseek is supported, it should not allow a seek to before the beginning of the file, but it should support seeking past the end of the file. Such seeks do not change the size of the file, but if it is followed by a write, the file size will increase.

Return Value If successful, this function returns the new value of the file position indicator.

This function must return -1 to indicate an error if for some reason no bytes could be written to the file. For many devices, the lseek operation is nonsensical (e.g. a computer monitor).

DEV_unlink ***Delete File***

Syntax **int DEV_unlink (const char * path);**

Description Remove the association of the pathname with the file. This means that the file may no longer be opened using this name, but the file may not actually be immediately removed.

Depending on the device, the file may be immediately removed, but for a device which allows open file descriptors to point to unlinked files, the file will not actually be deleted until the last file descriptor is closed. See [Section 8.2.3](#).

Return Value This function must return -1 to indicate an error if for some reason the file could not be unlinked (delayed removal does not count as a failure to unlink.)

If successful, this function returns 0.

DEV_rename ***Rename File***

Syntax **int DEV_rename (const char * old_name , const char * new_name);**

Description This function changes the name associated with the file.

- The *old_name* is the current name of the file.
- The *new_name* is the new name for the file.

Return Value This function must return -1 to indicate an error if for some reason the file could not be renamed, such as the file doesn't exist, or the new name already exists.

NOTE: It is inadvisable to allow renaming a file so that it is on a different device. In general this would require a whole file copy, which may be more expensive than you expect.

If successful, this function returns 0.

8.2.4 Adding a User-Defined Device Driver for C I/O

The function `add_device` allows you to add and use a device. When a device is registered with `add_device`, the high-level I/O routines can be used for I/O on that device.

You can use a different protocol to communicate with any desired device and install that protocol using `add_device`; however, the HOST functions should not be modified. The default streams `stdin`, `stdout`, and `stderr` can be remapped to a file on a user-defined device instead of HOST by using `freopen()` as in [Example 8-1](#). If the default streams are reopened in this way, the buffering mode will change to `_IOFBF` (fully buffered). To restore the default buffering behavior, call `setvbuf` on each reopened file with the appropriate value (`_IOLBF` for `stdin` and `stdout`, `_IONBF` for `stderr`).

The default streams `stdin`, `stdout`, and `stderr` can be mapped to a file on a user-defined device instead of HOST by using `freopen()` as shown in [Example 8-1](#). Each function must set up and maintain its own data structures as needed. Some function definitions perform no action and should just return.

Example 8-1. Mapping Default Streams to Device

```
#include <stdio.h>
#include <file.h>
#include "mydevice.h"

void main()
{
    add_device("mydevice", _MSA,
              MYDEVICE_open, MYDEVICE_close,
              MYDEVICE_read, MYDEVICE_write,
              MYDEVICE_lseek, MYDEVICE_unlink, MYDEVICE_rename);

    /*-----*/
    /* Re-open stderr as a MYDEVICE file */
    /*-----*/
    if (!freopen("mydevice:stderrfile", "w", stderr))
    {
        puts("Failed to freopen stderr");
        exit(EXIT_FAILURE);
    }

    /*-----*/
    /* stderr should not be fully buffered; we want errors to be seen as */
    /* soon as possible. Normally stderr is line-buffered, but this example */
    /* doesn't buffer stderr at all. This means that there will be one call */
    /* to write() for each character in the message. */
    /*-----*/
    if (setvbuf(stderr, NULL, _IONBF, 0))
    {
        puts("Failed to setvbuf stderr");
        exit(EXIT_FAILURE);
    }

    /*-----*/
    /* Try it out! */
    /*-----*/
    printf("This goes to stdout\n");
    fprintf(stderr, "This goes to stderr\n"); }
```

NOTE: Use Unique Function Names

The function names `open`, `read`, `write`, `close`, `lseek`, `rename`, and `unlink` are used by the low-level routines. Use other names for the device-level functions that you write.

Use the low-level function `add_device()` to add your device to the `device_table`. The device table is a statically defined array that supports n devices, where n is defined by the macro `_NDEVICE` found in `stdio.h/cstdio`.

The first entry in the device table is predefined to be the host device on which the debugger is running. The low-level routine `add_device()` finds the first empty position in the device table and initializes the device fields with the passed-in arguments. For a complete description, see [the `add_device` function](#).

