

SC100 C Compiler User's Manual



MNSC100CC/D
Rev. 1.6, 06/2000

SC100 C Compiler

User's Manual



microelectronics group

Lucent Technologies
Bell Labs Innovations



This document contains information on a new product. Specifications and information herein are subject to change without notice.

© Copyright Lucent Technologies Inc., 2000. All rights reserved.

© Copyright Motorola Inc., 2000. All rights reserved.

LICENSOR is defined as either Motorola, Inc. or Lucent Technologies Inc., whichever company distributed this document to LICENSEE. LICENSOR reserves the right to make changes without further notice to any products included and covered hereby. LICENSOR makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does LICENSOR assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation incidental, consequential, reliance, exemplary, or any other similar such damages, by way of illustration but not limitation, such as, loss of profits and loss of business opportunity. "Typical" parameters which may be provided in LICENSOR data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. LICENSOR does not convey any license under its patent rights nor the rights of others. LICENSOR products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support life, or for any other application in which the failure of the LICENSOR product could create a situation where personal injury or death may occur. Should Buyer purchase or use LICENSOR products for any such unintended or unauthorized application, Buyer shall indemnify and hold LICENSOR and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that LICENSOR was negligent regarding the design or manufacture of the part.

Motorola and  are registered trademarks of Motorola, Inc. Motorola, Inc. is an Equal Opportunity/Affirmative Action Employer. Lucent, Lucent Technologies, and (the Lucent Technologies insignia) are trademarks of Lucent Technologies Inc. Lucent Technologies Inc. is an Equal Opportunity/Affirmative Action Employer.

StarCore is a trademark of Motorola, Inc. It is used by Lucent Technologies with the authorization of Motorola.

All other tradenames, trademarks, and registered trademarks are the property of their respective owners.

Table of Contents

About This Book

Audience	xix
Organization	xix
Suggested Reading	xx
Conventions	xx
Acronyms and Abbreviations	xx

Chapter 1 Introduction

1.1	Overview of the SC100 C Compiler	1-1
1.2	The Global Optimization Approach	1-2
1.3	Compiling Applications	1-3
1.3.1	The Compiler Shell Program	1-3
1.3.2	Stages in the C Compilation Process	1-4

Chapter 2 Getting Started

2.1	A Quick Start	2-1
-----	---------------------	-----

Chapter 3 Using the SC100 C Compiler

3.1	The Shell Program	3-1
3.1.1	The C Compilation Process	3-2
3.1.2	Global Optimization	3-4
3.1.3	File Types and Extensions	3-7
3.1.4	Environment Variables	3-9
3.2	Invoking the Shell	3-10
3.2.1	Command Line Syntax	3-10
3.2.1.1	Using a Command File	3-10
3.3	Shell Control Options	3-11
3.3.1	Controlling the Behavior of the Shell	3-15
3.3.1.1	Controlling where the shell stops processing	3-15
3.3.1.2	Specifying a shell command file	3-16
3.3.1.3	Displaying the shell Help page	3-16
3.3.2	Specifying Preprocessing Options	3-17
3.3.2.1	Changing preprocessed output	3-17
3.3.2.2	Defining and undefining preprocessor macros	3-18
3.3.2.3	Adding directories to the #include file path	3-18

3.3.3	Overriding Input File Extensions	3-19
3.3.4	Output Filename and Location Options	3-20
3.3.5	Specifying C Language Options	3-21
3.3.5.1	Defining the language version	3-21
3.3.5.2	Adding debugging information to files	3-21
3.3.5.3	Changing the default char sign setting	3-21
3.3.6	Passing Options Through to Specific Tools	3-22
3.3.7	Setting the Options for Listings and Messages	3-23
3.3.7.1	Generating listing files	3-23
3.3.7.2	Controlling the type of information displayed	3-24
3.3.7.3	Suppressing warnings	3-24
3.3.7.4	Reporting all remarks and warnings	3-24
3.3.8	Specifying the Hardware Model and Configuration	3-25
3.3.8.1	Defining specific configuration and startup files	3-25
3.3.8.2	Defining memory mode	3-25
3.3.8.3	Copying initialized variables from ROM	3-26
3.3.8.4	Specifying the silicon revision	3-26
3.3.8.5	Defining the profiling information file	3-26
3.3.8.6	Specifying big-endian mode	3-26
3.4	Language Features	3-27
3.4.1	C Language Dialects	3-27
3.4.1.1	Standard Extensions	3-27
3.4.1.2	K&R/PCC mode	3-32
3.4.2	Types and Sizes	3-37
3.4.2.1	Characters	3-38
3.4.2.2	Integers	3-39
3.4.2.3	Floating point	3-40
3.4.2.4	Fractional representation	3-41
3.4.2.5	Pointers	3-41
3.4.2.6	Bit-fields	3-41
3.4.3	Fractional and Integer Arithmetic	3-43
3.4.4	Intrinsic Functions	3-46
3.4.4.1	Data types for intrinsic functions	3-46
3.4.4.2	Intrinsic function categories	3-47
3.4.4.3	Intrinsic functions examples	3-52
3.4.5	Pragmas	3-54
3.4.5.1	Syntax	3-54
3.4.5.2	Placement	3-54
3.4.5.3	Pragmas which apply to functions	3-56
3.4.5.4	Pragmas which apply to statements	3-58
3.4.5.5	Pragmas which apply to variables	3-61
3.4.6	Predefined Macros	3-63

Chapter 4 Interfacing C and Assembly Code

4.1	Inlining a Single Assembly Instruction	4-1
4.2	Inlining a Sequence of Assembly Instructions	4-2
4.2.1	Guidelines for Inlining Assembly Code Sequences	4-2
4.2.2	Defining an Inlined Sequence of Assembly Instructions	4-3
4.3	Calling an Assembly Function in a Separate File	4-7
4.3.1	Writing the Assembly Code	4-7
4.3.2	Calling the Assembly Function	4-8
4.3.3	Integrating the C and Assembly Files	4-8
4.4	Including Offset Labels in the Output File	4-9

Chapter 5 Optimization Techniques and Hints

5.1	Optimizer Overview	5-2
5.1.1	Code Transformations	5-2
5.1.2	Basic Blocks	5-2
5.1.3	Linear and Parallelized Code	5-3
5.1.4	Optimization Levels and Options	5-4
5.2	Using the Optimizer	5-6
5.2.1	Invoking the Optimizer	5-6
5.2.2	Optimizing for Space	5-7
5.2.3	Optimizing for Speed	5-7
5.2.4	Using Global Optimization	5-7
5.3	Optimization Types and Functions	5-8
5.3.1	Dependencies and Parallelization	5-8
5.3.2	High-Level Optimizations	5-9
5.3.2.1	Strength reduction (loop transformations)	5-10
5.3.2.2	Function inlining	5-16
5.3.2.3	Common subexpression elimination	5-17
5.3.2.4	Loop invariant code	5-17
5.3.2.5	Constant folding and propagation	5-17
5.3.2.6	Jump-to-jump elimination	5-18
5.3.2.7	Dead code elimination	5-18
5.3.2.8	Dead storage/assignment elimination	5-18
5.3.3	Low-Level Optimizations	5-19
5.3.3.1	Instruction scheduling	5-21
5.3.3.2	Low-level software pipelining	5-22
5.3.3.3	Conditional execution and predication	5-25
5.3.3.4	Speculative execution	5-26
5.3.3.5	Post-increment detection	5-27
5.3.3.6	Low-level peephole optimization	5-28
5.3.4	Time Optimizations	5-29
5.3.4.1	Time optimization levels and criteria	5-29
5.3.4.2	Prefix grouping	5-30
5.3.5	Space Optimizations	5-30
5.3.6	Global Optimizations	5-30

5.4	Guidelines for Using the Optimizer	5-31
5.4.1	Partial Summation Techniques	5-32
5.4.2	Multisample Techniques	5-35
5.4.2.1	Multisample implementation issues	5-37
5.4.2.2	Implementation example	5-39
5.4.3	General Hints	5-43
5.4.3.1	Software pipelining	5-43
5.4.3.2	Passing and returning large structs	5-43
5.4.3.3	Arithmetic operations	5-43
5.4.3.4	Local variables	5-43
5.4.3.5	Resource limitations	5-44
5.5	Optimizer Assumptions	5-44

Chapter 6 Runtime Environment

6.1	Startup Code	6-1
6.1.1	Bare Board Startup Code	6-2
6.1.2	C Environment Startup Code	6-3
6.1.2.1	C environment initialization code	6-3
6.1.2.2	Initialization of variables	6-3
6.1.2.3	C environment finalization code	6-3
6.1.2.4	Low level I/O services	6-3
6.1.3	Configuring Your Startup Code	6-4
6.2	Memory Models	6-5
6.2.1	Linker Command Files	6-5
6.3	Memory Layout and Configuration	6-6
6.3.1	Stack and Heap Configuration	6-7
6.3.1.1	Runtime stack	6-7
6.3.1.2	Dynamic memory allocation (heap)	6-8
6.3.2	Static Data Allocation	6-8
6.3.3	Configuring the Memory Map	6-9
6.3.3.1	Memory map configuration example	6-9
6.3.4	Machine Configuration File	6-10
6.3.4.1	Defining the memory configuration	6-10
6.3.5	Application Configuration File	6-12
6.3.5.1	File structure and syntax	6-13
6.3.5.2	Schedule section	6-13
6.3.5.3	Binding section	6-15
6.3.5.4	Overlay section	6-16
6.4	Calling Conventions	6-18
6.4.1	Stack Pointer	6-18
6.4.2	Stack-Based Calling Convention	6-18
6.4.3	Optimized Calling Sequences	6-20
6.4.4	Stack Frame Layout	6-21
6.4.5	Interrupt Handlers	6-22
6.4.6	Frame Pointer and Argument Pointer	6-22

6.4.7	Hardware Loops	6-22
6.4.8	Operating Modes.	6-23

Chapter 7 Runtime Libraries

7.1	Character Typing and Conversion (ctype.h).	7-2
7.1.1	Testing Functions	7-2
7.1.2	Conversion Functions	7-2
7.2	Floating Point Characteristics (float.h).	7-3
7.2.1	Floating Point Library Interface (fltmath.h)	7-4
7.2.1.1	Round_Mode	7-4
7.2.1.2	FLUSH_TO_ZERO	7-5
7.2.1.3	IEEE_Exceptions	7-5
7.2.1.4	EnableFPEExceptions	7-6
7.3	Integer Characteristics (limits.h).	7-7
7.4	Locales (locale.h)	7-7
7.5	Floating Point Math (math.h)	7-8
7.5.1	Trigonometric Functions.	7-8
7.5.2	Hyperbolic Functions	7-8
7.5.3	Exponential and Logarithmic Functions.	7-9
7.5.4	Power Functions	7-9
7.5.5	Other Functions.	7-9
7.6	Nonlocal Jumps (setjmp.h)	7-10
7.7	Signal Handling (signal.h)	7-10
7.8	Variable Arguments (stdarg.h)	7-10
7.9	Standard Definitions (stddef.h).	7-11
7.10	I/O Library (stdio.h)	7-11
7.10.1	Input Functions	7-11
7.10.2	Stream Functions	7-12
7.10.3	Output Functions.	7-13
7.10.4	Miscellaneous I/O Functions	7-13
7.11	General Utilities (stdlib.h)	7-14
7.11.1	Memory Allocation Functions	7-14
7.11.2	Integer Arithmetic Functions	7-14
7.11.3	String Conversion Functions	7-15
7.11.4	Searching and Sorting Functions	7-15
7.11.5	Pseudo Random Number Generation Functions.	7-15
7.11.6	Environment Functions.	7-16
7.11.7	Multibyte Character Functions	7-16
7.12	String Functions (string.h)	7-16
7.12.1	Copying Functions	7-17
7.12.2	Concatenation Functions.	7-17
7.12.3	Comparison Functions	7-17
7.12.4	Search Functions.	7-18
7.12.5	Other Functions.	7-18

7.13	Time Functions (time.h)	7-19
7.13.1	Time Constant.....	7-19
7.13.2	Process Time.....	7-19
7.14	Built-in Intrinsic Functions (prototype.h).....	7-20

Appendix A

Migrating from Other Environments

A.1	Code Migration Overview	A-1
A.2	Migrating Code Developed for DSP56600.....	A-2
A.3	Migrating Code Developed for TI6xx	A-6

List of Tables

3-1	File Types and Extensions	3-7
3-2	Shell Options Summary	3-12
3-3	Data Types and Sizes	3-37
3-4	Interpretation of 16-bit Data Values	3-44
3-5	Interpretation of 40-bit Data Values	3-44
3-6	Intrinsic Functions	3-48
3-7	Pragmas.	3-55
3-8	Predefined Macros	3-63
5-1	Optimization Options Summary	5-5
5-2	Summary of High-Level Optimizations	5-9
5-3	Summary of Low-Level Optimizations	5-19
5-4	Time Optimization Criteria.	5-29
6-1	Status Register Default Settings	6-2
6-2	Small Memory Model Default Values	6-7
6-3	Big Memory Model Default Values	6-7
6-4	Register Usage in the Stack-based Calling Convention	6-19
7-1	Supported ISO Libraries	7-1
7-2	Supported Non-ISO Libraries.	7-1
7-3	Testing Functions	7-2
7-4	Conversion Functions	7-2
7-5	Contents of File float.h	7-3
7-7	Locale Functions.	7-7
7-6	Contents of File limits.h	7-7
7-8	Trigonometric Functions.	7-8
7-9	Hyperbolic Functions	7-8
7-10	Exponential and Logarithmic Functions.	7-9
7-11	Power Functions	7-9
7-12	Other Functions.	7-9
7-13	Nonlocal Jumps.	7-10
7-14	Signal Handling (signal.h)	7-10
7-15	Variable Arguments (stdarg.h)	7-10
7-16	Standard Definitions (stddef.h).	7-11
7-17	Input Functions	7-11
7-18	Stream Functions	7-12

7-19	Output Functions.	7-13
7-20	Miscellaneous I/O Functions	7-13
7-21	Memory Allocation Functions	7-14
7-22	Integer Arithmetic Functions	7-14
7-23	String Conversion Functions	7-15
7-24	Searching and Sorting Functions	7-15
7-25	Pseudo Random Number Generation Functions.	7-15
7-26	Environment Functions.	7-16
7-27	Multibyte Character Functions	7-16
7-28	Copying Functions	7-17
7-29	Concatenation Functions.	7-17
7-30	Comparison Functions	7-17
7-31	Search Functions.	7-18
7-32	Other Functions.	7-18
7-33	Time Functions	7-19
7-34	Time Constant.	7-19
7-35	Built-in Intrinsic Functions.	7-20
A-1	DSP56600 Integer Data Type Differences	A-2
A-2	DSP56600 Fractional Data Type Differences	A-2
A-3	DSP56600 Pointer Size Differences.	A-3
A-4	DSP56600 Fractional Arithmetic Differences	A-3
A-5	DSP56600 Intrinsic Function Differences	A-4
A-6	DSP56600 Storage Specifiers.	A-5
A-7	DSP56600 Miscellaneous Differences	A-5

List of Figures

1-1	The SC100 C Compilation Process	1-5
3-1	The C Compilation Process	3-3
3-2	Non-Global Optimization	3-5
3-3	Global Optimization	3-6
3-4	File Extensions in the Shell Cycle	3-8
3-5	Characters—Memory Layout	3-38
3-6	Characters—Dn Register Layout	3-38
3-7	Characters—Rn Register Layout	3-38
3-8	Integers—Memory Layout	3-39
3-9	Integers—Alignment	3-39
3-10	Integers—Dn Register Layout	3-40
3-11	Integers—Rn Register Layout	3-40
3-12	Fractional Integers—Dn Register Layout	3-41
3-13	Extended Precision Fractional—Dn Register Layout	3-46
5-1	Linear and Parallelized Code	5-3
5-2	Square Loop	5-14
5-3	Triangular Loop	5-15
5-4	Sequence of Low-Level Transformation Optimizations	5-20
5-5	Single Sample and Multisample Kernels	5-35
5-6	Single ALU Operand and Memory Bandwidth	5-36
5-7	Quad ALU Operand and Memory Bandwidth	5-36
5-8	Options for Increasing Operand Bandwidth	5-37
5-9	Number of Samples and ALUs for Implementing DSP Algorithms	5-37
5-10	Quad Coefficient Loading from Memory	5-38
5-11	Misalignment when Loading Quad Operands	5-38
5-12	Quad ALU, Quad Sample FIR Filter Data Flow	5-39
5-13	FIR Filter Equations for Four Samples	5-40
5-14	Generic Kernel For Quad ALU FIR	5-40
6-1	SC100 Default Memory Layout	6-6
6-2	Stack Frame Layout	6-21

List of Examples

2-1	Sample source file - hello.c.	2-1
3-1	Invoking the shell	3-10
3-2	Defining a shell command file	3-16
3-3	Contents of a shell command file	3-16
3-4	Shell Help page (extract)	3-16
3-5	Overriding file extensions.	3-19
3-6	Specifying output files	3-20
3-7	Passing multiple options to the same tool.	3-22
3-8	Defining a profiling information file	3-26
3-9	Defining a predicate name	3-27
3-10	Deleting a predicate	3-28
3-11	Declaring an asm function	3-29
3-12	Returning the alignment requirement	3-30
3-13	Out of range warning	3-30
3-14	External entities in other scopes	3-30
3-15	Pointers to incomplete arrays	3-31
3-16	Prototyped parameter list	3-33
3-17	Omitting the declarator list	3-34
3-18	Keywords in type specifiers	3-34
3-19	Bit-field alignment to long word (1).	3-42
3-20	Bit-field alignment to character	3-42
3-21	Bit-field alignment to long word (2).	3-42
3-22	Bit-field offset.	3-42
3-23	Fractional arithmetic examples.	3-43
3-24	Integer arithmetic examples	3-43
3-25	Integer arithmetic computation.	3-45
3-26	Fractional arithmetic computation	3-45
3-27	Intrinsic functions	3-52
3-28	Intrinsic functions using extended precision	3-53
3-29	#pragma noline	3-56
3-30	#pragma save_ctxt	3-56
3-31	#pragma external.	3-57

3-32	#pragma interrupt	3-58
3-33	#pragma profile with constant value.	3-58
3-34	#pragma profile with frequency ratio	3-59
3-35	#pragma loop count	3-60
3-36	#pragma align	3-62
4-1	Inlining a single assembly instruction	4-2
4-2	Inlining syntax	4-4
4-3	Simple inlined assembly function.	4-4
4-4	Inlined assembly function with labels and hardware loops	4-5
4-5	Referencing global variables in an inlined assembly function	4-6
4-6	Assembly function in a separate file.	4-7
4-7	C code calling assembly function	4-8
4-8	Integrating C and assembly files.	4-8
4-9	Data structure shared between C and assembly	4-9
4-10	Specifying the output of offset information	4-9
4-11	Data structure offsets in the assembly output file.	4-9
4-12	Using symbolic offsets in assembly code.	4-10
5-1	Instruction transformation.	5-2
5-2	Invoking the optimizer with default settings	5-6
5-3	Invoking the optimizer for high-level optimizations only	5-6
5-4	Invoking the optimizer for time optimization.	5-6
5-5	Invoking the optimizer in global mode.	5-7
5-6	Simple instruction dependency.	5-8
5-7	Algorithm instruction dependency	5-9
5-8	Loop transformation - simple loop	5-10
5-9	Loop transformation - dynamic loop	5-11
5-10	Loop transformation - multi-step loop	5-12
5-11	Loop transformation - composed variables	5-13
5-12	Loop transformation - square loop	5-14
5-13	Loop transformation - triangular loop	5-15
5-14	Function inlining.	5-16
5-15	Common subexpression elimination.	5-17
5-16	Loop invariant code motion	5-17
5-17	Constant folding and propagation.	5-17
5-18	Jump-to-jump elimination.	5-18
5-19	Dead code elimination	5-18
5-20	Dead storage/assignment elimination	5-18

5-21	Instruction scheduling	5-21
5-22	Filling delay slots	5-21
5-23	Avoiding pipeline restrictions.	5-22
5-24	Software pipelining - complex FIR	5-23
5-25	Software pipelining - vector multiplication by a constant	5-24
5-26	Conditional execution and predication	5-25
5-27	Speculative execution	5-26
5-28	Post-increment detection.	5-27
5-29	Low-level peephole optimization	5-28
5-30	Combined pipelining and peephole optimizations	5-28
5-31	Simple and complex array accesses	5-31
5-32	MAC usage limited by dependency in loop	5-32
5-33	Partial summation for dual MAC usage	5-33
5-34	Alignment restrictions in algorithms	5-34
5-35	Single instruction quad ALU generic filter kernel	5-41
5-36	FIR_A4S4 quad ALU, quad sample C simulation code.	5-41
5-37	Avoiding software pipelining in source code	5-43
6-1	Creating a new startup file	6-4
6-2	Assembling the modified startup file	6-4
6-3	Using the modified startup file	6-4
6-4	Big and small memory models	6-5
6-5	Small memory mode instruction.	6-5
6-6	Allocating large arrays from the heap.	6-8
6-8	Modified memory configuration in the linker command file	6-9
6-7	Modified memory map configuration.	6-9
6-9	Defining a data memory space	6-11
6-10	Defining a program memory space.	6-11
6-11	Defining multiple memory spaces	6-12
6-12	Defining additional entry points for an application	6-15
6-13	Placing a variable at an absolute location.	6-16
6-14	Defining global variable overlays.	6-17
6-15	Function call and allocation of parameters.	6-19
7-1	Changing the round mode.	7-4
7-2	Disabling flushing to zero.	7-5
7-3	Using the exception status word	7-5
7-4	Setting a signal for exceptions	7-6
7-5	Timing an application	7-19

A-1 Migrating code from other environments A-2

About This Book

This manual describes the C compiler for the SC100 family of digital signal processor (DSP) cores, and provides detailed guidelines for its use.

Audience

This manual is intended for systems software developers, applications developers, system hardware developers, and microprocessor designers.

Organization

This manual is organized into seven chapters and one appendix, as follows:

- [Chapter 1, “Introduction,”](#) provides an overview of the SC100 C compiler and outlines the SC100 C compilation process.
- [Chapter 2, “Getting Started,”](#) provides the essential information and instructions that enable you to start using the SC100 C compiler.
- [Chapter 3, “Using the SC100 C Compiler,”](#) explains how to use the compiler, and describes the options and features that it supports.
- [Chapter 4, “Interfacing C and Assembly Code,”](#) describes the support provided by the compiler for interfacing between C source code and assembly code, and provides instructions for using this interface.
- [Chapter 5, “Optimization Techniques and Hints,”](#) explains how the SC100 optimizer operates, and describes the optimization levels and individual optimizations that can be applied.
- [Chapter 6, “Runtime Environment,”](#) describes the startup code used by the compiler, the layout and configuration of memory, and the calling conventions that the compiler supports.
- [Chapter 7, “Runtime Libraries,”](#) describes the C libraries and I/O libraries supported by the SC100 C compiler.
- [Appendix A, “Migrating from Other Environments,”](#) provides guidelines for migrating C code from other environments to the SC100 C compiler.

Suggested Reading

C, A Reference Manual, Samuel P. Harbison and Guy L. Steele Jr, Tartan Inc. (Prentice Hall: 1995).

SC100 Family of Digital Signal Processors, Specification Document Revision 0.3, Motorola Inc. (December 1998).

Conventions

This manual uses the following notational conventions:

- Courier monospaced type indicate commands, command parameters, code examples, expressions, datatypes, and directives.
- Italic type indicates replaceable command parameters.
- All source code examples are in C.

Acronyms and Abbreviations

The following list defines the acronyms and abbreviations used in this manual.

AGU	Address Generation Unit
ALU	Arithmetic Logic Unit
CFE	C Front End
DSP	Digital Signal Processor
FIR	Finite Impulse Response
IR files	Immediate Representation files
LLT	Low Level Transformations

Chapter 1

Introduction

This chapter provides an overview of the C compiler for the SC100 family of digital signal processor (DSP) cores, and outlines the SC100 C compilation process. It contains the following sections:

- [Section 1.1, “Overview of the SC100 C Compiler,”](#) describes the main features of the compiler.
- [Section 1.2, “The Global Optimization Approach,”](#) describes the benefits of the SC100 global optimization approach.
- [Section 1.3, “Compiling Applications,”](#) summarizes the main stages in the SC100 C compilation process.

1.1 Overview of the SC100 C Compiler

The SC100 C compiler accepts C source code files and produces optimized assembly code, which can be linked with assembly files and library object modules to form an executable program.

The compiler is supplied with a shell interface program, which enables you to invoke the development tools and phases that you require, in one simple step.

You can specify multiple source files of different types, together with a wide range of options, that allow you to choose the additional or specific functions to be performed by the compiler. The compiler also provides a choice of pragmas which you can use to embed optimization hints into your source code.

The compiler supports a comprehensive set of runtime libraries, a flexible runtime environment, and the facility for interfacing between C applications and assembly source code.

Using advanced optimization technology, the compiler is able to apply global optimization techniques to the source code, taking advantage of visibility of the entire application.

1.2 The Global Optimization Approach

The SC100 optimizer converts preprocessed source files into assembly output code, applying a range of code transformations which can significantly improve the efficiency of the executable program. The goal of the optimizer is to produce output code which is functionally equivalent to the original source code, while improving its performance in terms of execution time and/or code size.

The method used by traditional compilers is to optimize each source file individually, before compiling the optimized code and submitting all the compiled files to the linker. Because all the necessary information is not available when files are optimized individually, the compiler must make various assumptions, and is unable to produce the most efficient result.

To ensure optimal performance, the optimizer can take advantage of visibility of as much of the application as possible. The SC100 global binder links all modules into a single module on which all optimizations can be performed. As a result of this global approach, the performance of the optimizer is substantially improved, and the generated code is much more efficient than if produced without global optimization.

[Chapter 5, “Optimization Techniques and Hints,”](#) describes the optimization modes and functions in detail.

1.3 Compiling Applications

The SC100 compilation process consists of a series of steps, starting from the submission of source files and options to the C Front End (CFE), through the creation of Intermediate Representation (IR) files, the optimization of these files, and the output of optimized assembly code for linking into the final executable program.

You can perform all these processes in one single step, using the compiler shell program.

1.3.1 The Compiler Shell Program

The shell provides a one-step command-line interface, in which you specify the files to be processed for each compilation. At each stage, a different tool accepts the input files according to their file extensions, processes them, and outputs the transformed code for processing by the next development tool.

By default, the input files are progressed automatically through all the processing phases. The shell command line lets you select the exact development tools and processing stages that you require, and enables you to define any specific processing options, settings and/or default overrides that you need.

The options that you specify in the command line control the operation of the shell and of the tools used in the application development process. These options either affect the behavior of the shell itself or are dispatched to the different programs which the shell invokes.

The shell accepts a wide range of option types, including for example, those which perform specific actions, such as generating a list of included files, those which dictate how a source file should be treated, and those that control specific aspects of the C language features.

When you invoke the shell, the application development process is implemented automatically through all its various stages to the final production of the executable program.

[Chapter 3, “Using the SC100 C Compiler,”](#) provides a full description of the shell program and options.

1.3.2 Stages in the C Compilation Process

The following is an outline of the steps involved in compiling C source files into an executable program. These stages are illustrated in the flow diagram shown in Figure 1-1 on page 1-5:

1. The shell is invoked with the list of the C source files and assembly files to be processed, and the various options to be applied.
2. The C Front End (CFE) identifies each C source file by its file extension, preprocesses the source files, converts the files into Intermediate Representation (IR) files, and passes these to the optimizer.
3. The high-level phase of the optimizer translates each intermediate representation file into an assembly ASCII file, and performs a number of high-level optimizations. These include general (hardware-independent) and target-specific optimizations. Library files which have been created in IR form can be extracted by the optimizer, and included at this stage of processing. The optimization process also includes any relevant information contained in the application and machine configuration files.
4. The low level optimization phase carries out machine-specific, low level optimizations, and transforms the linear assembly code output by the previous phase into parallel assembly code.
5. At the end of the optimization, the optimized assembly files are output to the assembler, assembled together with any specified external assembly files, and from there output to the linker. The linker combines the assembly object files, extracts any required object modules from the library, and produces the executable application.

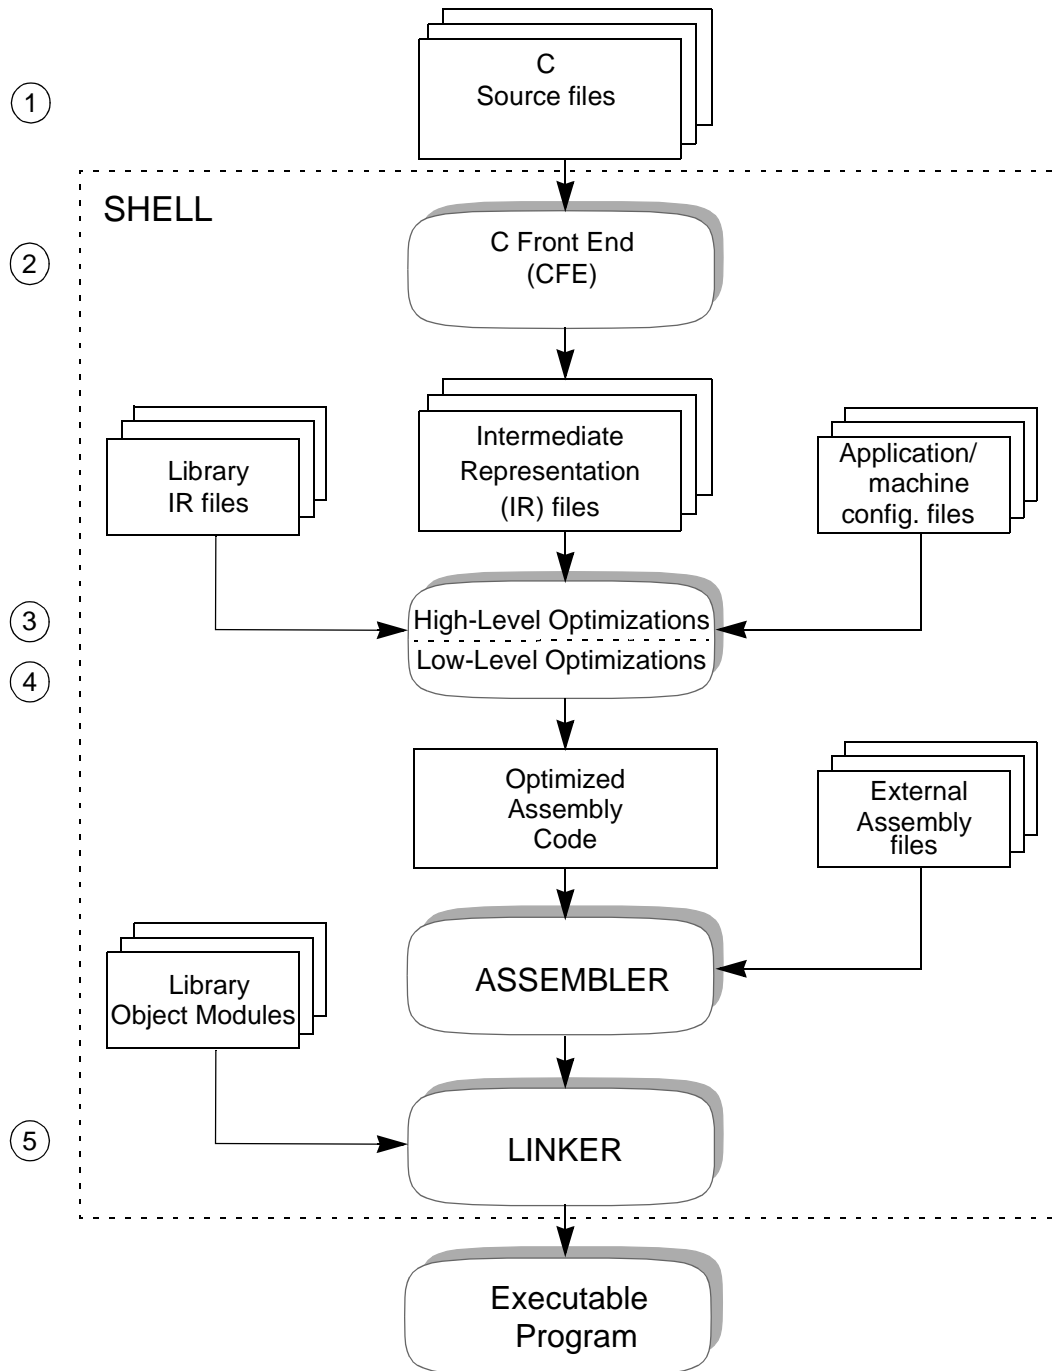


Figure 1-1. The SC100 C Compilation Process

Chapter 2

Getting Started

This chapter explains how to build and run a simple program using the SC100 C compiler.

2.1 A Quick Start

Creating and executing a program includes the following three steps:

1. Write the C source code, using the utility of your choice. In this example we will use a sample C source code file provided with your installation.
2. Compile and link the file, using the compiler shell.
3. Run the executable application that you have created.

These steps are described in greater detail below:

1. Locate the file `hello.c` in the `src/appnotes` subdirectory of your installation directory. Copy this file into your working directory.

Example 2-1 shows the C source code contained in the `hello.c` file:

Example 2-1. Sample source file - `hello.c`

```
/*hello.c*/
#include <stdio.h>
void main()
{
printf("Hello there!\n");
}
```

2. Enter the following command to instruct the shell program to compile and link the program:

```
ccsc100 -o hello.eld hello.c
```

3. Run the executable program, by entering the following:

```
sc100-sim -quiet -exec hello.eld
```

The message `Hello there!` is displayed.

Congratulations! You have successfully compiled, linked, and executed a program using the SC100 C compiler.

[Chapter 3, "Using the SC100 C Compiler,"](#) describes the shell program and the various file types in detail, and explains how to use the many options and language features that the compiler supports.

Chapter 3

Using the SC100 C Compiler

This chapter explains how to use the SC100 C compiler, and describes the options and features that the compiler supports. It contains the following sections:

- [Section 3.1, “The Shell Program,”](#) provides an overview of the compiler shell program, outlines the application development process and optimization modes, and lists the file types and environment variables recognized by the shell.
- [Section 3.2, “Invoking the Shell,”](#) explains how to initiate execution of the shell.
- [Section 3.3, “Shell Control Options,”](#) describes the options for controlling the operation of the shell and the development tools.
- [Section 3.4, “Language Features,”](#) describes the extensions and modes supported, as well as data types, intrinsic functions, pragmas, and predefined macros.

3.1 The Shell Program

The shell program controls the processing of C source files and other files into an executable application, through the preprocessing, compilation, optimization, assembly and linking stages.

The shell provides a one-step command line interface, in which you specify the files to be processed for each compilation. At each stage a different tool accepts the input files according to their file extensions, processes them, and outputs the transformed code for processing by the next development tool.

By default, the input files are progressed automatically through all the processing phases. The command line lets you select the exact development tools and processing stages that you require, and define any specific processing options, settings and/or default overrides that you need.

[Section 3.3, “Shell Control Options,”](#) describes in detail the options that you can include in the command line to change the behavior of the shell and the specific processing tools at the relevant stages of the development process.

3.1.1 The C Compilation Process

The following is an outline of the steps involved in compiling C source files into an executable program, as illustrated in the flow diagram shown in Figure 3-1 on page 3-3.

1. The shell is invoked with the list of the C source files and assembly files to be processed, and the various options to be applied.
2. The C Front End (CFE) identifies each C source file by its file extension, preprocesses the source files, converts the files into Intermediate Representation (IR) files, and passes these to the optimizer.
3. The high level phase of the optimizer translates each intermediate representation file into an assembly ASCII file, and performs a number of high level optimizations. These include general (hardware-independent) and target-specific optimizations. Library files which have been created in IR form can be extracted by the optimizer, and included at this stage of processing. The optimization process also includes any relevant information contained in the application and machine configuration files.
4. The low level optimization phase carries out machine-specific, low level optimizations, and transforms the linear assembly code output by the previous phase into parallel assembly code.
5. At the end of the optimization, the optimized assembly files are output to the assembler, assembled together with any specified external assembly files, and from there output to the linker. The linker combines the assembly object files, together with any specified external assembly files, extracts any required object modules from the library, and produces the executable application.

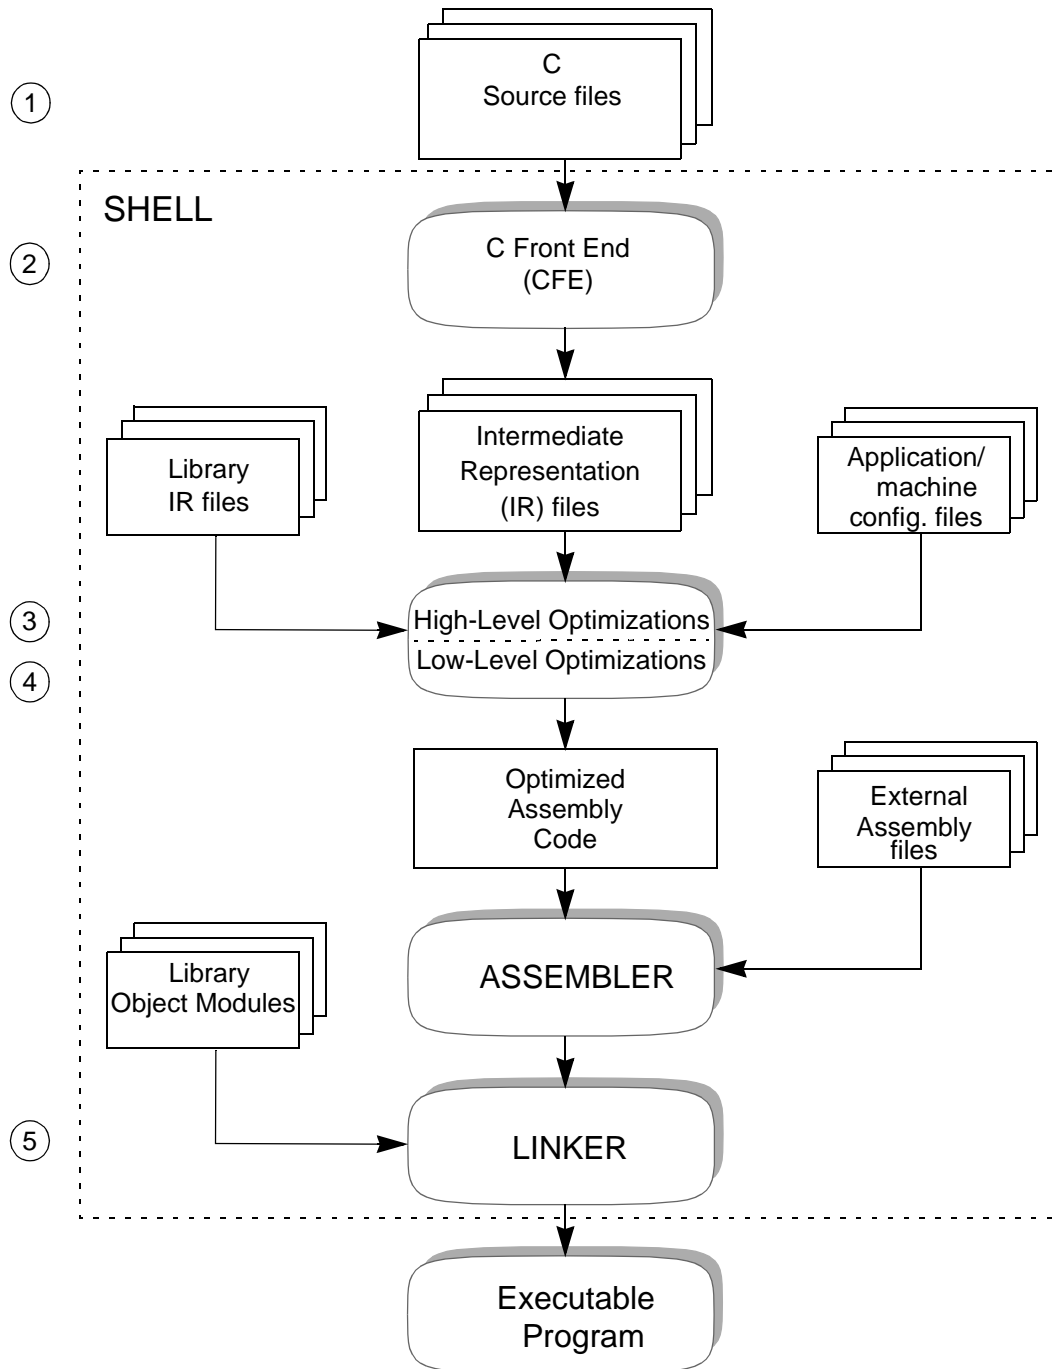


Figure 3-1. The C Compilation Process

3.1.2 Global Optimization

The method used by traditional compilers is to optimize each source file individually, before compiling the optimized code, and then submitting all the compiled files to the linker. Because not all the necessary information is available when files are optimized individually, the compiler must make various assumptions, and is unable to produce the most efficient result.

To ensure optimal performance, the optimizer can take advantage of visibility of as much of the application as possible. The SC100 global binder links all modules into a single module on which all optimizations can be performed. As a result of this global approach, the performance of the optimizer is substantially improved, and the generated code is much more efficient than if produced without global optimization.

Figure 3-2 on page 3-5, and Figure 3-3 on page 3-6, illustrate the different processing routes for non-global and global optimization.

Non-global optimization provides faster compilation, but produces less optimized code. This can be useful during the early stages of development, when you may need to compile different parts of the application separately.

Global optimization produces more efficient code, but the optimization process itself is slower than non-global optimization.

By default, the shell compiles source files without global optimization, for development purposes. You can choose to specify global optimization when you invoke the shell.

[Chapter 5, “Optimization Techniques and Hints,”](#) describes the optimization modes and functions in detail.