8.2.5 The device Prefix

A file can be opened to a user-defined device driver by using a device prefix in the pathname. The device prefix is the device name used in the call to `add_device` followed by a colon. For example:

```
FILE *fptr = fopen("mydevice:file1", "r");
int fd = open("mydevice:file2, O_RDONLY, 0);
```

If no device prefix is used, the HOST device will be used to open the file.

add_device	Add Device to Device Table				
<hr/>					
Syntax for C	<pre>#include <file.h> int add_device(char * name, unsigned flags , int (* dopen)(const char *path, unsigned flags, int llv_fd), int (* dclose)(int dev_fd), int (* dread)(int dev_fd, char *buf, unsigned count), int (* dwrite)(int dev_fd, const char *buf, unsigned count), off_t (* dlseek)(int dev_fd, off_t ioffset, int origin), int (* dunlink)(const char * path), int (* drename)(const char *old_name, const char *new_name));</pre>				
Defined in	lowlev.c in rtssrc.zip				
Description	<p>The <code>add_device</code> function adds a device record to the device table allowing that device to be used for I/O from C. The first entry in the device table is predefined to be the HOST device on which the debugger is running. The function <code>add_device()</code> finds the first empty position in the device table and initializes the fields of the structure that represent a device.</p> <p>To open a stream on a newly added device use <code>fopen()</code> with a string of the format <code>devicename : filename</code> as the first argument.</p> <ul style="list-style-type: none"> • The <i>name</i> is a character string denoting the device name. The name is limited to 8 characters. • The <i>flags</i> are device characteristics. The flags are as follows: <ul style="list-style-type: none"> <code>_SSA</code> Denotes that the device supports only one open stream at a time <code>_MSA</code> Denotes that the device supports multiple open streams More flags can be added by defining them in <code>file.h</code>. • The <i>dopen</i>, <i>dclose</i>, <i>dread</i>, <i>dwrite</i>, <i>dlseek</i>, <i>dunlink</i>, and <i>drename</i> specifiers are function pointers to the functions in the device driver that are called by the low-level functions to perform I/O on the specified device. You must declare these functions with the interface specified in Section 8.2.2. The device driver for the HOST that the TMS320C6000 debugger is run on are included in the C I/O library. 				
Return Value	<p>The function returns one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">0</td> <td>if successful</td> </tr> <tr> <td>-1</td> <td>on failure</td> </tr> </table>	0	if successful	-1	on failure
0	if successful				
-1	on failure				

Example

[Example 8-2](#) does the following:

- Adds the device *mydevice* to the device table
- Opens a file named *test* on that device and associates it with the FILE pointer *fid*
- Writes the string *Hello, world* into the file
- Closes the file

[Example 8-2](#) illustrates adding and using a device for C I/O:

Example 8-2. Program for C I/O Device

```
#include <file.h>
#include <stdio.h>
/*****
/* Declarations of the user-defined device drivers          */
*****/
extern int  MYDEVICE_open(const char *path, unsigned flags, int fno);
extern int  MYDEVICE_close(int fno);
extern int  MYDEVICE_read(int fno, char *buffer, unsigned count);
extern int  MYDEVICE_write(int fno, const char *buffer, unsigned count);
extern off_t MYDEVICE_lseek(int fno, off_t offset, int origin);
extern int  MYDEVICE_unlink(const char *path);
extern int  MYDEVICE_rename(const char *old_name, char *new_name);
main()
{
    FILE *fid;
    add_device("mydevice", _MSA, MYDEVICE_open, MYDEVICE_close, MYDEVICE_read,
              MYDEVICE_write, MYDEVICE_lseek, MYDEVICE_unlink, MYDEVICE_rename);
    fid = fopen("mydevice:test", "w");
    fprintf(fid, "Hello, world\n");

    fclose(fid);
}
```

8.3 Handling Reentrancy (_register_lock() and _register_unlock() Functions)

The C standard assumes only one thread of execution, with the only exception being extremely narrow support for signal handlers. The issue of reentrancy is avoided by not allowing you to do much of anything in a signal handler. However, BIOS applications have multiple threads which need to modify the same global program state, such as the CIO buffer, so reentrancy is a concern.

Part of the problem of reentrancy remains your responsibility, but the run-time-support environment does provide rudimentary support for multi-threaded reentrancy by providing support for critical sections. This implementation does not protect you from reentrancy issues such as calling run-time-support functions from inside interrupts; this remains your responsibility.

The run-time-support environment provides hooks to install critical section primitives. By default, a single-threaded model is assumed, and the critical section primitives are not employed. In a multi-threaded system such as BIOS, the kernel arranges to install semaphore lock primitive functions in these hooks, which are then called when the run-time-support enters code that needs to be protected by a critical section.

Throughout the run-time-support environment where a global state is accessed, and thus needs to be protected with a critical section, there are calls to the function `_lock()`. This calls the provided primitive, if installed, and acquires the semaphore before proceeding. Once the critical section is finished, `_unlock()` is called to release the semaphore.

Usually BIOS is responsible for creating and installing the primitives, so you do not need to take any action. However, this mechanism can be used in multi-threaded applications which do not use the BIOS LCK mechanism.

You should not define the functions `_lock()` and `_unlock()` functions directly; instead, the installation functions are called to instruct the run-time-support environment to use these new primitives:

```
void _register_lock (void ( *lock)());
```

```
void _register_unlock(void (*unlock)());
```

The arguments to `_register_lock()` and `_register_unlock()` should be functions which take no arguments and return no values, and which implement some sort of global semaphore locking:

```
extern volatile sig_atomic_t *sema = SHARED_SEMAPHORE_LOCATION;
static int sema_depth = 0;
static void my_lock(void)
{
    while (ATOMIC_TEST_AND_SET(sema, MY_UNIQUE_ID) != MY_UNIQUE_ID);
    sema_depth++;
}
static void my_unlock(void)
{
    if (--sema_depth) ATOMIC_CLEAR(sema);
}
```

The run-time-support nests calls to `_lock()`, so the primitives must keep track of the nesting level.

8.4 C6700 FastMath Library

The C6700 FastMath Library provides hand-coded assembly-optimized versions of certain math functions. These implementations are two to three times faster than those found in the standard run-time-support library. However, these functions gain speed improvements at the cost of accuracy in the result.

The C6700 FastMath library contains these files:

- `fastmath67x.lib`—object library for use with little-endian C/C++ code
- `fastmath67xe.lib`—object library for use with big-endian C/C++ code
- `fastmath67x.h`—header file to be included with C/C++ code

To use the C67x FastMath library, specify it before the standard run-time-support library when linking your program. For example:

```
cl6x -mv6700 --run_linker myprogram.obj --library=lnk.cmd --library=fastmath67x.lib --  
library=rts6700.lib
```

If you are using Code Composer Studio, include the C6700 FastMath library in your project, and ensure it appears before the standard run-time-support library in the Link Order tab in the Build Options dialog box.

For details, refer to the *TMS320C67x FastRTS Library Programmer's Reference* ([SPRU100](#)).

8.5 Library-Build Process

When using the C/C++ compiler, you can compile your code under a large number of different configurations and options that are not necessarily compatible with one another. Because it would be infeasible to include all possible run-time-support library variants, compiler releases pre-build only a small number of very commonly-used libraries such as `rts6200.lib`.

To provide maximum flexibility, the run-time-support source code is provided as part of each compiler release. You can build the missing libraries as desired. The linker can also automatically build missing libraries. This is accomplished with a new library build process, the core of which is the executable `mklib`, which is available beginning with CCS 5.1

8.5.1 Required Non-Texas Instruments Software

To use the self-contained run-time-support build process to rebuild a library with custom options, the following are required:

- `sh` (Bourne shell)
- `unzip` (InfoZIP `unzip` 5.51 or later, or equivalent)
You can download the software from <http://www.info-zip.org>.
- `gmake` (GNU make 3.81 or later)

More information is available from GNU at <http://www.gnu.org/software/make>. GNU make (`gmake`) is also available in earlier versions of Code Composer Studio. GNU make is also included in some Unix support packages for Windows, such as the MKS Toolkit, Cygwin, and Interix. The GNU make used on Windows platforms should explicitly report "This program build for Windows32" when the following is executed from the Command Prompt window:

```
gmake -h
```

All three of these programs are provided as a non-optional feature of CCS 5.1. They are also available as part of the optional XDC Tools feature if you are using an earlier version of CCS.

The `mklib` program looks for these executables in the following order:

1. in your `PATH`
2. in the directory `getenv("CCS_UTILS_DIR")/cygwin`
3. in the directory `getenv("CCS_UTILS_DIR")/bin`
4. in the directory `getenv("XDCROOT")`
5. in the directory `getenv("XDCROOT")/bin`

If you are invoking `mklib` from the command line, and these executables are not in your path, you must set the environment variable `CCS_UTILS_DIR` such that `getenv("CCS_UTILS_DIR")/bin` contains the correct programs.

8.5.2 Using the Library-Build Process

You should normally let the linker automatically rebuild libraries as needed. If necessary, you can run `mklib` directly to populate libraries. See [Section 8.5.2.2](#) for situations when you might want to do this.

8.5.2.1 Automatic Standard Library Rebuilding by the Linker

The linker looks for run-time-support libraries primarily through the `C6X_C_DIR` environment variable. Typically, one of the pathnames in `C6X_C_DIR` is *your install directory/lib*, which contains all of the pre-built libraries, as well as the index library `libc.a`. The linker looks in `C6X_C_DIR` to find a library that is the best match for the build attributes of the application. The build attributes are set indirectly according to the command-line options used to build the application. Build attributes include things like CPU revision. If the library is explicitly named (e.g. `rts6200.lib`), run-time support looks for that library exactly; otherwise, it uses the index library `libc.a` to pick an appropriate library.

The index library describes a set of libraries with different build attributes. The linker will compare the build attributes for each potential library with the build attributes of the application and will pick the best fit. For details on the index library, see the archiver chapter in the *TMS320C6000 Assembly Language Tools User's Guide*.