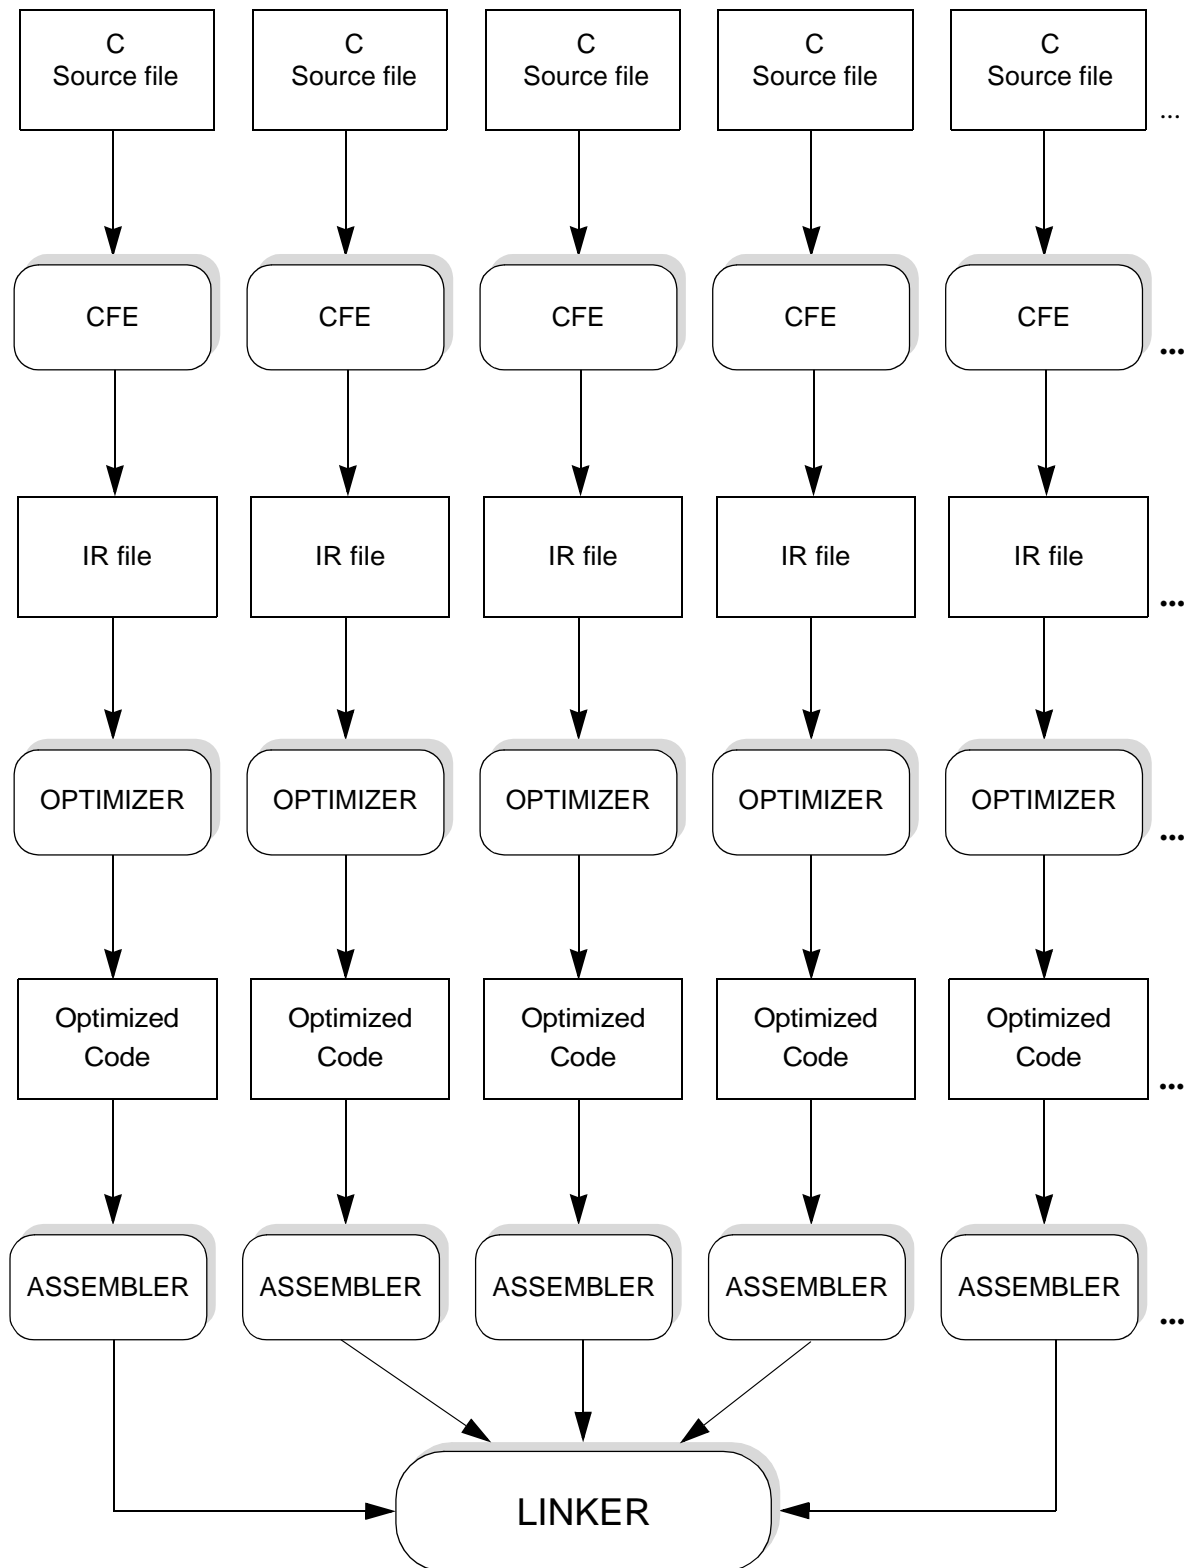


Figure 3-2. Non-Global Optimization

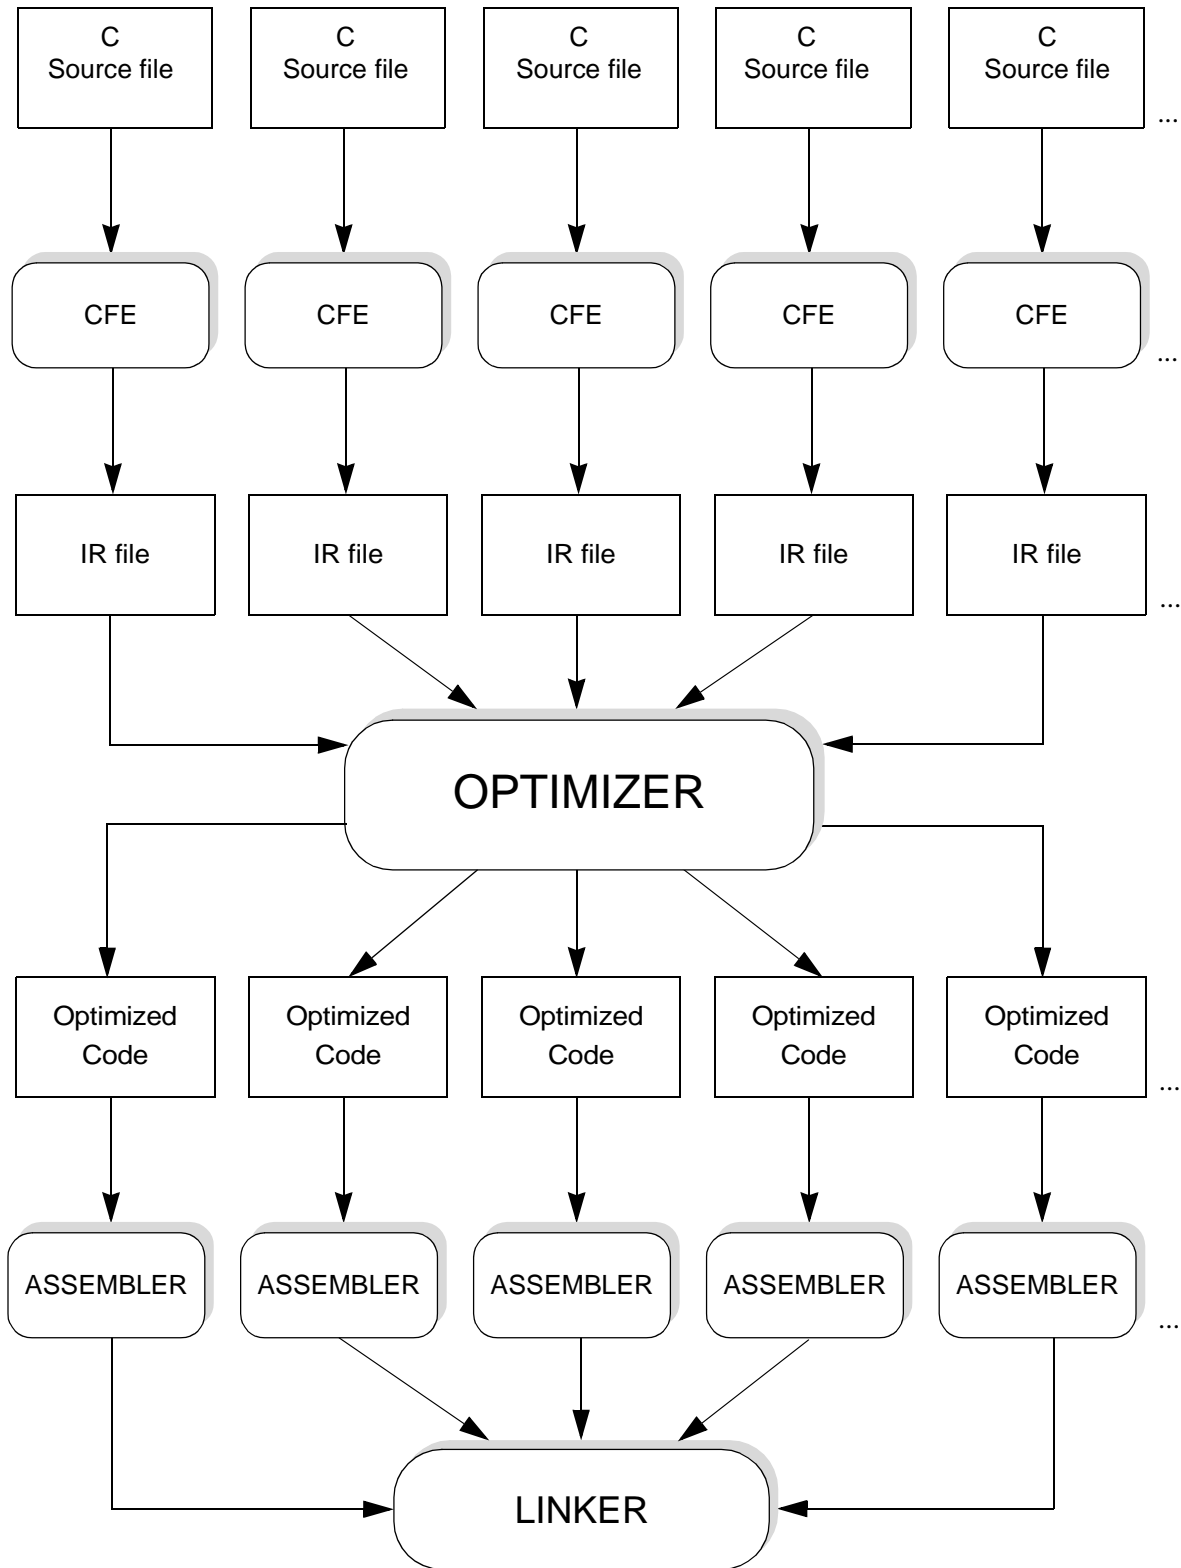


Figure 3-3. Global Optimization

3.1.3 File Types and Extensions

The shell program assumes that all items included in the command line that are not recognizable as options or option arguments are input file names. The extension for each file identifies the file type, and determines at which stage the shell will start processing the file. If the file extension is not recognized by any of the tools, the file will be treated as an input file to the linker.

The following table lists the file extensions and their corresponding file types, and shows which tool processes each file type.

Table 3-1. File Types and Extensions

Extension	File	Tool
.c	C source file	C Preprocessor
.h	C header file	
.i	Preprocessed C source	Front End
.obj	IR language file	Optimizer
.lib	IR library	Optimizer
.asm .sl	Assembly file	Assembler
.eln	Relocatable ELF object file	Linker
.elb	ELF library	Linker
.cmd	Linker command file	Linker

Note: It is possible to cause the shell to process a file as if it were a different file type, as described in [Section 3.3.3, “Overriding Input File Extensions.”](#)

The end result of the compilation process is an executable object file, with a file extension of `.eld`. Figure 3-4 on page 3-8 illustrates the assignment of file extensions at each stage of the shell processing cycle.

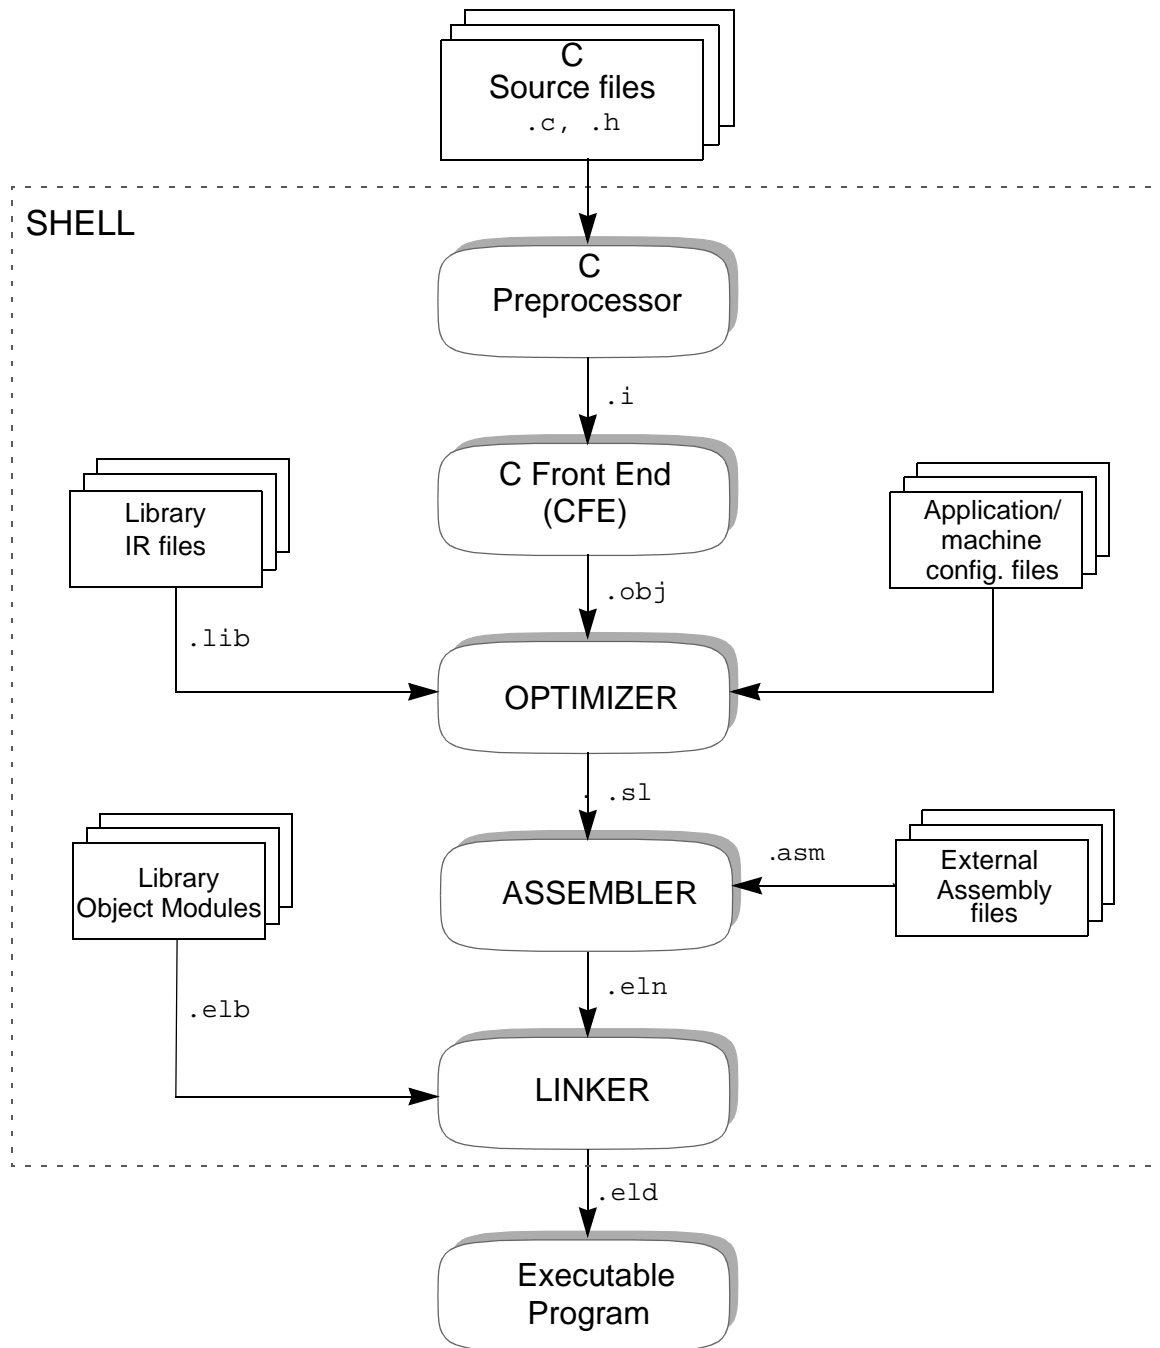


Figure 3-4. File Extensions in the Shell Cycle

3.1.4 Environment Variables

Each time the shell executes, it refers to certain environment variables which determine specific aspects of its behavior. These environment variables are defined during the installation process, and include the following:

<code>\$SC100_HOME</code>	Defines the root directory in which the executables, libraries, and tools are stored. This is set to the default location at installation. The compiler searches this directory for all the configuration and executable files that it requires.
---------------------------	--

3.2 Invoking the Shell

The shell is invoked using a single command line. This command line consists of the shell invocation command, one or more file names, and optionally, one or more shell options.

3.2.1 Command Line Syntax

The syntax of the shell command line is as follows:

```
ccsc100 [option...] file...
```

The three components of the command line are described below:

<code>ccsc100</code>	Invokes the compiler shell.
<code>option</code>	One or more options which control the way in which the shell and the various development tools operate. Section 3.3, “Shell Control Options,” describes the available options and their effect on the shell and the specified files. It is not mandatory to specify options in the command line.
<code>file</code>	The names (including extensions) of one or more files to be processed by the shell. These can be source, object, library and/or command files.

The following syntax rules apply:

- The command line must consist of only one single line.
- Individual options and files can be included in the command line in any order, and must be separated from each other by at least one space.
- Options may not be combined, and must be specified individually.
- Options which specify an argument, such as a file name or directory name, must be followed immediately by their argument(s), separated by at least one space.
- All file names, options, and arguments are case sensitive.

The shell command line shown in Example 3-1 specifies three C source files and the option `-c`, which instructs the shell to compile and assemble these files.

Example 3-1. Invoking the shell

```
ccsc100 -c one.c two.c three.c
```

3.2.1.1 Using a Command File

The command line can include one or more shell command files. These are files that you can create containing any number of options and arguments, which the shell will use as if they were part of the command line. [Section 3.3.1.2, “Specifying a shell command file,”](#) describes the use of shell command files in greater detail.

3.3 Shell Control Options

The options specified in the command line and command files control the operation of the shell, and of the tools used in the application development process.

The following categories of options are provided:

- Options that control the behavior of the shell
- Preprocessing options
- Options that override the file extension for input files
- Output filename and location options
- C language options
- Optimization pragma and code options
- Options that control the output of listing files and messages
- Pass-through options
- Hardware model and configuration options

Table 3-2 on page 3-12 provides a summary of the available options. This table is followed by detailed descriptions of each of the options, with the exception of the options relating to optimization, which are described in detail in [Chapter 5, “Optimization Techniques and Hints.”](#)

Table 3-2. Shell Options Summary*a) Options that control the behavior of the shell*

Option	Effect	Section	Page
-E	Stops after preprocessing source files. Removes comments.	3.3.1.1	3-15
-cfe	Stops after Front End. Does not invoke the optimizer. Enables the creation of libraries of object files for use with global optimization.	3.3.1.1	3-15
-S	Stops after compilation. Does not invoke the assembler.	3.3.1.1	3-15
-c	Compiles and assembles only. Does not invoke the linker.	3.3.1.1	3-15
-F <i>file</i>	Reads options from the specified file, and appends to command line.	3.3.1.2	3-16
-h or none	Displays the shell Help page, listing all available options.	3.3.1.3	3-16

b) Preprocessing options

Option	Effect	Section	Page
-C	Preserves comments in the preprocessing output.	3.3.2.1	3-17
-M <i>file</i>	Generates a MAKE file showing dependencies.	3.3.2.1	3-17
-MH <i>file</i>	Generates a list of #include files.	3.3.2.1	3-17
-D <i>mac [=def]</i>	Defines preprocessor macro.	3.3.2.2	3-18
-U <i>macro</i>	Undefines preprocessor macro.	3.3.2.2	3-18
-I <i>dir</i>	Adds directories to the #include file search path.	3.3.2.3	3-18
	Syntax note: The options -D, -U, and -I do not require a space before the argument.		

c) Options that override the file extension for input files

Option	Effect	Section	Page
-xc <i>file</i> [<i>file2 ...</i>]	Treats specified file(s) as C source file(s) (.c).	3.3.3	3-19
-xobj <i>file</i> [<i>file2 ...</i>]	Treats specified file(s) as IR language file(s) (.obj).	3.3.3	3-19
-xasm <i>file</i> [<i>file2 ...</i>]	Treats specified file(s) as assembler source file(s) (.asm or .s1).	3.3.3	3-19

d) Output filename and location options

Option	Effect	Section	Page
-o <i>file</i>	Assigns a filename (and extension) to the output file.	3.3.4	3-20
-r <i>dir</i>	Redirects all output to the specified directory.	3.3.4	3-20

Table 3-2. Shell Options Summary (Continued)

e) C language options

Option	Effect	Section	Page
-ansi	Strict ANSI mode. Assumes all C source files contain ANSI/ISO versions of the language, with no extensions. The default mode is the ANSI/ISO version with extensions.	3.3.5.1	3-21
-kr	K&R/pcc mode. Assumes all C source files contain K&R/pcc versions of the language. The default mode is the ANSI/ISO version with extensions.	3.3.5.1	3-21
-g	Adds debug information to generated files.	3.3.5.2	3-21
-sc (Default)	Makes char type variables signed.	3.3.5.3	3-21
-usc	Makes char type variables unsigned. The default setting is signed.	3.3.5.3	3-21

f) Optimization pragma and code options

Option	Effect	Section	Page
-O0	Disables all optimizations except constant folding. Outputs non-optimized assembly object code.	5.3.2.5	5-17
-O1	Omits all low level optimization steps. Performs all high level optimizations, and outputs optimized linear assembly object code. If the -Og option is also specified, high level optimizations are applied in global mode.	5.3.2	5-9
-O or -O2 (Default)	Performs all high level and low level optimizations. Outputs optimized assembly object code. If the -Og option is also specified, optimizations are applied in global mode.	5.3.3	5-19
-Os	Performs all optimizations, with more emphasis on space optimization than on time optimization. Outputs optimized assembly object code which is small as opposed to fast. Applies the optimizations in global mode if the -Og option is also specified.	5.3.5	5-30
-Ot or Ot1	Performs all optimizations, with more emphasis on time optimization than on space optimization. Outputs optimized assembly object code which is fast as opposed to small. If the -Og option is also specified, the optimizations are applied in global mode.	5.3.4	5-29
-Ot2	Performs all optimizations, with the most emphasis on time optimization. Outputs the fastest possible form of optimized assembly object code. Applies the optimizations in global mode if the -Og option is also specified.	5.3.4	5-29
-Og	Performs global optimization. The default is non-global optimization. This option can be specified together with any of the optimization options -O1 through -Ot2 listed above.	5.3.6	5-30

g) Pass-through options

Option	Effect	Section	Page
-Xasm <i>option</i>	Passes <i>option</i> to the assembler.	3.3.6	3-22
-Xlnk <i>option</i>	Passes <i>option</i> to the Linker.	3.3.6	3-22

Table 3-2. Shell Options Summary (Continued)*h) Options that control the output of listing files and messages*

Option	Effect	Section	Page
-de <i>[file]</i>	Retains generated error file.	3.3.7.1	3-23
-dm <i>[file]</i>	Generates a link map file.	3.3.7.1	3-23
-do	Adds to the assembly output file the offsets for C data structure field definitions.	3.3.7.1	3-23
-dL <i>[file]</i>	Generates a C list file.	3.3.7.1	3-23
-dL1 <i>[file]</i>	Generates a C list file including a list of #include files.	3.3.7.1	3-23
-dL2 <i>[file]</i>	Generates a C list file including expansions.	3.3.7.1	3-23
-dL3 <i>[file]</i>	Generates a C list file including both #include files and expansions.	3.3.7.1	3-23
-dx <i>[file]</i>	Generates a cross-reference information file.	3.3.7.1	3-23
-dc [0-4]	Generates a file showing calls in graphical tree form, in postscript. The number 0 to 4 specifies the paper size, A0 through A4.	3.3.7.1	3-23
-q or -w (Default)	Quiet mode. Displays errors only.	3.3.7.2	3-24
-v	Verbose mode. Displays full information.	3.3.7.2	3-24
-n	Displays command lines without executing.	3.3.7.2	3-24
-Wj	Suppresses warnings on local automatic variables that are used before their values are set.	3.3.7.3	3-24
-Wg	Suppresses global optimizer warnings.	3.3.7.3	3-24
-Wall	Reports all warnings and remarks.	3.3.7.4	3-24

i) Hardware model and configuration options

Option	Effect	Section	Page
-mc <i>file</i>	Specifies the file to be used as the machine configuration file, if different from the default file defined at installation.	3.3.8.1	3-25
-ma <i>file</i>	Specifies the file to be used as the application configuration file, if different from the default file defined at installation.	3.3.8.1	3-25
-crt <i>file</i>	Specifies the file to be used as the startup file, if different from the default file defined at installation.	3.3.8.1	3-25
-mb	Compiles in big-memory mode.	3.3.8.2	3-25
-mrom	Copies all initialized variables from ROM at startup.	3.3.8.3	3-26
-rev xx	Specifies the current revision of the chip.	3.3.8.4	3-26
-pro <i>file</i>	Reads profiling information from the specified file.	3.3.8.5	3-26
-be	Generates output for a big-endian target configuration. The default is a little-endian configuration.	3.3.8.6	3-26

3.3.1 Controlling the Behavior of the Shell

The options described in this section enable you to control the overall actions of the shell. You can specify the stage at which the shell program will stop processing, define files containing command line options, and display the invocation commands.

3.3.1.1 Controlling where the shell stops processing

By default, the shell will complete the entire processing cycle, from the input of source files through all the intermediate stages to the output of the final executable. If you want to stop the processing at a specific stage, you can use one of the options `-E`, `-cfe`, `-S`, or `-c`. In this way, you can process and check individual files or groups of files through different stages, until they are finally ready to be compiled and linked together.

Select one of the following options:

- `-E` The shell stops after preprocessing the C source files. Output files are assigned the extension `.i`. Comments are not preserved in the preprocessing output, unless the option `-C` is specified. See [Section 3.3.2.1, “Changing preprocessed output,”](#) for details of `-C` and other options that add specific features to the preprocessing output.
- `-cfe` The shell stops after processing the input source files through the Front End. You can use this option to check that the files are valid source files, which meet the essential requirements for processing by the shell, for example, they contain no syntax errors. This is primarily useful when preparing files for global optimization. Output files are IR files, assigned the extension `.-obj`.

The `-cfe` option enables you to create libraries of object files for use later when compiling in global optimization mode.
- `-S` The shell stops after compiling the source files to assembly files, and does not invoke the assembler. Output files are assigned the extension `.s1` by default.
- `-c` The shell stops after compiling C and assembly source files to object code, and does not invoke the linker. The object code output files are assigned the extension `.eln`.

Following processing with any of the above options, the output files are written to the current directory, or if the `-r` option has been included, to the specified directory. The output files are assigned the same names as the input files, with the extension for the selected option, as shown above. Any existing files in the directory with the same name and extension are overwritten.

The starting point for the processing of each input file is determined by its file extension. Refer to Table 3-1 on page 3-7 for an explanation of file extensions, and to [Section 3.3.3, “Overriding Input File Extensions,”](#) for a description of the options you can use to override these extensions.

3.3.1.2 Specifying a shell command file

You can create command files containing options and arguments, which the shell program will treat as if they were included on the command line.

Defining options and arguments within command files can save you input time when you invoke the shell program, and helps you overcome any imposed limitation on the length of the command line. Each time you invoke the shell, you can select the command file with the set of options that suit your specific requirements.

To specify a shell command file, specify the option `-F` followed by a filename. A command file can itself contain the option `-F` specifying another shell command file.

Example 3-2 illustrates the use of the `-F` option to specify the command file `proj.opt`.

Example 3-2. Defining a shell command file

```
ccsc100 -F proj.opt
```

Within the command file, each separate option (with or without an argument), file, or list of files must reside on a new line. You can specify as many lines as you wish, in any order. Comments can be included in the file using the `#` character. All characters between `#` and the end of the line are ignored by the shell.

The command file shown in Example 3-3 contains four lines which instruct the shell to invoke the linker with three application object files and one library file, generate a link map file, and output the executable program to a file named `appl.eld`.

Example 3-3. Contents of a shell command file

```
-o appl.eld           # output file name
-dm appl.map         # generate map file
file1.eln file2.eln file3.eln # object files
-l mylib.elb        # shared library
```

Note: If no map file is specified, the shell generates a file with the same file name as the specified `.eld` file, and the extension `.map`.

3.3.1.3 Displaying the shell Help page

You can display the shell Help page, which takes the form of a list of all the available shell options and arguments. Select the option `-h` to display this list.

Example 3-4 shows a section of the shell Help page:

Example 3-4. Shell Help page (extract)

```
-c    Compile and assemble only. Don't invoke the linker
-cfe  Stop after Front-End. (Used for global optimization)
-S    Generate assembly output file. Don't invoke assembler
-E    Preprocess only
-C    Preprocess only and keep comments
```

3.3.2 Specifying Preprocessing Options

The options described in this section enable you to control the preprocessing stage of the shell program, before the input files proceed through the Front End. You can change the output produced by the preprocessor, define one or more preprocessor macros, and define the directories to be searched for `#include` files.

3.3.2.1 Changing preprocessed output

You can specify any of the options listed below to change the format and content of the preprocessed output. These options can be specified in addition to the `-E` option, or instead of the `-E` option.

- `-C` Keeps all comments (preprocessor directives) in the preprocessing output. If you specify the `-E` option only, the preprocessed text is written to the output file with line control information only, and with all comments removed.
- `-M [file]` Instead of the normal preprocessing output, an output file is generated in `MAKE` format, containing a list showing the dependencies between the input source files.
If no file is specified, the output is sent to the standard output stream, `stdout`.
- `-MH [file]` Instead of the normal preprocessing output, an output file is generated containing a list of all the `#include` files used in the source. This list includes all levels of `#include` files, together with any nested files.
If no file is specified, the output is sent to the standard output stream, `stdout`.

3.3.2.2 Defining and undefining preprocessor macros

You can define one or more preprocessor macros, and you can remove the definition of a macro. [Section 3.4.6, “Predefined Macros,”](#) provides details of all the predefined macros supplied with the SC100 C compiler.

The following macro options can be specified more than once in the command line, to define and undefine different preprocessor macros:

`-D macro [=value]` Defines the named macro as a preprocessor macro, with the specified value. If *value* is omitted, the value 1 (one) is assumed. Once a preprocessor macro is defined with this option, it is passed by the shell to the preprocessor for all subsequent compilations until it is undefined with the `-U` option.

The space between the `-D` option and the named macro is optional.

`-U macro` Undefines the named macro by removing its previous definition. The macro will not be passed to the preprocessor unless it is redefined with the `-D` option. Any `-U` options in the command line are processed only *after* all `-D` options have been processed.

It is not necessary to enter a space between the `-U` option and the named macro.

3.3.2.3 Adding directories to the #include file path

The option `-I dir` adds the specified directory or directories to the path used by the shell to search for `#include` files. The string *dir* can be a list of directories.

To specify directory or directories for the `#include` file search path, specify the option `-I`, followed by a directory name or a list of directories. The space between the `-I` option and the *dir* string is optional. On UNIX hosts, separate the individual directories in the list with colons (:). On PC hosts, separate the individual directories with semicolons (;).

You can use this option more than once in a command line, and the directories or lists will be searched in the order in which the options are supplied.

3.3.3 Overriding Input File Extensions

You can change how the shell program treats a specific input file, by overriding the assumptions made by the shell based on the file's extension.

You can select any of the following options, as many times as required. After the selected option you can specify one or more filenames, separated by spaces.

`-xc file [file2 ...]` This option identifies the specified files as C language source files, as if they had the extension `.c`. The shell will process these files in exactly the same way as any other C source files specified in the command line, subject to any other processing options selected.

`-xobj file [file2 ...]` This option identifies the files as IR language files, as if they had been output by the Front End with the extension `.obj`. The files will be input for processing by the compiler.

`-xasm file [file2 ...]` This option instructs the shell program to identify the specified files as assembler source files, as if they had the extension `.asm` or `.s1`. The files will be assembled at the appropriate processing stage, and the object code will be made available to the linker.

For a list of the default extensions for each file type, see Table 3-1 on page 3-7.

These options can appear any number of times in the command line. Each option relates to one specified file or a list of files. The files identified by these options are processed normally in all other respects, and in the same relative order as other listed files.

In the following example, the input files `file1.ext` and `file2.bar`, specified after the option `-xc`, will be compiled as if they were C source files:

Example 3-5. Overriding file extensions

```
ccsc100 -c -xc file1.ext file2.bar
```

3.3.4 Output Filename and Location Options

These options let you specify the name and/or directory for the output files which the shell program will produce. By default, each output file is assigned the same name as the input file and is stored in the current directory.

The default file type and extension for the output files are determined by the stage at which the shell stops processing. For example, when the `-cfe` option has been selected, the output files produced by the Front End will have the extension `.obj`. If you wish, you can specify a different extension when you specify the file name. This will alter the way the shell will treat this file. For more information about overriding file extensions, refer to [Section 3.3.3, “Overriding Input File Extensions.”](#)

You can select either or both of the following options.

- `-o file` The output file is assigned the specified filename, and optionally the specified extension. Any existing file with the same name in the current directory, or in the specified directory, if the `-r` option is selected, is overwritten. You can specify this option more than once in the command line, for different files.

- `-r dir` All output files are redirected to the specified directory. This option can be specified only once in the command line.

In Example 3-6, the input file `file1.foo` will be treated as an input file to the linker (the default).

Example 3-6. Specifying output files

```
ccsc100 -o file.eld file1.foo
```

3.3.5 Specifying C Language Options

You can use the C language options described in this section to inform the shell of the language version being used in the source files, to add debugging information to generated files, and to define whether variables of type `char` should default to signed or unsigned.

3.3.5.1 Defining the language version

The default C language mode is the normal ANSI/ISO version with extensions, with all source files using the standard `.c` extension. You do not need to specify any language option if you use this mode. If, however, you use a different language version, you must select either the `-ansi` or the `-kr` option.

If you use the strict ANSI/ISO version of C, select the option `-ansi`. All input source files will be assumed by the Front End to be in the strict ANSI/ISO version of C, with no extensions. Any extensions found will be flagged with warnings.

If you use the K&R (Portable C Compiler, or PCC) dialect of C, select the option `-kr`. The shell program will assume that all source files are in this version of C.

See [Section 3.4, “Language Features,”](#) for details of the C language features supported in the default, strict ANSI and K&R modes.

You cannot compile source files in different C language versions at the same time. If you need to compile source files in different versions, you must use a separate shell command line for each version.

3.3.5.2 Adding debugging information to files

The option `-g` causes the shell program to include debugging information in the output files produced by all C compilations. The object files that are produced will be somewhat larger as they will contain source-level debugging information.

Note: Specifying this option does not affect the optimization level in any way.

3.3.5.3 Changing the default char sign setting

The default setting for all `char` type variables is signed. You can change this setting to make all `char` type variables default to unsigned using the `-usc` option. To change the setting back to make all `char` type variables default to signed, specify the `-sc` option.

3.3.6 Passing Options Through to Specific Tools

The options described in this section enable you to instruct the shell program to pass options to specific tools, such as the assembler or linker, as shown in Example 3-7.

Example 3-7. Passing multiple options to the same tool

```
-Xasm -occ
```

You can specify more than one option to be passed to the same tool in the same option statement, together with the arguments for each option. Multiple options, and their arguments where relevant, must be listed within quotation marks.

If a tool is invoked several times, the pass-through options are passed to the tool on each invocation, in addition to any other options passed directly by the shell program to the tool from the command line.

Specify either of the following options:

`-Xasm option` Passes the specified option(s) and arguments to the assembler.

`-Xlnk option` Passes the specified option(s) and arguments to the linker.

3.3.7 Setting the Options for Listings and Messages

The options in this section enable you to control the retention, display and printing of diagnostic and informational messages, and the generation of various listing and map files.

3.3.7.1 Generating listing files

By default the shell program does not retain the diagnostic and cross-reference information produced at different processing stages. You can select to retain one or more different types of information in listing files.

Use any combination of the following options to generate listing files containing the types of information you require. Each individual option can only be specified once in a shell command line.

- `-de [file]` The Front End creates a file containing all error messages generated during the compilation. The `-de` option retains this error file. If this option is not specified, the errors are displayed during processing, but not kept. If no file name is specified, a file is created with the same name as the source file, and the extension `.err`.
- `-dm [file]` Generates a link map file listing all the specific variables, applications and addresses used by the linker. If no file name is specified, a file is created with the same name as the executable, and the extension `.map`.
- `-do` Includes the details of C data structures in the output assembly file, showing the offsets for all field definitions in each data structure. Refer to [Chapter 4, “Interfacing C and Assembly Code,”](#) for a detailed description of the C/assembler interface.
- `-dL [file]` Generates a C list file, listing the entire contents of the source file. If no file name is specified, a file is created with the same name as the source file, and the extension `.lis`.
- `-dL1 [file]` Generates a C list file, listing the entire contents of the source file, with the addition of a list of `#include` files used by the source. If no file name is specified, a file is created with the same name as the source file, and the extension `.lis`.
- `-dL2 [file]` Generates a C list file, listing the entire contents of the source file, with the addition of expansions, such as macro expansions, line splices and trigraphs. If no file name is specified, a file is created with the same name as the source file, and the extension `.lis`.
- `-dL3 [file]` Generates a C list file, listing the entire contents of the source file, with the addition of a list of `#include` files, and expansions, such as macro expansions, line splices and trigraphs. If no file name is specified, a file is created with the same name as the source file, and the extension `.lis`.
- `-dx [file]` Generates a cross-reference information file, providing details of cross-references in the source file. If no file name is specified, a file is created with the same name as the source file, and the extension `.xrf`.
- `-dc [0-4]` Generates a file showing calls in graphical tree form, which can be printed using a postscript printer. Specify the size of the paper to be used for the printout: 0 for paper size A0, 1 for A1, and so on.

3.3.7.2 Controlling the type of information displayed

You can control the level and type of messages and information that the shell program displays using the following options:

- q or -w Quiet mode (the default). The shell program displays the minimum amount of information (errors only). Normal notices and banners are omitted. This option is useful when running the shell in batch mode or with the `MAKE` utility, when the display of normal progress information is not required.
- v Verbose mode. The shell program displays/prints all the commands and command line options being used, as it proceeds through the different processing stages and invokes the individual tools. The exact information output will depend on the processing stages performed by the shell.
- n Displays the specified shell processing actions without executing them. You can use this option before you invoke the shell, to check the actions the shell will take, based on the list of files and arguments specified in the command line.

3.3.7.3 Suppressing warnings

By default the shell reports all errors and warnings. You can use these options to reduce the number of messages that the shell program generates, by choosing to suppress specific types of warnings. This can be useful if, for example, you are testing incomplete sections of the program and you know in advance that certain warnings will be produced.

You can select either or both of the following options:

- wj This option suppresses warnings on local automatic variables that are used before their values are set. If you are testing partial code which you know does not assign values to all the local automatic variables, you can use this option to suppress all the "false" warnings that would otherwise be issued.
- wg By default, a warning is produced for each module identified as missing during the global optimization process. Use this option if you wish to suppress these warnings, for example when testing an incomplete application, or one that uses external modules.

3.3.7.4 Reporting all remarks and warnings

The shell reports all errors and warnings by default, but will not report remarks unless you specifically instruct it to do so. Select the option `-wall` to ensure that all remarks are reported, as well as all warnings and errors.

3.3.8 Specifying the Hardware Model and Configuration

The options in this category let you override some of the hardware and configuration settings defined as the default during installation.

3.3.8.1 Defining specific configuration and startup files

The default machine and application configuration files used by the compiler, and the startup file used by the linker, are defined during the installation process.

The machine configuration file includes information about the logical and physical memory maps. This information enables the global optimizer to dispatch variables to different memory areas in internal ROM or RAM.

The application configuration file contains information about how the application software and the hardware interact. The file includes sections about binding interrupt handlers, overlays, and application objects to specific addresses.

The startup file is used by the linker when it links the assembly code files with the standard libraries, and defines such items as the interrupt vector and set-up code executed upon system initialization.

These files, and their use in the run-time environment, are described in greater detail in [Chapter 6, “Runtime Environment.”](#)

You may wish to select other files to be used for configuration setup and initialization instead of the default files, for example, to specify certain devices which need to be initialized at startup.

To specify different files to be used at initialization, select one or all of the options listed below. For each option, specify the file name, and if the file is not in the current directory, specify the path.

<code>-mc <i>file</i></code>	The compiler reads the specified file instead of the default machine configuration file.
<code>-ma <i>file</i></code>	The compiler reads the specified file instead of the default application configuration file.
<code>-crt <i>file</i></code>	The linker links into the application the specified file instead of the default startup file.

For more detailed information, refer to [Chapter 6, “Runtime Environment.”](#)

3.3.8.2 Defining memory mode

The SC100 architecture instruction set supports both 16-bit and 32-bit addresses. If the application is small enough to allow all static data to fit into the lower 64K of the address space, then more efficient code can be generated. This mode (small memory mode) is the default, and assumes that all addresses are 16-bit.

If your application does not fit into 64K bytes, meaning that the use of 32-bit absolute addresses is required, you must instruct the shell to use the big memory model, by specifying the `-mb` option.

3.3.8.3 Copying initialized variables from ROM

During development you would normally use a loader to set the values for global variables, and to load these initialized variables into RAM at startup, together with the executable application.

When you have finished development, if your final application does not use a loader, you must ensure that when the completed application executes, the initialized variables will be copied from ROM into RAM. To do this, when you compile the final application version, specify the `-mrom` option.

Refer to [Chapter 6, “Runtime Environment,”](#) for more detailed information about the initialization of variables in the runtime environment.

3.3.8.4 Specifying the silicon revision

This option lets you specify the revision number of the silicon chip currently in use. You need to use this only if a new version has been issued since the original installation.

Select the option `-rev` followed by the revision number.

3.3.8.5 Defining the profiling information file

Profiling information describes the timing behavior of specific inputs for an application, meaning the exact number of times each specified source line executes. This information can be stored in a file and read by the compiler in order to enhance the efficiency of the optimization.

To instruct the compiler to read profiling information, select the option `-pro`, followed by a file name with one of the following:

- The name `tcovd`
- On a PC, any file name and the extension `.tdf`
- On a SPARCstation, any file name and the extension `.d`

The `-pro` option may be specified only once in the command line.

In Example 3-8, the compiler reads the profiling information file `coder.tdf`, and uses this information during optimization to determine the exact number of times that certain lines of code execute.

Example 3-8. Defining a profiling information file

```
ccsc100 -pro coder.tdf -o coder.eld coder.c
```

3.3.8.6 Specifying big-endian mode

By default, the compiler generates code based on the assumption that the architecture operates in little-endian mode, meaning the least significant bits in the lower address. If you want to run the application in an environment that operates in big-endian mode, meaning the most significant bits in the lower address, specify the option `-be`.

3.4 Language Features

This section describes the different language modes accepted by the SC100 C compiler. It also provides detailed information about the data types and sizes supported, fractional arithmetic representation, intrinsic functions, pragmas and predefined macros.