Now that the linker has decided which library to use, it checks whether the run-time-support library is present in C6X_C_DIR . The library must be in exactly the same directory as the index library libc.a. If the library is not present, the linker will invoke mklib to build it. This happens when the library is missing, regardless of whether the user specified the name of the library directly or allowed the linker to pick the best library from the index library.

The mklib program builds the requested library and places it in 'lib' directory part of C6X_C_DIR in the same directory as the index library, so it is available for subsequent compilations.

Things to watch out for:

- The linker invokes **mklib** and waits for it to finish before finishing the link, so you will experience a one-time delay when an uncommonly-used library is built for the first time. Build times of 1-5 minutes have been observed. This depends on the power of the host (number of CPUs, etc).
- In a shared installation, where an installation of the compiler is shared among more than one user, it is possible that two users might cause the linker to rebuild the same library at the same time. The **mklib** program tries to minimize the race condition, but it is possible one build will corrupt the other. In a shared environment, all libraries which might be needed should be built at install time; see [Section 8.5.2.2](#) for instructions on invoking **mklib** directly to avoid this problem.
- The index library must exist, or the linker is unable to rebuild libraries automatically.
- The index library must be in a user-writable directory, or the library is not built. If the compiler installation must be installed read-only (a good practice for shared installation), any missing libraries must be built at installation time by invoking **mklib** directly.
- The **mklib** program is specific to a certain version of a certain library; you cannot use one compiler version's run-time support's **mklib** to build a different compiler version's run-time support library.

8.5.2.2 Invoking mklib Manually

You may need to invoke **mklib** directly in special circumstances:

- The compiler installation directory is read-only or shared.
- You want to build a variant of the run-time-support library that is not pre-configured in the index library **libc.a** or known to mklib. (e.g. a variant with source-level debugging turned on.)

8.5.2.2.1 Building Standard Libraries

You can invoke mklib directly to build any or all of the libraries indexed in the index library **libc.a**. The libraries are built with the standard options for that library; the library names and the appropriate standard option sets are known to mklib.

This is most easily done by changing the working directory to be the compiler run-time-support library directory 'lib' and invoking the **mklib** executable there:

```
mklib --pattern=rts6200.lib
```

8.5.2.2.2 Shared or Read-Only Library Directory

If the compiler tools are to be installed in shared or read-only directory, mklib cannot build the standard libraries at link time; the libraries must be built before the library directory is made shared or read-only.

At installation time, the installing user must build all of the libraries which will be used by any user. To build all possible libraries, change the working directory to be the compiler RTS library directory 'lib' and invoke the mklib executable there:

```
mklib --all
```

Some targets have many libraries, so this step can take a long time. To build a subset of the libraries, invoke mklib individually for each desired library.

8.5.2.2.3 Building Libraries With Custom Options

You can build a library with any extra custom options desired. This is useful for building a debugging version of the library, or with silicon exception workarounds enabled. The generated library is not a standard library, and must not be placed in the 'lib' directory. It should be placed in a directory local to the project which needs it. To build a debugging version of the library rts6200.lib, change the working directory to the 'lib' directory and run the command:

```
mklib --pattern=rts6200.lib --name=rts6200_debug.lib --install_to=$Project/Debug --
extra_options="-g"
```

8.5.2.2.4 The mklib Program Option Summary

Run the following command to see the full list of options. These are described in [Table 8-1](#).

```
mklib --help
```

Table 8-1. The mklib Program Options

Option	Effect
<code>--index=filename</code>	The index library (libc.a) for this release. Used to find a template library for custom builds, and to find the source files (rtssrc.zip). REQUIRED.
<code>--pattern=filename</code>	Pattern for building a library. If neither <code>--extra_options</code> nor <code>--options</code> are specified, the library will be the standard library with the standard options for that library. If either <code>--extra_options</code> or <code>--options</code> are specified, the library is a custom library with custom options. REQUIRED unless <code>--all</code> is used.
<code>--all</code>	Build all standard libraries at once.
<code>--install_to=directory</code>	The directory into which to write the library. For a standard library, this defaults to the same directory as the index library (libc.a). For a custom library, this option is REQUIRED.
<code>--compiler_bin_dir=directory</code>	The directory where the compiler executables are. When invoking mklib directly, the executables should be in the path, but if they are not, this option must be used to tell mklib where they are. This option is primarily for use when mklib is invoked by the linker.
<code>--name=filename</code>	File name for the library with no directory part. Only useful for custom libraries.
<code>--options='str'</code>	Options to use when building the library. The default options (see below) are <i>replaced</i> by this string. If this option is used, the library will be a custom library.
<code>--extra_options='str'</code>	Options to use when building the library. The default options (see below) are also used. If this option is used, the library will be a custom library.
<code>--list_libraries</code>	List the libraries this script is capable of building and exit. ordinary system-specific directory.
<code>--log=filename</code>	Save the build log as <i>filename</i> .
<code>--tmpdir=directory</code>	Use <i>directory</i> for scratch space instead of the ordinary system-specific directory.
<code>--gmake=filename</code>	Gmake-compatible program to invoke instead of "gmake"
<code>--parallel=N</code>	Compile <i>N</i> files at once ("gmake -j <i>N</i> ").
<code>--query=filename</code>	Does this script know how to build FILENAME?
<code>--help</code> or <code>--h</code>	Display this help.
<code>--quiet</code> or <code>--q</code>	Operate silently.
<code>--verbose</code> or <code>--v</code>	Extra information to debug this executable.