3.4.1 C Language Dialects

The compiler accepts three variations of the C language. The standard C language mode is the normal ANSI/ISO version with extensions, as described in [Section 3.4.1.1, “Standard Extensions.”](#) This is the default mode.

The other two C language modes accepted are:

- **Strict ANSI/ISO mode:** Specified with the shell option `-ansi`. Any ISO C extensions are flagged with warnings.
- **K&R/PCC mode:** Specified with the shell option `-kr`. The compiler accepts the older K&R dialect of C, and provides almost complete compatibility with the widely used UNIX PCC (`pcc`) dialect. See [Section 3.4.1.2, “K&R/PCC mode,”](#) for details.

Source files of different C language types may not be compiled together, but once compiled they can be linked together into a single application.

3.4.1.1 Standard Extensions

This section lists the extensions which are normally accepted in standard C programs. When compiling in strict ANSI/ISO mode, the compiler issues warnings when these extensions are used.

3.4.1.1.1 Preprocessor extensions

The following preprocessor extensions are accepted:

- Comment text can appear at the end of preprocessing directives.
- Numbers are scanned according to the syntax for numbers. Thus, `0x123e+1` is scanned as three tokens instead of one invalid token.
- The `#assert` preprocessing extensions of AT&T System V release 4 are allowed. These enable the definition and testing of predicate names. Such names are in a name space distinct from all other names, including macro names. A predicate name can be defined by a preprocessing directive in one of two forms, as shown in Example 3-9:

Example 3-9. Defining a predicate name

```
#assert name
#assert name(token-sequence)
```

In the first form, the predicate is not given a value. In the second form, it is given the value *token-sequence*. Such a predicate can be tested in a `#if` expression, as follows: `#name(token-sequence)`. This expression has the value 1 if a `#assert` of that *name* with that *token-sequence* has appeared, otherwise it has the value 0. A predicate may be assigned more than one value at a given time.

- A predicate may be deleted by a preprocessing directive in one of two forms, as shown in Example 3-10:

Example 3-10. Deleting a predicate

```
#unassert name
#unassert name(token-sequence)
```

The first form removes all definitions of the indicated predicate name. The second form removes only the indicated definition, leaving any remaining definitions unchanged.

A number of predefined preprocessor macros are provided, as described in [Section 3.4.6, “Predefined Macros.”](#)

The pragmas described in [Section 3.4.5, “Pragmas,”](#) are available in all modes.

3.4.1.1.2 Syntax

The following syntax extensions are accepted:

- A translation unit (input file) can be empty, containing no declarations.
- An extra comma is allowed at the end of an `enum` list. Similarly, the final semicolon preceding the closing `}` of a struct or union specifier may be omitted. A remark is issued in both cases, except in `pcc` mode.
- A label definition may be followed immediately by a right brace. (Normally, a label definition must be followed by a statement.) A warning is issued.
- An empty declaration (a semicolon with nothing before it) is allowed. A remark is issued.
- An initializer expression that is a single value and is used to initialize an entire static array, `struct`, or `union` need not be enclosed in braces, except in strict ANSI C mode.
- A `struct` that has no named fields, but which has at least one unnamed field, is accepted by default. A diagnostic (a warning or error) is issued in strict ANSI C mode.

3.4.1.1.3 Declarations

The following declaration extensions are accepted:

- Static functions may be declared in function and block scopes. Their declarations are moved to the file scope.
- Benign redeclarations of `typedef` names are allowed, meaning that a `typedef` name may be redeclared in the same scope as the same type. A warning is issued.
- The compiler always accepts `asm` statements and declarations, with one exception, which is when compiling in strict ANSI C mode. The reason for this is that there would be a conflict with the ANSI C standard. For example, `asm("xyz");` would be interpreted by the Front End as an `asm` statement by default, while ANSI C would interpret this as a call of an implicitly-defined function `asm`.
- Functions declared as `asm` functions are accepted, and `__asm` is recognized as a synonym for `asm`. An `asm` function body is represented by an uninterpreted null-terminated string containing the text that appears in the source.

- An `asm` function must be declared with no storage class, with a prototyped parameter list, and with no omitted parameters, as shown in Example 3-11:

Example 3-11. Declaring an `asm` function

```
asm void f(int,int) {
    ...
}
```

- As an `asm` function must be output with a prototyped parameter list, these functions are valid for ANSI C modes only.

3.4.1.1.4 Types

The following type extensions are accepted:

- Bit-fields may have base types that are `enums` or integer types, as well as the types `int` and `unsigned int`. The use of any signed integer type is equivalent to using type `int`, and the use of any unsigned integer type is equivalent to using type `unsigned int`.
- The last member of a `struct` may have an incomplete array type. It may not be the only member of the `struct` (otherwise, the `struct` would have zero size).
- A file-scope array may have an incomplete `struct`, `union`, or `enum` type as its element type. The type must be completed before the array is subscripted (if it is subscripted), and by the end of the compilation if the array is not `extern`.
- The `enum` tags may be incomplete. The tag name can be defined and resolved later (by specifying the brace-enclosed list).
- Object pointer types and function parameter arrays that decay to pointers may use `restrict` as a type qualifier. Its presence is recorded in the compiler so that optimizations can be performed that would otherwise be prevented because of possible aliasing.
- The type `long float` is accepted as a synonym for `double`.
- Assignment of pointer types is allowed in cases where the destination type has added type qualifiers that are not at the top level (for example, `int **` to `const int **`).

3.4.1.1.5 Expressions and statements

The following extensions are accepted for expressions and statements:

- Assignment and pointer differences are allowed between pointers to types that are interchangeable, but not identical, for example, `unsigned char *` and `char *`. This includes pointers to same-sized integral types (e.g., typically, `int *` and `long *`). A warning is issued, except in `pcc` mode. A string constant may be assigned to a pointer to any kind of character, without a warning.
- In operations on pointers, a pointer to `void` is always implicitly converted to another type if necessary, and a null pointer constant is always implicitly converted to a null pointer of the right type if necessary. In ANSI C, some operators allow such conversions, while others do not, generally where such a conversion would not be logical.
- In an initializer, a pointer constant value may be cast to an integral type if the integral type is big enough to contain it.
- In an integral constant expression, an integer constant may be cast to a pointer type and then back to an integral type.
- In character and string escapes, if the character following the `\` has no special meaning, the value of the escape is the character itself. Thus `“\s” == “s”`. A warning is issued.

- Adjacent wide and non-wide string literals are not concatenated.
- In duplicate size and sign specifiers (e.g., `short short` or `unsigned unsigned`) the redundancy is ignored, and a warning is issued.
- `__ALIGNOF__` is similar to `sizeof`, but returns the alignment requirement value for a type, or 1 if there is no alignment requirement. It may be followed by a type or expression in parentheses, as shown in Example 3-12:

Example 3-12. Returning the alignment requirement

```
__ALIGNOF__(type)  
__ALIGNOF__(expression)
```

The expression in the second form is not evaluated.

- Identifiers may not contain dollar signs.
- `__INTADDR__(expression)` scans the enclosed expression as a constant expression, and converts it to an integer constant (it is used in the `offsetof` macro).
- The values of enumeration constants may be given by expressions that evaluate to unsigned quantities which fit in the `unsigned int` range but not in the `int` range. A warning is issued when such a result is possible, as shown in Example 3-13:

Example 3-13. Out of range warning

```
/* When ints are 32 bits: */  
enum a {w = -2147483648}; /* No warning */  
enum b {x = 0x80000000}; /* No warning */  
enum c {y = 0x80000001}; /* No warning */  
enum d {z = 2147483649}; /* Warning */
```

- The address of a variable with `register` storage class may be taken, and a warning is issued.
- The expression `& . . .` is accepted in the body of a function in which an ellipsis appears in the parameter list.
- An ellipsis may appear by itself in the parameter list of a function declaration, for example, `f (. . .)`. A diagnostic is issued in strict ANSI mode.
- External entities declared in other scopes are visible, as shown in Example 3-14. A warning is issued.

Example 3-14. External entities in other scopes

```
void f1(void) { extern void f(); }  
void f2() { f(); /* Using out of scope declaration */ }
```

- Pointers to incomplete arrays may be used in pointer addition, subtraction, and subscripting, as shown below in Example 3-15. A warning is issued if the value added or subtracted is anything other than a constant zero. Since the type pointed to by the pointer has zero size, the value added to or subtracted from the pointer is multiplied by zero and therefore has no effect on the result. Comparisons and pointer differences of such pairs of pointer types are also allowed. A warning is issued.

Example 3-15. Pointers to incomplete arrays

```
int (*p)[];  
...  
q = p[0];
```

- Pointers to different function types may be assigned or compared for equality (==) or inequality (!=) without an explicit type cast, and a warning is issued.
- A pointer to `void` may be implicitly converted to or from a pointer to a function type.
- Intrinsic functions are recognized as extensions only in the default C language mode (ANSI C with extensions). In all other modes they are treated as function calls.

3.4.1.2 K&R/PCC mode

When `pcc` mode is specified, the SC100 C compiler accepts the traditional C language defined by *The C Programming Language*, first edition, by Kernighan and Ritchie (K&R), Prentice-Hall, 1978. This mode provides almost complete compatibility with the Reiser CPP and Johnson PCC (`pcc`), both widely used as part of UNIX systems. Since there is no documentation of the exact behavior of those programs, complete compatibility cannot be guaranteed.

In general, when compiling in `pcc` mode, the compiler attempts to interpret a source program that is valid to `pcc` in the same way that `pcc` would. However, ANSI features that do not conflict with this behavior are not disabled.

In some cases where `pcc` allows a highly questionable construct, the compiler accepts it but gives a warning, where `pcc` would be silent. For example: `0x`, a degenerate hexadecimal number, is accepted as zero, but a warning is issued.

3.4.1.2.1 K&R/PCC mode preprocessor differences

The following are the preprocessor differences relative to the default standard mode:

- When preprocessing output is generated, the line-identifying directives will have the `pcc` form instead of the ANSI form.
- `__STDC__` is left undefined.
- Comments are deleted entirely (instead of being replaced by one space) in preprocessing output. Extra spaces are not generated in textual preprocessing output to prevent pasting of adjacent confusable tokens. As a result, the characters `a/**/b` will be `ab` in preprocessor output.
- The first directory searched for include files is the directory containing the file which contains the `#include` instead of the directory which contains the primary source file.
- Trigraphs are not recognized.
- Macro expansion is implemented differently. Arguments to macros are not macro-expanded before being inserted into the expansion of the macro. Any macro invocations in the argument text are expanded when the macro expansion is rescanned. With this method, macro recursion is possible and is checked for.
- Token pasting inside macro expansions is implemented differently. End-of-token markers are not maintained, so tokens that abut after macro substitution may be parsed as a single token.
- Macro parameter names inside character and string constants are recognized and are given substitutes.
- Macro invocations having too many arguments are flagged with a warning rather than an error. The extra arguments are ignored.
- Macro invocations having too few arguments are flagged with a warning rather than an error. A null string is used as the value of the missing parameters.
- Extra occurrences of `#else` (after the first has appeared in an `#if` block) are ignored, with a warning.

3.4.1.2.2 K&R/PCC mode syntax differences

The following are the syntax differences relative to the default standard mode:

- The keywords `signed`, `const`, and `volatile` are disabled, so that they can be user identifiers. The other non-K&R keywords (`enum` and `void`) are judged to have existed already in code and are not disabled.
- The `=` preceding an initializer may be omitted. A warning is issued. This was an anachronism even in K&R.
- `0x` is accepted as a hexadecimal 0, with a warning.
- `1E+` is accepted as a floating point constant with an exponent of 0, with a warning.
- The compound assignment operators may be written as two tokens (for example, `+=` may be written `+ =`).
- The compound assignment operators may be written in their old-fashioned reversed forms (for example, `-=` may be written `= -`). A warning is issued.
- The digits 8 and 9 are allowed in octal constants. (For example, the constant 099 has the value $9*8+9$, or 81.)
- The escape `\a` (alert) is not recognized in character and string constants.

3.4.1.2.3 K&R/PCC mode differences for declarations

The following are the declaration differences relative to the default ANSI mode:

- Declarations of the form `typedef some-type void;` are ignored.
- The names of functions and of external variables are always entered at the file scope.
- A function declared `static`, which is used and never defined, is treated as if its storage class were `extern` (instead of causing an error for being undefined).
- A file-scope array that has an unspecified storage class and remains incomplete at the end of the compilation will be treated as if its storage class is `extern`. In ANSI mode, the number of elements is changed to 1, and the storage class remains unspecified.
- When a function parameter list begins with a `typedef` identifier, the parameter list is considered prototyped only if the `typedef` identifier is followed by something other than a comma or right parenthesis, as shown below in Example 3-16. Function parameters are allowed to have the same names as `typedef` identifiers. In the normal ANSI mode, any parameter list that begins with a `typedef` identifier is considered prototyped, and Example 3-16 would produce an error.

Example 3-16. Prototyped parameter list

```
typedef int t;
int f(t) {}          /* Old-style list */
int g(t x) {}       /* Prototyped list, parameter x of type t */
```

- The empty declaration `struct x;` will not hide an outer-scope declaration of the same tag. It is taken to refer to the outer declaration.

- In a declaration of a member of a `struct` or union, the declarator list may be omitted entirely, to specify an unnamed field which requires padding, as shown in Example 3-17. Such a field may not be a bit-field.

Example 3-17. Omitting the declarator list

```
struct s {char a; int; char b[2];} v; /* sizeof(v) is 3 */
```

- No warning is generated for a storage specifier appearing in other than the first position in a list of specifiers (as in `int static`).
- The keywords `short`, `long`, and `unsigned` are treated as “adjectives” in type specifiers, and they may be used to modify a `typedef` type. For example, the declarations in Example 3-18 result in having type `unsigned long`:

Example 3-18. Keywords in type specifiers

```
typedef long size;  
unsigned size s;
```

- Free-standing tag declarations are allowed in the parameter declaration list for a function with old-style parameters.
- Declaration specifiers are allowed to be completely omitted in declarations. (ANSI C allows this only for function declarations.) Thus `i;` declares `i` as an `int` variable. A warning is issued.
- An identifier in a function is allowed to have the same name as a parameter of the function. A warning is issued.

3.4.1.2.4 K&R/PCC mode type differences

The following are the type differences relative to the default standard mode:

- Integral types with the same representation (size, signedness, and alignment) will be considered identical and may be used interchangeably. For example, this means that `int` and `long` will be interchangeable if they have the same size.
- All `enums` are given type `int`. In ANSI mode, smaller integral types will be used if possible.
- A “plain” `char` is considered to be the same as either `signed char` or `unsigned char`, depending on the command-line options. In ANSI C, “plain” `char` is a third type distinct from both `signed char` and `unsigned char`.
- All `float` functions are promoted to `double` functions, and any `float` function parameters are promoted to `double` function parameters.
- All `float` operations are executed as `double`.
- The types of large integer constants are determined according to the K&R rules. They will not be `unsigned` in some cases where ANSI C would define them that way.

3.4.1.2.5 K&R/PCC mode differences: expressions and statements

The following are the differences for expressions and statements relative to the default standard mode:

- Assignment is allowed between pointers and integers, and between incompatible pointer types, without an explicit cast. A warning is issued.
- A field selection of the form `p->field` is allowed even if `p` does not point to a `struct` or `union` that contains `field`. In this context, `p` must be a pointer or an integer. Similarly, `x.field` is allowed even if `x` is not a `struct` or `union` that contains `field`. In this case, `x` must be an lvalue. In both cases, if `field` is declared as a `field` in more than one `struct` or `union`, it must have the same offset in all instances.
- Overflows detected while folding signed integer operations on constants will cause warnings rather than errors.
- A warning will be issued for an `&` operator applied to an array. The type of such an operation is “address of array element” rather than “address of array”.
- For the shift operators `<<` and `>>`, the usual arithmetic conversions are done on the operands as they would be for other binary operators. The right operand is then converted to `int`, and the result type is the type of the left operand. In ANSI C, the integral promotions are done on the two operands separately, and the result type is the type of the left operand. The effect of this difference is that, in `pcc` mode, a `long` shift count will force the shift to be done as `long`.
- String literals will not be shared. Identical string literals will cause multiple copies of the string to be allocated.
- The expression `sizeof` may be applied to bit-fields. The size is that of the underlying type (for example `unsigned int`).
- Any lvalues cast to a type of the same size remain lvalues, except when they involve a floating point conversion.
- A warning rather than an error is issued for integer constants that are larger than can be accommodated in an `unsigned long`. The value is truncated to an acceptable number of low-order bits.
- Expressions in a `switch` statement are cast to `int`. This differs from the ANSI C definition in that a `long` expression may be truncated.
- The promotion rules for integers are different: `unsigned char` and `unsigned short` are promoted to `unsigned int`.

3.4.1.2.6 K&R/PCC differences: remaining incompatibilities

The additional known cases where the compiler is not compatible with `pcc` are as follows:

- Token pasting is not implemented outside of macro expansions (meaning, in the primary source line) when two tokens are separated only by a comment. That is, `a/**/b` is not considered to be `ab`. The `pcc` compiler's behavior in such a case can be obtained by preprocessing to a text file and then compiling that file.

The textual output from preprocessing is also equivalent but not identical. The blank lines and white space will not be exactly the same as those produced in `pcc`.

- The `pcc` compiler considers the result of a `?:` operator to be an `lvalue` if the first operand is constant and the second and third operands are compatible `lvalues`. The compiler never treats the result of the `?:` operator as an `lvalue`.
- The `pcc` compiler misparses the third operand of a `?:` operator in a way that some programs exploit, as follows:

`i ? j : k += 1` is parsed by `pcc` as `i ? j : (k += 1)`

This is not correct, since the precedence of the `+=` operator is lower than the precedence of the `?:` operator. The compiler will generate an error in such a case.

- The `lint` utility recognizes the keywords for its special comments anywhere in a comment, regardless of whether they are preceded by other text in the comment. The compiler only recognizes the keywords when they are the first identifier following an optional initial series of blanks and/or horizontal tabs. In addition, `lint` recognizes only a single digit of the `VARARGS` count. The compiler accumulates as many digits as appear in the count.

3.4.2 Types and Sizes

The data types supported by the compiler are summarized in Table 3-3 below. The table shows the size for each data type in memory and in the two register types, the 40-bit data register (Dn), and the 32-bit address register (Rn). Table 3-3 also shows the required alignment and the value range for each data type, together with a reference to the section in this chapter which provides greater detail about the data type.

Table 3-3. Data Types and Sizes

Type	Size (in bits)				Range		Details	
	Memory	Dn	Rn	Align	Minimum	Maximum	Section	Page
char	8	40	32	8	-128	127	3.4.2.1	3-38
unsigned char	8	40	32	8	0	255	3.4.2.1	3-38
short	16	40	32	16	-32,768	32,767	3.4.2.2	3-39
unsigned short	16	40	32	16	0	65,535	3.4.2.2	3-39
int	32	40	32	32	-2,147,483,648	2,147,483,647	3.4.2.2	3-39
unsigned int	32	40	32	32	0	4,294,967,295	3.4.2.2	3-39
long	32	40	32	32	-2,147,483,648	2,147,483,647	3.4.2.2	3-39
unsigned long	32	40	32	32	0	4,294,967,295	3.4.2.2	3-39
float, double, and long double	32	40	32	32	-1.17E-38	1.17E+38	3.4.2.3	3-40
fractional short ¹	16	40	-	16	-1	0.99969842	3.4.2.4	3-41
fractional long / int ²	32	40	-	32	-1	0.9999999953	3.4.2.4	3-41
pointer	32	40	32	32	0	0xFFFFFFFF	3.4.2.5	3-41

1. Fractional short is not a language type. It can be used with intrinsic functions only, and maps to the predefined type short.
2. Fractional long/int is not a language type. It can be used with intrinsic functions only, and maps to the predefined type long/int.

3.4.2.1 Characters

A character, whether signed or unsigned, is stored in memory in one byte (8 bits), and is always aligned on an 8-bit boundary. Arrays of characters occupy one byte per character. Figure 3-5 shows the memory layout for characters.




Figure 3-5. Characters—Memory Layout

When loaded into registers, signed characters are signed extended, while unsigned characters are zero extended. Figure 3-6 illustrates the layout for signed and unsigned characters in the Dn (40-bit) data register. “S” indicates the signed extension of the value.




Figure 3-6. Characters—Dn Register Layout

Figure 3-7 shows the layout for signed and unsigned characters in the Rn (32-bit) address register.




Figure 3-7. Characters—Rn Register Layout

3.4.2.2 Integers

Integer arithmetic is performed using data sizes appropriate to the arithmetic operation. Short integers use at least 16-bit wide operations (single-precision integer arithmetic), and long integers use at least 32-bit (double-precision integer arithmetic).

Short and long integers are stored in memory using little-endian representation (the least significant bits in the lower address), unless the option `-be` is specified.

Integer arithmetic overflow wraps around and does not result in any additional side effects.

Figure 3-8 shows the memory layout for short and long integers.



Figure 3-8. Integers—Memory Layout

Short integers must be aligned on 2-byte (16-bit) boundaries, while long integers must be aligned on a 4-byte (32-bit) boundary. Figure 3-9 illustrates the alignment of short and long integers, in conjunction with characters.




Figure 3-9. Integers—Alignment

As with characters, when loaded into registers, signed integers are signed extended, while unsigned integers are zero extended.

Figure 3-10 illustrates the layout for signed and unsigned short and long integers in the Dn (40-bit) data register. “S” indicates the signed extension of the value.



Figure 3-10. Integers—Dn Register Layout

Figure 3-11 shows the layout for signed and unsigned short and long integers in the Rn (32-bit) address register.

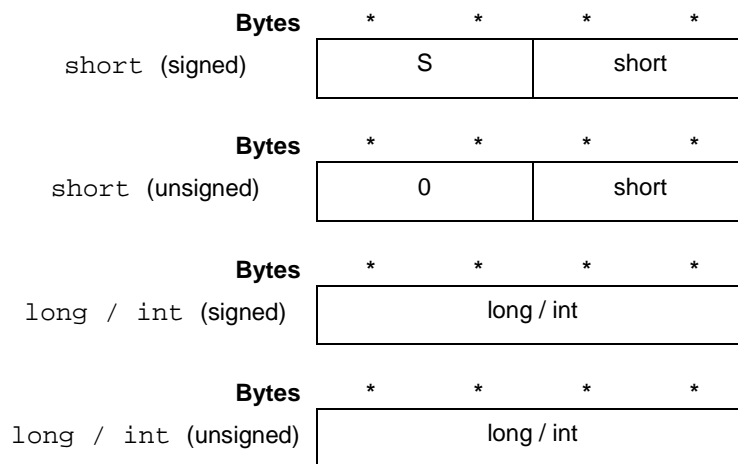


Figure 3-11. Integers—Rn Register Layout

3.4.2.3 Floating point

Floating point, double, and long double type integers are mapped to a single precision IEEE-754 type, using 32 bits (4 bytes). The compiler generates calls for library functions to evaluate floating point expressions. The representation of these integers in memory and in the registers is exactly the same as for long integers.

3.4.2.4 Fractional representation

Since C does not provide built-in support for fractional types, the syntactic representation of fractional types and operations is implemented by intrinsic functions using integer data types. See [Section 3.4.4, “Intrinsic Functions,”](#) for details of the intrinsic functions supported.

Fixed -point arithmetic is performed using 16-bit, 32-bit, 40-bit, and 64-bit operations. Fractional integers are stored in memory using little-endian representation, meaning the least significant bits in the lower address, unless the option `-be` is specified.

Fractional type overflows may saturate and do not result in any additional side effect. Rounding and saturation modes are determined as part of the startup code, or with optional intrinsic function calls.

Operations on double and extended precision type objects are limited to assignments and fractional arithmetic using intrinsic functions only. See [Section 3.4.3, “Fractional and Integer Arithmetic,”](#) for further information. Integer operations on extended precision types are not supported.

Fractional types are mapped to their corresponding predefined types. A fractional `short` maps to the predefined type `short`, a fractional `long` maps to the predefined type `long`, and a fractional `int` maps to the predefined type `int`.

Figure 3-12 illustrates the layout for fractional `short` and `long` integers in the Dn (40-bit) data register, which is the only register used for fractional integer types. “S” indicates the signed extension of the value.




Figure 3-12. Fractional Integers—Dn Register Layout

When loading data from memory into data registers, the compiler aligns the data in the registers according to the context in which the data is used.

3.4.2.5 Pointers

Pointers contain addresses of data objects or functions. Pointers are represented in memory using 32 bits (4 bytes). In the small memory model, although pointers are represented in memory using 32 bits, only 16 bits are meaningful. The representation of pointers in memory and in the registers is exactly the same as for unsigned long integers, as shown in [Section 3.4.2.2, “Integers.”](#)

3.4.2.6 Bit-fields

Members of structures are always allocated on byte boundaries, and are aligned according to their fundamental base type. However, bit-fields in structures can be allocated at any bit and of any length not exceeding the size of a long word (32 bits). Signed and unsigned bit-fields are permitted and are sign extended when fetched. A bit-field of type `int` is considered signed.

Bit-fields are always allocated from the low-order end of a word (right to left or little-endian), even if the option `-be` is specified. Bit-field sizes are not allowed to cross a long word boundary.

In the following example, the structure `more` has 4-byte alignment and will have a size of 4 bytes. This is because the bit-fields in the structure are governed by the fundamental type `long` which requires a 4-byte alignment.

Example 3-19. Bit-field alignment to long word (1)

```
struct more {
    long first : 3;
    unsigned int second : 8;
};
```

The structure `less` shown in Example 3-20 requires only a one byte alignment because this is the requirement of the fundamental type `char` used in this structure.

Example 3-20. Bit-field alignment to character

```
struct less {
    unsigned char third : 3;
    unsigned char fourth : 8;
};
```

The alignments are driven by the underlying type, not the width of the fields. These alignments are to be considered along with any other structure members.

In Example 3-21 below, the structure `careful` requires a 4-byte alignment; its bit-fields require only a one byte alignment, but the field `fluffy` requires a 4-byte alignment because its fundamental type is `long`.

Example 3-21. Bit-field alignment to long word (2)

```
struct careful {
    unsigned char third : 3;
    unsigned char fourth : 8;
    long fluffy;
};
```

Fields within structures and unions begin on the next possible suitably aligned boundary for their data type. For fields that are not bit-fields, this is a suitable byte alignment. Bit-fields begin at the next available bit offset, with the following exception: the first bit-field after a member that is not a bit-field will be allocated on the next available byte boundary.

In the following example, the offset of the field `c` is one byte. The structure itself has 4-byte alignment and is four bytes in size because of the alignment restrictions introduced by using the `long` underlying data type for the bit-field.

Example 3-22. Bit-field offset

```
struct s {
    int bf: 5;
    char c;
}
```

3.4.3 Fractional and Integer Arithmetic

The ability to perform both integer and fractional arithmetic is one of the strengths of the SC100 C compiler.

Fractional arithmetic is typically required for computation-intensive algorithms such as digital filters, speech coders, vector and array processing, digital control, or other signal processing tasks. In this mode, the data is interpreted as fractional values, and the computations are performed interpreting the data as fractional. Fractional arithmetic examples are shown in Example 3-23.

Often, saturation is used when performing calculations in this mode to prevent the severe distortion that occurs in an output signal generated from a result where a computation overflows without saturation. Saturation can be selectively enabled or disabled so that intermediate calculations can be performed without limiting, and limiting is only done on final results.

Example 3-23. Fractional arithmetic examples

$0.5 * 0.25 \rightarrow 0.125$

$0.625 + 0.25 \rightarrow 0.875$

$0.125 / 0.5 \rightarrow 0.25$

$0.5 \gg 1 \rightarrow 0.25$

It is important to note that the notation used in Example 3-23 is for illustration purposes only, since C does not support the specification of fractional constants using floating-point notation. The compiler implements fractional arithmetic using intrinsic functions based on integer data types. For more information, see [Section 3.4.2.4, “Fractional representation,”](#) and [Section 3.4.4, “Intrinsic Functions.”](#)

Integer arithmetic is invaluable for controller code, array indexing and address computations, peripheral setup and handling, bit manipulation, and other general purpose tasks, as shown in Example 3-24.

Example 3-24. Integer arithmetic examples

$4 * 3 \rightarrow 12$

$1201 + 79 \rightarrow 1280$

$63 / 9 \rightarrow 7$

$100 \ll 1 \rightarrow 200$

Data in a memory location or register can be interpreted as fractional or integer, depending on the needs of a user's program. Table 3-4 shows how a 16-bit value can be interpreted as either a fractional or integer value, depending on the location of the binary point.

Table 3-4. Interpretation of 16-bit Data Values

Binary Representation ¹	Hexadecimal Representation	Integer Value (decimal)	Fractional Value (decimal)
0.100 0000 0000 0000	0x4000	16384	0.5
0.010 0000 0000 0000	0x2000	8192	0.25
0.001 0000 0000 0000	0x1000	4096	0.125
0.111 0000 0000 0000	0x7000	28672	0.875
0.000 0000 0000 0000	0x0000	0	0.0
1.100 0000 0000 0000	0xC000	-16384	-0.5
1.110 0000 0000 0000	0xE000	-8192	-0.25
1.111 0000 0000 0000	0xF000	-4096	-0.125
1.001 0000 0000 0000	0x9000	-28672	-0.875

1. Note: This corresponds to the location of the binary point when interpreting the data as fractional. If the data is interpreted as integer, the binary point is located immediately to the right of the LSB.

The following equation shows the relationship between a 16-bit integer and a fractional value:

$$\text{Fractional Value} = \text{Integer Value} / (2^{15})$$

There is a similar equation relating 40-bit integers and fractional values:

$$\text{Fractional Value} = \text{Integer Value} / (2^{31})$$

Table 3-5 shows how a 40-bit value can be interpreted as either an integer or fractional value, depending on the location of the binary point.

Table 3-5. Interpretation of 40-bit Data Values

Hexadecimal Representation	40-bit Integer in Entire Accumulator (decimal)	16-bit Integer in MSP (decimal)	Fractional Value (decimal)
0x0 4000 0000	1073741824	16384	0.5
0x0 2000 0000	536870912	8192	0.25
0x0 0000 0000	0	0	0.0
0xF C000 0000	-1073741824	-16384	-0.5
0xF E000 0000	-536870912	-8192	-0.25

The following code fragment illustrates the use of integer arithmetic:

Example 3-25. Integer arithmetic computation

```
a = a + b*c;
```

Example 3-26 provides an example of the use of an intrinsic function to implement fractional arithmetic.

Example 3-26. Fractional arithmetic computation

```
a = I_mac(a,b,c);
```

[Section 3.4.4, “Intrinsic Functions,”](#) describes the use of intrinsic functions in greater detail.

3.4.4 Intrinsic Functions

The compiler supports a large number of intrinsic (built-in) functions that map directly to SC100 assembly instructions. As C does not support fractional types and operations, these intrinsic functions enable fractional operations to be implemented using integer data types.

The syntax of the compiler group of intrinsic functions is structured for full compatibility with the ETSI and ITU reference implementations of bit-exact standards.

3.4.4.1 Data types for intrinsic functions

The following four data types are defined for specific use with intrinsic functions:

- Fractional short, a 16-bit fractional value mapped to a short, as described in [Section 3.4.2.4, “Fractional representation.”](#)
- Fractional long, a 32-bit fractional value mapped to a long, as described in [Section 3.4.2.4, “Fractional representation.”](#)
- Extended precision fractional, a 40-bit value which can only be used in intrinsic functions. See [Section 3.4.4.1.1, “Extended precision fractional,”](#) for details.
- Double precision fractional, a 64-bit value which can only be used in intrinsic functions. See [Section 3.4.4.1.2, “Double precision fractional,”](#) for details.

Extended and double precision fractional types enable algorithms to be defined which require precision larger than 32 bits. These data types can be used only with intrinsic functions and with assignments. Variables defined as extended and double precision fractionals cannot be used for standard arithmetical or other operations.

3.4.4.1.1 Extended precision fractional

The extended precision fractional (`word40`) is a 40-bit data type which occupies the entire Dn (40-bit) register, as shown in Figure 3-13:




Figure 3-13. Extended Precision Fractional—Dn Register Layout

This data type is mapped in the compiler as a structure containing two elements:

- A 32-bit integer placed to the right of the binary point.
- An 8-bit integer placed to the left of the binary point. These “guard bits” can be used to ensure a more accurate result when an overflow occurs.

When stored in memory, an extended precision fractional variable occupies 64 bits. The least significant 32 bits are stored in the first 32-bit word, and the 8 most significant guard bits are stored in the second 32-bit word in an undefined position.

See Table 3-6 on page 3-48 for a list of intrinsic functions for fractional arithmetic using guard bits.

See Example 3-28 on page 3-53 for an illustration of the use of intrinsic functions with extended precision fractional variables.

3.4.4.1.2 Double precision fractional

The double precision fractional data type (`word64`) consists of 64 bits, all of which are assumed to be to the right of the binary point. This data type is mapped in the compiler as a structure containing two 32-bit elements.

See Table 3-6 on page 3-48 for a list of intrinsic functions for double precision data types.

3.4.4.1.3 Fractional constants

Fractional constants require integer notation, since floating point notation is not supported. For example, to express the value `0.5` as a fractional constant, the integer representation in hexadecimal must be used in the source code, in this case `0x4000`. For further examples of fractional values and their corresponding hexadecimal representations, see Table 3-4 on page 3-44.

3.4.4.1.4 Initializing variables with fractional values

Variables can be initialized as fractional values, using the following macros:

- `WORD16` initializes a value as a fractional short.
- `WORD32` initializes a value as a fractional long.

For example, `short x = WORD16(0.5)` initializes `x` as a fractional short with the value `0x4000`.

3.4.4.2 Intrinsic function categories

The following categories of intrinsic functions are provided:

- Fractional arithmetic
- Long fractional arithmetic
- Double precision fractional arithmetic
- Extended precision fractional arithmetic, with guard bits
- Architecture primitives
- Architecture primitives that generate identical assembly instructions
- Bit reverse addressing

Table 3-6 on page 3-48 lists all the supported intrinsic functions by category, with a brief description of each function. See [Section 7.14, “Built-in Intrinsic Functions \(prototype.h\)”](#) on page 7-20, for more detailed information about each of the fractional arithmetic and long fractional arithmetic intrinsic functions.

Section 3.4.4.3, “[Intrinsic functions examples](#),” which follows Table 3-6, contains example code segments illustrating the use of a number of intrinsic functions.

Table 3-6. Intrinsic Functions

a) Fractional arithmetic

Intrinsic Function	Declaration	Description
add	short add(short,short)	Short add
sub	short sub(short,short)	Short subtract
mult	short mult(short,short)	Short multiply
div_s	short div_s(short,short)	Short divide
mult_r	short mult_r(short,short)	Multiply and round
L_mac	long L_mac(long,short,short)	Multiply accumulate
mac_r	short mac_r(long,short,short)	Multiply accumulate and round
L_msu	long L_msu(long,short,short)	Multiply subtract
msu_r	short msu_r(long,short,short)	Multiply subtract and round
abs_s	short abs_s(short)	Short absolute value
negate	short negate(short)	Short negate
round	short round(long)	Round
shl	short shl(short,short)	Short shift left
shr	short shr(short,short)	Short shift right
shr_r	short shr_r(short,short)	Short shift right and round
norm_s	short norm_s(short)	Normalize any fractional value
max	short max(short,short)	Maximum value of any two short fractional values
min	short min(short,short)	Minimum value of any two short fractional values
saturate	short saturate(short)	Short saturation

Table 3-6. Intrinsic Functions (Continued)*b) Long fractional arithmetic*

Intrinsic Function	Declaration	Description
L_add	long L_add(long,long)	Long add
L_sub	long L_sub(long,long)	Long subtract
L_mult	long L_mult(short,short)	Long multiply
extract_h	short extract_h(long)	Extract 16 MSB of long word
extract_l	short extract_l(long)	Extract 16 LSB of long word
L_deposit_h	long L_deposit_h(short)	Deposit short in MSB
L_deposit_l	long L_deposit_l(short)	Deposit short in LSB
L_abs	long L_abs(long)	Long absolute value
L_negate	long L_negate(long)	Long negate
norm_l	short norm_l(long)	Normalize any long fractional value
L_max	long L_max(long,long)	Maximum value of any two long fractional values
L_min	long L_min(long,long)	Minimum value of any two long fractional values
L_shl	long L_shl(long,short)	Long shift left
L_shr	long L_shr(long,short)	Long shift right
L_shr_r	long L_shr_r(long,short)	Long shift right and round
L_sat	long L_sat(long)	Long saturation

c) Double precision fractional arithmetic

Intrinsic Function	Declaration	Description
D_mult	Word64 D_mult(long,long)	Double precision multiply
D_mac	Word64 D_mac(Word64,long,long)	Double precision multiply accumulate
D_msu	Word64 D_msu(Word64,long,long)	Double precision multiply subtract
D_add	Word64 D_add(Word64,Word64)	Double precision add
D_sub	Word64 D_sub(Word64,Word64)	Double precision subtract
D_cmpeq	short D_cmpeq(Word64,Word64)	Double precision compare equal
D_cmpgt	short D_cmpgt(Word64,Word64)	Double precision compare greater than
D_sat	Word64 D_sat(Word64)	Double precision saturation
D_round	long D_round(Word64)	Double precision round
D_set	Word64 D_set(long,unsigned long)	Concatenate two longs into one double precision value
D_extract_l	unsigned long D_extract_l(Word64)	Extract 32 LSB of double precision value
D_extract_h	long D_extract_h(Word64)	Extract 32 MSB of double precision value

Table 3-6. Intrinsic Functions (Continued)*d) Extended precision fractional arithmetic (with guard bits)*

Intrinsic Function	Declaration	Description
X_mult	Word40 X_mult(short,short)	Short multiply to long long word
X_mac	Word40 X_mac(Word40,short,short)	Short multiply accumulate to long long word
X_msu	Word40 X_msu(Word40,short,short)	Short multiply subtract to long long word
X_set	Word40 X_set(char,unsigned long)	Concatenate char and unsigned long into one long long word
X_add	Word40 X_add(Word40,Word40)	Long add including guard bits
X_sub	Word40 X_sub(Word40,Word40)	Long subtract including guard bits
X_shl	Word40 X_shl(Word40,short)	Long shift left with guard bits
X_shr	Word40 X_shr(Word40,short)	Long shift right with guard bits
X_extract_h	short X_extract_h(Word40)	Extract 16 MSB of long long word
X_extract_l	short X_extract_l(Word40)	Extract 16 LSB of long long word
X_round	short X_round(Word40)	Round long long value
X_norm	short X_norm(Word40)	Normalize any long long fractional value
X_rol	Word40 X_rol(Word40)	Rotate left a long long word
X_ror	Word40 X_ror(Word40)	Rotate right a long long word
X_abs	Word40 X_abs(Word40)	Long absolute value with guard bits
X_sat	long X_sat(Word40)	Long saturation including guard bits
X_or	Word40 X_or(Word40,Word40)	Logical OR two long values with guard bits
X_trunc	long X_trunc(Word40)	Truncate guard bits
X_extend	Word40 X_extend(long)	Sign extend long value to include guard bits
X_cmpeq	short X_cmpeq(Word40,Word40)	Fractional compare equal with guard bits
X_cmpgt	short X_cmpgt(Word40,Word40)	Fractional compare greater than with guard bits

Table 3-6. Intrinsic Functions (Continued)

e) Architecture primitives

Intrinsic Function	Declaration	Description
L_rol	long L_rol(long)	Rotate left a long
L_ror	long L_ror(long)	Rotate right a long
mpyuu	long mpyuu(long, long)	Long multiply 16 LSB of two long words, treating both words as unsigned values
mpyus	long mpyus(long, long)	Long multiply 16 LSB of the first long word, treated as an unsigned value, by 16 MSB of the second long word, treated as signed
mpysu	long mpysu(long, long)	Long multiply 16 MSB of the first long word, treated as a signed value, by 16 LSB of the second long word, treated as unsigned
setnosat	setnosat()	Set saturation mode off
setsat32	setsat32()	Set saturation mode on
set2crm	set2crm()	Set rounding mode to two's-complement rounding mode
setcnvrm	setcnvrm()	Set rounding mode to convergent rounding mode

f) Architecture primitives that generate identical assembly instructions

Intrinsic Function	Declaration	Description
debug	void debug()	Enter <code>Debug</code> mode
debugev	void debugev()	Generate <code>Debug</code> event
mark	void mark()	If trace buffer enabled, write program counter to trace buffer
stop	void stop()	Enter <code>Stop</code> low power mode
trap	void trap()	Execute <code>Trap</code> exception
wait	void wait()	Enter <code>Wait</code> low power mode
ei	void ei()	Enable interrupts
di	void di()	Disable interrupts
illegal	void illegal()	Execute <code>illegal</code> exception

g) Bit reverse addressing

Intrinsic Function	Declaration	Description
InitBitReverse	InitBitReverse	Allocate a bit reverse iterator
BitReverseUpdate	BitReverseUpdate	Increment the iterator with bit reverse
EndBitReverse	EndBitReverse	Free bit reverse iterator

3.4.4.3 Intrinsic functions examples

The following example illustrates the use of a number of intrinsic functions.

Example 3-27. Intrinsic functions

```
#include <prototype.h>
void Iir(short Input[],
         short Coef[],
         short FiltOut[])
{
    long L_Sum;
    short int Stage, Smp;

    FiltOut[0] = Input[0];
    for (Smp = 1; Smp < S_LEN; Smp++)
    {
        L_Sum = L_msu(LPC_ROUND, FiltOut[Smp - 1], Coef[0]);
        for (Stage = 1; ((0 < (Smp - Stage)) && Stage < NP); Stage++)
            L_Sum = L_msu(L_Sum, FiltOut[Smp - Stage - 1], Coef[Stage]);

        L_Sum = L_shl(L_Sum, ASHIFT);
        L_Sum = L_msu(L_Sum, Input[Smp], 0x8000);
        FiltOut[Smp] = extract_h(L_Sum);
    }
}
```

Example 3-28 illustrates the use of extended precision variables and intrinsic functions using guard bits:

Example 3-28. Intrinsic functions using extended precision

```
#include <prototype.h>
docorr()
{
    int j, i;
    int shift_val;
    short corr_0;

    Word40 E_acc, E_sum;

    E_acc = X_extend(0);      E_sum = X_extend(0);

    for (i = 0; i < M1; i++)
    {
        for (j = 0; j < M2; j++)
            E_acc = X_mac (E_acc, sample[j], coeff[j] );

        L_sample[i] = X_sat(E_acc);

        E_acc = X_abs(E_acc);
        E_sum = X_add(E_sum, E_acc);
    }

    shift_val = X_norm(E_sum);
    corr_0 = 0;

    for (i = 0; i < M1; i++)
    {
        sample[i] = round (L_shr (L_sample[i], shift_val));
        corr_0 = sub (corr_0, sample[i]);
    }

    corr = corr_0;
}
```

3.4.5 Pragas

Pragas allow you greater control over your application, enabling you to give the compiler specific additional information about how to process certain statements. The pragmas that you specify in your code provide the compiler with context-specific hints which can save the compiler unnecessary operations, and help to further enhance the optimization process.

You can include as many pragmas as necessary in your source code. The sections that follow describe the syntax and placement rules for pragmas.

3.4.5.1 Syntax

The pragmas supported by the compiler have the following general syntax:

```
#pragma pragma-name [argument(s)]
```

One or more of the arguments may be optional. Arguments are comma-delimited.

Each pragma must fit into one line.

3.4.5.2 Placement

Each pragma is applicable only in a certain context, and must be placed accordingly. Four categories of pragmas can be defined according to the placement rules, as follows:

- Pragmas which apply to functions can appear only in the scope of the function, after the opening “{”.
- Pragmas which apply to statements must be placed immediately before the relevant statement, or immediately before any comment lines which precede the statement.
- Pragmas which apply to variables must follow the object definition, or any comment lines which follow that definition. Objects referred to by pragmas must be explicitly defined.

The pragmas supported by the compiler are listed in Table 3-7 on page 3-55. The sections that follow the table provide a brief summary and example of the syntax and use of each pragma. The detailed functioning of each pragma is described in [Chapter 5, “Optimization Techniques and Hints.”](#)

Table 3-7. Pragmas

<i>a) Function pragmas</i>	Description	Section	Page
<code>#pragma inline</code>	Forces function inlining.	3.4.5.3.1	3-56
<code>#pragma noinline</code>	Disables function inlining.	3.4.5.3.1	3-56
<code>#pragma save_ctxt</code>	Forces save and restore of all registers that are used in this procedure.	3.4.5.3.2	3-56
<code>#pragma external func</code> [<i>name = string,</i> <i>convention = number,</i> <i>nosideeffects</i>]	Defines a function as external to the C application, or as a function that can be called from outside the application.	3.4.5.3.3	3-57
<code>#pragma interrupt func</code>	Defines the specified function as an interrupt handler.	3.4.5.3.4	3-58
<i>b) Pragmas which apply to statements</i>	Description	Section	Page
<code>#pragma profile value</code>	Sets profiling information for a statement.	3.4.5.4.1	3-58
<code>#pragma loop_count</code> (<i>lower_bound,</i> <i>upper_bound,</i> <i>{2/4},</i> <i>remainder</i>)	Specifies the minimum and maximum limits for a loop, the loop count divider (2 or 4), and the use of the remainder.	3.4.5.4.2	3-59
<i>c) Pragmas which apply to variables</i>	Description	Section	Page
<code>#pragma align var_name {4/8}</code>	Forces stricter alignment on an object. Needed for paired moves.	3.4.5.5.1	3-61
<code>#pragma align *var_name {4/8}</code>	Indicates that the address of the variable referenced by a pointer is aligned as specified.	3.4.5.5.1	3-61

3.4.5.3 Pragas which apply to functions

The pragmas in this category provide additional information about specific functions, and are defined in the scope of the function to which they apply, directly after the “{” which marks the start of the scope.

3.4.5.3.1 Forcing or disabling function inlining

Inlining enables the compiler to improve optimization by replacing a function call by the entire function. For very small functions, for example, where the overhead of the function call is greater than the size of the function itself, this can be very efficient. For more information about function inlining, refer to [Chapter 5, “Optimization Techniques and Hints.”](#)

You can use `#pragma inline` to force the compiler to inline a specific function, or `#pragma noline` to prevent the compiler from inlining a certain function. In the code segment shown in Example 3-29, any calls to the function which follows `#pragma noline` will not be inlined.

Example 3-29. #pragma noline

```
static int proc_30(int a)
{
#pragma noline
    int tab_30[1000];

    tab_30[0] = 4*a;
    return(tab_30[0]);
}
```

3.4.5.3.2 Saving the entire context of the system

During normal processing, the compiler saves the contents of registers that have been changed, and any other essential data. You can force the compiler to save the entire context of the machine, including all registers that are used in this procedure, so that it can be restored if necessary to its previous state, at the exact point at which the specific function started to execute.

Using `#pragma save_ctxt` to save the entire system status can incur a large overhead, and should only be used where absolutely necessary.

The following example illustrates the use of `#pragma save_ctxt` to force the compiler to save the complete machine context upon entry to the specified function.

Example 3-30. #pragma save_ctxt

```
void EntryPoint()
{
#pragma save_ctxt
...
}
```

3.4.5.3.3 Defining a function as external

When the compiler encounters an unresolved function call, it assumes by default that this is a call to an external function that exists outside the application. The pragma `#pragma external` enables you to:

- Confirm this assumption, by informing the compiler that the call is to an external function defined outside the application
- Define the function as an internal function that can be called from outside the application

The effect of the pragma depends on its placement, as described below:

- If `#pragma external` is specified in the global scope, the compiler does not expect to find the body of the function within the current application. The compiler uses standard calling conventions to call the function, and does not issue warnings for unresolved references. Specifying `#pragma external` in the global scope is valid only in global optimization mode.
- If `#pragma external` is specified within the function scope, followed by the body of the defined function, the compiler recognizes this as an internal function that can be called from outside the application.

The following optional parameters can be specified with `#pragma external`:

- Specify `name = string` to provide a specific function name, to override the default linkage name allocated to the function.
- Define `convention = number` to select the calling convention to be used instead of the default standard convention. See [Chapter 6, “Runtime Environment,”](#) for further information about calling conventions.
- Specify `nosideeffects` if the function does not change any variable values in the application, and can be moved or duplicated in other parts of the application without making any changes.

When `nosideeffects` is specified, the compiler does not need to make worst case assumptions about any possible impact that the function may have within the application.

In the first part of Example 3-31, `printf` is defined as an external function that does not exist within the application, and that has no effect on any variables in the application. In the second part of the example, the function `ICanBeCalled` is defined inside the application and may be called by external function calls. This function therefore has to obey the standard calling conventions.

Example 3-31. `#pragma external`

```
extern void printf();
#pragma external printf [nosideeffects]

void main()
{
    printf("Hello there\n");
}

void ICanBeCalled(int X, int Y)
{
    #pragma external ICanBeCalled [name = "xyz"]
    ...
}
```

3.4.5.3.4 Defining a function as an interrupt handler

A function that operates as an interrupt handler differs from other functions in three basic respects:

- It must save and restore all resources that it uses, as it can be called at any time an interrupt occurs, and cannot assume any conventions.
- It runs in “exception” mode, which forces the compiler to generate instructions that are slightly different from the instructions issued in normal mode.
- It cannot be passed parameters nor return a value.

You can use `#pragma interrupt` to define a function as an interrupt handler, as shown in the following example.

Example 3-32. #pragma interrupt

```
void IntHandler();
#pragma interrupt IntHandler
extern long Counter;
void IntHandler()
{
Counter++;
}
```

3.4.5.4 Pragmas which apply to statements

Pragmas which apply to statements are placed immediately before the relevant statement.

3.4.5.4.1 Specifying a profile value

By default, the profiler provided with the compiler enables it to make the necessary assumptions about the number of times to execute a given statement. You can specify `#pragma profile`, followed by a value and immediately preceding a statement, to specify to the compiler the exact number of times that the statement executes.

In Example 3-33, the value following `#pragma profile` notifies the compiler that the loop executes only 10 times. If `#pragma profile` is not specified, the compiler assumes that, since this is a loop with dynamic bounds, the loop executes 25 times (the default). It is important to note that this assumption affects the optimization of the program, and not its correctness.

Example 3-33. #pragma profile with constant value

```
#include <prototype.h>
int energy (short block[], int N)
{
    int i;

    long int L_tmp = 0;

    for (i = 0; i < N; i++)
#pragma profile 10
        L_tmp = L_mac (L_tmp, block[i], block[i]);

    return round (L_tmp);
}
```

With `if-then-else` constructs, `#pragma profile` can be used to inform the compiler which branch executes more frequently, and the frequency ratio between the two branches, meaning the number of times one branch executes in relation to the other.

In Example 3-34, the two `#pragma profile` statements have the values 5 and 50. These values notify the compiler that the `else` branch section executes 10 times more frequently than the first (implied `then`) section. When used in this way, the exact `#pragma profile` values are not significant, since they indicate the frequency ratio, and not the absolute values. In this example, the values 1 and 10 would convey the same information.

Example 3-34. `#pragma profile` with frequency ratio

```
#include <prototype.h>
int energy (short block[], int N)
{
    int i;
    long int L_tmp = 0;

    if ( N>50)
#pragma profile 5
        for (i = 0; i < 50; i++)
            L_tmp = L_mac (L_tmp, block[i], block[i]);
    else
#pragma profile 50
        for (i = 0; i < N; i++)
            L_tmp = L_mac (L_tmp, block[i], block[i]);

    return round (L_tmp);
}
```

3.4.5.4.2 Defining a loop count

The compiler tries to evaluate the number of times a loop iterates using the static information available. In cases where this static information is not supplied to the compiler, if you know the upper and lower limits of a loop, you can use `#pragma loop_count` to provide these values. Supplying such information, which cannot always be discerned automatically by the compiler, enables generation of more efficient code.

Similarly, specifying a divider for the loop count enables the optimizer to unroll loops in the most efficient way. The loop count can be divided by either 2 or 4, corresponding to the number of execution units. You can also instruct the compiler whether to use the remainder, if there is one following division of the loop count, to execute the loop an additional number of times.

The syntax of `#pragma loop_count` is:

```
#pragma loop_count (lower_bound, upper_bound, [{2/4}, [remainder]])
```

Define a value for `lower_bound` for the minimum number of times the loop will iterate, and a value for `upper_bound` for the maximum number of times.

The divider parameter is optional. Only the values 2 or 4 may be specified as the divider.

To specify that a remainder should be used for the loop count, specify a value for `remainder`. The `remainder` argument is only valid if a value has been specified for the divider.

The pragma `#pragma loop_count` must be placed inside the loop to which it relates, and outside any nested loops which the loop contains.

In Example 3-35, the loop will always iterate at least 4 times and at most 512 times. The iteration count will always be divisible by 4. As no remainder is specified, any remainder from the division will be disregarded.

Example 3-35. #pragma loop count

```
void correlation2 (short vec1[], short vec2[], int N, short *result)
{
    long int L_tmp = 0;
    int i;

    for (i = 0; i < N; i++)
        #pragma loop_count (4,512,4)
        L_tmp = L_mac (L_tmp, vec1[i], vec2[i]);

    *result = round (L_tmp);
}
```


3.4.5.5 Pragmas which apply to variables

These pragmas are placed immediately after the definition of the object(s) to which they refer. Objects referred to by pragmas must first be explicitly defined.

3.4.5.5.1 Alignment of variables

Objects are usually aligned according to their size, as described in [Section 3.4.2, “Types and Sizes.”](#) The default alignment for arrays is determined by their base type.

An array may need to be aligned to a specified value before it can be passed to an external function. The pragma `#pragma align` can be used to force the alignment of arrays passed to an external function, to meet the specific alignment requirements of the function.

To force the alignment of an array before passing it to an external function, specify `#pragma align`, followed by the defined array object, and either the value 4 for 4-byte (32-bit double word) alignment or 8 for 8-byte (64-bit quad word) alignment.

Certain instructions, such as `move .2w` and `move .4w`, which move words in pairs, may require alignment to be applied that is stricter than the alignment defined for the data types involved.

In certain cases, the compiler cannot assess the alignment for dynamic objects and has to assume that the objects have the alignment requirements for their base type. As a result, the compiler cannot use the multiword move instructions for these objects. By specifying the exact alignment for one or more objects, you can enable the compiler to use these multiword moves and generate more efficient code.

You can use the pragma `#pragma align` to provide the compiler with specific alignment information about pointers to arrays, in order to enable the compiler to use multiword move instructions.

To inform the compiler that the address of an array is aligned as required for multiword moves, specify `#pragma align`, followed by the pointer to the array object, and either the value 4 for 4-byte alignment or 8 for 8-byte alignment. When using `#pragma align` in this way, you should ensure that the object is in fact aligned as required, since this form of the pragma does not force the alignment.

In the first part of Example 3-36, array `a` is forced to 8-byte alignment before being passed to the external function `Energy`. The second part of the example informs the compiler that both input vectors are aligned to 32 bits. The instruction `move.2f` may be used here.

Example 3-36. #pragma align

```
#include <prototype.h>
short a[10];
#pragma align a 8

extern int Energy( short a[] );
int foo()
{
    return Energy(a);
}
short Cor(short vec1[], short vec2[], int N)
{
#pragma align *vec1 4
#pragma align *vec2 4

    long int L_tmp = 0;
    long int L_tmp2 = 0;
    int i;

    for (i = 0; i < N; i += 2)
        L_tmp = L_mac(L_tmp, vec1[i], vec2[i]);
        L_tmp2 = L_mac(L_tmp2, vec1[i+1], vec2[i+1]);

    return round(L_tmp + L_tmp2);
}
```

3.4.6 Predefined Macros

The compiler shell maintains a number of predefined macros, including standard C macros, and additional macros which are specific to the SC100 C compiler and the SC100 architecture. Table 3-8 lists these predefined macros.

Table 3-8. Predefined Macros

Macro Name	Description
<code>__LINE__</code>	The line number of the current source line.
<code>__FILE__</code>	The name of the current source file.
<code>__DATE__</code>	The compilation date, as a character string in the form Mmm dd yyyy e.g. Jan 23 1999.
<code>__TIME__</code>	The compilation time, as a character string in the form hh:mm:ss.t
<code>__STDC__</code>	Decimal constant 1, indicating ANSI conformance.
<code>__STDC_VERSION__</code>	Defined in ANSI C mode as 199409L.
<code>__SIGNED_CHARS__</code>	Defined when char is signed by default
<code>__VERSION__</code>	The version number of the compiler, as a character string in the form nn.nn.
<code>__INCLUDE_LEVEL__</code>	Decimal constant, indicating the current depth of file inclusion.
<code>__ENTERPRISE_C__</code>	Defined for use with the Enterprise compiler. If your source file may be compiled with other compilers apart from the Enterprise, this macro should be included in a conditional statement to ensure that the appropriate commands are activated, for example: <pre>#ifdef __ENTERPRISE_C_ (Enterprise-specific commands) #else #endif</pre>
<code>__SC100__</code>	Defined for use with all compilers based on the SC100 architecture. If your source file may be compiled with other compilers apart from those based on the SC100 architecture, this macro should be included in a conditional statement to ensure that the appropriate commands are activated, as shown in the following example: <pre>#ifdef __SC100_ (SC100-specific commands) #else #endif</pre>
<code>__SC140__</code>	The architecture variant, which specifies the number of MAC units to be used by the compiler.

Chapter 4

Interfacing C and Assembly Code

The SC100 C compiler supports interfacing between C source code and assembly code, enabling access to functionality not provided by C. This chapter describes the features of this interface and provides instructions, guidelines, and examples.

The following sections are contained in this chapter:

- [Section 4.1, “Inlining a Single Assembly Instruction,”](#) explains how to use individual assembly instructions in your C source code.
- [Section 4.2, “Inlining a Sequence of Assembly Instructions,”](#) describes how to embed an assembly function consisting of a sequence of assembly instructions into your C code.
- [Section 4.3, “Calling an Assembly Function in a Separate File,”](#) explains how an assembly function that is contained in a separate file can be used in conjunction with your C source files.
- [Section 4.4, “Including Offset Labels in the Output File,”](#) describes the use of symbolic offsets for C data structures in the assembly output file.

4.1 Inlining a Single Assembly Instruction

A single assembly instruction can be inlined in a sequence of C statements and compiled by the compiler. To ensure successful compilation of an inlined assembly instruction, note the following guidelines:

- The compiler passes an inlined instruction to the assembly output file in the form of text, and therefore has no knowledge of the contents or side effects of the instruction. It is important that you ensure that there is no risk of the instruction affecting the C and/or assembly environment and producing unpredictable results. For example, you should not use an inlined assembly instruction to change the contents of registers, as the compiler has no knowledge of such changes. Similarly, you should not include any jumps or labels which access the C code and may affect the correctness of the tracking algorithms.
- Inlined assembly code instructions are ignored by the optimizer.
- Since the compiler treats the assembly instruction as a string of text, it cannot perform any error checking on the instruction. Check the syntax and text of the instruction carefully prior to compilation. Errors in assembly code are identified only at the assembly stage of the compilation process.
- A single inlined assembly instruction cannot reference a C object. The only way to reference a C object in assembly code is by inlining a sequence of assembly instructions, as described in [Section 4.2, “Inlining a Sequence of Assembly Instructions.”](#)

To inline a single assembly instruction, use the `asm` statement. The syntax is as for a standard function call, with one argument enclosed in double quotation marks, as shown below in Example 4-1.

Example 4-1. Inlining a single assembly instruction

```
asm("wait");
```

4.2 Inlining a Sequence of Assembly Instructions

It is possible to use assembly code that references C objects, by defining a separate function that consists of a sequence of assembly instructions, and inlining this in your C code. Such a function is implemented entirely in assembly and may not include C statements, but can accept parameters referenced by the assembly code.

4.2.1 Guidelines for Inlining Assembly Code Sequences

The following guidelines are similar to those for the inlining of individual assembly instructions, described in [Section 4.1, “Inlining a Single Assembly Instruction,”](#) and apply also to the use of inlined sequences of assembly code:

- The compiler passes a sequence of inlined instructions to the assembly output file as a string of text, and therefore has no knowledge of the contents or side effects of the instructions. It is important that you ensure that the assembly function does not affect the C and/or assembly environment and does not produce unpredictable results. For example, do not use inlined assembly instructions to change the contents of registers, and do not alter the sequence of C code instructions by specifying jumps, as the compiler has no knowledge of such changes.
- Functions based on inlined sequences of assembly code cannot be used by the optimizer, and are ignored during optimization. Avoid using assembly-based functions if a C alternative is available, in order to ensure maximum optimization of the code.
- The compiler performs no error checking on the sequence of assembly instructions. Assembly code errors are identified only at the assembly stage of the compilation process.

The guidelines listed below apply specifically to the use of inlined sequences of assembly code:

- When passing parameters to an inlined sequence of assembly instructions, registers are not automatically allocated. You must specify for each parameter the register in which the parameter enters or exits the function. There is no need to save and restore the registers before and after the function.
- The compiler is unable to deduce whether an inlined function is likely to affect the application, for example, if it modifies global variables. It is important that you provide the compiler with this information if there is a possibility that the function may have any side effects.
- A function that is initially defined as stand-alone may in certain circumstances be included in another sequence of instructions. Inlined functions should therefore not use statements such as `RTS`. If the function is used in a sequence of instructions, the compiler adds the necessary return statements automatically.

- Local variables are not automatically allocated by the compiler for use by assembly functions. If the function requires the use of local variables, you must allocate these specifically on the stack or define them as static variables.
- Assembly functions defined as a sequence of instructions can access global variables in the C source code, since these are static by definition.

4.2.2 Defining an Inlined Sequence of Assembly Instructions

When defining a sequence of inlined assembly instructions, you define the header for the function before the body of the instructions, and you specify the registers to be used by each parameter. You can define a list of read parameters, a list of write parameters, and/or a list of modified registers, as appropriate.

The syntax for inlining a sequence of assembly instructions is as follows:

```
asm <func prototype>
{
asm_header
    optional arg binding
    optional return value
    optional read list
    optional write list
    optional modified reg list
asm_body
    <asm code>
asm_end
}

optional arg binding
    .arg
        <ident> in <reg>;
        <ident> in <reg>;
        ...
optional return value
    return in <reg>

optional read list:
    .read <ident>,<ident>,...;

optional write list:
    .write <ident>,<ident>,...;

optional modified reg list:
    .reg <reg>, <reg>, ...;
```

The following syntax conventions apply:

- Identifiers must have the prefix `_` (underscore).
- Registers must have the prefix `$` (dollar sign).
- Labels must have the suffix `.` (period).

Example 4-2 shows the syntax for an inlined assembly function that takes two arguments as input parameters and returns one value. The first argument is passed in the register `d0`, and the second parameter in the register `r1`. The result is returned in `d0`.

Example 4-2. Inlining syntax

```
asm int t6( int param1, int *param2)
{
asm_header
.arg
    _param1 in $d0;
    _param2 in $r1;
return in $d0;
.reg $d0,$d1,$r1;
asm_body
    move.l (r1),d1
    add    d0,d1,d0
asm_end
}
```

In Example 4-3, the function `t6` accepts two parameters, an integer `p1` passed in register `d14`, and a pointer `p2` passed in `r7`. The result of the function is returned in `d14`.

Example 4-3. Simple inlined assembly function

```
#include <stdio.h>
int A[10] = {1,2,3,4,5,6,7,8,9,0};
asm int t6(int p1, int *p2)
{
asm_header
.arg
    _p1 in $d14;
    _p2 in $r7;
return in $d14;
.reg $d14,$d1,$r7;
asm_body
    move.l (r7),d1
    add    d14,d1,d14
asm_end
}
int main()
{
    int k = 8;
    int s;

    s = t6(k,&A[3]);

    printf("S= %d\n",s);

return s;
}
```

Example 4-4 shows the use of labels and hardware loops within inlined assembly functions. You should use hardware loops within assembly functions only if you know that the loop nesting is legal. In this example, the function is called from outside a loop, and the use of hardware loops is therefore allowed.

Example 4-4. Inlined assembly function with labels and hardware loops

```
#include <stdio.h>
char sample[10] = {9,6,7,1,0,5,1,8,2,6};
int status;

asm char t7(int p)
{
    asm_header
    .arg
    _p in $d7;
    return in $d8;
    .reg $d7,$d8,$r1;
    asm_body

    clr    d8                move.l    #_sample,r1                doen3
d7
    dosetup3 _L10
    loopstart3
_L10:
    move.b (r1),d1
    add    d8,d1,d8
    inc    d1
    move.b d1,(r1)+
    loopend3

    asm_end
}

int main()
{
    int m = 8;
    int s,i;

    for(i=0;i < 10;i++) {
        sample[i] *= 2;
        printf("%d ",sample[i]);
    }
    printf("\n");
    s = (int)t7(m);
    printf("S= %d\n",s);

    for(i=0;i < 10;i++)
        printf("%d ",sample[i]);
    printf("\n");
    return 1;
}
```

Example 4-5 shows how global variables are referenced within an inlined assembly function. Global variables are accessed using their linkage name, which is by default the variable name prefixed by the character `_` (underscore). The variables `vector1` and `vector2` are therefore accessed within the function as `_vector1` and `_vector2` respectively.

Example 4-5. Referencing global variables in an inlined assembly function

```
#include <stdio.h>

short vector1[] = {1,2,3,4,5,6,7,8,9,10,11,12,13,14,15};
int vector2[] = {11,12,13,14,15,16,17,18,19,1,2,3,4,5,6};
short result_1=0;
int result_2=0;

asm void test(int n, short *r1,int *r2)
{
    asm_header
    .arg
        _n in $r1;
        _r1 in $r3;
        _r2 in $r7;
    .reg $d0,$r1,$r6,$r11,$r3,$r7;
asm_body
    move.l    #_vector1,r6
    move.l    #_vector2,r11
    addl1a   r1,r6
    addl2a   r1,r11
    move.w   (r6),d0
    asrr     #<2,d0
    move.w   d0,(r3)           move.l           (r11),d1
    asl     d1,d2
    move.l   d2,(r7)
asm_end
}

int main(void)
{
    test(12,&result_1,&result_2);

    printf("Status = %d %d\n",(int)result_1, result_2);

    return (int)result_2;
}
```

4.3 Calling an Assembly Function in a Separate File

The compiler supports calls to assembly functions that are contained in separate files, and enables you to integrate these files with your C application.

To include a call to an assembly function in your program, follow the steps described below:

1. Write the assembly function in a separate file from your C source files. Use the standard calling conventions, as described in [Chapter 6, “Runtime Environment.”](#)
2. If required, assemble the file. This step is optional.
3. In your C source file, define the assembly function as an external function.
4. Specify both the C source file and the assembly file as input files in the shell command line to integrate the files during compilation.

The following examples show how a segment of C code calls a function that performs an FFT algorithm implemented in assembly.

4.3.1 Writing the Assembly Code

Example 4-6 shows the assembly code for the FFT algorithm, in the file `fft.s1`.

Example 4-6. Assembly function in a separate file

```

;
; extern void fft(short *, short*);
;
; Parameters:  pointer to input buffer in r0
;             pointer to output buffer in r1
;
_fft:
    push        d6            push        d7            ;Save and restore
d6, d7, r6, r7, according to
    push        r6            push        r7            ;calling
conventions.

    < implementation of FFT algorithm >

    pop         r6            pop         r7
    pop         d6            pop         d7

    rts

```

4.3.2 Calling the Assembly Function

The C code that calls the FFT function is shown in Example 4-7. This source code is saved in the file `test_fft.c`.

Example 4-7. C code calling assembly function

```
#include <stdio.h>
extern void fft(short *, short*);
#pragma external fft

short in_block[512];
short out_block[512];

int in_block_length, out_block_length;

void main()
{
    int i;
    FILE *fp;
    int status;

    in_block_length=512;
    out_block_length=512;

    fp=fopen("in.dat","rb");
    if( fp== 0 )
    {
        printf("Can't open parameter file: input_file.dat\n");
        exit(-1);
    }

    printf("Processing function fft \n");

    while ((status=fread(in_block, sizeof(short), in_block_length, fp)) ==
in_block_length)
    {
        fft(in_block,out_block);
    }
}
```

4.3.3 Integrating the C and Assembly Files

Example 4-8 shows how the two input files are specified in the shell command line:

Example 4-8. Integrating C and assembly files

```
ccsc100 -o test_fft.eld test_fft.c fft.sl
```

4.4 Including Offset Labels in the Output File

In some cases when assembly functions are called, data structures need to be shared between the C source code and the assembly code. In the following example, the layout of the structure `complex` needs to be used by the assembly code.

Example 4-9. Data structure shared between C and assembly

```
struct complex
{
    short r;
    short i;
};

struct complex CVEC1, CVEC2;
volatile struct complex res;

void main()
{
    cmpy (&CVEC1, &CVEC2, &res);
}
```

The `-do` option in the shell command line instructs the compiler to include the details of C data structures in the output assembly file. You can specify this as an additional option in the command line, as shown in Example 4-10:

Example 4-10. Specifying the output of offset information

```
ccsc100 -o test.eld test.c cmpy.sl -do
```

When the `-do` option is specified, the output file shows the offsets for all field definitions in each data structure defined in the C source code. The symbolic label is composed of:

`<module name>_<structure name>_<field name>`, as shown in the following example:

Example 4-11. Data structure offsets in the assembly output file

```
test_complex_r    equ        0
test_complex_i    equ        2
```

The symbolic labels in the output file can be used in the assembly code, making the code more readable, as illustrated in Example 4-12. Using these symbolic labels also makes maintenance of the assembly code easier when changes are made to the C code.

Example 4-12. Using symbolic offsets in assembly code

```
=====
; Function cmpy
;
; Parameter x      passed in r0
; Parameter y      passed in r1
; Parameter result passed in (sp-12)
;=====

        global _cmpy
        align   2

_cmpy
[
    move.2f (r0),d0d1
    move.2f (r1),d2d3
]
[
    mpy d0,d2,d5          mpy d0,d3,d7
    macr -d1,d3,d5        macr d1,d2,d7          move.l(sp-12),r2
]
    rtsd
    moves.f d5,(r2+test_complex_r)  moves.f d7,(r2+test_complex_i)

```

Chapter 5

Optimization Techniques and Hints

This chapter explains how the SC100 optimizer operates, and describes the optimization levels and individual optimizations which can be applied. The following sections are included in this chapter:

- [Section 5.1, “Optimizer Overview,”](#) provides a general description of the optimizer, illustrates how the optimizer transforms code, and outlines the available optimization levels, modes and options.
- [Section 5.2, “Using the Optimizer,”](#) explains how to invoke the optimizer, and how to achieve the required results for your application.
- [Section 5.3, “Optimization Types and Functions,”](#) describes the individual optimizations in detail.
- [Section 5.4, “Guidelines for Using the Optimizer,”](#) provides advice on ways to write source code that can make the best use of the optimizer’s capabilities.
- [Section 5.5, “Optimizer Assumptions,”](#) describes the assumptions made by the optimizer in different circumstances.

5.1 Optimizer Overview

The SC100 optimizer converts preprocessed source files into assembly output code, applying a range of code transformations which can significantly improve the efficiency of the executable program. The goal of the optimizer is to produce output code which is functionally equivalent to the original source code, while improving its performance in terms of execution time and/or code size.

5.1.1 Code Transformations

The optimizer is extremely versatile, and can transform the code in a number of ways to achieve optimal results. These code transformations include:

- Substituting instructions with more efficient code
- Removing redundant instructions
- Inserting instructions to simplify operations
- Hoisting or lowering instructions to reduce unnecessary operations

Example 5-1 below illustrates each of these instruction transformations:

Example 5-1. Instruction transformation

<u>Transformation type</u>	<u>Before optimization</u>	<u>After optimization</u>
Substitution:	<code>move.w #0,d0</code>	<code>clr d0</code>
Removal:	<code>move.w #0,d0</code> <code>move.w (r1),d0</code>	<code>move.w (r1),d0</code>
Insertion and removal:	<code>for(i=0; i<3; i++)</code> <code>foo();</code>	<code>foo();</code> <code>foo();</code> <code>foo();</code>
Hoisting:	<code>if(a<3)</code> <code> Tbit = TRUE;</code> <code>else</code> <code> Tbit = FALSE;</code>	<code>Tbit = FALSE;</code> <code>if(a<3)</code> <code> Tbit = TRUE;</code>

5.1.2 Basic Blocks

The majority of the code transformations operate on basic blocks of code. A basic block of code is a linear sequence of instructions for which there is only one entry point and one exit point. There are no branches in a basic block. In general, bigger basic blocks enable better optimization, since the scope for further optimization is increased.

5.1.3 Linear and Parallelized Code

The optimizer can produce code that takes full advantage of the multiple execution units provided by the SC100 architecture.

Executable programs process instructions in the form of execution sets, with one execution set per cycle. The optimizer can increase the number of instructions in an execution set, enabling two or more execution units to process instructions in parallel, in the same cycle. In this way, linear code is transformed into parallelized code:

- **Linear code** uses only one execution unit, regardless of the number of units available. Each execution set consists of one instruction only.
- **Parallelized code** execution sets can comprise multiple instructions which execute in parallel using the available number of execution units. Parallelized code executes faster and more efficiently than linear code.

Figure 5-1 illustrates the transformation of linear code, comprising a series of single instruction execution sets, into parallelized code, which consists of execution sets containing one or more instructions each:

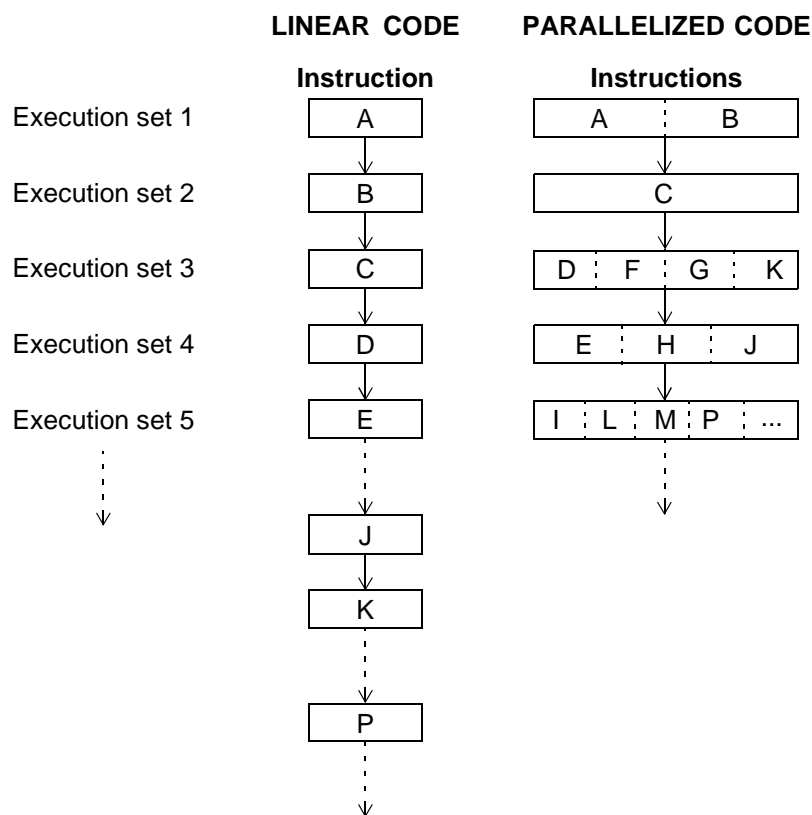


Figure 5-1. Linear and Parallelized Code

Dependencies between instructions can restrict the level of parallelization that the optimizer can achieve. For more information, see [Section 5.3.1, “Dependencies and Parallelization,”](#) and [Section 5.4, “Guidelines for Using the Optimizer.”](#)

5.1.4 Optimization Levels and Options

Three basic optimization levels are provided, all of which maintain a balance between space and time optimizations:

- Level 0 compiles the fastest and produces the slowest output as linear code.
- Level 1 takes longer to compile, applies high-level, machine independent optimizations and produces optimized linear code.
- Level 2 (the default) compiles more slowly than Level 1, applies all high-level optimizations, as well as all low-level, machine dependent optimizations, and can produce faster, parallelized code.

Two additional options enable you to apply both high and low-level optimizations, while weighting the optimization process in favor of either program speed or size:

- The speed optimization option maximizes execution speed at the expense of size. Programs or modules that have been optimized for speed execute faster, but consume more memory.
- The space optimization option minimizes size at the expense of speed. Programs or modules that have been optimized for space require a smaller amount of memory but run comparatively slower.

Only one of the above six optimization options may be selected for each compilation.

The optimizer can apply any one of the above options in either global or non-global optimization mode.

Global optimization is a complex process which requires significantly more compilation time than non-global optimization. In global mode the optimizer applies the required level of optimization across all the files in the application at the same time, and as a result produces the most efficient program code.

Global optimization is generally applied at the end of the development cycle, after all source files have been compiled and optimized individually or in groups. By default, the optimizer operates in non-global optimization mode.

Table 5-1 below summarizes the optimization options. This table also includes a cross-reference to the sections containing detailed descriptions of the individual optimizations applied by each of the options.

Table 5-1. Optimization Options Summary

Option	Description	Benefits	Section	Page
-O0 (Level 0)	<ul style="list-style-type: none"> Performs constant folding only. Outputs non-optimized linear code. 	<ul style="list-style-type: none"> Compiles fastest. 	5.3.2.5	5-17
-O1 (Level 1)	<ul style="list-style-type: none"> Performs all high-level (non-parallelized) optimizations, such as function inlining. Omits all low-level optimization steps. Outputs optimized linear code. 	<ul style="list-style-type: none"> Compiles faster than option -O2 (the default). Produces faster programs than option -O0. Generates assembly code which correlates clearly with the C source code, and can assist debugging. 	5.3.2	5-9
-O or -O2 (Level 2) (Default)	<ul style="list-style-type: none"> Performs all high-level and low-level optimizations. Outputs optimized non-linear assembly object code. 	<ul style="list-style-type: none"> Takes advantage of parallel execution units and can produce faster programs than option -O1. Compiles faster than option -Og (global optimization). 	5.3.3	5-19
-Os	<ul style="list-style-type: none"> Performs almost all optimizations, with more emphasis on reducing code size than on optimization for speed. 	<ul style="list-style-type: none"> Produces optimized assembly object code which is small as opposed to fast. 	5.3.5	5-30
-Og	<ul style="list-style-type: none"> Performs global optimization. Can be specified together with any of the other optimization options. Produces the most efficient results when specified together with the -O2 (default) option. Compiles significantly slower than the other options. 	<ul style="list-style-type: none"> Takes advantage of the visibility of all files to implement the specified optimization level across the entire application. Produces the fastest runtime code. 	5.3.6	5-30

5.2 Using the Optimizer

By default, the compiler optimizes all source code files using Level 2 optimization in non-global mode. You can choose to optimize your source code at the level that you require at each stage of program development, and you can optimize individual sections of the program according to their purpose in the application. For example, you may wish to prepare your application as follows:

- **During initial development stages:** Use the default Level 2 optimization to compile your source code files, individually or in groups. If required, optimize certain sections of the application for maximum speed, and optimize other sections for size, to reduce the memory space they occupy.
- **During final development stages:** Select Level 2 and global mode, in order to apply all optimizations globally across the entire application. The compilation is slower, but produces the most effective optimization results.

You select the optimization level and mode to be applied by specifying one or more options in the shell command line, as described below in [Section 5.2.1, “Invoking the Optimizer.”](#)

5.2.1 Invoking the Optimizer

The optimizer can be invoked by including the required option(s) in the shell command line or command file, as illustrated in the examples that follow. For more detailed information about the use and syntax of the shell command line, refer to [Section 3.2, “Invoking the Shell,”](#) on page 3-10.

The command line shown in Example 5-2 invokes the optimizer with one input source file, and the default optimization settings. The optimizer applies Level 2 optimizations in non-global mode, with a balance between space and time optimizations.

Example 5-2. Invoking the optimizer with default settings

```
ccsc100 -o file.eld file.c
```

Example 5-3 shows how to invoke the optimizer with the Level 1 option, to apply high-level optimizations only. The optimizer maintains a balance between time and space optimizations, and operates in non-global mode.

Example 5-3. Invoking the optimizer for high-level optimizations only

```
ccsc100 -O1 -o file.eld file.c
```

In Example 5-4, the optimizer applies all optimizations in non-global mode, with the emphasis on increasing the speed of the program at the expense of its size.

Example 5-4. Invoking the optimizer for time optimization

```
ccsc100 -Ot -o file.eld file.c
```

The command line shown in Example 5-5 invokes the optimizer in global mode. The optimizer processes all the specified source files together, applying the default Level 2 optimizations to all the modules in the application.

Example 5-5. Invoking the optimizer in global mode

```
ccsc100 -Og -o file.eld file1.c file2.c file3.c
```

5.2.2 Optimizing for Space

Your application, or specific parts of it, may require code that occupies the least possible space in memory. You can optimize the file(s) for space at the expense of speed. All optimizations are applied, with the emphasis on functions which reduce code size.

To activate space optimization, specify the `-Os` option in the shell command line. See [Section 5.3.5, “Space Optimizations,”](#) for details of the thresholds and optimization functions for this option.

5.2.3 Optimizing for Speed

You can optimize your application, or the parts of it which are speed-sensitive, to produce code that is fast as opposed to small.

To optimize for speed, specify option `-Ot`, `-Ot1` or `-Ot2` in the shell command line, as illustrated in Example 5-4 on page 5-6. The `-Ot2` option produces the fastest possible form of assembly object code, with the least saving in memory size.

Each of these options will activate most of the main optimizations, with the emphasis on optimization for speed rather than space. The difference between the options is in the thresholds at which the individual optimizations are activated. [Section 5.3.4, “Time Optimizations,”](#) lists the thresholds and describes how the optimizations are activated for each option.

5.2.4 Using Global Optimization

Once you have optimized your individual source files and groups of files, you can invoke the optimizer in global mode to ensure maximum optimization across the entire application, in order to produce the most efficient code.

In global optimization mode, all the code in the application is processed by the compiler at the same time. The optimizer has no need to make worst case assumptions since all the necessary information is available. This enables the optimizer to achieve an extremely powerful level of optimization.

The main disadvantages of compiling in global optimization mode are the high consumption of resources required, and the slow compilation time. In addition, because of the interdependency that global optimization creates between all segments of the application, the entire application needs to be recompiled if any one source code file is changed. For these reasons global optimization is generally used at the final stage of development.

For a graphic representation of how the compiler operates in non-global and global optimization modes, see [Section 3.2, “Invoking the Shell,”](#) on page 3-10.

To activate global optimization, specify the `-Og` option in the shell command line, as shown in Example 5-5 on page 5-7. While you can specify this option together with any of the other optimization level options, global optimization is generally recommended with optimization level 2. The `-O2` option is the default and may be omitted.

5.3 Optimization Types and Functions

The optimizer implements two main types of optimization:

- High-level optimizations improve the output code without taking into account the properties of the target machine. These optimizations are described in detail in [Section 5.3.2, “High-Level Optimizations.”](#)
- Low-level optimizations achieve code improvements by exploiting the architecture features of the target machine. [Section 5.3.3, “Low-Level Optimizations,”](#) provides a description of these optimizations.

These two sets of optimizations can be applied with the emphasis on reducing the memory size required by the program, or on achieving maximum program execution speed. See [Section 5.3.4, “Time Optimizations,”](#) and [Section 5.3.5, “Space Optimizations,”](#) for further detail.

Both sets of optimizations can be applied to individual files and groups of files, in non-global and in global mode. Refer to [Section 5.2.4, “Using Global Optimization,”](#) and [Section 5.3.6, “Global Optimizations,”](#) for more information.

Changes in the code as a result of one optimization may enable another optimization to be applied, producing an accumulative effect.

5.3.1 Dependencies and Parallelization

Dependency between instructions directly limits how successfully the optimizer can apply the various optimizations. An instruction is considered to be dependent on another if a change in their order of execution influences the result of the operation.

The optimizer can group instructions into parallelized execution sets only if these instructions do not contain dependencies. For a description of parallelized execution sets, refer to [Section 5.1.3, “Linear and Parallelized Code.”](#) Parallelization of different parts of the program, or of iterations of the same loop, can significantly increase the speed of the executable application.

Example 5-6 illustrates a simple dependency between two instructions. The value of `d0` is entirely different when the order of these instructions is reversed. These instructions cannot be executed in parallel.

Example 5-6. Simple instruction dependency

```
move.w    #5,d0                ; Sets register d0 to 5
add       d0,d1,d2             ; Adds the values in d0 and d1 into register d2
```

An example of dependency arising from an algorithm is shown in Example 5-7. The value of the variable `sum` must be calculated before it can be used in the `L_mac` instruction.

Example 5-7. Algorithm instruction dependency

```
sum = mpy(a,b);
result = L_mac(sum,c,d);
```

The optimizer can operate most effectively with code which contains as few dependencies as possible. [Section 5.4, “Guidelines for Using the Optimizer,”](#) provides more detailed advice for writing code that avoids dependencies and makes the best use of the optimizations.

The sections that follow describe the operation of individual optimizations in detail, and are intended for advanced users of the SC100 C compiler. If you do not require this level of detail, you may wish to skip these sections, and turn directly to [Section 5.4, “Guidelines for Using the Optimizer.”](#)

5.3.2 High-Level Optimizations

In the High-Level optimization phase, a number of general, hardware-independent optimizations are implemented. All high-level optimizations are applied when either optimization level 1 (option `-O1`) or the default optimization level 2 (option `-O2`) is selected. These high-level optimizations are summarized in Table 5-2 below, and are described in detail in the sections that follow.

Table 5-2. Summary of High-Level Optimizations

Optimization	Description	Section	Page
Strength reduction (loop transformations)	Transforms array access patterns and induction variables in loops, and replaces them with pointer accesses	5.3.2.1	5-10
Function inlining	Substitutes a function call with the code of the function	5.3.2.2	5-16
Common subexpression elimination	Replaces an expression with its value if it occurs more than once	5.3.2.3	5-17
Loop invariant code	Moves code outside a loop if its value is unchanged by the loop	5.3.2.4	5-17
Constant folding and propagation	Calculates the value of an expression at compilation time if it contains known static constants	5.3.2.5	5-17
Jump-to-jump elimination	Combines jump instructions	5.3.2.6	5-18
Dead code elimination	Removes code that is never executed	5.3.2.7	5-18
Dead storage/assignment elimination	Removes redundant variables and value assignments	5.3.2.8	5-18

The output from the high-level optimization phase is in the form of linear assembly code.