Examples:

To build all standard libraries and place them in the compiler's library directory:

```
mklib --all --index=$C_DIR/lib
```

To build one standard library and place it in the compiler's library directory:

```
mklib --pattern=rts6200.lib --index=$C_DIR/lib
```

To build a custom library that is just like rts6200.lib, but has symbolic debugging support enabled:

```
mklib --pattern=rts6200.lib --extra_options="-g" --index=$C_DIR/lib --install_to=$Project/Debug --
name=rts6200_debug.lib
```

8.5.3 Extending mklib

The **mklib** API is a uniform interface that allows Code Composer Studio to build libraries without needing to know exactly what underlying mechanism is used to build it. Each library vendor (e.g. the TI compiler) provides a library-specific copy of 'mklib' in the library directory that can be invoked, which understands a standardized set of options, and understands how to build the library. This allows the linker to automatically build application-compatible versions of any vendor's library without needing to register the library in advance, as long as the vendor supports mklib.

8.5.3.1 Underlying Mechanism

The underlying mechanism can be anything the vendor desires. For the compiler run-time-support libraries, mklib is just a wrapper which knows how to unpack Makefile from rtsrc.zip and invoke gmake with the appropriate options to build each library. If necessary, mklib can be bypassed and Makefile used directly, but this mode of operation is not supported by TI, and the you are responsible for any changes to Makefile. The format of the Makefile and the interface between mklib and the Makefile is subject to change without notice. The mklib program is the forward-compatible path.

8.5.3.2 Libraries From Other Vendors

Any vendor who wishes to distribute a library that can be rebuilt automatically by the linker must provide:

- An index library (like 'libc.a', but with a different name)
- A copy of mklib specific to that library
- A copy of the library source code (in whatever format is convenient)

These things must be placed together in one directory that is part of the linker's library search path (specified either in C6X_C_DIR or with the linker --search_path option).

If mklib needs extra information that is not possible to pass as command-line options to the compiler, the vendor will need to provide some other means of discovering the information (such as a configuration file written by a wizard run from inside CCS).

The vendor-supplied mklib must at least accept all of the options listed in [Table 8-1](#) without error, even if they do not do anything.

C++ Name Demangler

The C++ compiler implements function overloading, operator overloading, and type-safe linking by encoding a function's prototype and namespace in its link-level name. The process of encoding the prototype into the linkname is often referred to as name mangling. When you inspect mangled names, such as in assembly files, disassembler output, or compiler or linker diagnostics, it can be difficult to associate a mangled name with its corresponding name in the C++ source code. The C++ name demangler is a debugging aid that translates each mangled name it detects to its original name found in the C++ source code.

These topics tell you how to invoke and use the C++ name demangler. The C++ name demangler reads in input, looking for mangled names. All unmangled text is copied to output unaltered. All mangled names are demangled before being copied to output.

Topic	Page
9.1 Invoking the C++ Name Demangler	269
9.2 C++ Name Demangler Options	269
9.3 Sample Usage of the C++ Name Demangler	270

9.1 Invoking the C++ Name Demangler

The syntax for invoking the C++ name demangler is:

```
dem6x [options] [filenames]
```

dem6x	Command that invokes the C++ name demangler.
<i>options</i>	Options affect how the name demangler behaves. Options can appear anywhere on the command line. (Options are discussed in Section 9.2.)
<i>filenames</i>	Text input files, such as the assembly file output by the compiler, the assembler listing file, the disassembly file, and the linker map file. If no filenames are specified on the command line, dem6x uses standard input.

By default, the C++ name demangler outputs to standard output. You can use the `-o` file option if you want to output to a file.

9.2 C++ Name Demangler Options

The following options apply only to the C++ name demangler:

--abi=eabi	Demangles EABI identifiers
-h	Prints a help screen that provides an online summary of the C++ name demangler options
-o file	Outputs to the given file rather than to standard out
-u	Specifies that external names do not have a C++ prefix
-v	Enables verbose mode (outputs a banner)

9.3 Sample Usage of the C++ Name Demangler

The examples in this section illustrate the demangling process. [Example 9-1](#) shows a sample C++ program. [Example 9-2](#) shows the resulting assembly that is output by the compiler. In this example, the linknames of all the functions are mangled; that is, their signature information is encoded into their names.