5.3.2.1 Strength reduction (loop transformations)

The purpose of strength reduction is to increase the effectiveness of the code by transforming operations which are “expensive” in terms of resources, into less expensive, linear operations. For example, addition and subtraction are linear functions which require less operation cycles than multiplication and division.

When an address calculation that contains multiplication is replaced by one containing addition, the amount of resources required by the code is significantly reduced, since addition can be implemented using the complex addressing mode of the Address Generation Unit (AGU). Where the multiplication appears within a loop, the benefit of the replacement is further increased.

The strength reduction optimization identifies and transforms induction variables, meaning variables whose successive values form an arithmetic progression, usually within a loop. An example of an induction variable is a subscript which points to the addresses of array elements, and increases with each iteration of the loop. The computation of such a variable can be moved to a position outside the loop to avoid repeated operations, and/or transformed for use with linear operations.

Simple and complex loops and array access patterns are transformed where possible into simpler, linear forms, as described in the sections that follow.

5.3.2.1.1 Simple loops

Example 5-8 shows the generated pseudocode and output assembly code for a simple loop which initializes an array. The loop structure is static, meaning that its induction variables, the loop counter *i* and the array offset *t1*, both increase by increments of known constant values.

Example 5-8. Loop transformation - simple loop

C source code

```
int table[100];
step = 1;

for(i=0; i<100; i+=step)
    table[i] = 0;
```

Pseudocode before optimization

```

    i = 0;
L1  t1 = i * 4;
    table[t1] = 0;
    i++;
    if(i<100) goto L1
```

Pseudocode after optimization

```

    i = 0;
L1  t1 = i * 4;
    table[t1] = 0;
    t1 = t1 + 4;
    i++;
    if(i<100) goto L1
```

Assembly code output

```
move.l #_table,r0   clr d2
loopstart3
    move.l d2,(r0)+
loopend3
```

Before optimization, the calculation of the value of *t1* is within the loop, and is incremented by multiplication. After optimization, the initial value of *t1* is set outside the loop, and its value is incremented inside the loop by addition. The resulting values are identical for both forms, but in the optimized version the resource overhead is considerably lower.

The same principles also apply to more complex loop structures and array access patterns, as described in the sections that follow:

- Dynamic loops, in which increments are based on a variable whose value is not known at compilation time
- Multi-step loops, in which the loop iterator increments more than once in each iteration of the loop
- Composed variable loops, in which one or more variables or iterators are linked to each other in a linear relationship
- Square loops, which access elements in a two-dimensional array as in a matrix, on a row-by-row basis
- Triangular loops, which are similar to square loops, but which access each row in the matrix from an incremented starting position in each subsequent row

5.3.2.1.2 Dynamic loops

In a dynamic loop, one or more increments are based on variables whose values are not known at compilation time.

Example 5-9 shows the generated code for a dynamic loop in which the value of the loop increment and its upper limit are not known at the time of compilation. The optimization removes the initial multiplication instruction from the body of the loop, and inside the loop the multiplication increment instruction is replaced by an addition instruction.

Example 5-9. Loop transformation - dynamic loop

C source code

```
step = step_table[1];
for(i=0; i<MAX; i+=step)
    table[i] = 0;
```

Pseudocode before optimization

```
    step = step_table[1];
    i = 0;
L1 t1 = i * 2;
    table[t1] = 0;
    i = i + step;
    if(i<MAX) goto L1
```

Assembly code output

```
L2
    clr     d3          add     d1,d0,d1
    move.l d3,(r1)    cmpge.w #100,d1
    adda   r2,r1
    jf     L2
```

Pseudocode after optimization

```
    i = 0;
    step = step_table[1];
    t1 = i * 2;
    t2 = step * 2;
L1 table[t1] = 0;
    t1 = t1 + t2;
    i = i + step;
    if(i<MAX) goto L1
```

5.3.2.1.3 Multi-step loops

Loops in which the loop iterator increments more than once in each iteration of the loop are defined as multi-step loops.

In the multi-step loop shown in Example 5-10, the loop iterator *i* increments twice within the loop. In this case, *i* is transformed into an induction variable which increments in linear progression in three stages.

Example 5-10. Loop transformation - multi-step loop

C source code

```
int table[10];
for(i=0; i<10; i++)
    table[i] = i;
    i++;
    table[i] = 0;
```

Pseudocode before optimization

```
    i = 0;
L1 t1 = i * 2;
    table[t1] = i;
    i = i + 1;
    t2 = i * 2;
    table[t2] = i;
    i = i + 1;
    if(i<10) goto L1
```

Pseudocode after optimization

```
    i = 0;
    t1 = i * 2;
    t2 = i * 2 + 2;
    t3 = i;
Repeat 10 times:
    table[t1] = t3;
    table[t2] = 0;
    t1 = t1 + 4;
    t2 = t2 + 4;
    t3 = t3 + 2;
```

Assembly code output

```
    loopstart3
L93

    move.l  d0,(r0)+n3  add#<2,d0

    move.l  d2,(r1)+n3
    loopend3
```

5.3.2.1.4 Composed variable loops

A composed variable loop incorporates one or more variables or iterators which have a linear relationship between them. The loop transformation optimizes such loops by moving the multiplication instruction to a position outside the loop, and by substituting one of the variables with a constant.

This optimization can be applied only when the variables are linked by linear arithmetic functions, meaning those calculations involving addition or subtraction of the variables, or multiplication of a variable by a constant. Functions which include non-linear operations, such as multiplication of two induction variables, cannot be optimized in this way.

Example 5-11 illustrates the generated code for a composed variables loop. In this example the increment is the result of a linear calculation using the two induction variables *i* and *j*.

Example 5-11. Loop transformation - composed variables

C source code

```
int table[100];
for(i=0, j=0; i<10; i++)
    table[10 * i + j] = i;
    j++;
```

Pseudocode before optimization

```
    i = 0;
    j = 0;
    t1 = i * 10;
L1 t2 = t1 + j;
    t3 = t2 * 2; /* address */
    table[t3] = i;
    i = i + 1;
    t1 = t1 + 10;
    j = j + 1;
    if(j < 10) goto L1
```

Pseudocode after optimization

```
    i = 0; j = 0;
    t1 = 1 * 10;
    t2 = t1 + j;
    t3 = t2 * 2;
    Repeat 10 times:
        table[t3] = i;
        i = i + 1;
        t3 = t3 + 22;
```

Assembly code output

```
    loopstart3
L93
    move.l    d0,(r0)+n3    inc    d0
    loopend3
```

5.3.2.1.5 Square loops

A square loop is a two-dimensional array access pattern which is similar to a matrix in which cells are accessed horizontally in rows, starting at the first cell in each row.

The code that is initially generated for a square loop uses a doubly-nested loop with two induction variables. These variables are incremented by multiplication, as the loop progresses through the array elements in each row, and at the start of each new row, as shown in Figure 5-2.




Figure 5-2. Square Loop

The loop transformation changes such a two-dimensional array into one row containing all the elements in one straight string. The multiplication instructions are replaced by additions, as the progression can now be performed on a linear basis. An example of the transformation of a square loop is shown below in Example 5-12.

Example 5-12. Loop transformation - square loop

C source code

```
int table[70][70];
int i, j;
for(i=0; i<35; i++)
    for(j=0; j<70; j++)
        c+=table[i][j];
```

Pseudocode before optimization

```
    i = 0;
L1 j = 0;
L2  tmp1 = i * 140;
    tmp2 = j * 2;
    tmp3 = tmp1 + tmp2;
    tmp4 = table[tmp3];
    c = c + tmp4;
    j++;
    if(j < 70) goto L2
    i++;
    if(i<35) goto L1
```

Pseudocode after optimization

```
tmp1 = 0;
Repeat 35 times
    tmp2 = table + tmp1;
    /* row base address */
    Repeat 70 times
        tmp3 = *tmp2;
        /* pointer to cell */
        c = c + tmp3;
        tmp2 = tmp2 + 2;
    tmp1 = tmp1 + 140;
    /* next matrix row */
```

Assembly code output

```
    loopstart2
L98
    doensh3    d0
    loopstart3
L97
    move.l    d1,(r0)+
    loopend3
L94
    loopend2
```

5.3.2.1.6 Triangular loops

A triangular loop array access pattern is similar to the square loop described above, except that the pointer moves to an incremented starting position in each row. The starting position pointer increments by linear progression, as shown in Figure 5-3:




Figure 5-3. Triangular Loop

A triangular loop is transformed into a mainly linear based loop, incorporating the offset increment as an addition operation. Example 5-13 illustrates the transformation of a triangular loop.

Example 5-13. Loop transformation - triangular loop

C source code

```
int table[70][70];
int i, j;
for(i=0; i<70; i++)
    for(j=i+3; j<70; j++)
        table[i][j] = 0;
```

Pseudocode before optimization

```
    i = 0;
L1 j = i
    if(j>=70) goto L3
L2    tmp1 = i * 140;
        tmp2 = j * 2;
        tmp3 = tmp1 + tmp2;
        table[tmp3] = 0;
        j++;
        if(j < 70) goto L2
L3 i++;
    if(i<70) goto L1
```

Pseudocode after optimization

```
    i = 0;
    tmp1 = 0;
    Repeat 70 times
        j = 1 + 3;
        if(j>=70) goto L3
        tmp2 = j * 2;
        tmp3 = tmp1 + tmp2; /* offset */
        tmp4 = table + tmp3;
        /* base + offset */
        Repeat (70-j)
            *tmp4 = 0;
            *tmp4 = tmp4 + 2;
L3 i++;
    tmp1 = tmp1 + 140;
    /* next matrix row */
```

Example 5-13. Loop transformation - triangular loop

Assembly code output

```
    loopstart2
L98
    cmpgt    d0,d2
    jf      L3
    tfr     d0,d4 sub    d0,d2,d3
    doensh3  d3
    asll    #<2,d4
    add     d4,d1,d5
    move.l  d5,r1
    adda    #>_table,r1,r0
    loopstart3
L97
    move.l  d6,r0)+
    loopend3
L3
    add     #280,d1,d1
    inc     d0
    loopend2
```

5.3.2.2 Function inlining

Inlining replaces a call to a function with a copy of the code for the function. In cases where the procedure call and return may be more time-consuming than the function itself, function inlining can significantly increase the speed of the program. Function inlining generates larger executable code.

The function inlining optimization can be particularly effective in global optimization mode, as it can be applied across all available files, and operates in conjunction with other global optimizations.

The following example shows how the operation executed by the function `Check` is incorporated into the code itself, removing the call to the function.

Example 5-14. Function inlining

Before optimization

```
int Check(int x);
{
    return (x>10);
}
void main()
{
    if (Check(y))
        a = 5;
}
```

After optimization

```
void main()
{
    if (y>10)
        a = 5;
}
```

You can force or suppress function inlining at specific points in the code, using the pragmas `#pragma inline` and `#pragma noinline`. Refer to [Section 3.4.5, “Pragmas,”](#) on page 3-54, for further details.

5.3.2.3 Common subexpression elimination

Where an expression appears in more than one place in the code and has the same computed value in each instance, this optimization replaces the expression itself with its result. Values loaded from memory can be included in this process, as well as values based on arithmetic computations. In the example shown below, the variable `x` replaces the repeated subexpression `e + f`.

Example 5-15. Common subexpression elimination

Before optimization

```
d = e + f + g;
y = e + f + z;
```

After optimization

```
x = e + f;
d = x + g;
y = x + z;
```

5.3.2.4 Loop invariant code

The term “invariant code” refers to an instruction which appears inside a loop, but whose value is not directly affected by the execution of the loop. This optimization moves such an instruction to a position outside the loop, with the result that the instruction is not repeated each time the loop executes. In Example 5-16, the variable `z` is set to the computed value of `2 * b + 1` before the loop executes, and this calculation is removed from the iteration.

Example 5-16. Loop invariant code motion

Before optimization

```
b = c;
for(i=0; i<3; i++)
    d[i] = 2 * b + 1;
```

After optimization

```
b = c;
z = 2 * b + 1;
for(i=0; i<3; i++)
    d[i] = z;
```

5.3.2.5 Constant folding and propagation

This optimization identifies expressions which contain `int` values known to be constants and calculates their value at compilation time. The value of the expression then replaces the expression itself, as shown in Example 5-17 below. This is the only optimization implemented if optimization level 0 (option `-O0`) is selected.

Example 5-17. Constant folding and propagation

Before optimization

```
X = 2;
Y = X + 10;
Z = 2 * Y;
```

After optimization

```
X = 2;
Y = 12;
Z = 24;
```

5.3.2.6 Jump-to-jump elimination

This optimization combines two jump operations into one, in cases where the code executes a jump to an address, and at that address immediately jumps to a different address.

In Example 5-18, the two jump instructions `goto J1;` and `goto J2;` are replaced by a direct jump to `J2`.

Example 5-18. Jump-to-jump elimination

<u>Before optimization</u>	<u>After optimization</u>
<code>if(x)</code>	<code>if(x)</code>
<code>...</code>	<code>...</code>
<code>else</code>	<code>else</code>
<code>goto J1;</code>	<code>goto J2;</code>
<code>J1:</code>	
<code>goto J2;</code>	

5.3.2.7 Dead code elimination

This optimization removes segments of “dead” code, meaning code that cannot possibly be executed. The code may be dead from the start, or it may become dead as a result of other optimizations. For example, the code may specify a condition which can never be true. In the example shown below, the variable `c` is type `char`, which can never have a value greater than 255, and therefore the `if` condition will never be met.

Example 5-19. Dead code elimination

<u>Before optimization</u>	<u>After optimization</u>
<code>char c;</code>	<code>a = 2;</code>
<code>if c > 300</code>	
<code>a = 1;</code>	
<code>else</code>	
<code>a = 2;</code>	

5.3.2.8 Dead storage/assignment elimination

Dead storage or assignment occurs when a variable is assigned a value, either directly or as a result of an expression, and is not used again anywhere in the code, or receives another value before being used. This optimization removes any unnecessary instructions and unused memory locations which may result from such cases. This redundancy may arise as a result of other optimizations.

In Example 5-20, before optimization the variable `a` is assigned the value 5, and is not used before it is reassigned the value 7. The dead storage/assignment elimination optimization removes the redundant instruction `a = 5`. If the variable `a` was not used at all after being assigned a value, it would be removed completely.

Example 5-20. Dead storage/assignment elimination

<u>Before optimization</u>	<u>After optimization</u>
<code>a = 5;</code>	<code>a = 7;</code>
<code>..</code>	
<code>a = 7;</code>	

5.3.3 Low-Level Optimizations

The Low-Level Transformations (LLT) phase is a separate modular stage of the optimization process which implements a number of low-level optimizations. This phase transforms the linear code generated by the high-level optimization phase into parallel assembly code, which can take advantage of the parallel execution units of the SC100 architecture.

The degree of parallelization that the optimizer is able to achieve is limited by the number and type of dependencies within the source code. See [Section 5.3.1, “Dependencies and Parallelization,”](#) for a summary of these issues.

[Section 5.4, “Guidelines for Using the Optimizer,”](#) provides detailed advice about preparing your source code with a view to reducing dependencies and realizing the maximum potential for optimization.

All low-level optimizations are applied when the level 2 optimization (option -O2) is selected. The speed and time optimization levels (options -Os, -Ot, -Ot1 and -Ot2) apply most of the low-level optimizations, with different thresholds. See [Section 5.3.4, “Time Optimizations,”](#) and [Section 5.3.5, “Space Optimizations,”](#) for more detail.

Low-level optimizations are not activated at all when either option -O0 or option -O1 is selected.

The major low-level optimizations are summarized below in Table 5-3, and are described in detail in the sections that follow.

Table 5-3. Summary of Low-Level Optimizations

Optimization	Description	Section	Page
Instruction scheduling	Executes multiple instructions in the same cycle, fills delay slots associated with a branch operation, and avoids pipeline restrictions	5.3.3.1	5-21
Low-level software pipelining	Rearranges instructions in a loop to minimize dependencies	5.3.3.2	5-22
Conditional execution and predication	Transforms a branch into a sequence of conditional actions	5.3.3.3	5-25
Speculative execution	Moves instructions from conditional to unconditional paths	5.3.3.4	5-26
Post-increment detection	Combines the functions of incrementing (or decrementing) a pointer and accessing the computed address into one instruction	5.3.3.5	5-27
Low-level peephole optimization	Merges a sequence of instructions into a single instruction	5.3.3.6	5-28

The optimizer applies the low-level optimizations in a predefined sequence, and invokes some of the optimizations more than once, as illustrated in Figure 5-4. Each optimization is directly affected by the result of the preceding optimization.

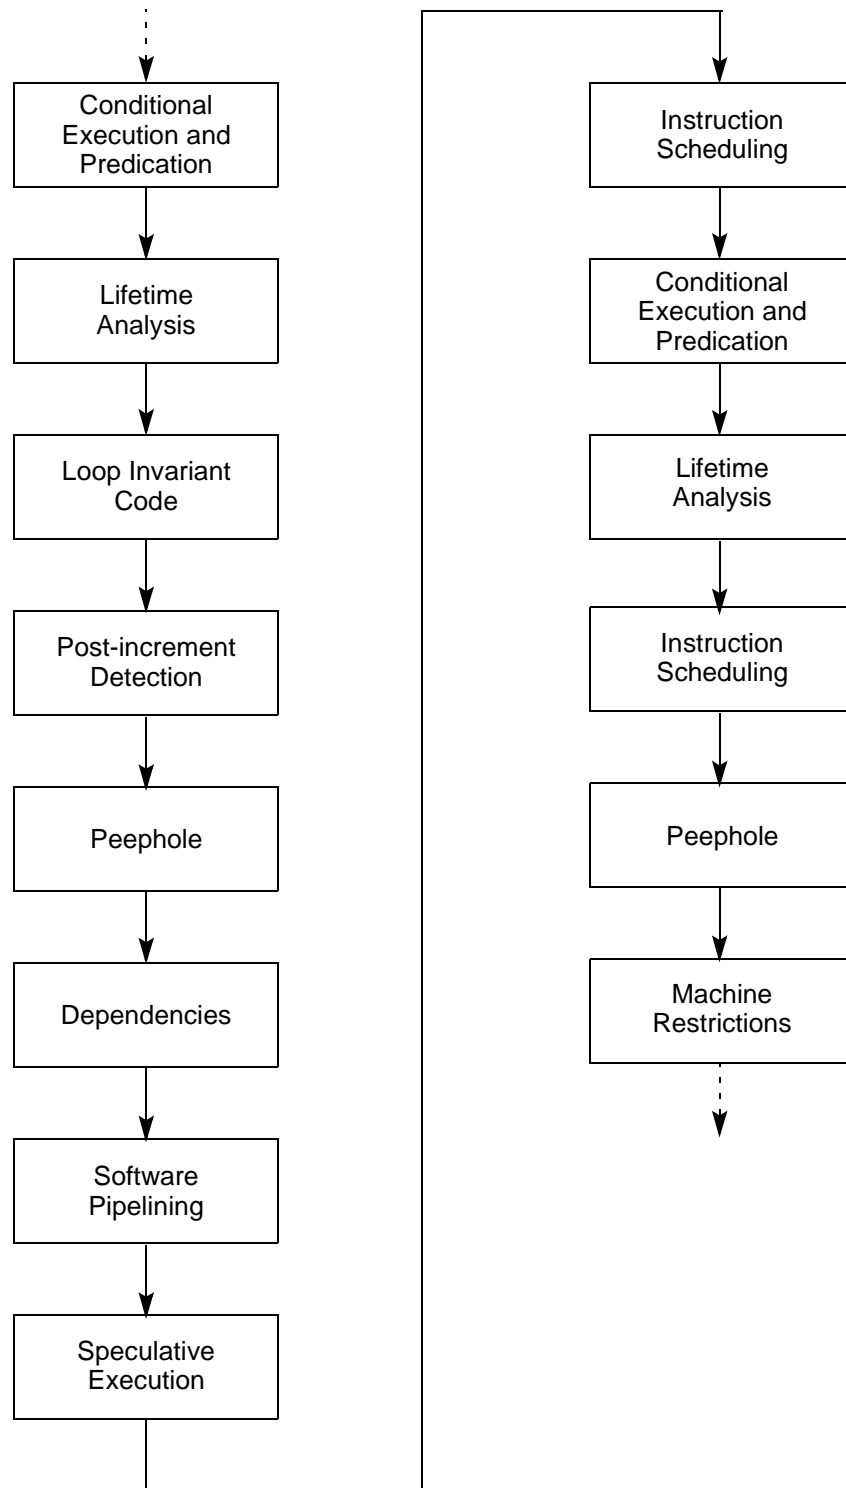


Figure 5-4. Sequence of Low-Level Transformation Optimizations

5.3.3.1 Instruction scheduling

The main purpose of this optimization is to execute as many instructions as possible from the same instruction stream in the same cycle. The amount of dependency between the instructions limits the extent to which this can be achieved.

The instruction scheduling optimization organizes instructions into execution sets wherever it is possible to do so, making best use of the Data Arithmetic Units and Address Generation Units provided by the SC100 architecture.

Example 5-21 illustrates the use of instruction scheduling:

Example 5-21. Instruction scheduling

Before optimization

```
move.l d0,(r0)
inc    d0
tfra   r3,r0
adda   #12,r3
move   (r1)+,d1
```

After optimization

```
move.l d0,(r0)  inc    d0
tfra   r3,r0    move   (r1)+,d1
adda   #12,r3
```

Instruction scheduling serves two further purposes:

- Filling delay slots when branch instructions are executed, as described below
- Rescheduling operations that are not dependent on pipeline restricted instructions, as described in [Section 5.3.3.1.2, “Avoiding pipeline restrictions,”](#) on page 5-22

5.3.3.1.1 Filling delay slots

A branch instruction requires three cycles to execute if the branch is taken. When a branch executes, the prefetch queue is lost, and the cycles used for the other instructions are wasted, since they cannot execute until the branch instruction has completed. The wasted cycles are termed *delay slots*.

The instruction scheduling optimization checks whether other operations can be executed at the same time as the branch instruction. This is not possible if there are limiting factors, for example:

- The branch instruction is directly affected by the instructions which precede it.
- There are specific dependencies between the branch and the other instructions.

If there are no limiting factors, the scheduler rearranges the code, in order to use the delay slots efficiently. In the following example, the code has been reorganized to enable three instructions to execute during the time that the branch requires to complete its operation.

Example 5-22. Filling delay slots

Before optimization

```
move.l d0,(r0)
inc    d0
tfr    d5,d2
rts
```

After optimization

```
rtsd
move.l d0,(r0)  inc    d0  tfr d5,d2
```

5.3.3.1.2 Avoiding pipeline restrictions

Certain instructions, for example, a move to an Rn register, are subject to pipeline restrictions. The effect of these instructions may not be implemented until two or more cycles after the instruction executes. In such cases, an operation which is dependent on the result of such an instruction, and which follows it immediately, must wait until the result is available.

The instruction scheduling optimization rearranges the sequence of such instructions where possible, using the cycle(s) which would otherwise be wasted to implement one or more operations that are not dependent on the restricted instruction.

In Example 5-23, the `clr` instruction has been rescheduled, since it can execute before the effect of the `move.l` instruction is implemented, whereas the `move.w` instruction must wait for the results of the `move.l` operation.

Example 5-23. Avoiding pipeline restrictions

<u>Before optimization</u>	<u>After optimization</u>
<code>move.l d0, (r0)</code>	<code>move.l d0, (r0)</code>
<code>nop</code>	<code>clr d0</code>
<code>move.w (r0), dl</code>	<code>move.w (r0), dl</code>
<code>clr d0</code>	

5.3.3.2 Low-level software pipelining

Software pipelining provides a further level of loop optimization, in addition to the high-level optimizations which operate on loops.

The software pipelining optimization attempts to rearrange the sequence of instructions inside a loop, in order to minimize dependencies between such instructions, and thus increase the level of parallelization.

For example, a segment of code may consist of three instructions, A, B and C, within a loop which iterates 4 times. In some cases, the code may be reorganized into a different sequence without affecting its result, for example:

1. Instruction A
2. Instructions B, C, A, in a loop which iterates 3 times
3. Instruction B
4. Instruction C

The revised arrangement of the instructions results in fewer dependencies than in the original code.

This optimization is applied only to innermost loops of small or moderate size, which contain no branches or function calls within the loop. It is most effective when applied to loops that execute a large number of times.

Each iteration of a software pipelined loop may contain instructions from a different iteration of the original loop.

Software pipelining increases code size in almost all circumstances. When optimization for size is specified, software pipelining is suppressed entirely. See [Section 5.3.4, “Time Optimizations,”](#) for further details.

The following example shows how the software pipelining optimization reduces the number of iterations and rearranges instructions both within and outside the loop, thus enabling the maximum number of instructions that are not dependent on each other to execute in parallel.

Example 5-24. Software pipelining - complex FIR

C source code

```
for (i = 0; i < N; i++)
{
L_tmpr = L_mac (L_tmpr, sample[i].r, coeff[N - i - 1].r;
L_tmpr = L_msu (L_tmpr, sample[i].i, coeff[N - i - 1].i;

L_tmpr = L_mac (L_tmpr, sample[i].i, coeff[N - i - 1].r;
L_tmpr = L_mac (L_tmpr, sample[i].r, coeff[N - i - 1].i;
}
```

Before optimization

```
loop n times:
  move.w (r0)+,d4
  move.w (r1)-,d3
  mac    d3,d4,d5
  move.w (r0)+,d1
  move.w (r1)-,d2
  mac    -d1,d2,d5
  mac    d3,d1,d6
  mac    d2,d4,d6
```

After optimization

```
/* Prolog */
move.w (r0)+,d4      move.w (r1)-,d3
mac    d3,d4,d5      move.w (r0)+,d1      move.w(r1)-,d2

loop n-1 times:
  /*start loop*/
  [
    mac    d3,d1,d6      mac    -d1,d2,d5
    move.w (r0)+,d4      move.w (r1)-,d3
  ]
  [
    mac    d3,d4,d5      mac    d2,d4,d6
    move.w (r0)+,d1      move.w (r1)-,d2
  ]
  /*endloop*/
  /* Epilog */
mac    d3,d1,d6      mac    -d1,d2,d5
mac    d2,d4,d6
```

In the following example, the loop iterates only 8 times, instead of the 10 in the original code, since two iterations have been unrolled. The loop executes in a single cycle. During this cycle the loop:

- Loads a value from iteration $i+2$
- Multiplies the value from iteration $i+1$
- Stores the result value from iteration i

Example 5-25. Software pipelining - vector multiplication by a constant

C source code

```
for (i=0; i<10; i++)
    b[i] = mult(a[i], 0x4000);
```

Assembly code after optimization

```
doensh3  #<8 ; Pipelining loop twice
move.l   #_a,r1
move.f   #16384,d1
move.f   (r1)+,d0  move.l #_b,r0
mpy      d0,d1,d2  move.f (r1)+,d0
loopstart3
```

L93

```
[
moves.f  d2,(r0)+  ; *
mpy      d0,d1,d2  ; **
move.f   (r1)+,d0  ; ***
]
loopend3
```

L92

```
moves.f  d2,(r0)+  mpy d0,d1,d2
moves.f  d2,(r0)+
```

5.3.3.3 Conditional execution and predication

The conditional execution and predication optimization simplifies small conditional structures and transforms the branch into one sequence.

An example of this transformation is shown in Example 5-26, in which two branches are removed.

Example 5-26. Conditional execution and predication

C source code

```
If(a < 0){
    lower_bound = 0;
    i = 0;
}else
    lower_bound = a;
```

Generated code before optimization

```
    move.w    a,d0
    tstgt    d0
    bf       L_False
    clr      d2
    clr      d3
    bra      L_AfterIf
L_False
    tfr      d0,d2
L_AfterIf
    move.w    d2,lower_bound
```

Generated code after optimization

```
    move.w    a,d0
    tstgt    d0
    ift      clr d2      clr d3
    iff      tfr d0,d2
    move.w    d2,lower_bound
```

An additional advantage of this optimization is that it increases the size of the basic blocks in the optimized code segment, making further optimization more effective.

It is important to note, however, that the conditional execution optimization adds one word for each branch that it replaces (*ift* and *iff* in the above example). As a result, the impact on the size of the program can be considerable. Generally, this optimization is only activated for small structures where the number of instructions added is less or equal to the number of instructions saved. The optimization levels which specify size as an important consideration apply specific thresholds for this optimization. For details, see [Section 5.3.4, “Time Optimizations.”](#)

5.3.3.4 Speculative execution

The speculative execution optimization moves instructions from conditional to unconditional paths, in order to fill execution slots that would not otherwise be used.

If an empty execution slot is available when a condition statement is encountered, the instructions are rearranged so that the conditional instructions execute unconditionally in previous cycles to the condition. If the condition is true and the `ift` instruction has been executed, or if the condition is false and the `iff` instruction has been executed, a cycle has been gained. If the condition result does not match the moved instruction, the appropriate instruction is executed as normal, with no loss of cycles.

Example 5-26 shows an example of this transformation. In this example, the first `iff` instruction is moved so that it executes in the same cycle as the `cmpgt` instruction. If the result of the conditional operation is true, the `ift` instruction is executed in the next cycle. If the result is false, the instruction that was previously the second `iff` is executed, with the result that only one cycle is used instead of two.

Example 5-27. Speculative execution

C source code

```
If(var > 5)
    x[3] = a;
else
    y = b;
```

Generated code before optimization

```
cmpgt #5,d1
nop
iff    move.l x+6,r0
iff    move.l d3,_y
ift    move.l d2,(r0)
```

Generated code after optimization

```
move.l x+6,r0    cmpgt #5,d1
nop
iff    move.l d3,_y
ift    move.l d2,(r0)
```

This optimization can be implemented successfully for one or more instructions if:

- Sufficient slots are available.
- There are no dependencies between the instruction in the conditional path and other instructions.
- The conditional instruction does not have any specific side effects.

5.3.3.5 Post-increment detection

This optimization exploits the features of the SC100 architecture, and increases code efficiency in terms of both size and speed. It identifies the instructions which use arithmetic functions to modify pointers, and which access the computed addresses, and replaces them with special post-increment or post-decrement address mode instructions which combine both functions.

The increment (or decrement) factor is not limited to the values 2 or 4, since any one of the four index registers (n0 through n3) may be used, as illustrated in Example 5-28.

Example 5-28. Post-increment detection

Generated code before optimization

```
L150
  move.l  #_L_R,r4
  move.l  #_CGUpdates,r5
  doen3   #<8
  dosetup3      L183
  loopstart3
L183
  move.l  (r4),d0
  move.l  (r5),d1
  mac     d0,d1,d2
  adda   #<4,r4
  adda   #<12,r5
  loopend3
L152
```

Generated code after optimization

```
L150
  doensh3  #<7      ;
  Pipelining loop once
  move.w   #3,n3
  move.l   #_L_R,r4
  move.l   #_CGUpdates,r5
  move.l   (r4)+,d0
  move.l   (r5)+n3,d1
  loopstart3
L183
  [
    mac     d0,d1,d2  ; *
    move.l  (r5)+n3,d1 ; **
    move.l  (r4)+,d0  ; **
  ]
  loopend3
L152
  mac     d0,d1,d2
```

5.3.3.6 Low-level peephole optimization

The low-level peephole optimization identifies sequences of instructions that can be merged into a single instruction, and implements this transformation, as shown in Example 5-29.

Example 5-29. Low-level peephole optimization

<u>Generated code before optimization</u>		<u>Generated code after optimization</u>	
deca	r0	decgea	r0
move.w	#33,d0	move.w	#33,d0
tstgea.l	r0		

Example 5-30 illustrates a combination of pipelining and peephole optimizations. After pipelining, the final `mac` instruction, which has been moved outside the loop, is merged with the `rnd` instruction to form a `macr` instruction.

Example 5-30. Combined pipelining and peephole optimizations

<u>Generated code before optimization</u>		<u>Generated code after optimization</u>	
	doen #9	doen #8	;Pipelining loop once
	dosetup0 L1	dosetup0 L1	
	loopstart0	move.w (r0)+,d3	
L1		move.w (r1)+,d2	
	move.w (r0)+,d3	loopstart0	
	move.w (r1)+,d2		
	mac d2,d3,d7	L1	mac d2,d3,d7 ; *
			move.w (r1)+,d2 ; **
	loopend0		move.w (r0)+,d3 ; **
			loopend0
	rnd d7		
			macr d2,d3,d7

5.3.4 Time Optimizations

The following three options are provided to enable you to optimize your code with the emphasis on maximizing program execution speed:

- The `-Ot` option produces code that is faster, and possibly larger, than the code output by the default `-O2` option.
- The `-Ot1` option results in yet faster program execution, and a potentially larger program than the `-Ot` option.
- The `-Ot2` option produces the fastest code of the three. Programs optimized using this option are likely to occupy more memory size than those produced with either option `-Ot` or `-Ot1`.

5.3.4.1 Time optimization levels and criteria

When optimizing for speed, the compiler includes all optimizations in both the high-level and low-level optimization phases, subject to the specific criteria for the selected option, as shown in Table 5-4.

Table 5-4. Time Optimization Criteria

Criteria	-Ot	-Ot1	-Ot2	Description
Small loop size.	25	25	50	Maximum size of a loop (in instructions) to which the following optimizations are applied: <ul style="list-style-type: none"> • Software pipelining • Prefix grouping (See Section 5.3.4.2, "Prefix grouping.")
Conditional execution block size.	5	5	10	Maximum size of a basic block (in instructions) which is converted to use the <code>ift/iff</code> instructions.
Loop bound type for software pipelining.	Static	Static	Dynamic	Determines whether software pipelining is applied to loops: <ul style="list-style-type: none"> • Static: Software pipelining is applied only to loops whose bounds (upper and lower limits) are constants or static variables known at compilation time. • Dynamic: Software pipelining is applied to loops with bounds that are not known at compilation time.
Loop alignment.	No	No	Yes	Forces the loop start address to 16-byte alignment. Non-short loops perform faster if aligned in this way. This may increase the size of the code.

All the optimizations included in the previous table increase the size of the code.

Further information about software pipelining is provided in [Section 5.3.3.2, "Low-level software pipelining."](#) For details of the conditional execution optimization, see [Section 5.3.3.3, "Conditional execution and predication."](#)

5.3.4.2 Prefix grouping

Instruction grouping is applied by the optimizer wherever possible, in order to make best use of the available multiple execution units. In addition to "natural" grouping of instructions, which increases efficiency and does not increase code size, the optimizer can implement prefix grouping. Prefix grouping is a mechanism whereby an additional word is introduced into the code in order to force more than one instruction to execute in the same cycle.

Prefix grouping improves performance in terms of speed, but increases the size of the code. According to the optimization option selected, the optimizer activates prefix grouping as follows:

- Option `-Ot2`: on the entire code
- Option `-Ot1`: on all loops
- Options `-Ot` and `-O2`: on small loops only (the default)

5.3.5 Space Optimizations

When you select the `-Os` option, the optimizer aims to produce code that occupies as little memory space as possible. In certain cases, this may be at the expense of program speed, resulting in a program that is small, but which executes relatively slower than a program optimized with the default option `-O2`.

- The compiler executes both high-level and low-level optimization phases when compiling with the `-Os` option. In order to minimize the size of the code, the optimizer does not apply the high-level optimization, function inlining, since this always increases code size.

All the low-level optimizations are applied, except for conditional execution, which is not used if the basic block involved contains more than five instructions.

In this mode, the optimizer uses only serial grouping when encoding assembly instructions, since code size is increased when prefixes are added, as described in [Section 5.3.4.2, "Prefix grouping."](#) The optimizer applies conditional execution and predication to small structures only, because this optimization adds to the size of the code. See [Section 5.3.3.3, "Conditional execution and predication,"](#) for further details.

5.3.6 Global Optimizations

Global optimization mode produces the most effective form of optimization, since optimizations are applied across all the files in the application. The option `-Og` can be specified in the command line together with any of the optimization options except the `-O0` option, and is most effective when used with the default level `-O2`.

In addition to implementing the selected level of optimization across all the files, global optimization also applies two specific optimizations:

- Function inlining across multiple files, which applies the optimization described in [Section 5.3.2.2, "Function inlining,"](#) to the whole program. As with function inlining for individual files, this increases the size of the code, but can considerably increase execution speed.
- Optimization of access to global and static variables.

5.4 Guidelines for Using the Optimizer

The optimizer produces the best possible results when the source code is written in a simple and straightforward manner. Complex structures and algorithms should be avoided wherever possible, since these can reduce the effectiveness of many of the optimizations.

During the various optimization phases, the compiler attempts to convert all the structures in the code into a form that is independent of the style of an individual user, and that can be processed efficiently by the individual optimizations. By following the basic rules of clarity and simplicity when writing your code, you help the optimizer to retrieve the specific information it needs, and to apply the maximum amount of optimization.

For example, when accessing arrays you should use simple access instructions wherever possible, and avoid using complex access instructions which use pointers, as shown in Example 5-31:

Example 5-31. Simple and complex array accesses

a) Simple array access (recommended)

```
a[i]
```

b) Complex array access (not recommended)

```
p = &a[0]  
*p++;
```

[Section 5.4.3, “General Hints,”](#) provides further general guidelines for writing simple code structures to assist optimization.

You can further enhance the results of the optimization by applying two specific techniques that help the optimizer take full advantage of the multiple execution units of the SC100 architecture:

- Partial summation, which reduces dependencies in a loop, enabling multiple iterations of a loop in parallel
- Multisample processing, a programming technique which processes multiple samples simultaneously

These techniques are described in the sections that follow.

5.4.1 Partial Summation Techniques

One of the optimizer's major functions is to produce parallelized code that fully utilizes the available number of multiply-accumulate (MAC) units. The number of MAC units that can be used in an execution set, meaning the number of instructions executed in the same cycle, is usually limited by the degree of dependency within the code.

The partial summation programming technique helps you reduce the dependencies in the loops in your source code, in such a way that the iterations can execute in parallel. By structuring your source code using partial summation techniques wherever possible, you enable the optimizer to further reduce dependencies and increase parallelization.

In Example 5-32, the inner loop can use only a single MAC per cycle, because of the inner dependency within the algorithm. The same output code is generated when compiling for a single, dual, or quad MAC StarCore system.

Example 5-32. MAC usage limited by dependency in loop

Source code

```
void Iir(short Input[], short Coef[], short FiltOut[])
{
    long L_Sum = 0;  short int Stage, Smp;  int LoopCount;

    FiltOut[0] = Input[0];
    for (Smp = 1; Smp < S_LEN; Smp++)
    {
        L_Sum = LPC_ROUND;  LoopCount = (Smp<NP ? Smp : NP );

        for (Stage = 0; Stage < LoopCount; Stage++)
            L_Sum = L_msu(L_Sum, FiltOut[Smp - Stage -1],Coef[Stage]);

        L_Sum = L_shl(L_Sum, ASHIFT);
        L_Sum = L_msu(L_Sum, Input[Smp], 0x8000);
        FiltOut[Smp] = extract_h(L_Sum);
    }
}
```

Generated code

```
doenshl d0
move.f r2)+,d0      move.f (r0)-,d1
loopstart1
PL001
mac    -d0,d1,d2    move.f (r0)-,d1      move.f (r2)+,d0
loopend1
PL000
mac    -d0,d1,d2
```

Example 5-33 illustrates how you can use partial summation to split the inner loop in the above example to enable two parallel iterations. The loop iterates half the number of times. The sum is accumulated using two variables, which are combined outside the loop.

Example 5-33. Partial summation for dual MAC usage

Source code

```
for (Stage = 0; Stage < (LoopCount>>1); Stage++)
{
    L_Sum = L_msu(L_Sum, FiltOut[Smp - 2*Stage -1], Coef[2*Stage]);
    L_Sum1 = L_msu(L_Sum1, FiltOut[Smp - 2*Stage -2], Coef[2*Stage+1]);
}

L_Sum = L_shl(L_Sum+L_Sum1, ASHIFT);
L_Sum = L_msu(L_Sum, Input[Smp], 0x8000);
```

Generated code

```
doenshl d0
    move.2f    (r2)+,d0d1    move.2f    (r0)-,d6d7
    loopstart0
PL001
    [
    mac        -d0,d6,d2
    mac        -d1,d7,d5
    move.2f    (r0)-,d6d7
    move.2f    (r2)+,d0d1
    ]
    loopend0
PL000
    mac        -d0,d6,d2    mac        -d1,d7,d5
```

The same technique can be used for compiling with a quad MAC system, by splitting the loop into four iterations, using four variables and one quarter the number of iterations.

It is important to note that partial summation is not suitable for algorithms with bit-exact requirements. This technique changes the order of the calculation, and may affect the value of the result in cases where statements must be executed in the exact order specified.

In certain algorithms the effectiveness of the partial summation technique may be limited because of alignment restrictions. For example, the `move.2f` instruction, which is required for partial summation, must be used on a long word boundary.

In Example 5-33, this restriction is satisfied, and the partial summation technique can be used successfully. Example 5-34 shows an algorithm for which partial summation cannot be used. This is because the second iteration produces an odd value for the variable `i`, with the result that the `move.2f` instruction violates the alignment requirement.

Example 5-34. Alignment restrictions in algorithms

```
for (i = 0; i < DataBlockSize; i++)
{
    Delay[(DataBlockSize-i)] = DataIn[i];
    sum1 = 0; sum2 = 0;
    for (j = 0; j < FirSize/2; j++)
    {
        sum = L_mac(sum, Coef[2*j], Delay[2*j-i]);
        sum = L_mac(sum, Coef[2*j+1], Delay[2*j-i+1]);
    }
    Result = round(sum);
}
```

The multisample techniques described in the following section help you write source code which enables the optimizer to take further advantage of multiple execution units. You can apply multisample techniques even if you cannot use partial summation for certain algorithms because of alignment restrictions or bit-exact requirements.

5.4.2 Multisample Techniques

To obtain high performance, a pipelining technique called “multisample” programming is used to process multiple samples simultaneously. The multisample programming techniques enable you to obtain high performance by taking full advantage of the SC100 multiple-ALU architecture.

The following terminology is used throughout this section:

- **Generic Kernel:** The minimum required operations of the algorithm. The generic kernel is the theoretical minimum size of the kernel without considering implementation constraints.
- **Basic Kernel:** The inner loop of a DSP algorithm. This may contain several replications of the generic kernel or additional code for pipelining. The basic kernel is actually implemented on the DSP and is subject to implementation constraints.
- **Operand:** A value used as an input to an ALU.
- **Delays:** Values stored as a delay line for referencing past values.
- **Iteration:** The complete execution of a basic kernel.
- **Loop pass:** The execution of the instructions within the basic kernel. Many loop passes may be needed to complete a single iteration of the kernel.

To process several samples simultaneously, operands (both coefficients and variables) are reused within the kernel. Although a coefficient or operand is loaded once from memory, multiple ALUs may use the value, or the value may be used in a later step of the kernel.

Figure 5-5 illustrates the structure of a single sample and multisample algorithm.




Figure 5-5. Single Sample and Multisample Kernels

In a single sample algorithm (Figure 5-5 A), samples are processed by the algorithm serially. The kernel processes a single input sample and generates a single output sample. For an algorithm such as an FIR, samples are input to the FIR kernel one at a time. The FIR kernel generates a single output for each input sample. Blocks of samples are processed using loops and executing the FIR kernel several times.

In contrast, the multisample algorithm (Figure 5-5 B) takes multiple samples at the input in parallel and generates multiple samples at the output simultaneously. The multisample algorithm operates on data in small blocks. Operands and coefficients are held in registers, and applied to both samples simultaneously, resulting in fewer memory accesses.