Example 9-1. C++ Code for `calories_in_a_banana`

```
class banana {
public:
    int calories(void);

    banana();

    ~banana();
};

int calories_in_a_banana(void)
{
    banana x;

    return x.calories();
}
```

Example 9-2. Resulting Assembly for `calories_in_a_banana`

```
_calories_in_a_banana__Fv:
; ** -----*
        CALL    .S1    ___ct__6bananaFv ; |10|
        STW    .D2T2  B3,*SP--(16)    ; |9|
        MVKL   .S2    RL0,B3          ; |10|
        MVKH   .S2    RL0,B3          ; |10|
        ADD    .S1X   8,SP,A4         ; |10|
        NOP
RL0:    ; CALL OCCURS                ; |10|
        CALL    .S1    _calories__6bananaFv ; |12|
        MVKL   .S2    RL1,B3          ; |12|
        ADD    .S1X   8,SP,A4         ; |12|
        MVKH   .S2    RL1,B3          ; |12|
        NOP
RL1:    ; CALL OCCURS                ; |12|
        CALL    .S1    ___dt__6bananaFv ; |13|
        STW    .D2T1  A4,*+SP(4)      ; |12|
        ADD    .S1X   8,SP,A4         ; |13|
        MVKL   .S2    RL2,B3          ; |13|
        MVK    .S2    0x2,B4          ; |13|
        MVKH   .S2    RL2,B3          ; |13|
RL2:    ; CALL OCCURS                ; |13|
        LDW    .D2T1  *+SP(4),A4      ; |12|
        LDW    .D2T2  *++SP(16),B3    ; |13|
        NOP
        RET    .S2    B3              ; |13|
        NOP
        ; BRANCH OCCURS                ; |13|
```

Executing the C++ name demangler demangles all names that it believes to be mangled. Enter:

```
dem6x calories_in_a_banana.asm
```

The result is shown in [Example 9-3](#). The linknames in [Example 9-2](#) `__ct__6bananaFv`, `__calories__6bananaFv`, and `__dt__6bananaFv` are demangled.

Example 9-3. Result After Running the C++ Name Demangler

```

calories_in_a_banana():
; ** -----*
    CALL    .S1 banana::banana()      ; |10|
    STW     .D2T2 B3,*SP--(16)        ; |9|
    MVKL    .S2 RL0,B3                 ; |10|
    MVKH    .S2 RL0,B3                 ; |10|
    ADD     .S1X 8,SP,A4                ; |10|
    NOP 1
RL0:      ; CALL OCCURS                 ; |10|
    CALL    .S1 banana::calories()    ; |12|
    MVKL    .S2 RL1,B3                 ; |12|
    ADD     .S1X 8,SP,A4                ; |12|
    MVKH    .S2 RL1,B3                 ; |12|
    NOP 2
RL1:      ; CALL OCCURS                 ; |12|
    CALL    .S1 banana::~~banana()    ; |13|
    STW     .D2T1 A4,*+SP(4)           ; |12|
    ADD     .S1X 8,SP,A4                ; |13|
    MVKL    .S2 RL2,B3                 ; |13|
    MVK     .S2 0x2,B4                 ; |13|
    MVKH    .S2 RL2,B3                 ; |13|
RL2:      ; CALL OCCURS                 ; |13|
    LDW     .D2T1 *+SP(4),A4           ; |12|
    LDW     .D2T2 *++SP(16),B3        ; |13|
    NOP 4
    RET     .S2 B3                      ; |13|
    NOP 5
    ; BRANCH OCCURS                     ; |13|

```

Glossary

absolute lister— A debugging tool that allows you to create assembler listings that contain absolute addresses.

assignment statement— A statement that initializes a variable with a value.

autoinitialization— The process of initializing global C variables (contained in the .cinit section) before program execution begins.

autoinitialization at run time— An autoinitialization method used by the linker when linking C code. The linker uses this method when you invoke it with the `--rom_model` link option. The linker loads the .cinit section of data tables into memory, and variables are initialized at run time.

alias disambiguation— A technique that determines when two pointer expressions cannot point to the same location, allowing the compiler to freely optimize such expressions.

aliasing— The ability for a single object to be accessed in more than one way, such as when two pointers point to a single object. It can disrupt optimization, because any indirect reference could refer to any other object.

allocation— A process in which the linker calculates the final memory addresses of output sections.

ANSI— American National Standards Institute; an organization that establishes standards voluntarily followed by industries.

archive library— A collection of individual files grouped into a single file by the archiver.

archiver— A software program that collects several individual files into a single file called an archive library. With the archiver, you can add, delete, extract, or replace members of the archive library.

assembler— A software program that creates a machine-language program from a source file that contains assembly language instructions, directives, and macro definitions. The assembler substitutes absolute operation codes for symbolic operation codes and absolute or relocatable addresses for symbolic addresses.

assignment statement— A statement that initializes a variable with a value.

autoinitialization— The process of initializing global C variables (contained in the .cinit section) before program execution begins.

autoinitialization at run time— An autoinitialization method used by the linker when linking C code. The linker uses this method when you invoke it with the `--rom_model` link option. The linker loads the .cinit section of data tables into memory, and variables are initialized at run time.