Multisample algorithms are ideal for block processing algorithms where data is buffered and processed in groups (such as speech coders). Figure 5-5 B shows two samples being processed simultaneously. However, the number of simultaneous samples depends on the processor architecture and type of algorithm.

Most DSP algorithms have a multiply-accumulate (MAC) at their core. On a load/store machine, the register file is the source/destination of operands to/from memory. For the ALU, the register file is the source/destination of operands. On a single sample, single ALU algorithm, the memory bandwidth is typically equal to the operand bandwidth, as shown in Figure 5-6.




Figure 5-6. Single ALU Operand and Memory Bandwidth

When increasing the number of ALUs to four, the bandwidth increases as shown in Figure 5-7.

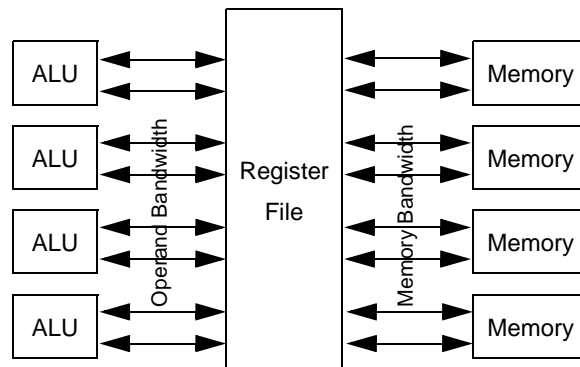


Figure 5-7. Quad ALU Operand and Memory Bandwidth

Quadrupling the number of ALUs quadruples the operand bandwidth. If there is one address generator per operand, this results in eight address generators. This is undesirable because it requires an 8 port memory and a significant amount of address generation hardware.

The SC140 DSP core solves this problem by providing up to a quad operand load/store over a single bus. With two quad operand loads, eight operands can be loaded using two address generators.

Although quad operand loading provides the proper memory bandwidth, some algorithms have special memory alignment requirements. These alignment requirements make it difficult to use multiple operand load/stores.

Multisample algorithms provide a solution for implementing algorithms with memory alignment requirements. By reusing previously loaded values, the number of operands loaded from memory is reduced, which relaxes the alignment constraints.

Both techniques for increasing operand bandwidth, by using wider data buses or by reusing operands, are shown in Figure 5-8.

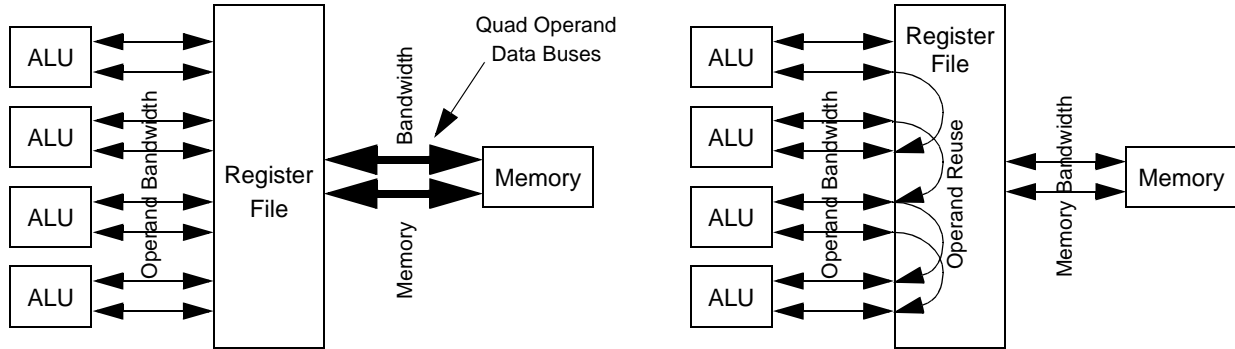


Figure 5-8. Options for Increasing Operand Bandwidth

To introduce the multisample technique, four example DSP kernels are written in multisample form. The DSP kernels presented are direct form FIR filter, direct form IIR filter, correlation and biquad filter.

5.4.2.1 Multisample implementation issues

When implementing a DSP algorithm such as an FIR filter, trade-offs are made between the number of samples processed and the number of ALUs as shown in Figure 5-9.

		Number of ALUs		
		1	2	4
Number of Samples	1	1 sample, 1 ALU	1 sample, 2 ALUs	1 sample, 4 ALUs
	2	2 samples, 1 ALU	2 samples, 2 ALUs	2 samples, 4 ALUs
	4	4 samples, 1 ALU	4 samples, 2 ALUs	4 samples, 4 ALUs

Figure 5-9. Number of Samples and ALUs for Implementing DSP Algorithms

As the kernel computes more samples simultaneously, the number of memory loads decreases because data and coefficient values are being reused. However, to obtain this reuse, more intermediate results are required, which typically requires more registers in the processor architecture.

If the operand memory requires wait states, this technique improves the speed of the algorithm. If the operand memory is full speed, then the algorithm does not execute any faster, but may reduce power consumption because the number of memory accesses has been reduced.

By using more ALUs, it is theoretically possible to compute an algorithm faster. Moving across the row theoretically applies 1, 2 or 4 ALUs to the algorithm. To apply multiple ALUs, some degree of parallelism is required in the algorithm to partition the computations.

Although computing a single sample with multiple ALUs is theoretically possible, limitations in the DSP hardware may not allow this style of algorithm to be implemented. In particular, most processors typically require operands to be aligned in memory and multiple operand load/stores to be aligned.

For example, a double operand load requires an even address and a quad operand load requires a double even address. These types of restrictions are typical to reduce the complexity of the address generation hardware (particularly for modulo addressing).

Restricting the boundaries of the load makes implementing some algorithms very difficult or impossible. This is easiest to explain by way of example. Consider a series of (aligned) quad operand loads from memory, as shown in Figure 5-10.




Figure 5-10. Quad Coefficient Loading from Memory

The loads in Figure 5-10 do not have a problem with alignment because loads occur from double even addresses.

Alignment problems typically occur with algorithms implementing delay lines in memory. These algorithms delete the oldest delay and replace it with the newest sample. This is typically done by using modulo addressing and “backing up” the pointer after the sample is processed. This leads to an addressing alignment problem as shown in Figure 5-11.




Figure 5-11. Misalignment when Loading Quad Operands

On the first iteration of the kernel, quad data values are loaded starting from a double even address. This does not create an alignment problem. However, at the end of the first iteration, the pointer is backed up by one to delete the oldest sample. On the next iteration, the pointer is not at a double even address and the quad data load is not aligned.

A solution to the alignment problem is to reduce the number of operands moved on each data bus. This relaxes the alignment issue. However, in order to maintain the same operand bandwidth, each loaded operand must be used multiple times. This is a situation where multisample processing is useful.

As the number of samples per iteration increases, more operands are reused and the number of moves per sample is reduced. With fewer moves per sample, the number of memory loads is decreased allowing fewer operands per bus. Fewer operands per bus allows the data to be loaded with fewer restrictions on alignment.

5.4.2.2 Implementation example

The FIR_A4S4 Quad ALU, quad sample, is the highest performance implementation on a quad ALU SC100 DSP.

To further increase the performance of the FIR filter, four ALUs may be used. To avoid misalignment, four samples are processed simultaneously. The quad ALU, quad sample FIR data flow is shown in Figure 5-12.




Figure 5-12. Quad ALU, Quad Sample FIR Filter Data Flow

Input samples are grouped together four at a time. Coefficients and delays are loaded and applied to all four input values to compute four output values. By using four ALUs, the execution time of the filter is only one quarter the execution time of a single ALU filter.

To develop the FIR filter equations for processing four samples simultaneously, the equations for the current sample $y(n)$ and the next three output samples $y(n+1)$, $y(n+2)$ and $y(n+3)$ are as shown in Figure 5-13.

$$\begin{array}{l}
 y(n) = x(n)C_0 + x(n-1)C_1 + x(n-2)C_2 + x(n-3)C_3 + x(n-4)C_4 + x(n-5)C_5 + x(n-6)C_6 + x(n-7)C_7 \\
 y(n+1) = x(n+1)C_0 + x(n)C_1 + x(n-1)C_2 + x(n-2)C_3 + x(n-3)C_4 + x(n-4)C_5 + x(n-5)C_6 + x(n-6)C_7 \\
 y(n+2) = x(n+2)C_0 + x(n+1)C_1 + x(n)C_2 + x(n-1)C_3 + x(n-2)C_4 + x(n-3)C_5 + x(n-4)C_6 + x(n-5)C_7 \\
 y(n+3) = x(n+3)C_0 + x(n+2)C_1 + x(n+1)C_2 + x(n)C_3 + x(n-1)C_4 + x(n-2)C_5 + x(n-3)C_6 + x(n-4)C_7
 \end{array}$$

Figure 5-13. FIR Filter Equations for Four Samples

The generic kernel has the following characteristics:

- Four parallel MACs.
- One coefficient is loaded and used by all four MACs in the same generic kernel.
- One delay value is loaded, used by the generic kernel and saved for the next three generic kernels.
- Three delays are reused from the previous generic kernel.

To develop the structure of the quad ALU kernel, the filter operations are written in parallel and the loads are moved ahead of where they are first used. This creates the generic kernel shown in Figure 5-14.

				Generic Kernel	
$y(n) = 0$	$y(n+1) = 0$	$y(n+2) = 0$	$y(n+3) = 0$		load $x(n+3)$
$y(n) += C_0 * x(n)$	$y(n+1) += C_0 * x(n+1)$	$y(n+2) += C_0 * x(n+2)$	$y(n+3) += C_0 * x(n+3)$	} Generic Kernel	load $x(n+2)$
$y(n) += C_1 * x(n-1)$	$y(n+1) += C_1 * x(n)$	$y(n+2) += C_1 * x(n+1)$	$y(n+3) += C_1 * x(n+2)$		load $x(n+1)$
$y(n) += C_2 * x(n-2)$	$y(n+1) += C_2 * x(n-1)$	$y(n+2) += C_2 * x(n)$	$y(n+3) += C_2 * x(n+1)$		load C_0 , load $x(n)$
$y(n) += C_3 * x(n-3)$	$y(n+1) += C_3 * x(n-2)$	$y(n+2) += C_3 * x(n-1)$	$y(n+3) += C_3 * x(n)$		load C_1 , load $x(n-1)$
$y(n) += C_4 * x(n-4)$	$y(n+1) += C_4 * x(n-3)$	$y(n+2) += C_4 * x(n-2)$	$y(n+3) += C_4 * x(n-1)$		load C_2 , load $x(n-2)$
$y(n) += C_5 * x(n-5)$	$y(n+1) += C_5 * x(n-4)$	$y(n+2) += C_5 * x(n-3)$	$y(n+3) += C_5 * x(n-2)$		load C_3 , load $x(n-3)$
$y(n) += C_6 * x(n-6)$	$y(n+1) += C_6 * x(n-5)$	$y(n+2) += C_6 * x(n-4)$	$y(n+3) += C_6 * x(n-3)$		load C_4 , load $x(n-4)$
$y(n) += C_7 * x(n-7)$	$y(n+1) += C_7 * x(n-6)$	$y(n+2) += C_7 * x(n-5)$	$y(n+3) += C_7 * x(n-4)$	load C_5 , load $x(n-5)$	
					load C_6 , load $x(n-6)$
					load C_7 , load $x(n-7)$

Figure 5-14. Generic Kernel For Quad ALU FIR

The generic kernel requires four MACs and two parallel loads. Example 5-35 illustrates how the kernel in Figure 5-14 is implemented in a single instruction.

Example 5-35. Single instruction quad ALU generic filter kernel

```
y(n) += C * D1      y(n+1) += C * D2      y(n+2) += C * D3      y(n+3) += C * D4
Load C, Copy D3 to D4, Copy D2 to D3, Copy D1 to D2, Load D1
```

To provide delay reuse, the delays are copied by using temporary variables D1, D2, D3 and D4 as a delay line. This imposes a requirement on the kernel to perform two MACs and five move operations (two loads and three copies) in a single instruction.

Example 5-36 contains an example of C simulation code which implements the generic kernel shown in Figure 5-14 on page 5-40.

Example 5-36. FIR_A4S4 quad ALU, quad sample C simulation code

```
#include <prototype.h>
#include <stdio.h>

#define DataBlockSize40    // size of data block to process
#define FirSize           8    // number of coefficients in FIR

Word16 DataIn[DataBlockSize] = {
    328, 9830, 8192, -6553, -3277, 3277, 3277, -6553, -9830, 4915,
    8192, -6553, 328, 9830, 4915, -6553, -3277, 3277, 3277, -9830,
    4915, -3277, -9830, 8192, -6553, 328, 9830, -6553, 3277, 3277,
    3277, 328, 9830, 4915, -3277, -9830, 8192, -6553, -6553, 3277
};

Word16 Coef[FirSize] = {
    3277, 6553, -9830, -6553, -4915, 3277, 8192, -6553
};

Word16 Delay[FirSize+3];

#define IncMod(a) (a=((a+1)%(FirSize+3)))
#define DecMod(a) (a=((a+FirSize+2)%(FirSize+3)))

int main()
{
    int DelayPtr;
    Word32 sum1,sum2,sum3,sum4;
    Word16 D1,D2,D3,D4;
    int i,j;

    DelayPtr = 0; // init delay ptr

    for (i = 0; i < DataBlockSize; i += 4) { // do 4 samples at a time

        Delay[DelayPtr] = DataIn[i];    DecMod(DelayPtr);
        Delay[DelayPtr] = DataIn[i+1]; DecMod(DelayPtr);
        Delay[DelayPtr] = DataIn[i+2]; DecMod(DelayPtr);
```

```
Delay[DelayPtr] = DataIn[i+3];

sum1 = 0;           // init sum to zero
sum2 = 0;           // init sum to zero
sum3 = 0;           // init sum to zero
sum4 = 0;           // init sum to zero

D4 = Delay[DelayPtr];   IncMod(DelayPtr);
    D3 = Delay[DelayPtr];   IncMod(DelayPtr);
    D2 = Delay[DelayPtr];   IncMod(DelayPtr);

for (j = 0; j < FirSize / 4 ; j++) { // evaluate FIR
    D1 = Delay[DelayPtr];   // get delay
    IncMod(DelayPtr);

        sum1 = L_mac ( sum1, Coef[4*j], D1 );
        sum2 = L_mac ( sum2, Coef[4*j], D2 );
        sum3 = L_mac ( sum3, Coef[4*j], D3 );
        sum4 = L_mac ( sum4, Coef[4*j], D4 );

    D4 = Delay[DelayPtr];   // get delay
    IncMod(DelayPtr);

        sum1 = L_mac ( sum1, Coef[4*j+1], D4 );
        sum2 = L_mac ( sum2, Coef[4*j+1], D1 );
        sum3 = L_mac ( sum3, Coef[4*j+1], D2 );
        sum4 = L_mac ( sum4, Coef[4*j+1], D3 );

    D3 = Delay[DelayPtr];   // get next delay
    IncMod(DelayPtr);

        sum1 = L_mac ( sum1, Coef[4*j+2], D3 );
        sum2 = L_mac ( sum2, Coef[4*j+2], D4 );
        sum3 = L_mac ( sum3, Coef[4*j+2], D1 );
        sum4 = L_mac ( sum4, Coef[4*j+2], D2 );

    D2 = Delay[DelayPtr];   // get next delay
    IncMod(DelayPtr);

        sum1 = L_mac ( sum1, Coef[4*j+3], D2 );
        sum2 = L_mac ( sum2, Coef[4*j+3], D3 );
        sum3 = L_mac ( sum3, Coef[4*j+3], D4 );
        sum4 = L_mac ( sum4, Coef[4*j+3], D1 );
}
DecMod(DelayPtr);

printf("Index: %d, output: %d\n",i,round(sum1));
printf("Index: %d, output: %d\n",i+1,round(sum2));
printf("Index: %d, output: %d\n",i+2,round(sum3));
printf("Index: %d, output: %d\n",i+3,round(sum4));
```


5.4.3 General Hints

In addition to the specific techniques described in the previous sections, there are a number of general guidelines that you should follow when writing source code, in order to assist the optimizer to produce the most efficient results. These guidelines are described in the sections that follow.

5.4.3.1 Software pipelining

The optimizer implements sophisticated levels of software pipelining, saving you the need to introduce software pipelining into your source code. It is important that you do not include any manual form of software pipelining into your source code, as this can conflict with the algorithms used by the optimizer, resulting ultimately in less efficient optimization.

Example 5-37 shows two forms of source code for the same loop. The first version contains no pipelining, and is the recommended source code form. This will generate more efficient and smaller code than the second version, which pipelines the first iteration at the C level outside the loop. The type of manual pipelining shown in the second version should be avoided.

Example 5-37. Avoiding software pipelining in source code

1. No pipelining (recommended)

```
L_R = 0;
for (J = 0; J < S_LEN; J++)
    L_R = L_mac(L_R, WBasisVecs[J + (I * S_LEN)], WInput[J]);
```

2. Manual pipelining (not recommended)

```
L_R = L_mult(WBasisVecs[I * S_LEN], WInput[0]);
for (J = 1; J < S_LEN; J++)
    L_R = L_mac(L_R, WBasisVecs[J + (I * S_LEN)], WInput[J]);
```

5.4.3.2 Passing and returning large structs

Instead of passing and returning large structs using their value, use pointers to large structs wherever possible.

5.4.3.3 Arithmetic operations

Whenever you can, use constants instead of variables for shift, division, or remainder operations.

5.4.3.4 Local variables

Any local variable that you specify should be initialized before it is used.

5.4.3.5 Resource limitations

The SC100 architecture provides a total of 16 Dn registers and 16 Rn registers. If the number of active variables is greater than the number that the registers can accommodate, the compiler maps the extra variables to memory, resulting in less efficient code.

For best results, you should take account of these physical limitations when writing your source code. For example, when preparing a set of instructions to execute in one cycle, remember that there is a restriction on the number of operands that can be used in a single cycle.

5.5 Optimizer Assumptions

The optimizer uses the information passed to it by the compiler, in order to ensure that the optimizations applied during the various optimization stages do not affect the original accuracy of the program.

At the time that the compiler accumulates this information, it assumes that only two types of variables can be accessed while inside a function, either indirectly through a pointer or by another function call:

- Global variables, meaning all variables within the file scope or application scope
- Local variables, whose addresses are retrieved implicitly by the automatic conversion of array references to pointers, or explicitly by the `&` operator

If your programs conform to the standard ANSI/ISO version of C, this assumption does not affect your code. If the code that you are compiling is not standard, and it violates this assumption, the optimization process may adversely affect the behavior of the program.

To avoid unexpected results, and to ensure that your program executes correctly once optimized, follow the coding guidelines listed below:

- Don't make assumptions based on memory layout when using pointers. For example, if `x` points to the first member of a structure, `x+1` may not necessarily point to the second member of the same structure. Similarly, if `y` is defined as a pointer to the first declared variable in a list, do not assume that `y+1` points to the second variable in the list.
- When referencing an array, keep the references inside the array bounds.
- Ensure that all the required arguments are passed to functions.
- When subscribing one array, don't access another array indirectly. For example, if in the construct `x[y-x]`, `x` and `y` are the same type of array, the construct is equivalent to `*(x+(y-x))`, which is equivalent to `*y`. Thus the construct actually references the array `y`.
- When pointing to objects, don't reference outside the bounds of these objects. The optimizer assumes that all references of the form `*(p+i)` apply within the bounds of the variable(s) to which `p` points.
- When the need arises for variables that are accessed by external processes, be sure to declare the variables as `volatile`. Use this keyword judiciously, as it may have adverse effects on optimization.

Chapter 6

Runtime Environment

This chapter describes the startup code used by the SC100 C compiler, the layout and configuration of memory, and the calling conventions which the compiler supports.

It contains the following sections:

- [Section 6.1, “Startup Code,”](#) provides details of the runtime code used for system initialization and finalization.
- [Section 6.2, “Memory Models,”](#) describes the two memory models supported by the compiler.
- [Section 6.3, “Memory Layout and Configuration,”](#) describes the way that the compiler maps memory, and explains how you can change this configuration to suit your requirements.
- [Section 6.4, “Calling Conventions,”](#) describes the stack-based and other calling conventions that the compiler supports.

6.1 Startup Code

The compiler runtime startup code consists of the following components:

- Initialization code, which is executed when the program is initiated and before its main function is called
- Finalization code, which controls the closedown of the application after the program’s main function terminates
- Entry points for low level I/O services
- The interrupt vector table
- Support for debugging tools

The entire startup code is contained in a single assembly code file, named `crtsc100.asm`, which is located in the directory `install-dir/src/rtlib`. The object module for this file is named `crtsc100.e1n`, and is located in the directory `install-dir/lib`.

The compiler startup code contains two phases:

- **Bare board startup code**, which is used for programs which execute without the support of any runtime executive or operating system. This phase resets the interrupt vector and initializes all necessary hardware registers.
- **C environment startup code**, which is a mandatory phase for all configurations. This phase initializes the runtime structure of the application for the C environment, and includes the finalization code used following termination of the program.

6.1.1 Bare Board Startup Code

The bare board startup phase assumes that no operating system or runtime executive is running. It performs the various actions which are normally carried out automatically by the operating system or runtime executive, as follows:

1. The reset interrupt vector is set to point to the system entry point `__crt0_start`, as if the system has just been reset. The interrupt vector table holds the addresses of all interrupt handlers. The first entry in this table is the system entry point. All other entries in the interrupt vector table point by default to the `abort` function. See Chapter 7 for more information about `abort` and other runtime functions.
2. The hardware registers are initialized as follows:
 - The four modulo (M) registers (m0–m3) are initialized to linear addressing.
 - The status register is set to an initial value taken from the linker command file used at link time. This file includes a label `SR_setting`, which defines the initial value to be assigned to the status register following system reset. Table 6-1 shows the default status register settings.

Table 6-1. Status Register Default Settings

Setting Type	Value
Mode:	Exception mode
Interrupt level:	7
Saturation:	On
Rounding mode:	NEAREST_EVEN

3. If the system includes a timer, the timer is activated.
4. The bare board startup phase terminates by jumping to the C environment startup code entry point, `__start`.

6.1.2 C Environment Startup Code

The C environment startup phase is applicable to all programs. The entry point for this phase is `__start`. This phase includes initialization code used prior to program start, and finalization code used after the application terminates.

6.1.2.1 C environment initialization code

The following initialization actions are executed before the application starts:

1. The memory map is set up and initialized. The stack pointer (SP) value is loaded into memory by the stack start address, located at `StackStart`. This label is defined in the linker command file and used by the linker at link time. For further information about the memory map, see [Section 6.2, “Memory Models.”](#)
2. If the `-mrom` option has been specified in the shell command line, initialized variables are copied from ROM into RAM. This option is required for applications which do not use a loader.
3. The `argv` and `argc` arguments are set up.
4. Interrupts are enabled. Until this point, interrupts have been disabled.
5. The application main procedure entry point is called using the function `main`.

6.1.2.2 Initialization of variables

If your system uses a loader, this will by default initialize all variables. In systems that do not include a loader, it is important that you specify the `-mrom` option when you compile the final version of your application, to ensure that the initialized variables are copied from ROM into RAM at startup.

Note: Before a C program executes, certain global variables may assume the assignment of an initial value of zero. The compiler does not preinitialize variables automatically. You must ensure that your code includes explicit initialization of any variable that must have an initial value of zero.

6.1.2.3 C environment finalization code

The following steps are activated on return from the application main function:

- The runtime function `exit` is called. This terminates any I/O services which have not yet terminated, and stops the processor by issuing the `stop` instruction.

Note: Certain embedded real time applications never terminate. Such termination activities do not usually pertain to embedded applications, but may be of use during early development and debugging stages.

6.1.2.4 Low level I/O services

The C environment startup code includes the input and output of low level, buffered I/O services. The code uses calls to `__send` and `__receive` in order to interface with debugging tools and/or runtime systems.

6.1.3 Configuring Your Startup Code

If the default runtime setup does not match your configuration, you need to modify your startup code accordingly.

To create your own runtime configuration code, follow the steps described below:

1. Make your own copy of the default startup file, `crtsc100.asm`, with a name of your choice, as shown in the following example:

Example 6-1. Creating a new startup file

```
cp install-dir/src/rtlib/crtsc100.asm mysc100.asm
```

2. Make the required changes to the new file.
3. Assemble the modified file, as shown in Example 6-2.

Example 6-2. Assembling the modified startup file

```
asmsc100 -b -l mysc100.asm
```

The generated object file has the same file name as the source file, and the extension `.eln`. In this example, the object file generated is `mysc100.eln`.

4. Use the modified file by specifying the `-crt` option in the shell command line, as shown in Example 6-3, to ensure that the modified startup file is used at link time.

Example 6-3. Using the modified startup file

```
ccsc100 -crt mysc100.eln my-object-files.eln
```

6.2 Memory Models

The compiler architecture supports both 16-bit and 32-bit addresses. If the application is small enough to allow all static data to fit into the lower 64K of the address space, then more efficient code can be generated. This mode, the small memory model, is the default, and assumes that all addresses are 16-bit.

The big memory model does not restrict the amount of space allocated to addresses. This model is selected with the option `-mb` when the shell is invoked.

When the compiler uses the big memory model to access a data object, whether static or global, it must use a longer instruction that includes a 32-bit address. This operation requires an additional word, and as a result it produces code that is larger, and in some cases, slower, than a similar operation using the small memory model.

Example 6-4 illustrates the use of the `move.l` instruction in the big and small memory models. In this example, the assembler interprets the address in the first instruction as a 32-bit address, and allocates the maximum space for it. In the second instruction, the `<` symbol indicates to the assembler that this address fits into a 16-bit space, thus preventing the allocation of unnecessary program memory.

Example 6-4. Big and small memory models

Big memory model (3 16-bit words):

```
move.l      address,d0
```

Small memory model (2 16-bit words):

```
move.l      <address,d0
```

Certain instructions can be used only in small memory mode. If `<` is omitted in conjunction with these instructions, an error results. Example 6-5 shows the instruction `bmset.w`, which sets bit `#zero` in the specified address, and is valid only in small memory mode.

Example 6-5. Small memory mode instruction

```
bmset.w #0001,<address
```

For maximum efficiency, it is recommended that you place data in the lower 64K of the memory map, in order to enable the compiler to use small memory mode.

6.2.1 Linker Command Files

The SC100 linker refers to a linker command file at link time, for various runtime values, addresses and labels. Two linker command files are provided, one for each memory mode. These files are `crtscsmm.cmd`, used in small memory mode, and `crtscbmm.cmd`, used when big memory mode is selected. Both files are located in the `install-dir/etc` directory.

6.3 Memory Layout and Configuration

The SC100 default memory layout is a single linear block which is divided into data and code areas. C programs generate code and data in sections. The compiler places each of these sections in its own continuous space in memory.

The default layout of the SC100 memory is illustrated in Figure 6-1.

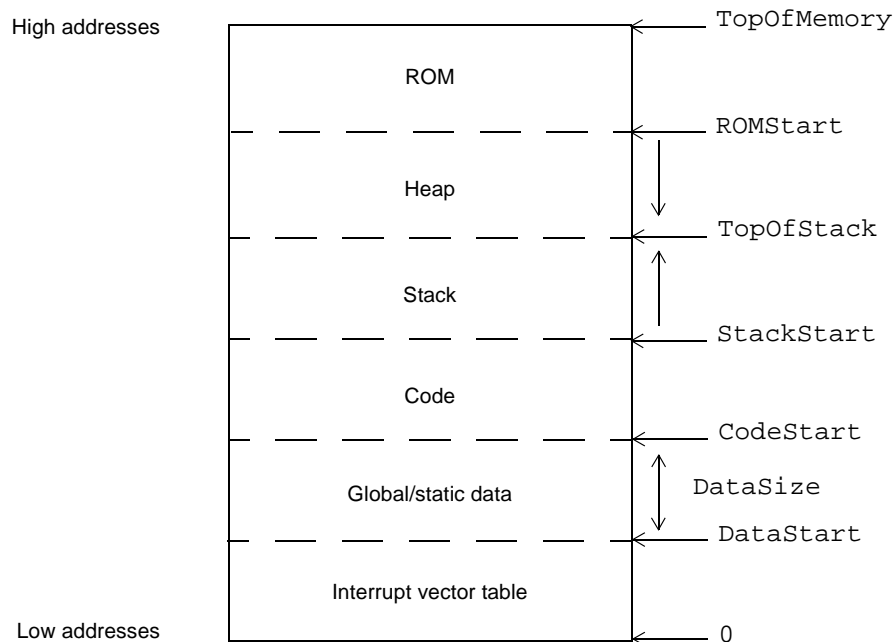


Figure 6-1. SC100 Default Memory Layout

Both memory models use the same default layout, but with different default values that define the distribution of the memory areas, as shown in Table 6-2 and Table 6-3 on page 6-7. These default values can be changed, and the memory map can be configured to meet your specific requirements, as described in [Section 6.3.3, “Configuring the Memory Map.”](#)

The layout and functionality of the stack and heap are common to both memory models, and are described in the sections that follow.

The default memory map values for the small memory model are listed in Table 6-2. These values are held in the file `crtscsmm.cmd`.

Table 6-2. Small Memory Model Default Values

From	Default value	To	Default value	Contents
0		0x1ff		Interrupt vector table
0x200		DataSize-1	0xffffd	Global and static variables
CodeStart	0x10000	StackStart-1	0x27fff	Program code
StackStart	0x28000	TopOfStack	0x7fff0	Stack and heap
ROMStart	0x80000	TopOfMemory	0xfffff	ROM

Table 6-3 lists the default memory map values for the big memory model. These values are held in the file `crtscbmm.cmd`.

Table 6-3. Big Memory Model Default Values

From	Default value	To	Default value	Contents
0		0x1ff		Interrupt vector table
0x200		DataSize-1	0x17fff	Global and static variables
CodeStart	0x18000	StackStart-1	0x3ffff	Program code
StackStart	0x40000	TopOfStack	0x7fff0	Stack and heap
ROMStart	0x80000	TopOfMemory	0xfffff	ROM

6.3.1 Stack and Heap Configuration

The heap and stack are allocated from the same area of memory and must be contiguous. The compiler always treats the stack and heap as a continuous area of memory. The other sections of memory can be distributed, and there are no restrictions relating to their location.

6.3.1.1 Runtime stack

The compiler allocates an area of memory to the runtime stack, which is used for the following purposes:

- Allocation of local variables
- Passing arguments to functions
- Saving function return addresses
- Saving temporary results

The stack is allocated in the area above the space used for code, and grows in an upward direction toward the top of memory. The compiler uses the SP register to manage this stack.

The SC100 architecture includes two stack pointers:

- NSP, used when the processor is running in Normal mode
- ESP, used when the processor is running in Exception mode

As shown in Table 6-1 on page 6-2, the default mode at initialization is Exception mode.

The compiler makes no assumptions about which stack pointer to use, and uses the pointer for the current processor mode to point to the address at the top of the stack.

When the system is initialized, the stack pointer for the current mode is set by default to the address of the location directly after the code area, as defined in `StackStart` in the linker command file. The actual address of the stack is determined at link time.

The stack pointer for the current processor mode is automatically incremented by the C environment at the entry to a function. This ensures that sufficient space is reserved for the execution of the function. At the function exit, the stack pointer is decremented, and the stack is restored to its previous size prior to function entry. If your application includes assembly language routines and C code, you must ensure at the end of each assembly routine that the current stack pointer is restored to its pre-routine entry state.

Note: If you change the default memory configuration, remember to allow sufficient space for the stack to grow. If a stack overflow occurs at runtime, this will cause your program to fail. The compiler does not check for stack overflow during compilation or at runtime.

6.3.1.2 Dynamic memory allocation (heap)

The runtime libraries supported by the compiler include a number of functions which enable you to allocate memory dynamically for variables. See Chapter 7 for details of the runtime libraries supported. Since C does not support the dynamic allocation of memory, the compiler assigns an area of memory as a heap for this purpose.

The compiler allocates memory from a global pool for the stack and the heap together. The lower address of the area assigned to the stack and heap is defined in `StackStart`, in the linker command file. The heap starts at the top of memory, and is allocated in a downward direction toward the stack.

Objects that are dynamically allocated are addressed only with pointers, and not directly. The amount of space that can be allocated to the heap is limited by the amount of available memory in your system.

To make more efficient use of the space allocated to data, you can use the heap to allocate large arrays, instead of defining them as static or global.

For example, a definition such as `struct large array1[80];` can be defined using a pointer and the `malloc` function, as illustrated in Example 6-6.

Example 6-6. Allocating large arrays from the heap

```
struct large *array1;  
array1 = (struct large *)malloc(80*sizeof(struct large));
```

6.3.2 Static Data Allocation

When you compile your application in non-global mode, the allocations for each file are assigned to different sections of data memory. At link time these are dispatched to different addresses.

When compiling in global mode, the compiler uses the same data section for all allocations. If you want to override this and to instruct the compiler to use non-contiguous data blocks, you can edit the machine configuration file to define the exact memory map of the system that you want to use. For more details, see [Section 6.3.4, “Machine Configuration File.”](#)

6.3.3 Configuring the Memory Map

The default values in the SC100 memory map are easily configurable, by modifying the linker command file. When making such changes, it is important that you ensure that the code size and data size values that you specify do not overlap.

The stack and the heap must be always be located together in one contiguous area of memory. The compiler makes no assumptions about the layout of the other sections of memory, which can be split and distributed over non-contiguous parts of memory, as required.

Section 6.3.3.1, “[Memory map configuration example](#),” provides an example of a requirement for a modified memory map configuration, and describes the changes to be defined in the linker command file for this sample configuration.

6.3.3.1 Memory map configuration example

This example assumes that you have a system with non-contiguous memory, and would like to configure the memory as follows:

- All code placed in external memory (addresses 0x10000000 through 0x10100000)
- All data placed in internal memory
- Some local memory reserved for the most frequently used functions and overlays (addresses 0x10000 through 0x20000)
- All data placed in the lower 64K addresses, in order to be able to use the small memory model compilation mode

The memory map that meets these requirements is shown in Example 6-7:

Example 6-7. Modified memory map configuration

From	To	Contents
0	0x1fff	Interrupt vectors
0x200	0xffffd	Global and static variables
0x10000	0x1fffff	Local code
0x20000	0x7fff0	Stack and heap
0x80000	0xfffff	ROM
0x10000000	0x100fffff	External code

Example 6-8 shows the definitions in the `crtscsmm.cmd` file that specify this memory map configuration:

Example 6-8. Modified memory configuration in the linker command file

```
.provide _DataSize,    0x10000      ; Sets the data size.
.provide _CodeStart,  0x10000000   ; Sets the loader code start address.
.provide _StackStart, 0x20000      ; Sets the stack start address;
                                ; the stack grows upwards.
.provide _TopOfStack, 0x7fff0      ; The heap start address;
                                ; the heap grows downwards.
.provide _ROMStart,   0x80000      ; Sets the ROM start address.
```

6.3.4 Machine Configuration File

The machine configuration file contains the following:

- Information about data types and alignment requirements, used by the compiler for reference. This data must **not** be changed.
- Memory structure information, used by the compiler to allocate variables in the data sections of memory. This information can be modified if required.

By default, the compiler uses the file `proc.config`, located in the `install-dir/etc` directory. A different machine configuration file can be specified using the `-mc` option in the shell command line.

The SC100 memory structure consists of physical and logical memory maps, as follows:

- Physical memory is divided into several memory spaces. Each memory space is a physical entity consisting of a data bus and an address bus. A physical memory space is defined in terms of its size in words and the width of its address bus, and comprises blocks of words with contiguous addresses, described as physical memory areas.
- Logical memory areas are defined as blocks of memory words with contiguous addresses. These words are used by the compiler as if they were in physical memory areas. The addresses of the logical areas are mapped as offsets to physical memory addresses at link time.

This dual memory map structure provides a high degree of flexibility during the loading of application code.

6.3.4.1 Defining the memory configuration

Each memory space is defined individually in the machine configuration file, by specifying a space identifier and a description, comprising:

- Memory space type: program or data.
- Word size, in bytes.
- Area list, defining one or more logical areas.
- The addresses in the logical areas, as positive integers, used as offsets to physical memory areas.
- Physical area type: single-port RAM (`ramsp`), dual-port RAM (`ramdp`) or ROM (`rom`).
- Attached spaces (optional). This is used for dual-port RAM only, when `ramdp` is the defined area type, to specify the two memory spaces. It is important that the code ensures address consistency between the corresponding spaces.

The syntax for defining a memory space is as follows:

```
space definition:
  define space <space identifier>:
    space_type;
    word_size;
    area_list;
  end define
;
space_type:
  program | data
;
word_size:
  word : byte_number
;
area_list:
  area | area_list area
;
area:
  address_value .. address_value : area_type opt_attached_spaces ;
;
area_type:
  ramsp | ramdp | rom
;
opt_attached_spaces:
  [ space_number , space_number ] |
;
```

In Example 6-9, a one-word data space is defined, providing one logical area that can be used for the allocation of variables.

Example 6-9. Defining a data memory space

```
define space data_0 :
  data;
  word : 2;
  0x0000 .. 0xffff : ramsp;
end define
```

Example 6-10 shows the definition of a 2-word program space in ROM.

Example 6-10. Defining a program memory space

```
define space pgm :
  program;
  word : 4;
  0x0000 .. 0x3fff : rom;
end define
```

At link time, these areas are mapped to the relevant physical memory space, and the actual addresses are calculated as offsets to the physical space starting address.

A data space can be divided into multiple logical areas, as shown in Example 6-11. When the compiler executes in global mode, it divides memory into these logical areas, and allocates variables accordingly.

Example 6-11. Defining multiple memory spaces

```
define space data_1 :
  data;
  word : 2;
  0x0000 .. 0x3fff : ramsp;
  0x0800 .. 0xffff : ramdp [data_0,data_1];
  0x10000 .. 0x13fff : ramsp;
  0x40000 .. 0x47fff : ramsp;
end define
```

Note: If you define new memory spaces in the machine configuration file, it is important that you also add these space definitions in the linker command file, to enable the linker to locate them at link time.

6.3.5 Application Configuration File

The application configuration file contains information about the interaction between the application software and the hardware. This file indicates to the compiler how to compile specific software units in order to ensure efficient sharing of hardware resources, in particular memory space. This information can be modified, to suit the requirements of your application.

The default application configuration file is named `minimal.appli`, and is located in the `install-dir/etc` directory. A different application configuration file can be specified, using the `-ma` option.

This file contains the following functional section types:

- Schedule section, which defines the entry points for the software units in the application, and their overlay capabilities for local variables. See [Section 6.3.5.2, “Schedule section,”](#) for details.
- Binding section, which specifies the links between software interrupt routines and hardware interrupt vectors, and between software-defined variables and fixed memory addresses. See [Section 6.3.5.3, “Binding section,”](#) for details.
- Overlay section, which specifies the overlay capabilities of global variables. See [Section 6.3.5.4, “Overlay section,”](#) for details.

6.3.5.1 File structure and syntax

More than one section of each type can be included in the file. The order in which the sections are defined in the file is unimportant. Each of the section types is optional and can be omitted.

The syntax of the application configuration file is as follows:

```
translation_unit:
    header_section
    configuration section_list
end configuration
;
header_section:
    opt_version
;
opt_version:
    version string_content |
;
section_list:
    section | section_list section
;
section:
    schedule_section | binding_section | overlay_section
;
```

6.3.5.2 Schedule section

The schedule section defines the entry point structure of an application, by specifying a "call tree". The call tree root is a C function name that defines the starting entry point for an application. Each node in the call tree is the name of an entry point of a unit that can be called during the execution of the application.

Each call tree node is defined as a call tree item, and is given a `ct` number that is unique for the application. A call tree item can be one of three types:

- Background task, identifying the main entry point, defined as `main`
- Interrupt handler, identifying an interrupt routine entry point, defined as `it_entry`, with a number that is used by the binding section to link to the associated hardware interrupt vector
- Task entry point, defined as `task_entry`, for example, an operating system task

The schedule section can optionally include an overlay specification, which informs the compiler which groups of local variables can use the same memory location during execution of the application. The compiler is able to overlay groups of local variables automatically, but only when it is clear that the two sets of variables do not share the same lifetime, and are therefore not active simultaneously. By specifying overlays in this file, you provide the necessary information in advance to help the compiler make more efficient use of memory space.

The overlay specification in the schedule section relates to local variables only. Overlays for global variables are specified in the overlay section, as described in [Section 6.3.5.4, "Overlay section."](#)

The syntax of the schedule section is as follows:

```
schedule_list:
    schedule_elmt | schedule_list schedule_elmt
;
schedule_elmt:
    call_tree_list ; opt_overlay_spec
;
call_tree_list:
    call_tree_item | call_tree_list call_tree_item
;
call_tree_item:
    ct [int_constant] : main = ident ; |
    ct [int_constant] : it_entry int_constant = ident ; |
    ct [int_constant] : task_entry = ident ;
;
opt_overlay_spec:
    overlay = entry_overlay_list ; |
;
entry_overlay_list:
    [group_list]
;
group_list:
    group | group_list, group
;
group:
    [entry_number_list]
;
entry_number_list:
    entry_number | entry_number_list, entry_number
;
entry_number:
    ct[int_constant] | int_constant
;
```


Example 6-12 defines two entry points, in addition to `main`. The function `task1()` is defined as a task entry point and the function `int_entry()` is defined as an interrupt handler.

Note that defining a function as an interrupt handler in the application configuration file is equivalent to using `#pragma interrupt` in the source file. For more details, see [Section 6.4.5, “Interrupt Handlers.”](#)

Example 6-12. Defining additional entry points for an application

```
configuration
schedule
    ct[0] : main = _main;
    ct[1] : task_entry = _task1;
    ct[2] : it_entry 0 = _int_entry;
end schedule
binding
    place __stackX on space 0 at 1;
end binding
end configuration
```

6.3.5.3 Binding section

The binding section performs the following functions:

- Assignment of fixed memory addresses to variables. A full memory address is specified with a memory binding directive, using the following syntax:

```
memory_binding_directive:
    place full_ident on space_identifier at number
```

- Specification of the links between fixed interrupt entries and hardware interrupt vector addresses. An interrupt binding directive is used to specify an interrupt entry number, in the range 0-15, and the corresponding hardware vector number, in the range 1-16, using the following syntax:

```
it_binding_directive:
    place it_vector interrupt_number on space_identifier at vector_number
```

The syntax of the binding section is as follows:

```
binding_directive:
    memory_binding_directive | it_binding_directive
;
binding_directive_list:
    binding_directive binding_directive_list ; binding_directive
;
binding section:
    binding
        binding_directive_list
    end binding
;
```

In Example 6-13, the location of global variable `mem` is fixed at absolute address `0x2000`:

Example 6-13. Placing a variable at an absolute location

```
configuration
schedule
  ct[0] : main = _main;
  ct[1] : it_entry 0 = _int_entry;
end schedule
binding
  place ___stackX on space 0 at 1;
  place _mem on space 0 at 0x2000;
end binding
end configuration
```

6.3.5.4 Overlay section

The overlay section specifies how the compiler should overlay global variables in order to further reduce the amount of memory required for data. As with local variables, in many cases the compiler can automatically detect that two data objects do not share the same lifetime and as a result, the memory allocated to these objects can be shared. This feature is needed for cases where the compiler cannot identify statically that the object lifetimes of global variables do not conflict.