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 statements that are grouped together within braces and treated as an entity.

.bss section— One of the default object file sections. You use the assembler `.bss` directive to reserve a specified amount of space in the memory map that you can use later for storing data. The `.bss` section is uninitialized.

byte— Per ANSI/ISO C, the smallest addressable unit that can hold a character.

- C/C++ compiler**— A software program that translates C source statements into assembly language source statements.
- code generator**— A compiler tool that takes the file produced by the parser or the optimizer and produces an assembly language source file.
- COFF**— Common object file format; a system of object files configured according to a standard developed by AT&T. These files are relocatable in memory space.
- 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 documents or improves readability of a source file. Comments are not compiled, assembled, or linked; they have no effect on the object file.
- compiler program**— A utility that lets you compile, assemble, and optionally link in one step. The compiler runs one or more source modules through the compiler (including the parser, optimizer, and code generator), the assembler, and the linker.
- configured memory**— Memory that the linker has specified for allocation.
- constant**— A type whose value cannot change.
- cross-reference listing**— An output file created by the assembler that lists the symbols that were defined, what line they were defined on, which lines referenced them, and their final values.
- .data section**— One of the default object file 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.
- direct call**— A function call where one function calls another using the function's name.
- 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).
- disambiguation**— See *alias disambiguation*
- dynamic memory allocation**— A technique used by several functions (such as malloc, calloc, and realloc) to dynamically allocate memory for variables at run time. This is accomplished by defining a large memory pool (heap) and using the functions to allocate memory from the heap.
- ELF**— Executable and Linkable Format; a system of object files configured according to the System V Application Binary Interface specification.
- emulator**— A hardware development system that duplicates the TMS320C6000 operation.
- entry point**— A point in target memory where execution starts.
- environment variable**— A system symbol that you define and assign to a string. Environmental variables are often included in Windows batch files or UNIX shell scripts such as .cshrc or .profile.
- epilog**— The portion of code in a function that restores the stack and returns. See also *pipelined-loop epilog*.
- executable object file**— A linked, executable object file that is downloaded and executed on a target 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 defined or declared in a different program module.

- file-level optimization**— A level of optimization where the compiler uses the information that it has about the entire file to optimize your code (as opposed to program-level optimization, where the compiler uses information that it has about the entire program to optimize your code).
- function inlining**— The process of inserting code for a function at the point of call. This saves the overhead of a function call and allows the optimizer to optimize the function in the context of the surrounding code.
- global symbol**— A symbol that is either defined in the current module and accessed in another, or accessed in the current module but defined in another.
- 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.
- indirect call**— A function call where one function calls another function by giving the address of the called function.
- initialization at load time**— An autoinitialization method used by the linker when linking C/C++ code. The linker uses this method when you invoke it with the `--ram_model` link option. This method initializes variables at load time instead of run time.
- initialized section**— A section from an object file that will be linked into an executable object file.
- input section**— A section from an object file that will be linked into an executable object file.
- integrated preprocessor**— A C/C++ preprocessor that is merged with the parser, allowing for faster compilation. Stand-alone preprocessing or preprocessed listing is also available.
- interlist feature**— A feature that inserts as comments your original C/C++ source statements into the assembly language output from the assembler. The C/C++ statements are inserted next to the equivalent assembly instructions.
- intrinsics**— Operators that are used like functions and produce assembly language code that would otherwise be inexpressible in C, or would take greater time and effort to code.
- ISO**— International Organization for Standardization; a worldwide federation of national standards bodies, which establishes international standards voluntarily followed by industries.
- kernel**— The body of a software-pipelined loop between the pipelined-loop prolog and the pipelined-loop epilog.
- K&R C**— Kernighan and Ritchie C, the de facto standard as defined in the first edition of *The C Programming Language* (K&R). Most K&R C programs written for earlier, non-ISO C compilers should correctly compile and run without modification.
- label**— A symbol that begins in column 1 of an assembler source statement and corresponds to the address of that statement. A label is the only assembler statement that can begin in column 1.
- linker**— A software program that combines object files to form an executable object file that can be allocated into 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 section program counter (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 places an executable object file into system memory.
- loop unrolling**— An optimization that expands small loops so that each iteration of the loop appears in your code. Although loop unrolling increases code size, it can improve the performance of your code.