Defining the overlay specification for global variables includes the following:

- Grouping the global variables into sets that can share the same memory space. In the overlay section syntax, the full identity is specified for each global variable, or list of variables, and defined as `symbol_list`.
- Defining each set of global variables as a `symbol_group`, associated with a `symbol_list` and an identifying group number.
- Specifying compatibility clauses that define which symbol groups can be overlaid, using the keyword `discern`.
- Specifying a list of compatibility clauses to indicate which symbol groups in the application can share the same memory space.

The syntax of the overlay section is as follows:

```
overlay section:
  overlay
    opt_overlay_spec
    compatibility_list
  end overlay
;
symbol_list:
  full_ident | symbol_list, full_ident
;
symbol_group:
  SG [number] = [symbol_list] ;
;
```

```
symbol_group_list:
    symbol_group | symbol_group_list symbol_group
;
sg_ref:
    SG [number]
;
sg_list:
    sg_ref | sg_list, sg_ref
;
compatibility_clause:
    discern_sg_ref : sg_list ;
;
compatibility_list:
    compatibility_clause | compatibility_list compatibility_clause
;
```

Example 6-14 shows an overlay section that specifies that the the application will never access the two global arrays, `arr1` and `arr2`, at the same time, and they can therefore share the same physical memory location.

Example 6-14. Defining global variable overlays

```
configuration
schedule
    ct[0] : main = _main;
end schedule
binding
    place ___stackX on space 0 at 1;
end binding
overlay
    sg[0] = [_arr1];
    sg[1] = [_arr2];
    discern sg[0] : sg[1];
end overlay
end configuration
```

6.4 Calling Conventions

The compiler supports a stack-based calling convention. Additional calling conventions are also supported. Calling conventions can be mixed within the same application.

Specific calling conventions can be enforced using pragmas. For further information about the use of pragmas for this purpose, refer to [Section 3.4.5, “Pragmas,”](#) on page 3-54.

When compiling in separate compilation mode, non-static functions use the stack-based calling convention.

6.4.1 Stack Pointer

The SP register serves as the stack pointer, which points to the first available location. The stack direction is toward higher addresses, meaning that a push is implemented as `(sp)+`. The stack pointer must always be 8-byte aligned.

6.4.2 Stack-Based Calling Convention

The following calling conventions are supported:

- The first (left-most) function parameter is passed in `d0` if it is a numeric scalar or in `r0` if it is an address parameter, regardless of its size. The second function parameter is passed in `d1` if it is a numeric scalar, or in `r1` if it is an address parameter, regardless of its size. The remaining parameters are pushed on the stack. Long parameters are pushed on the stack using little endian mode, with the least significant bits in the lower addresses.
- Structures and union objects that can fit in a register are treated as numeric parameters, and are therefore candidates to be passed in a register.
- Numeric return values are returned in `d0`. Numeric address return values are returned in `r0`. Functions returning large structures, meaning structures that do not fit in a single register, receive and return the returned structure address in `r2`. The space for the returned object is allocated by the caller.
- Functions with a variable number of parameters allocate all parameters on the stack.
- Parameters are aligned in memory according to the base parameter type, with the exception of characters and unsigned characters that have a 32-bit alignment.

The following registers are saved by the caller: `d0-d5`, `r0-r5`, `n0-n3`.

The following registers are saved by the callee, if actually used: `d6-d7`, `r6-r7`.

The compiler assumes that the current settings of the following operating control bits are correct:

- Saturation mode
- Round mode
- Scale bits

The application is responsible for setting these mode bits correctly.

Example 6-15 shows two function calls and the parameters that are allocated for each call.

Example 6-15. Function call and allocation of parameters

Function call:

```
foo(int a1, struct fourbytes a2, struct eightbytes a3, int *a4)
```

Parameters:

```
a1 - in d0
a2 - in d1
a3 - in stack
a4 - in stack
```

Function call:

```
bar(long *b1, int b2, int b3[])
```

Parameters:

```
b1 - in r0
b2 - in d1
b3 - in stack.
```

The stack-based calling convention must be used when calling functions that are required to maintain a calling stack.

The compiler is able to use optimized calling sequences for functions that are not exposed to external calls.

Local and formal variables are allocated on the stack and in registers.

Table 6-4 summarizes register usage in the stack-based calling convention.

Table 6-4. Register Usage in the Stack-based Calling Convention

Register	Used as	Caller Saved	Callee Saved
d0	First numeric parameter Return numeric value	+	
d1	Second numeric parameter	+	
d2-d5		+	
d6-d7			+
d8-d15		+	
r0	First address parameter Return address value	+	
r1	Second address parameter	+	
r2	Big object return address	+	
r3-r5		+	
r6	Optional argument pointer		+
r7	Optional frame pointer		+
n0-n3, m0-m3		+	

6.4.3 Optimized Calling Sequences

A stack-less convention may be used when calling functions that are not reentrant, if this technique generates more efficient code than other conventions.

This convention will be used only if the function is not visible to external code.

When using this calling convention, local variables may be allocated statically, meaning not on a stack. Functions with mutually exclusive lifetimes may share space for their local variables.

Actual parameters are placed by the calling function at the locations allocated for the formal parameters in the called function. The compiler may use registers and memory locations as required when allocating locations for the formal parameters.

Under this calling convention, all registers are classified as caller-saved.

Return values from functions are placed in the space allocated for the function return value in the calling function. The compiler may use a register or a memory location as the space for the function return value.

6.4.4 Stack Frame Layout

The stack pointer points to the top (high address) of the stack frame. Space at higher addresses than the stack pointer is considered invalid and may actually be unaddressable. The stack pointer value must always be a multiple of eight.

Figure 6-2 shows typical stack frames for a function, indicating the relative position of local variables, parameters and return addresses.

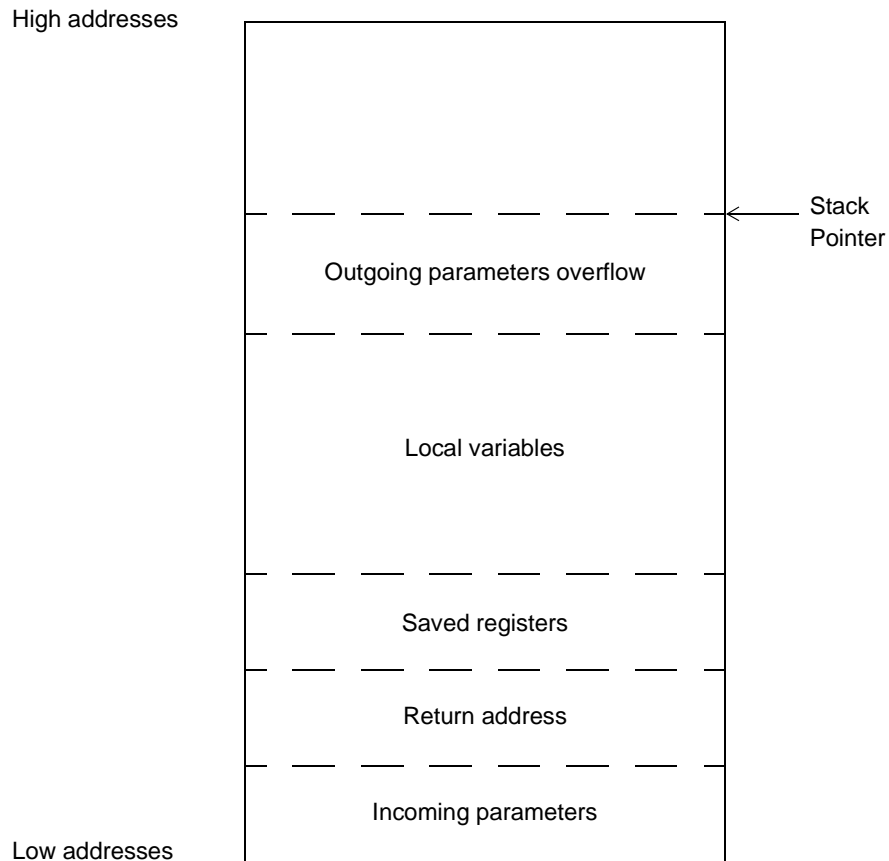


Figure 6-2. Stack Frame Layout

The caller must reserve stack space for return variables that do not fit in registers. This return buffer area is typically located with the local variables. This space is typically allocated only for functions that make calls that return structures. Beyond these requirements, a function is free to manage its stack frame as necessary.

The outgoing parameter overflow block is located at the top (higher addresses) of the frame. Any incoming argument spill generated for `varargs` and `stdargs` processing must be at the bottom (low addresses) of the frame.

The caller puts argument variables that do not fit in registers into the outgoing parameter overflow area. If all arguments fit in registers, this area is not required. A caller has the option to allocate argument overflow space sufficient for the worst case call, use portions of this space as necessary, and/or leave the stack pointer unchanged between calls.

Local variables that do not fit into the local registers are allocated space in the local variables area of the stack. If there are no such variables, this area is not required.

6.4.5 Interrupt Handlers

Functions which require no parameters and return no result can be designated as interrupt handler functions. The process of creating an interrupt handler function includes:

- Defining the function as an interrupt handler
- Linking the function to the appropriate interrupt vector entry

An interrupt handler can be defined in one of two ways:

- Using `#pragma interrupt` in the source code. For more detail about this pragma, refer to [Section 3.4.5.3.4, “Defining a function as an interrupt handler,”](#) on page 3-58.
- Defining an interrupt entry point in the application, by editing the schedule section of the application configuration file, as described in [Section 6.3.5, “Application Configuration File.”](#)

To create the link between the function and the interrupt vector entry, you can use any one of the following options:

- In the code that calls the function, place a call to the handler function in the interrupt vector entry.
- Use the `signal.h` library function to insert a call to the interrupt handler function into the required interrupt vector entry. For syntax details, see [Section 7.7, “Signal Handling \(signal.h\),”](#) on page 7-10.
- If the function is very small, you can embed it in the interrupt vector entry, by modifying the startup code file, `crtsc100.asm`. The size of each interrupt vector entry is 64 bytes. With this option, there is no need for an explicit call from the vector to the function.

Interrupt handler functions always follow the stack-based calling convention. When an interrupt function is called, the interrupt handler saves all registers and all other resources that are modified by the function. Upon returning from the function all registers and hardware loop state saved at entry are restored to their original state.

Local variables are saved on the stack. Interrupt handlers that are known to be non-interruptible may also allocate data statically.

Return from interrupt is implemented using an RTE instruction.

6.4.6 Frame Pointer and Argument Pointer

The compiler does not use a frame pointer or an argument pointer.

If, however, the use of a frame pointer or an argument pointer is required by external code, `r7` may be allocated as a frame pointer and `r6` as an argument pointer. When these registers are allocated as frame pointer and/or argument pointer they should be saved and restored as part of the function prolog/epilog code.

6.4.7 Hardware Loops

All hardware loop resources are available for use by the compiler. It is assumed that no nesting occurs when entering a function. As a result, a function may use all 4 nesting levels for its own use. An additional side effect of this assumption is that loops that include a function call as part of the loop code cannot be implemented using hardware loops, unless the compiler can infer the nesting level of the called function from static variables known at compilation time.

Loops are nested beginning with loop counter 3 at the innermost nesting level.

6.4.8 Operating Modes

The compiler makes the following assumptions regarding runtime operating modes and the machine state:

- All modulo (M) registers ($m0-m3$) are assumed to contain the value -1 (linear addressing). If the use of an M register is required, the using function must restore the M register to the value -1 before returning or before calling another function.
- No specific settings are assumed for the operating mode settings in the OMR register. The compiler assumes that the default settings in the startup code, including saturation modes, rounding mode and scale bits, are set by the user. You can control and change these operating modes during execution of the application. Refer to the SC100 architecture documentation for further details.

Chapter 7

Runtime Libraries

This chapter describes the C libraries and I/O libraries supported by the SC100 C compiler. Each table in this chapter is organized in alphabetical order, according to the file, function, or constant name in the first column in the table.

Table 7-2 summarizes the ISO standard C libraries that the compiler supports.

Table 7-1. Supported ISO Libraries

Header file	Description	Section	Page
<code>ctype.h</code>	Character Typing and Conversion	7.1	7-2
<code>float.h</code>	Floating Point Characteristics	7.2	7-3
<code>limits.h</code>	Integer Characteristics	7.3	7-7
<code>locale.h</code>	Locales	7.4	7-7
<code>math.h</code>	Floating Point Math	7.5	7-8
<code>setjmp.h</code>	Nonlocal Jumps	7.6	7-10
<code>signal.h</code>	Signal Handling	7.7	7-10
<code>stdarg.h</code>	Variable Arguments	7.8	7-10
<code>stddef.h</code>	Standard Definitions	7.9	7-11
<code>stdio.h</code>	I/O Library	7.10	7-11
<code>stdlib.h</code>	General Utilities	7.11	7-14
<code>string.h</code>	String Functions	7.12	7-16
<code>time.h</code>	Time Functions	7.13	7-19

The non-ISO C library supported by the compiler is shown in Table 7-2. This library contains the built-in intrinsic functions supplied with the compiler.

Table 7-2. Supported Non-ISO Libraries

Header file	Description	Section	Page
<code>prototype.h</code>	Built-in Intrinsic Functions	7.14	7-20

7.1 Character Typing and Conversion (ctype.h)

The `ctype.h` library contains the following function types:

- Testing functions
- Conversion functions

7.1.1 Testing Functions

Table 7-3 lists the testing functions that the compiler supports.

Table 7-3. Testing Functions

Function	Purpose
<code>int isalnum(int)</code>	Tests for <code>isalpha</code> or <code>isdigit</code>
<code>int isalpha(int)</code>	Tests for <code>isupper</code> or <code>islower</code>
<code>int iscntrl(int)</code>	Tests for any control character
<code>int isdigit(int)</code>	Tests for decimal digit character
<code>int isgraph(int)</code>	Tests for any printing character except space
<code>int islower(int)</code>	Tests for lowercase alphabetic character
<code>int isprint(int)</code>	Tests for any printing character including space
<code>int ispunct(int)</code>	Tests for any printing character not space and not <code>isalnum</code>
<code>int isspace(int)</code>	Tests for white-space characters
<code>int isupper(int)</code>	Tests for uppercase alphabetic character
<code>int isxdigit(int)</code>	Tests for hexadecimal digit character

7.1.2 Conversion Functions

Table 7-4 lists the conversion functions that the compiler supports.

Table 7-4. Conversion Functions

Function	Purpose
<code>int tolower(int)</code>	Converts uppercase alphabetic character to the equivalent lower case character
<code>int toupper(int)</code>	Converts lowercase alphabetic character to the equivalent uppercase character

7.2 Floating Point Characteristics (float.h)

The compiler represents floating point numbers using IEEE format (ANSI/IEEE Std 754-1985). Only single precision floating point format is supported.

The contents of `float.h` are listed in Table 7-5.

Table 7-5. Contents of File float.h

Constant	Value	Purpose
FLT_DIG		6 Number of decimal digits of precision
DBL_DIG		6
LDBL_DIG		6
FLT_EPSILON	1.1920929E-07	Minimum positive number χ such that $1.0 + \chi$ does not equal 1.0
DBL_EPSILON	1.1920929E-07	
LDBL_EPSILON	1.1920929E-07	
FLT_MANT_DIG		24 Number of base-2 digits in the mantissa
DBL_MANT_DIG		24
LDBL_MANT_DIG		24
FLT_MAX_10_EXP		38 Maximum positive integers n such that 10^n is
DBL_MAX_10_EXP		38 representable
LDBL_MAX_10_EXP		38
FLT_MAX_EXP		128 Maximum positive integer n such that 2^{n-1} is
DBL_MAX_EXP		128 representable
LDBL_MAX_EXP		128
FLT_MAX	3.4028235E+38	Maximum positive floating point number
DBL_MAX	3.4028235E+38	
LDBL_MAX	3.4028235E+38	
FLT_MIN_10_EXP		-39 Minimum negative integer n such that 10^n is
DBL_MIN_10_EXP		-39 representable
LDBL_MIN_10_EXP		-39
FLT_MIN_EXP		-126 Minimum negative integer n such that 2^{n-1} is
DBL_MIN_EXP		-126 representable
LDBL_MIN_EXP		-126
FLT_MIN	5.8774817E-39	Minimum positive number
DBL_MIN	5.8774817E-39	
LDBL_MIN	5.8774817E-39	
FLT_RADIX		2 Floating point exponent is expressed n radix 2.
FLT_ROUNDS		1 Floating point rounding is to nearest even number.

7.2.1 Floating Point Library Interface (fltmath.h)

This header file defines the software floating point library interface. Most of these functions are called by the code generator of the compiler for floating point expression evaluation. They may also be called directly by user code.

The floating point library supports the full IEEE-754 single-precision floating point standard.

Three configuration parameters and one status word can be used. Each of these is described in the following sections.

- Round_Mode
- FLUSH_TO_ZERO
- IEEE_Exceptions
- EnableFPExceptions

7.2.1.1 Round_Mode

Four rounding modes are supported:

- ROUND_TO_NEAREST_EVEN. The representable value nearest to the infinitely precise intermediate value is the result. If the two nearest representable values are equally near (tie), then the one with the least significant bit equal to zero (even) is the result.
- ROUND_TOWARDS_ZERO. The result is the value closest to, and no greater in magnitude than, the infinitely precise intermediate result.
- ROUND_TOWARDS_MINUS_INF. The result is the value closest to and no greater than the infinitely precise intermediate result (possibly minus infinity).
- ROUND_TOWARDS_PLUS_INF. The result is the value closest to and no less than the infinitely precise intermediate result (possibly plus infinity).

By default, the rounding mode is set to ROUND_TO_NEAREST_EVEN.

Following is an example of changing the round mode to ROUND_TOWARDS_MINUS_INF:

Example 7-1. Changing the round mode

```
#include <fltmath.h>
. . .
Round_Mode = ROUND_TOWARDS_MINUS_INF.
```

7.2.1.2 FLUSH_TO_ZERO

This is a boolean configuration item that sets the behavior of un-normalized numbers. When set to true (default) all un-normalized values are flushed to zero. This leads to better performance, but a smaller dynamic range.

For example, to disable the `FLUSH_TO_ZERO` option, you would specify the following:

Example 7-2. Disabling flushing to zero

```
#include <fltmath.h>
. . .
FLUSH_TO_ZERO = 0;
```

7.2.1.3 IEEE_Exceptions

This is a status word that represents the IEEE exceptions that were raised during the last floating point operation. By default, the floating point library sets these values but does not handle any of these exceptions.

The following exceptions are supported:

- `IEEE_Inexact`
- `IEEE_Divide_By_Zero`
- `IEEE_Underflow`
- `IEEE_Overflow`
- `IEEE_Signaling_Nan`

See the IEEE standard for the exact description of these exceptions.

Following is an example of how to use the exception status word:

Example 7-3. Using the exception status word

```
#include <fltmath.h>
float x,y;
. . .
x = x*y;
if (IEEE_Exceptions & IEEE_Overflow)
{
<handle overflow>
}
```

7.2.1.4 EnableFPExceptions

This is a bit field mask. Setting a flag enables raising an SIGFPE signal if the last FP operation raised this exception. For example:

Example 7-4. Setting a signal for exceptions

```
#include <fltmath.h>
#include <signal.h>
void SigFPHandler(int x)
{
    switch (IEEE_Exceptions)
    {
        case IEEE_Overflow:
            . . .
        case IEEE_Divide_by_zero:
            . . .
    }
}
float x,y;
. . .
EnableFPExceptions = IEEE_Overflow | IEEE_Divide_by_zero;
signal(SIGFPE, SigFPHandler)
x = x*y /*This will raise SIGFPE if overflow or divide by zero occur */
```

This example installs a signal for handling overflow and divide by zero exceptions.

Note: Because the signal handling installs the handler address into the interrupt table, this example works only if the interrupt vector table is located in RAM. If the call to SIGNAL is not able to install the new handler, SIG_ERR is returned.

7.3 Integer Characteristics (limits.h)

The contents of `limits.h` are listed in Table 7-6.

Table 7-6. Contents of File `limits.h`

Constant	Value	Purpose
<code>CHAR_BIT</code>		8 Width of <code>char</code> type, in bits
<code>CHAR_MAX</code>		127 Maximum value for <code>char</code>
<code>CHAR_MIN</code>		-128 Minimum value for <code>char</code>
<code>INT_MAX</code>	2147483647	Maximum value for <code>int</code>
<code>INT_MIN</code>	(-2147483647-1)	Minimum value for <code>int</code>
<code>UINT_MAX</code>	429496729u	Maximum value for unsigned <code>int</code>
<code>LONG_MAX</code>	2147483647	Maximum value for long <code>int</code>
<code>LONG_MIN</code>	(-2147483647-1)	Minimum value for long <code>int</code>
<code>ULONG_MAX</code>	429496729uL	Maximum value for unsigned long <code>int</code>
<code>MB_LEN_MAX</code>		2 Maximum number of bytes in a multibyte character
<code>SCHAR_MAX</code>		127 Maximum value for signed <code>char</code>
<code>SCHAR_MIN</code>		-128 Minimum value for signed <code>char</code>
<code>UCHAR_MAX</code>		255 Maximum value for unsigned <code>char</code>
<code>SHRT_MAX</code>	32767	Maximum value for short <code>int</code>
<code>SHRT_MIN</code>	-32768	Minimum value for short <code>int</code>
<code>USHRT_MAX</code>	65536u	Maximum value for unsigned short <code>int</code>

7.4 Locales (locale.h)

Table 7-7 lists the locales functions that the compiler supports.

Table 7-7. Locale Functions

Function	Purpose
<code>localeconv(void)</code>	
<code>setlocale(int category, const char* locale)</code>	

Note: These functions are supported for compatibility purposes, and have no effect.

7.5 Floating Point Math (math.h)

The `math.h` library contains the following function types:

- Trigonometric functions
- Hyperbolic functions
- Exponential and logarithmic functions
- Power functions
- Other functions

The compiler runtime environment fully implements the `math.h` library using floating point emulation.

7.5.1 Trigonometric Functions

Table 7-8 lists the trigonometric functions that the compiler supports.

Table 7-8. Trigonometric Functions

Function	Purpose
<code>double acos(double)</code>	arc cosine
<code>double asin(double)</code>	arc sine
<code>double atan(double)</code>	arc tangent
<code>double atan2(double, double)</code>	arc tangent2
<code>double cos(double)</code>	cosine
<code>double sin(double)</code>	sine
<code>double tan(double)</code>	tangent

7.5.2 Hyperbolic Functions

Table 7-9 lists the hyperbolic functions that the compiler supports.

Table 7-9. Hyperbolic Functions

Function	Purpose
<code>double cosh(double)</code>	Hyperbolic cosine
<code>double sinh(double)</code>	Hyperbolic sine
<code>double tanh(double)</code>	Hyperbolic tangent

7.5.3 Exponential and Logarithmic Functions

Table 7-10 lists the exponential and logarithmic functions that the compiler supports.

Table 7-10. Exponential and Logarithmic Functions

Function	Purpose
<code>double exp(double)</code>	Exponential
<code>double frexp(double, int*)</code>	Splits floating point into fraction and exponent
<code>double ldexp(double, int)</code>	Computes value raised to a power
<code>double log(double)</code>	Natural logarithm
<code>double log10(double)</code>	Base ten (10) logarithm
<code>double modf(double, double*)</code>	Splits floating point into fraction and integer

7.5.4 Power Functions

Table 7-11 lists the power functions that the compiler supports.

Table 7-11. Power Functions

Function	Purpose
<code>double pow(double, double)</code>	Raises value to a power
<code>double sqrt(double)</code>	Square root

7.5.5 Other Functions

Table 7-12 lists the other functions that the compiler supports.

Table 7-12. Other Functions

Function	Purpose
<code>double ceil(double)</code>	Ceiling
<code>double fabs(double)</code>	Floating point absolute number
<code>double floor(double)</code>	Floor
<code>double fmod(double, double)</code>	Floating point remainder

7.6 Nonlocal Jumps (setjmp.h)

Table 7-13 lists the nonlocal jumps that the compiler supports.

Table 7-13. Nonlocal Jumps

Function	Purpose
<code>typedef unsigned int jmp_buf[32]</code>	Buffer used to save the execution context
<code>void longjmp(jmp_buf, int)</code>	Nonlocal jump
<code>int setjmp(jmp_buf)</code>	Nonlocal return

7.7 Signal Handling (signal.h)

Table 7-14 lists the signal handling that the compiler supports.

Table 7-14. Signal Handling (signal.h)

Function	Purpose
<code>int raise(int)</code>	Raises a signal
<code>void(*signal(int, void (*)(int)))(int)</code>	Installs a signal handler

7.8 Variable Arguments (stdarg.h)

Table 7-15 lists the variable arguments that the compiler supports.

Table 7-15. Variable Arguments (stdarg.h)

Function	Purpose
<code>va_arg(_ap, _type) (*(_type*)((_ap) -= sizeof(_type)))</code>	Returns next parameter in argument list
<code>va_end(_ap) (void)0</code>	Performs cleanup of argument list
<code>va_list</code>	Type declaration of variable argument list
<code>va_start(_ap, _parmN) (void)(_ap = (char*)&_parmN)</code>	Performs initialization of argument list

7.9 Standard Definitions (stddef.h)

Table 7-16 lists the standard definitions that the compiler supports.

Table 7-16. Standard Definitions (stddef.h)

Function	Purpose
<code>NULL((void*)0)</code>	Null pointer constant
<code>offsetof(type, member)</code>	Field offset in bytes from start of structure
<code>typedef int ptrdiff_t</code>	Signed integer type resulting from the subtraction of two pointers
<code>typedef int size_t</code>	Unsigned integer type that is the data type of the <code>sizeof</code> operator
<code>typedef short wchar_t</code>	Wide character type, as defined in ISO C

7.10 I/O Library (stdio.h)

The `stdio.h` library contains the following function types:

- z functions
- Stream functions
- Output functions
- Miscellaneous I/O functions

7.10.1 Input Functions

Table 7-17 lists the input functions that the compiler supports.

Table 7-17. Input Functions

Function	Purpose
<code>int fgetc(FILE*)</code>	Inputs a single character if available from specified stream
<code>size_t fread(void*, size_t, size_t, FILE*)</code>	Inputs a size number of characters from <code>stdin</code>
<code>int fscanf(FILE*, const char*, ...)</code>	Inputs text from the specified stream
<code>int getc(FILE*)</code>	Inputs a single character if available from specified stream
<code>int getchar(void)</code>	Inputs a single character if available from <code>stdin</code>
<code>int scanf(const char*, ...)</code>	Inputs text from <code>stdin</code>
<code>int sscanf(const char*, const char*, ...)</code>	Inputs text from specified string

7.10.2 Stream Functions

Table 7-18 lists the stream functions that the compiler supports.

Table 7-18. Stream Functions

Function	Purpose
<code>void clearerr(FILE*)</code>	Clears the EOF and error indicators for the specified stream
<code>int fclose(FILE*)</code>	Flushes the specified stream and closes the file associated with it
<code>int feof(FILE*)</code>	Tests the EOF indicator for the specified stream
<code>int ferror(FILE*)</code>	Tests the error indicator for the specified stream
<code>int fgetpos(FILE*, fpos_t*)</code>	Stores the current value of the file position indicator for the specified stream
<code>FILE* freopen(const char*, const char*, FILE*)</code>	Opens the specified file in the specified mode, using the specified stream
<code>int fseek(FILE*, long int, int)</code>	Sets the file position indicator for the specified stream
<code>int fsetpos(FILE*, const fpos_t*)</code>	Sets the file position indicator for the specified stream to the specified value
<code>long int ftell(FILE*)</code>	Retrieves the current value of the file position indicator for the current stream
<code>int remove(const char*)</code>	Makes the specified file unavailable by its defined name
<code>int rename(const char*, const char*)</code>	Assigns to the specified file a new filename
<code>void rewind(FILE*)</code>	Sets the file position indicator for the specified stream to the beginning of the file
<code>void setbuf(FILE*, char*)</code>	Defines a buffer and associates it with the specified stream. A restricted version of <code>setvbuf()</code>
<code>int setvbuf(FILE*, char*, int, size_t)</code>	Defines a buffer and associates it with the specified stream
<code>stderr</code>	Standard error stream (Value = 3)
<code>stdin</code>	Standard input stream (Value = 1)
<code>stdout</code>	Standard output stream (Value = 2)
<code>FILE* tmpfile(void)</code>	Creates a temporary file
<code>char* tmpnam(char*)</code>	Generates a valid filename, meaning a filename that is not in use, as a string

7.10.3 Output Functions

Table 7-19 lists the output functions that the compiler supports.

Table 7-19. Output Functions

Function	Purpose
<code>char* fgets(char*, int, FILE*)</code>	Outputs characters to the specified stream
<code>int fprintf(FILE*, const char*, ...)</code>	Outputs the specified text to the specified stream
<code>int fputc(int, FILE*)</code>	Outputs a single character to the specified stream
<code>int fputs(const char*, FILE*)</code>	Outputs a string to the specified stream
<code>size_t fwrite(const void*, size_t, size_t, FILE*)</code>	Outputs a size number of characters to <code>stdout</code>
<code>char* gets(char*)</code>	Outputs characters into the user's buffer
<code>void perror(const char*)</code>	Outputs an error message
<code>int printf(const char*, ...)</code>	Outputs the specified text to <code>stdout</code>
<code>int putc(int, FILE*)</code>	Outputs a single character to the specified stream
<code>int putchar(int)</code>	Outputs a single character
<code>int puts (const char*)</code>	Outputs the string to <code>stdout</code> , followed by a newline
<code>int sprintf(char*, const char*, ...)</code>	Outputs the specified text to the specified buffer
<code>int vfprintf(FILE*, const char*, va_list)</code>	Outputs the variable arguments to the specified stream
<code>int vprintf(const char*, va_list)</code>	Outputs the variable arguments to <code>stdout</code>
<code>int vsprintf(char*, const char*, va_list)</code>	Outputs the variable arguments to the specified buffer

7.10.4 Miscellaneous I/O Functions

Table 7-20 lists the miscellaneous I/O functions that the compiler supports.

Table 7-20. Miscellaneous I/O Functions

Function	Purpose
<code>int fflush(FILE*)</code>	Causes the output buffers to be emptied to their destinations
<code>FILE* fopen(const char*, const char*)</code>	Associates a stream with a file
<code>int ungetc(int, FILE*)</code>	Moves the character back to the head of the input stream

7.11 General Utilities (stdlib.h)

The `stdlib.h` library contains the following function types:

- Memory allocation functions
- Integer arithmetic functions
- String conversion functions
- Searching and sorting functions
- Pseudo random number generation functions
- Environment functions
- Multibyte functions

7.11.1 Memory Allocation Functions

Table 7-21 lists the memory allocation functions that the compiler supports.

Table 7-21. Memory Allocation Functions

Function	Purpose
<code>void free(void*)</code>	Returns allocated space to heap
<code>void* calloc(size_t, size_t)</code>	Allocates heap space initialized to zero
<code>void* malloc(size_t)</code>	Allocates heap space
<code>void* realloc(void*, size_t)</code>	Allocates a larger heap space and returns previous space to heap

7.11.2 Integer Arithmetic Functions

Table 7-22 lists the integer arithmetic functions that the compiler supports.

Table 7-22. Integer Arithmetic Functions

Function	Purpose
<code>int abs(int)</code>	Absolute value
<code>div_t div(int, int)</code>	Quotient and remainder
<code>long int labs(long int)</code>	Computes absolute value and returns as long int
<code>ldiv_t ldiv(long int, long int)</code>	Quotient and remainder of long int

7.11.3 String Conversion Functions

Table 7-23 lists the string conversion functions that the compiler supports.

Table 7-23. String Conversion Functions

Function	Purpose
<code>double atof(const char*)</code>	String to float
<code>int atoi(const char*)</code>	String to int
<code>long int atol(const char*)</code>	Long
<code>double strtod(const char*, char**)</code>	Double
<code>long int strtol(const char*, char**, int)</code>	Long
<code>unsigned long int strtoul(const char*, char**, int)</code>	Unsigned long

7.11.4 Searching and Sorting Functions

Table 7-24 lists the searching and sorting functions that the compiler supports.

Table 7-24. Searching and Sorting Functions

Function	Purpose
<code>void *bsearch(const void*, const void*, size_t, size_t, int(*)(const void*, const void*))</code>	Binary search
<code>void *qsort(void*, size_t, size_t, int(*)(const void*, const void*))</code>	Quick sort

7.11.5 Pseudo Random Number Generation Functions

Table 7-25 lists the pseudo random number generation functions that the compiler supports.

Table 7-25. Pseudo Random Number Generation Functions

Function	Purpose
<code>int rand(void)</code>	Random number generator
<code>void srand(unsigned int)</code>	Initializes the random number generator

7.11.6 Environment Functions

Table 7-26 lists the environment functions that the compiler supports.

Table 7-26. Environment Functions

Function	Purpose
<code>void abort(void)</code>	Causes an abnormal termination
<code>int atexit(void (*)(void))</code>	Registers a function to be called at normal termination
<code>void exit(int)</code>	Causes a normal termination
<code>char *getenv(const char *name)¹</code>	Gets environment variable
<code>int system(const char *string)¹</code>	Passes command to host environment

1. This function is supported for compatibility purposes and has no effect.

7.11.7 Multibyte Character Functions

Table 7-27 lists the multibyte character functions that the compiler supports.

Table 7-27. Multibyte Character Functions

Function	Purpose
<code>int mblen(const char*, size_t)</code>	Multibyte string length
<code>size_t mbstowcs(wchar_t*, const char*, size_t)</code>	Converts multibyte string to wide character string
<code>int mbtowc(wchar_t*, const char*, size_t)</code>	Converts multibyte to wide character
<code>int wctomb(char*, wchar_t)</code>	Converts wide character to multibyte
<code>size_t wcstombs(char*, const wchar_t*, size_t)</code>	Converts wide character string to multibyte string

7.12 String Functions (string.h)

The `string.h` library contains the following function types:

- Copying functions
- Concatenation functions
- Comparison functions
- Search functions
- Other functions

7.12.1 Copying Functions

Table 7-28 lists the copying functions that the compiler supports.

Table 7-28. Copying Functions

Function	Purpose
<code>void* memcpy(void*, const void*, size_t)</code>	Copies data
<code>void* memmove(void*, const void*, size_t)</code>	Swaps data
<code>char* strcpy(char*, const char*)</code>	Copies a string
<code>char* strncpy(char*, const char*, size_t)</code>	Copies a string of a maximum length

7.12.2 Concatenation Functions

Table 7-29 lists the concatenation functions that the compiler supports.

Table 7-29. Concatenation Functions

Function	Purpose
<code>char* strcat(char*, const char*)</code>	Concatenates a string to the end of another string
<code>char* strncat(char*, const char*, size_t)</code>	Concatenates a string of specified maximum length to the end of another string

7.12.3 Comparison Functions

Table 7-30 lists the comparison functions that the compiler supports.

Table 7-30. Comparison Functions

Function	Purpose
<code>int memcmp(const void*, const void*, size_t)</code>	Compares data
<code>int strcmp(const char*, const char*)</code>	Compares strings
<code>int strcoll(const char*, const char*)</code>	Compares strings based on locale
<code>int strncmp(const char*, const char*, size_t)</code>	Compares strings of maximum length
<code>size_t strxfrm(char*, const char*, size_t)</code>	Transforms a string into a second string of the specified size

7.12.4 Search Functions

Table 7-31 lists the search functions that the compiler supports.

Table 7-31. Search Functions

Function	Purpose
<code>void* memchr(const void*, int, size_t)</code>	Searches for a value in the first number of characters
<code>char* strchr(const char*, int)</code>	Searches a string for the first occurrence of char
<code>size_t strcspn(const char*, const char*)</code>	Searches a string for the first occurrence of char in string set and returns the number of characters skipped
<code>char strpbrk(const char*, const char*)</code>	Searches a string for the first occurrences of char in string set and returns a pointer to that location
<code>char* strrchr(const char*, int)</code>	Searches a string for the last occurrence of char
<code>size_t strspn(const char*, const char*)</code>	Searches a string for the first occurrence of char not in string set.
<code>char* strstr(const char*, const char*)</code>	Searches a string for the first occurrence of string
<code>char* strtok(char*, const char*)</code>	Separates a string into tokens

7.12.5 Other Functions

Table 7-32 lists the other functions that the compiler supports.

Table 7-32. Other Functions

Function	Purpose
<code>void* memset(void*, int, size_t)</code>	Copies a value into each number of characters
<code>char* strerror(int)</code>	Returns string for associated error condition
<code>size_t strlen(const char*)</code>	Returns size of string

7.13 Time Functions (time.h)

Table 7-33 lists the time functions that the compiler supports.

Table 7-33. Time Functions

Function	Purpose
<code>char *asctime(const struct tm *timeptr)</code>	Converts time to ASCII representation
<code>clock_t clock()</code>	Returns processor time
<code>typedef unsigned long clock_t</code>	Type used for measuring time
<code>char *ctime (const time_t *timer)</code>	Converts time to ASCII representation
<code>double difftime(time_t time1, time_t time0)</code>	Returns difference in seconds
<code>time_t mktime(struct tm *timeptr)</code>	Converts struct tm to time_t
<code>size_t strftime (char *s, size_t maxsize, const char *format, const struct tm *timeptr)</code>	Converts an ASCII string to time_t
<code>time_t time(time_t *timer)</code>	Returns processor time (same as <code>clock</code>)
<code>typedef unsigned long time_t</code>	Type used for measuring time
<code>struct tm *gmtime(const time_t *timer)</code>	Returns time in GMT time zone
<code>struct tm *localtime(const time_t *timer)</code>	Returns time in local time zone

7.13.1 Time Constant

Table 7-34 shows the time constant that the compiler supports.

Table 7-34. Time Constant

Constant	Value	Purpose
<code>CLOCKS_PER_SEC</code>		TBD

7.13.2 Process Time

The `clock` function returns the current value of the system timer. This function must be configured to match the actual system timer configuration. The timer is started and set for a maximum period during the initialization of any C program that references the `clock` function, and is used only by this function. The return value of `clock` has type `clock_t`, which is `unsigned long`.

The following example shows how to use the `clock` function to time your application:

Example 7-5. Timing an application

```
#include <time.h>
clock_t start, end, elapsed;
/* . . . application setup . . . */
start = clock( );
/* . . . application processing . . . */
end = clock( );
elapsed = end - start; /* Assumes no wrap-around */
printf("Elapsed time: %Lu * 2 cycles. \n", elapsed);
```

7.14 Built-in Intrinsic Functions (prototype.h)

The compiler supports a set of built-in intrinsic functions that enable fractional operations to be implemented using integer data types, by mapping directly to SC100 assembly instructions.

Table 7-35 lists these built-in intrinsic functions.

Table 7-35. Built-in Intrinsic Functions

Function	Purpose
short abs_s(short var1)	Short absolute value of var1. For example, the result of abs_s(-32768) is +32767.
short add(short var1, short var2)	Short add. Performs the addition var1+var2 with overflow control and saturation. The 16-bit result is set at +32767 when overflow occurs, or at -32768 when underflow occurs.
BitReverseUpdate	Increments the iterator with bit reverse.
Word64 D_add(Word64 D_var1, Word64 D_var2)	Double precision add. Performs the addition D_var1+D_var2 with overflow control and saturation.
short D_cmpeq(Word64 D_var1, Word64 D_var2)	Double precision compare equal. Compares two 64-bit values and returns a 16-bit result containing the value ??? if the values are equal, or ??? if they are not.
short D_cmpgt(Word64 D_var1, Word64 D_var2)	Double precision compare greater than. Compares two 64-bit values and returns a 16-bit result containing ???.
long D_extract_h(Word64 D_var1)	Double precision extract high. Returns the 32 MSB of the 64-bit value D_var1.
unsigned long D_extract_l(Word64 D_var1)	Double precision extract low. Returns the 32 LSB of the 64-bit value D_var1 as an unsigned 32-bit value.
Word64 D_mac(Word64 D_var3, long L_var1, long L_var2)	Double precision multiply accumulate. Multiplies L_var1 by L_var2 and shifts the result left by 1. Adds the 64-bit result to L_var3 with saturation, and returns a 64-bit result. For example: D_mac(D_var3, L_var1, L_var2) = D_add(D_var3, D_mult(L_var1, L_var2)).
Word64 D_msu(Word64 D_var3, long L_var1, long L_var2)	Double precision multiply subtract. Multiplies L_var1 by L_var2 and shifts the result left by 1. Subtracts the 64-bit result from D_var3 with saturation, and returns a 64-bit result. For example: D_msu(D_var3, L_var1, L_var2) = D_sub(D_var3, D_mult(L_var1, L_var2)).
Word64 D_mult(long L_var1, long L_var2)	Double precision multiply. The 64-bit result of the multiplication of L_var1 by L_var2 with one shift left, for example: D_mult(L_var1, L_var2) = D_shl((L_var1*L_var2), 1).
long D_round(Word64 D_var1)	Double precision round. Rounds the lower 32 bits of the 64-bit D_var1 into the MS 32 bits with saturation. Shifts the resulting bits right by 32 and returns the 32-bit value.
Word64 D_sat(Word64 D_var1)	Double precision saturation. Saturates a 64-bit value.

Table 7-35. Built-in Intrinsic Functions (Continued)

Function	Purpose
Word64 D_set(long L_var1, unsigned long L_var2)	Double precision concatenation. Concatenates two 32-bit values, L_var1 and unsigned L_var2, into one 64-bit value.
Word64 D_sub(Word64 D_var1, Word64 D_var2)	Double precision subtract. 64-bit subtraction of the two 64-bit variables (D_var1-D_var2) with overflow control and saturation.
void debug()	Generates assembly instruction to enter Debug mode.
void debugev()	Generates assembly instruction to issue Debug event.
void di()	Generates assembly instruction to disable interrupts.
short div_s(short var1,short var2)	Short divide. Produces a result which is the fractional integer division of var1 by var2; var1 and var2 must be positive, and var2 must be greater or equal to var1. The result is positive (leading bit equal to 0) and truncated to 16 bits. If var1 = var2 then div(var1,var2) = 32767.
void ei()	Generates assembly instruction to enable interrupts.
EndBitReverse	Frees bit reverse iterator.
short extract_h(long L_var1)	Long extract high. Returns the 16 MSB of L_var1.
short extract_l(long L_var1)	Long extract low. Returns the 16 LSB of L_var1.
void illegal()	Generates assembly instruction to execute illegal exception.
InitBitReverse	Allocates a bit reverse iterator.
long L_abs(long L_var1)	Long absolute value of L_var1. Saturates in cases where the value is -214783648.
long L_add(long L_var1,long L_var2)	Long add. 32-bit addition of the two 32-bit variables (L_var1+L_var2) with overflow control and saturation. The result is set at +2147483647 when overflow occurs, or at -2147483648 when underflow occurs.
long L_deposit_h(short var1)	Deposit short in MSB. Deposits the 16-bit var1 into the 16 MS bits of the 32-bit output. The 16 LS bits of the output are zeroed.
long L_deposit_l(short var1)	Deposit short in LSB. Deposits the 16-bit var1 into the 16 LS bits of the 32-bit output. The 16 MS bits of the output are sign extended.
long L_mac(long L_var3,short var1, short var2)	Multiply accumulate. Multiplies var1 by var2 and shifts the result left by 1. Adds the 32-bit result to L_var3 with saturation, and returns a 32-bit result. For example: L_mac(L_var3,var1,var2) = L_add(L_var3,L_mult(var1,var2)).
long L_max(long L_var1,long L_var2)	Compares the values of two 32-bit variables and returns the higher value of the two.
long L_min(long L_var1,long L_var2)	Compares the values of two 32-bit variables and returns the lower value of the two.