- 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 process of inserting source statements into your code in place of a macro call.
- map file**— An output file, created by the linker, that shows the memory configuration, section composition, section allocation, symbol definitions and the addresses at which the symbols were defined for your program.
- memory map**— A map of target system memory space that is partitioned into functional blocks.
- name mangling**— A compiler-specific feature that encodes a function name with information regarding the function's arguments return types.
- object file**— An assembled or linked file that contains machine-language object code.
- object library**— An archive library made up of individual object files.
- operand**— An argument of an assembly language instruction, assembler directive, or macro directive that supplies information to the operation performed by the instruction or directive.
- optimizer**— A software tool that improves the execution speed and reduces the size of C programs. See also *assembly optimizer*.
- options**— Command-line parameters that allow you to request additional or specific functions when you invoke a software tool.
- output section**— A final, allocated section in a linked, executable module.
- parser**— A software tool that reads the source file, performs preprocessing functions, checks the syntax, and produces an intermediate file used as input for the optimizer or code generator.
- partitioning**— The process of assigning a data path to each instruction.
- pipelining**— A technique where a second instruction begins executing before the first instruction has been completed. You can have several instructions in the pipeline, each at a different processing stage.
- pop**— An operation that retrieves a data object from a stack.
- pragma**— A preprocessor directive that provides directions to the compiler about how to treat a particular statement.
- preprocessor**— A software tool that interprets macro definitions, expands macros, interprets header files, interprets conditional compilation, and acts upon preprocessor directives.
- program-level optimization**— An aggressive level of optimization where all of the source files are compiled into one intermediate file. Because the compiler can see the entire program, several optimizations are performed with program-level optimization that are rarely applied during file-level optimization.
- prolog**— The portion of code in a function that sets up the stack. See also *pipelined-loop prolog*.
- push**— An operation that places a data object on a stack for temporary storage.
- quiet run**— An option that suppresses the normal banner and the progress information.
- 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-time environment**— The run time parameters in which your program must function. These parameters are defined by the memory and register conventions, stack organization, function call conventions, and system initialization.
- run-time-support functions**— Standard ISO functions that perform tasks that are not part of the C language (such as memory allocation, string conversion, and string searches).
- run-time-support library**— A library file, `rts.src`, that contains the source for the run time-support functions.
- section**— A relocatable block of code or data that ultimately will be contiguous with other sections in the memory map.
- sign extend**— A process that fills 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/C++ code or assembly language code that is compiled or assembled to form an object file.
- stand-alone preprocessor**— A software tool that expands macros, `#include` files, and conditional compilation as an independent program. It also performs integrated preprocessing, which includes parsing of instructions.
- static variable**— A variable 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; their previous value is resumed when the function or program is reentered.
- storage class**— An entry in the symbol table that indicates how to access a symbol.
- 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 ultimately will occupy continuous space in the 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 that can be used by a debugging tool such as a simulator or an emulator.
- target system**— The system on which the object code you have developed is executed.
- .text section**— One of the default object file sections. The `.text` section is initialized and contains executable code. You can use the `.text` directive to assemble code into the `.text` section.
- trigraph sequence**— A 3-character sequence that has a meaning (as defined by the ISO 646-1983 Invariant Code Set). These characters cannot be represented in the C character set and are expanded to one character. For example, the trigraph `??'` is expanded to `^`.
- trip count**— The number of times that a loop executes before it terminates.
- unconfigured memory**— Memory that is not defined as part of the memory map and cannot be loaded with code or data.
- uninitialized section**— A object file section that reserves space in the memory map but that has no actual contents. These sections are built with the `.bss` and `.usect` directives.
- unsigned value**— A value that is treated as a nonnegative number, regardless of its actual sign.
- variable**— A symbol representing a quantity that can assume any of a set of values.

word— A 32-bit addressable location in target memory

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46C and to discontinue any product or service per JESD48B. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as "components") are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

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

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

TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license 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 significant portions of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

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

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI's goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or "enhanced plastic" are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have **not** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components which meet ISO/TS16949 requirements, mainly for automotive use. Components which have not been so designated are neither designed nor intended for automotive use; and TI will not be responsible for any failure of such components to meet such requirements.

Products

Audio	www.ti.com/audio
Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DLP® Products	www.dlp.com
DSP	dsp.ti.com
Clocks and Timers	www.ti.com/clocks
Interface	interface.ti.com
Logic	logic.ti.com
Power Mgmt	power.ti.com
Microcontrollers	microcontroller.ti.com
RFID	www.ti-rfid.com
OMAP Mobile Processors	www.ti.com/omap
Wireless Connectivity	www.ti.com/wirelessconnectivity

Applications

Automotive and Transportation	www.ti.com/automotive
Communications and Telecom	www.ti.com/communications
Computers and Peripherals	www.ti.com/computers
Consumer Electronics	www.ti.com/consumer-apps
Energy and Lighting	www.ti.com/energy
Industrial	www.ti.com/industrial
Medical	www.ti.com/medical
Security	www.ti.com/security
Space, Avionics and Defense	www.ti.com/space-avionics-defense
Video and Imaging	www.ti.com/video

TI E2E Community e2e.ti.com