Table 7-35. Built-in Intrinsic Functions (Continued)

Function	Purpose
<code>long L_msu(long L_var3,short var1, short var2)</code>	Multiply subtract. Multiplies <code>var1</code> by <code>var2</code> and shifts the result left by 1. Subtracts the 32-bit result from <code>L_var3</code> with saturation, and returns a 32-bit result. For example: <code>L_msu(L_var3,var1,var2) = L_sub(L_var3,L_mult(var1,var2)).</code>
<code>long L_mult(short var1,short var2)</code>	Long multiply. The 32-bit result of the multiplication of <code>var1</code> by <code>var2</code> with one shift left, for example: <code>L_mult(var1,var2)= L_shl((var1*var2),1)</code> and <code>L_mult(-32768,-32768) = 2147483647.</code>
<code>long L_negate(long L_var1)</code>	Long negate. Negates the 32-bit variable <code>L_var1</code> with saturation. Saturates in cases where the value is <code>-2147483648(0x8000 0000).</code>
<code>long L_rol(long L_var1)</code>	Long rotate left. Rotates the 32-bit variable <code>L_var1</code> left into a 40-bit value, and returns a 32-bit result..
<code>long L_ror(long L_var1)</code>	Long rotate right. Rotates the 32-bit variable <code>L_var1</code> right into a 40-bit value, and returns a 32-bit result.
<code>long L_sat(long L_var1)</code>	Saturates a 32-bit value.
<code>long L_shl(long L_var1,short var2)</code>	Long shift left. Arithmetically shifts the 32-bit <code>L_var1</code> left <code>var2</code> positions. Zero fills the <code>var2</code> LSB of the result. If <code>var2</code> is negative, arithmetically shifts <code>L_var1</code> right by <code>var2</code> with sign extension. Saturates the result in cases where underflow or overflow occurs.
<code>long L_shr(long L_var1,short var2)</code>	Long shift right. Arithmetically shifts the 32-bit <code>L_var1</code> right <code>var2</code> positions with sign extension. If <code>var2</code> is negative, arithmetically shifts <code>L_var1</code> left by <code>var2</code> and zero fills the <code>var2</code> LSB of the result. Saturates the result in cases where underflow or overflow occurs.
<code>long L_shr_r(long L_var1,short var2)</code>	Long shift right and round. Same as <code>L_shr(L_var1,var2)</code> but with rounding. Saturates the result in cases where underflow or overflow occurs.
<code>long L_sub(long L_var1,long L_var2)</code>	Long subtract. 32-bit subtraction of the two 32-bit variables (<code>L_var1-L_var2</code>) with overflow control and saturation. The result is set at <code>+2147483647</code> when overflow occurs or at <code>-2147483648</code> when underflow occurs.
<code>short mac_r(long L_var3,short var1, short var2)</code>	Multiply accumulate and round. Multiplies <code>var1</code> by <code>var2</code> and shifts the result left by 1. Adds the 32-bit result to <code>L_var3</code> with saturation. Rounds the LS 16 bits of the result into the MS 16 bits with saturation and shifts the result right by 16. Returns a 16-bit result.
<code>void mark()</code>	Generates assembly instruction to write program counter to trace buffer, if trace buffer enabled.
<code>short max(short var1, short var2)</code>	Compares the values of two 16-bit variables and returns the higher value of the two.
<code>short min(short var1, short var2)</code>	Compares the values of two 16-bit variables and returns the lower value of the two.
<code>long mpyuu(long L_var1,long L_var2)</code>	Multiplies the 16 LSB of two 32-bit variables, treating both variables as unsigned values, and returns a 32-bit result.

Table 7-35. Built-in Intrinsic Functions (Continued)

Function	Purpose
<code>long mpyus(long L_var1, long L_var2)</code>	Multiplies the 16 LSB of the 32-bit variable <code>L_var1</code> , treated as an unsigned value, by the 16 MSB of the 32-bit variable <code>L_var2</code> , treated as a signed value. Returns a 32-bit result.
<code>long mpysu(long L_var1, long L_var2)</code>	Multiplies the 16 MSB of the 32-bit variable <code>L_var1</code> , treated as a signed value, by the 16 LSB of the 32-bit variable <code>L_var2</code> , treated as an unsigned value. Returns a 32-bit result.
<code>short msu_r(long L_var3, short var1, short var2)</code>	Multiply subtract and round. Multiplies <code>var1</code> by <code>var2</code> and shifts the result left by 1. Subtracts the 32-bit result from <code>L_var3</code> with saturation. Rounds the LS 16 bits of the result into the MS 16 bits with saturation and shifts the result right by 16. Returns a 16-bit result.
<code>short mult(short var1, short var2)</code>	Short multiply. Performs the multiplication of <code>var1</code> by <code>var2</code> and gives a 16-bit result which is scaled, for example: <code>mult(var1, var2) = extract_l(L_shr((var1 * var2), 15))</code> and <code>mult(-32768, -32768) = 32767</code> .
<code>short mult_r(short var1, short var2)</code>	Multiply and round. Same as <code>mult</code> with rounding, for example: <code>mult_r(var1, var2) = extract_l(L_shr((var1*var2)+16384), 15)</code> and <code>mult_r(-32768, -32768) = 32767</code> .
<code>short negate(short var1)</code>	Short negate. Negates <code>var1</code> with saturation. Saturates in cases where the value is <code>-32768</code> , for example: <code>negate(var1) = sub(0, var1)</code> .
<code>short norm_l(long L_var1)</code>	Normalizes any long fractional value. Produces the number of left shifts needed to normalize the 32-bit variable <code>L_var1</code> for positive values on the interval with minimum of <code>1073741824</code> and maximum of <code>2147483647</code> , and for negative values on the interval with minimum of <code>-2147483648</code> and maximum of <code>-1073741824</code> . In order to normalize the result, the following operation must be executed: <code>norm_L_var1 = L_shl(L_var1, norm_l(L_var1))</code> .
<code>short norm_s(short var1)</code>	Normalizes any fractional value. Produces the number of left shifts needed to normalize the 16-bit variable <code>var1</code> for positive values on the interval with minimum of <code>16384</code> and maximum of <code>32767</code> , and for negative values on the interval with minimum of <code>-32768</code> and maximum of <code>-16384</code> . In order to normalize the result, the following operation must be executed: <code>norm_var1 = shl(var1, norm_s(var1))</code> .
<code>short round(long var1)</code>	Round. Rounds the lower 16 bits of the 32-bit number into the MS 16 bits with saturation. Shifts the resulting bits right by 16 and returns the 16-bit number, for example: <code>round(L_var1) = extract_h(L_add(L_var1, 32768))</code> .
<code>short saturate(short var1)</code>	Saturates a 16-bit value.
<code>setcnvrm()</code>	Sets rounding mode to convergent rounding mode.

Table 7-35. Built-in Intrinsic Functions (Continued)

Function	Purpose
<code>set2crm()</code>	Sets rounding mode to two's-complement rounding mode.
<code>void setnosat()</code>	Clears saturation mode bit in status register.
<code>void setsat32()</code>	Sets saturation mode bit in status register.
<code>short shl(short var1,short var2)</code>	Short shift left. Arithmetically shifts the 16-bit <code>var1</code> left <code>var2</code> positions. Zero fills the <code>var2</code> LSB of the result. If <code>var2</code> is negative, arithmetically shifts <code>var1</code> right by <code>var2</code> with sign extension. Saturates the result in cases where underflow or overflow occurs.
<code>short shr(short var1,short var2)</code>	Short shift right. Arithmetically shifts the 16-bit <code>var1</code> right <code>var2</code> positions with sign extension. If <code>var2</code> is negative, arithmetically shifts <code>var1</code> left by <code>var2</code> with sign extension. Saturates the result in cases where underflow or overflow occurs.
<code>short shr_r(short var1,short var2)</code>	Short shift right and round. Same as <code>shr(var1, var2)</code> but with rounding. Saturates the result in cases where underflow or overflow occurs.
<code>void stop()</code>	Generates assembly instruction to enter <code>Stop</code> low power mode.
<code>short sub(short var1,short var2)</code>	Performs the subtraction with overflow control and saturation. The 16-bit result is set at +32767 when overflow occurs or at -32768 when underflow occurs.
<code>void trap()</code>	Generates assembly instruction to execute <code>Trap</code> exception.
<code>void wait()</code>	Generates assembly instruction to enter <code>wait</code> low power mode.
<code>Word40 X_abs(Word40 X_var1)</code>	40-bit absolute value of <code>X_var1</code> .
<code>Word40 X_add(Word40 X_var1, Word40 X_var2)</code>	Extended precision add. Performs the addition <code>X_var1+X_var2</code> without saturation. .
<code>short X_cmpeq(Word40 X_var1, Word40 X_var2)</code>	Extended precision compare equal. Compares two 40-bit values and returns a 16-bit result containing the value <code>???</code> if the values are equal, or <code>???</code> if they are not.
<code>short X_cmpgt(Word40 X_var1, Word40 X_var2)</code>	Extended precision compare greater than. Compares two 40-bit values and returns a 16-bit result containing <code>???</code> .
<code>Word40 X_extend(long L_var1)</code>	Sign extend 32-bit value to 40-bit value.
<code>short X_extract_h(Word40 X_var1)</code>	Extended precision extract high. Returns the 16 MSB of the 40-bit value <code>X_var1</code> .
<code>short X_extract_l(Word40 X_var1)</code>	Extended precision extract low. Returns the 16 LSB of the 40-bit value <code>X_var1</code> .
<code>Word40 X_mac(Word40 X_var3, short var1,short var2)</code>	Extended precision multiply accumulate. Multiplies <code>var1</code> by <code>var2</code> and shifts the result left by 1. Adds the 40-bit result to <code>X_var3</code> without saturation, and returns a 40-bit result. For example: <code>X_mac(X_var3, var1, var2) = X_add(X_var3, X_mult(var1, var2))</code> .

Table 7-35. Built-in Intrinsic Functions (Continued)

Function	Purpose
Word40 X_msu(Word40 X_var3, short var1,short var2)	Extended precision multiply subtract. Multiplies var1 by var2 and shifts the result left by 1. Subtracts the 40-bit result from var3 without saturation, and returns a 40-bit result. For example: X_msu(X_var3,var1,var2) = X_sub(X_var3,X_mult(var1,var2)).
Word40 X_mult(short var_1,short var_2)	Extended precision multiply. The 40-bit result of the multiplication of var1 by var2 with one shift left, for example: X_mult(var1,var2) = X_shl((var1*var2),1).
short X_norm(Word40 X_var1)	Normalizes a 40-bit fractional value.
Word40 X_or(Word40 X_var1, Word40 X_var2)	Performs logical OR on two 40-bit values.
Word40 X_rol(Word40)	Rotates left a 40-bit value.
Word40 X_ror(Word40)	Rotates right a 40-bit value.
short X_round(Word40 X_var1)	Extended precision round. Rounds the lower 16 bits of the 40-bit number into the MS 16 bits without saturation. Shifts the resulting bits right by 16 and returns the 16-bit number.
long X_sat(Word40 X_var1)	Extended precision saturation. Saturates a 40-bit value.
Word40 X_set(char var1, unsigned long L_var2)	Extended precision concatenation. Concatenates an 8-bit character value and an unsigned 32-bit value into one 40-bit value.
Word40 X_shl(Word40 X_var1,short var2)	Extended shift left. Arithmetically shifts the 40-bit X_var1 left var2 positions. Zero fills the var2 LSB of the result. If var2 is negative, arithmetically shifts X_var1 right by var2 with sign extension.
Word40 X_shr(Word40 X_var1,short var2)	Extended shift right. Arithmetically shifts the 40-bit X_var1 right var2 positions with sign extension. If var2 is negative, arithmetically shifts X_var1 left by var2 and zero fills the var2 LSB of the result.
Word40 X_sub(Word40 X_var1, Word40 X_var1)	Extended precision subtract. 40-bit subtraction of the two 40-bit variables (X_var1-X_var2) without saturation.
long X_trunc(Word40 X_var1)	Truncates 40-bit value into 32-bit value.

Appendix A

Migrating from Other Environments

The SC100 C compiler provides header files that make it easy to migrate C code developed for certain other compilers. The compilation and its results may be affected in various ways by the differences between specific compiler environments and the compiler. The effects may include, for example, assembler errors for inlined code that is not supported, or loss of efficiency for functions that are supported, but implemented in a different way.

This Appendix contains the following sections:

- [Section A.1, “Code Migration Overview,”](#) provides general guidelines for migrating code from another environment to the compiler.
- [Section A.2, “Migrating Code Developed for DSP56600,”](#) describes the issues to be considered when migrating code developed for the DSP56600 compiler family.
- [Section A.3, “Migrating Code Developed for TI6xx,”](#) describes the differences to take into account when migrating code developed for the TI6xx family of compilers.

A.1 Code Migration Overview

In most circumstances, the compiler can successfully compile standard ANSI code that:

- Does not use compiler-specific extensions
- Does not rely implicitly on the sizes of data types
- Does not rely on system-specific features, such as memory maps or peripherals
- Does not rely on undefined compiler behavior

The compiler runtime libraries include a header file for each environment for which code is accepted, as follows:

- DSP56600 compilers: `port566toSC1.h` header file
- TI6xx compilers: `portc6xttoSC1.h` header file

The features used in the specified environment are defined in the relevant header file with correct values, to ensure that the code is not affected and compiles successfully.

To use these definitions, just include the appropriate header file to your source code. For example, when migrating code from the DSP56600 compiler environment, include the `port566toSC1.h` header file, as shown in Example A-1.

Example A-1. Migrating code from other environments

```
#include <port566toSC1.h>
void main()
{
}
```

A.2 Migrating Code Developed for DSP56600

When using the SC100 C compiler with code developed for the DSP56600 family of compilers, the following differences should be taken into account:

- **Integer data types:** The DSP56600 and SC100 compilers map certain integer data types to different sizes. Table A-1 lists the data type size discrepancies that relate to integers:

Table A-1. DSP56600 Integer Data Type Differences

Data Type	DSP56600 Compiler	SC100 C Compiler
char	Saved in memory as 16 bits. Some operations are performed with 16 bits, others with 8.	8 bits
unsigned char		
packed char	8 bits	Not supported
int	16 bits	32 bits
unsigned int		
enum	16 bits	32 bits

- **Fractional data types:** DSP56600 compilers use built-in data types for declaring fractional variables. The SC100 C compiler uses standard integer types for both fractional and integer values. Table A-2 lists the fractional data type differences:

Table A-2. DSP56600 Fractional Data Type Differences

Data Type	DSP56600 Compiler	SC100 C Compiler
16-bit fraction	<code>_fract</code>	Word16
32-bit fraction	<code>long_fract</code>	Word32
40-bit accumulator	<code>long_fract</code>	Word40
64-bit fraction	Not supported	Word64
Complex fractions	<code>_complex</code>	Not supported directly

- **Floating point data types:** DSP56600 compilers represent floating point data types according to a 32-bit proprietary format. The SC100 C compiler maps fractional data types to a single-precision IEEE-754 type, using 32 bits. As a result, there may be differences in the numerical accuracy of floating point calculations.

- **Pointers:** The difference in pointer size between the two compilers is shown in Table A-3:

Table A-3. DSP56600 Pointer Size Differences

Data Type	DSP56600 Compiler	SC100 C Compiler
pointer to char	16 bits	32 bits
pointer to short	16 bits	32 bits, even addresses only
pointer to long	16 bits	32 bits, quad addresses only

In most circumstances, the difference in pointer size is unlikely to have any impact, since the relevant addresses are usually mapped to different numerical values on different processors.

- **Fractional arithmetic:** DSP56600 compilers support fractional arithmetic using integer-like operators, such as + and *. The SC100 C compiler implements fractional operations through the use of intrinsic functions. Table A-4 lists the DSP56600 fractional operations and shows the equivalent SC100 C compiler intrinsic functions:

Table A-4. DSP56600 Fractional Arithmetic Differences

Fractional Operation	DSP56600 Compiler	SC100C Compiler
Addition	+	Word16 add Word32 L_add
Subtraction	-	Word16 sub Word32 L_sub
Absolute value	_fabs _lfabs	Word16 abs_s Word32 L_abs
Multiplication	*	Word16 mult Word32 L_mult Word16 mult_r
Shift right	>>	Word16 shr Word32 L_shr
Shift left	<<	Word16 shl Word32 L_shl
Negate	-	Word16 negate Word32 L_negate
Round	_fract_round	Word16 round
Divide	_pdiv	Word16 div_s
Normalize	Can be implemented using _asm	Word16 norm_s Word16 norm_l
Saturation control	Can be implemented using _asm	void setnosat void setsat32

The SC100 C compiler supports many more fractional operations, including 40-bit and 64-bit fractional functions, which are not supported in the DSP56600 environment.

- **Inlined assembly and C code:** DSP56600 compilers use `_inline` and `_asm` to designate a C routine for inlining, and to define the instructions, operands and modifiers for inlined assembly statements. The SC100 C compiler uses the pragma `#pragma inline` to specify an inlined function. See [Chapter 4, “Interfacing C and Assembly Code,”](#) for more information.

- **Intrinsic functions:** The SC100 C compiler library routines support a number of DSP56600 intrinsic functions, as shown in Table A-5:

Table A-5. DSP56600 Intrinsic Function Differences

Description	DSP56600 Compiler	SC100 C Compiler
Bit field operations	<code>_bfchg()</code> <code>_bfclr()</code> <code>_bfset()</code> <code>_bftsth()</code> <code>_bftstl()</code>	Can be implemented by library routines
Cache control	<code>_cache_get_start()</code> <code>_cache_get_end()</code> <code>_pflush()</code> <code>_pflushun()</code> <code>_pfree()</code> <code>_plock()</code> <code>_punlock()</code>	Not available
Fraction to integer coercion	<code>_fract2int()</code> <code>_lfract2long</code>	Not needed (both represented by integers)
Integer to fraction coercion	<code>_intt2fract()</code> <code>_long2lfract()</code>	Not needed (both represented by integers)
Extend byte in accumulator	<code>_ext()</code>	Not applicable
Fractional square root	<code>_fsqrt()</code>	Can be implemented by a library routine
String copy (inlined)	<code>_strcmp()</code>	Supported as a library routine (<code>strcmp</code>)
Absolute of long integer	<code>_labs</code>	<code>labs()</code>
Insert NOP instruction	<code>_nop()</code>	<code>_asm("nop")</code>
STOP instruction	<code>_stop()</code>	<code>stop()</code>
Software interrupt	<code>_swi()</code>	<code>trap()</code>
WAIT instruction	<code>_wait()</code>	<code>wait()</code>
Viterbi operation	<code>_vsl</code>	Can be implemented by a library routine

- **Pragmas:** The functions of the DSP56600 inlined assembly pragmas `asm`, `asm_noflush` and `endasm` are supported by the SC100 C compiler using a function qualifier. The SC100 C compiler loop optimization pragma `#pragma loop_count` is the equivalent of the DSP56600 pragmas `iterate_at_least_once` and `no_iterate_at_least_once`.

The following DSP56600 pragmas have no equivalent in the SC100 C compiler environment:

- `cache_align_now`
- `cache_sector_size`
- `cache_region_start`
- `cache_region_endpack_strings`
- `nopack_strings`
- `source`
- `nosource`
- `jumptable_memory`

- **Interrupt handlers:** The SC100 C compiler pragma `interrupt` performs the function of both `_fast_interrupt` and `_long_interrupt` in the DSP56600 environment.
- **Storage specifiers:** The DSP56600 compilers support a number of storage specifiers, which are either not used in the SC100 environment, or are specified at link time, as shown in Table A-6:

Table A-6. DSP56600 Storage Specifiers

Storage	DSP56600 Compiler	SC100 C Compiler
X memory	<code>_X</code>	Not applicable
Y memory	<code>_Y</code>	Not applicable
Program memory	<code>_P</code>	Not applicable
L memory	<code>_L</code>	Not applicable
Lowest 64 words in data memory	<code>_near</code>	Not applicable
Internal memory	<code>_internal</code>	Specified at link time
External memory	<code>_external</code>	Specified at link time
Absolute address for global variable	<code>_at</code>	Specified at link time in the application configuration file

- **Miscellaneous:** Table A-7 outlines some further differences between the two compilers:

Table A-7. DSP56600 Miscellaneous Differences

Description	DSP56600 Compiler	SC100 C Compiler
Wrap-around semantics for fractional data	<code>_nosat</code>	Not applicable
Force DSP56300 GNU calling convention	<code>_compatible</code>	Not applicable
Circular buffer support	<code>_circ</code>	Addressing calculations using the C modulo (%) operator

A.3 Migrating Code Developed for TI6xx

The following differences should be considered when using the compiler with code developed for the TI6xx family of compilers:

- **Data Types:** TI6xx compilers map the integer type `long` to 40 bits. The compiler defines the integer type `long` as 32 bits. C code that relies on the fact that type `long` is 40 bits wide must be modified before it can be migrated.
- **Keywords:** The TI6xx keywords `register`, `near` and `far` are not supported by the compiler. When including the migration header file, these keywords are accepted but have no effect on the compilation results.

The TI6xx keywords `interrupt` and `inline` are supported, but are implemented differently, using `#pragma inline` and `#pragma interrupt`. As a result, no automatic translation is provided. The code must be modified to use the pragmas supported by the compiler. For further information, see [Section 3.4.5, “Pragmas,”](#) on page 3-54.

- **Pragmas:** TI6xx pragmas are ignored. Warnings are issued, but the correctness of the compilation is not affected.
- **Inlined assembly code:** By definition, inlined assembly code is not portable from one environment to another. The assembler is unable to recognize inlined TI6xx assembly code, and issues errors.
- **Intrinsic functions:** The TI6xx intrinsic functions listed in the `portc6xtosc1.h` header file are supported. These are functionally equivalent to their corresponding TI6xx intrinsic functions, but their performance may be significantly affected.

Index

A

- abort environment function 7-16
- abs integer arithmetic function 7-14
- abs_s intrinsic function 3-48, 7-20
- acos trigonometric function 7-8
- add intrinsic function 3-48, 7-20
- align #pragma 3-55
- Alignment
 - bit-fields 3-41
 - variables 3-61
- ansi shell option 3-13, 3-21
- Application configuration file 6-12
 - binding section 6-15
 - overlay section 6-16
 - schedule section 6-13
- Application entry point 6-3
- Arithmetic
 - fixed point 3-41
 - floating point 3-41
 - fractional 3-43
 - integer 3-43
- asctime time function 7-19
- asin trigonometric function 7-8
- asm statement 4-2
- Assembly functions 4-7
- Assembly instruction inlining
 - asm statement 4-2
- Assembly instructions
 - inlining sequence 4-2
 - inlining single instruction 4-1
- atan trigonometric function 7-8
- atan2 trigonometric function 7-8
- atexit environment function 7-16
- atof string conversion function 7-15
- atoi string conversion function 7-15
- atol string conversion function 7-15

B

- Bare board startup 6-2
- Bare board startup code 6-1
- Basic block 5-2, 5-25
- be shell option 3-14, 3-26, 3-39, 3-41, 3-42
- Big memory model 6-5
- Big-endian mode 3-26
- Binding section
 - application configuration file 6-15
- Bit-fields 3-41

- BitReverseUpdate intrinsic function 3-51, 7-20
- Built-in intrinsic functions 7-20

C

- C environment startup 6-3
- C environment startup code 6-1
- C language
 - dialects 3-27
 - extensions 3-27
 - K&R 3-32
 - PCC 3-32
- C language options 3-21
- C shell option 3-12, 3-17
- c shell option 3-12, 3-15
- Call tree 6-13
- Calling convention
 - stack-based 6-18
 - stack-less 6-20
- calloc memory allocation function 7-14
- ceil function 7-9
- cfe shell option 3-12, 3-15
- Character typing 7-2
- clearerr stream function 7-12
- clock time function 7-19
- clock_t time function 7-19
- Code
 - linear 5-3
 - migrating from other environments A-1
 - parallelized 5-3
 - transformations 5-2
- Command file 3-10, 3-16
- Command line 3-10
 - syntax 3-10
- Common subexpression elimination 5-17
- Comparison functions 7-17
- Compatibility clause 6-16
- Compilation process 1-4, 3-2
- Composed variable loop 5-13
- Concatenation functions 7-17
- Conditional execution 5-25
- Configuration
 - memory map 6-9
 - startup code 6-4
- Constant folding 5-17
- Control options 3-11
- Conversion functions 7-2
- Copying functions 7-17
- cos trigonometric function 7-8

cosh hyperbolic function 7-8
-crt shell option 3-14, 3-25
ctime time function 7-19
ctype.h library 7-1, 7-2

D

-D shell option 3-12, 3-18
D_add intrinsic function 3-49, 7-20
D_cmpeq intrinsic function 3-49, 7-20
D_cmpgt intrinsic function 3-49, 7-20
D_extract_h intrinsic function 3-49, 7-20
D_extract_l intrinsic function 3-49, 7-20
D_mac intrinsic function 3-49, 7-20
D_msu intrinsic function 3-49, 7-20
D_mult intrinsic function 3-49, 7-20
D_round intrinsic function 3-49, 7-20
D_sat intrinsic function 3-49, 7-20
D_set intrinsic function 3-49, 7-21
D_sub intrinsic function 3-49, 7-21
Data allocation
 static 6-8
Data types 3-37
 bit-fields 3-41
 character 3-38
 double precision fractional 3-46
 extended precision fractional 3-46
 floating point 3-40
 fractional long 3-46
 fractional representation 3-41
 fractional short 3-46
 integer 3-39
 pointers 3-41
-dc shell option 3-14, 3-23
-de shell option 3-14, 3-23
Dead assignment
 elimination 5-18
Dead code
 elimination 5-18
Dead storage
 elimination 5-18
debug intrinsic function 3-51, 7-21
debugev intrinsic function 3-51, 7-21
Delay slots 5-21
Dependencies
 between instructions 5-3
Dependency 5-8
di intrinsic function 3-51, 7-21
Dialects
 C language 3-27
difftime time function 7-19
div integer arithmetic function 7-14
div_s intrinsic function 3-48, 7-21
-dL shell option 3-14, 3-23
-dL1 shell option 3-14, 3-23

-dL2 shell option 3-14, 3-23
-dL3 shell option 3-14, 3-23
-dm shell option 3-14, 3-23
-do shell option 3-14, 3-23
Double precision 3-47
DSP56600 compiler
 differences A-2
 header file A-1, A-2
 migrating code A-1, A-2
-dx shell option 3-14, 3-23
Dynamic loop 5-11
Dynamic memory allocation 6-8

E

-E shell option 3-12, 3-15
ei intrinsic function 3-51, 7-21
Elimination
 dead assignment 5-18
 dead code 5-18
 dead storage 5-18
 jump-to-jump 5-18
 subexpression 5-17
EndBitReverse intrinsic function 3-51, 7-21
Entry points 6-13
Environment functions 7-16
Environment variables 3-9
Execution sets 5-3
 parallelized 5-8
Execution units 5-3
exit environment function 7-16
exp function 7-9
Exponential functions 7-9
Extended 3-50
Extended precision 3-46
Extensions 3-7, 3-19
 C language 3-27
external #pragma 3-55
External function 3-57
extract_h intrinsic function 3-49, 7-21
extract_l intrinsic function 3-49, 7-21

F

-F shell option 3-12, 3-16
fabs function 7-9
fclose stream function 7-12
feof stream function 7-12
ferror stream function 7-12
fflush I/O function 7-13
fgetc input function 7-11
fgetpos stream function 7-12
fgets output function 7-13
File extensions 3-7, 3-19
File types 3-7

- Finalization code 6-1, 6-3
- float.h library 7-1, 7-3
- Floating point arithmetic 3-41
- Floating point characteristics 7-3
- Floating point math 7-8
- floor function 7-9
- fmod function 7-9
- fopen I/O function 7-13
- fprintf output function 7-13
- fputc output function 7-13
- fputs output function 7-13
- Fractional
 - arithmetic 3-43
 - constants 3-47
 - representation 3-41
 - values 3-47
- fread input function 7-11
- free memory allocation function 7-14
- freopen stream function 7-12
- frexp function 7-9
- fscanf input function 7-11
- fseek stream function 7-12
- fsetpos stream function 7-12
- ftell stream function 7-12
- Function inlining 5-16
- Functions
 - built-in intrinsic 7-20
 - comparison 7-17
 - concatenation 7-17
 - conversion 7-2
 - copying 7-17
 - environment 7-16
 - exponential 7-9
 - external 3-57
 - hyperbolic 7-8
 - I/O 7-13
 - input 7-11
 - integer arithmetic 7-14
 - intrinsic 3-46, 3-48, 7-20
 - logarithmic 7-9
 - memory allocation 7-14
 - multibyte character 7-16
 - output 7-13
 - power 7-9
 - pseudo random number generation 7-15
 - search 7-18
 - searching 7-15
 - sorting 7-15
 - stream 7-12
 - string 7-16
 - string conversion 7-15
 - testing 7-2
 - time 7-19
 - trigonometric 7-8

- fwrite output function 7-13

G

- g shell option 3-13, 3-21
- General utilities 7-14
- getc input function 7-11
- getchar input function 7-11
- getenv environment function 7-16
- gets output function 7-13
- Global
 - mode 5-4, 5-7, 5-30
 - optimization 1-2, 3-4, 5-4, 5-7, 5-30
- Global variables 6-16
- gmtime time function 7-19
- Guidelines
 - optimizer 5-31

H

- h shell option 3-12, 3-16
- Hardware loops 6-22
- Hardware registers
 - initialization 6-2
- Header file
 - Tl6xx compiler A-1
- Heap 6-8
- High level optimizations 5-9
- Hyperbolic functions 7-8

I

- I shell option 3-12, 3-18
- I/O functions 7-13
- I/O services
 - low level 6-3
 - termination 6-3
- illegal intrinsic function 3-51, 7-21
- Include files 3-18
- InitBitReverse intrinsic function 3-51, 7-21
- Initialization
 - M registers 6-2
 - status registers 6-2
 - variables 3-26, 6-3
- Initialization code 6-1, 6-3
- Initializing variables with fractional values 3-47
- inline #pragma 3-55
- Inlining 3-56, 5-16
 - sequence of assembly instructions 4-2
 - single assembly instruction 4-1
- Input file extension 3-19
- Input functions 7-11
- Instruction scheduling 5-21
- Instruction transformations 5-2
- Integer arithmetic 3-43
- Integer arithmetic functions 7-14

- Integer characteristics 7-7
- interrupt #pragma 3-55
- Interrupt entry 6-15, 6-22
- Interrupt handler 3-58, 6-13, 6-22
- Interrupt vector 6-2, 6-15, 6-22
- Interrupts 6-3
- Intrinsic functions 3-46
 - architecture primitives 3-51
 - assembly instruction architecture primitives 3-51
 - bit reverse addressing 3-51
 - double precision fractional arithmetic 3-49
 - fractional arithmetic 3-48
 - fractional arithmetic with guard bits 3-50
 - long fractional arithmetic 3-49
- Invariant code loop 5-17
- isalnum testing function 7-2
- isalpha testing function 7-2
- iscentrl testing function 7-2
- isdigit testing function 7-2
- isgraph testing function 7-2
- islower testing function 7-2
- ISO libraries 7-1
- isprint testing function 7-2
- ispunct testing function 7-2
- isspace testing function 7-2
- isupper testing function 7-2
- isxdigit testing function 7-2

J

- Jump-to-jump elimination 5-18

K

- K&R mode 3-32
- kr shell option 3-13, 3-21

L

- L_abs intrinsic function 3-49, 7-21
- L_add intrinsic function 7-21
- L_deposit_h intrinsic function 3-49, 7-21
- L_deposit_l intrinsic function 3-49, 7-21
- L_mac intrinsic function 3-48, 7-21
- L_max intrinsic function 3-49, 7-21
- L_min intrinsic function 3-49, 7-21
- L_msu intrinsic function 3-48, 7-22
- L_mult intrinsic function 7-22
- L_negate intrinsic function 3-49, 7-22
- L_rol intrinsic function 3-51, 7-22
- L_ror intrinsic function 3-51, 7-22
- L_sat intrinsic function 3-49, 7-22
- L_shl intrinsic function 3-49, 7-22
- L_shr intrinsic function 3-49, 7-22
- L_shr_r intrinsic function 3-49, 7-22
- L_sub intrinsic function 3-49, 7-22

- labs integer arithmetic function 7-14
- L-add intrinsic function 3-49
- ldexp function 7-9
- ldiv integer arithmetic function 7-14
- Libraries
 - ISO 7-1
 - non-ISO 7-1
- limits.h library 7-1, 7-7
- Linear code 5-3
- Linker command file 6-3, 6-5
- Listing files 3-23
- Little-endian 3-26
- Little-endian mode 3-42
- Little-endian representation 3-39, 3-41
- L-mult intrinsic function 3-49
- locale.h library 7-1, 7-7
- localeconv locales function 7-7
- Locales functions 7-7
- localtime time function 7-19
- log function 7-9
- log10 function 7-9
- Logarithmic functions 7-9
- Logical memory 6-10
- Loop
 - composed variable 5-13
 - dynamic 5-11
 - multi-step 5-12
 - simple 5-10
 - square 5-13
 - transformations 5-10
- Loop count 3-59
- loop_count #pragma 3-55
- Loops
 - hardware 6-22
- Low level optimizations 5-8, 5-19
- Low level peephole 5-28
- Low level transformations (LLT) 5-19

M

- M registers
 - initialization 6-2
 - value 6-23
- M shell option 3-12, 3-17
- ma shell option 3-14, 3-25
- mac_r intrinsic function 3-48, 7-22
- Machine configuration file 6-10
- Macros 3-18
 - fractional values 3-47
 - predefined 3-63
 - preprocessor 3-18
- Main entry point 6-13
- malloc memory allocation function 7-14
- mark intrinsic function 3-51, 7-22
- math.h library 7-1, 7-8

- max intrinsic function 3-48, 7-22
- mb shell option 3-14, 3-25
- mc shell option 3-14, 3-25
- memchr search function 7-18
- memcmp comparison function 7-17
- memcpy copying function 7-17
- memmove copying function 7-17
- Memory
 - logical 6-10
 - mode 3-25, 6-5
 - physical 6-10
- Memory allocation
 - dynamic 6-8
 - functions 7-14
- Memory layout
 - default 6-6
- Memory map
 - configuration 6-9
 - default values 6-7
 - initialization 6-3
- Memory model
 - big 6-5
 - small 6-5
- Memory space 6-10
- memset function 7-18
- Messages 3-23
- MH shell option 3-12, 3-17
- Migrating code A-1
- min intrinsic function 3-48, 7-22
- mktime time function 7-19
- Mode
 - global 5-30
 - K&R/PCC 3-32
 - non-global optimization 5-4
 - optimization 5-4
- modf function 7-9
- mpysu intrinsic function 3-51, 7-23
- mpyus intrinsic function 3-51, 7-23
- mpyuu intrinsic function 3-51, 7-22
- mrom shell option 3-14
- msu_r intrinsic function 3-48, 7-23
- mult intrinsic function 3-48, 7-23
- mult_r intrinsic function 3-48, 7-23
- Multibyte character functions 7-16
- Multiple execution units 5-3
- Multi-step loop 5-12

N

- n shell option 3-14, 3-24
- negate intrinsic function 3-48, 7-23
- noinline #pragma 3-55
- Non-global optimization 3-4, 5-4

- Non-ISO libraries 7-1
- Nonlocal jumps 7-10
- norm_l intrinsic function 3-49, 7-23
- norm_s intrinsic function 3-48, 7-23

O

- O shell option 3-13, 5-5, 5-19
- o shell option 3-12, 3-20
- O0 shell option 3-13, 5-5, 5-17
- O1 shell option 3-13, 5-5, 5-9
- O2 shell option 3-13, 5-5, 5-19
- Og shell option 3-13, 5-5, 5-30
- Optimization
 - for size 5-7, 5-30
 - for speed 5-7, 5-29
 - global 3-4, 5-4, 5-7, 5-30
 - high level 5-8, 5-9
 - levels 5-4
 - low level 5-8, 5-19
 - non-global 3-4, 5-4
 - options 5-4, 5-5
- Optimizer
 - guidelines 5-31
 - invoking 5-6
- Options
 - C language 3-21
 - control 3-11
 - extensions 3-19
 - messages 3-23
 - output files 3-20
 - shell 3-12
- Os shell option 3-13, 5-5, 5-30
- Ot shell option 3-13, 5-29
- Ot1 shell option 3-13, 5-29
- Ot2 shell option 3-13, 5-29
- Output files 3-20
- Output functions 7-13
- Overlay section
 - application configuration file 6-16
- Overlay specification 6-13

P

- Parallelized code 5-3
- Parallelized execution sets 5-8
- Passing options 3-22
- perror output function 7-13
- Physical memory 6-10
- Pipeline restrictions 5-22
- Pointers 3-41
- Post-increment detection 5-27
- pow function 7-9
- Power functions 7-9

Pragmas

- #pragma align 3-55, 3-62
- #pragma external 3-55, 3-57
- #pragma inline 3-55
- #pragma interrupt 3-58
- #pragma loop count 3-55, 3-60
- #pragma noinline 3-55, 3-56
- #pragma profile 3-55, 3-58, 3-59
- #pragma save 3-55
- #pragma save_ctxt 3-56
- placement 3-54
- syntax 3-54

Predefined macros 3-63

Prefix grouping 5-30

Preprocessing options 3-17

Preprocessor macros 3-18

printf output function 7-13

-pro shell option 3-14, 3-26

Process time 7-19

profile #pragma 3-55

Profile information file 3-26

Profile value 3-58

prototype.h 7-1

prototype.h library 7-1, 7-20

Pseudo random number generation functions 7-15

putc output function 7-13

putchar output function 7-13

puts output function 7-13

Q

-q shell option 3-14, 3-24

R

-r shell option 3-12, 3-20

rand pseudo random number generation function 7-15

realloc memory allocation function 7-14

remove stream function 7-12

rename stream function 7-12

Reporting 3-24

Reset interrupt vector 6-2

-rev shell option 3-14, 3-26

rewind stream function 7-12

round intrinsic function 3-48, 7-23

Runtime

- environment 6-1

- startup code 6-1

S

-S shell option 3-12, 3-15

saturate intrinsic function 3-48, 7-23

save_ctxt #pragma 3-55

-sc shell option 3-13, 3-21

scanf input function 7-11

Schedule section

- application configuration file 6-13

Search functions 7-18

Searching functions 7-15

set2cnvrm intrinsic function 3-51, 7-23

set2crm intrinsic function 3-51, 7-24

setbuf stream function 7-12

setjmp.h library 7-1, 7-10

setlocale locales function 7-7

setnosat intrinsic function 3-51, 7-24

setsat32 intrinsic function 3-51, 7-24

setvbuf stream function 7-12

Shell 1-3, 3-1

Shell command file 3-16

Shell command line 3-10

Shell options

- behavior control 3-12

- C language 3-13

- file extension override 3-12

- hardware model and configuration 3-14

- optimization pragma and code 3-13

- output filename and location 3-12

- output of listing files and messages control 3-14

- pass-through 3-13

- preprocessing 3-12

- stop processing 3-15

- summary 3-12

shl intrinsic function 3-48, 7-24

shr intrinsic function 3-48, 7-24

shr_r intrinsic function 3-48, 7-24

Signal handling 7-10

signal.h library 6-22, 7-1, 7-10

Silicon revision 3-26

Simple loop 5-10

sin trigonometric function 7-8

sinh hyperbolic function 7-8

Small memory model 6-5

Software pipelining 5-22

Sorting functions 7-15

Space optimization 5-7, 5-30

Speculative execution 5-26

Speed

- optimization 5-7

sprintf output function 7-13

sqrt function 7-9

Square loop 5-13

srand pseudo random number generation function 7-15

sscanf input function 7-11

Stack

- frame 6-21

- memory allocation 6-7

- pointer 6-8, 6-18

- space 6-21

- start address 6-3, 6-8

- Stack-based
 - calling convention 6-18
- Stack-less
 - calling convention 6-20
- Standard definitions 7-11
- Startup code 6-1
 - bare board 6-1
 - C environment 6-1
 - configuration 6-4
- Static data allocation 6-8
- Status registers
 - default settings 6-2
 - initialization 6-2
- stdarg.h library 7-1, 7-10
- stddef.h library 7-1, 7-11
- stderr stream function 7-12
- stdin stream function 7-12
- stdio.h library 7-1, 7-11
- stdlib.h library 7-1, 7-14
- stdout stream function 7-12
- stop 3-15
- stop intrinsic function 3-51, 7-24
- strcat concatenation function 7-17
- strchr search function 7-18
- strcmp comparison function 7-17
- strcoll comparison function 7-17
- strcpy copying function 7-17
- strcspn search function 7-18
- Stream functions 7-12
- Strength reduction 5-10
- strerror function 7-18
- strftime time function 7-19
- String conversion functions 7-15
- String functions 7-16
- string.h library 7-1, 7-16
- strlen function 7-18
- strncat concatenation function 7-17
- strncmp comparison function 7-17
- strncpy copying function 7-17
- strpbrk search function 7-18
- strrchr search function 7-18
- strspn search function 7-18
- strstr search function 7-18
- strtod string conversion function 7-15
- strtok search function 7-18
- strtol string conversion function 7-15
- strtoul string conversion function 7-15
- strxfrm comparison function 7-17
- sub intrinsic function 3-48, 7-24
- Subexpression elimination 5-17
- Symbolic labels 4-10
- System context 3-56

T

- tan trigonometric function 7-8
- tanh hyperbolic function 7-8
- Task entry point 6-13
- Termination
 - I/O services 6-3
- Testing functions 7-2
- TI6xx compiler
 - header file A-1
 - migrating code A-1
- Time constant 7-19
- time function 7-19
- Time functions 7-19
- Time optimization 5-29
- time.h library 7-1, 7-19
- time_t time function 7-19
- Timer 6-2
- tolower conversion function 7-2
- toupper conversion function 7-2
- Transformations
 - loop 5-10
- trap intrinsic function 3-51, 7-24
- Trigonometric functions 7-8

U

- U shell option 3-12, 3-18
- ungetc I/O function 7-13
- usc shell option 3-13, 3-21

V

- v shell option 3-14, 3-24
- Variable arguments 7-10
- Variables
 - alignment 3-61
 - initialization 6-3
- vfprintf output function 7-13
- vprintf output function 7-13
- vsprintf output function 7-13

W

- w shell option 3-14, 3-24
- wait intrinsic function 3-51, 7-24
- Wall shell option 3-14, 3-24
- Warnings 3-24
- Wg shell option 3-14, 3-24
- Wj shell option 3-14, 3-24
- WORD16 macro 3-47
- WORD32 macro 3-47
- Word40 extended precision fractional 3-46
- Word64 double precision fractional 3-47

X

X_abs intrinsic function 3-50, 7-24
X_add intrinsic function 3-50, 7-24
X_cmpeq intrinsic function 3-50, 7-24
X_cmpgt intrinsic function 3-50, 7-24
X_extend intrinsic function 3-50, 7-24
X_extract_h intrinsic function 3-50, 7-24
X_extract_l intrinsic function 3-50, 7-24
X_mac intrinsic function 3-50, 7-24
X_msu intrinsic function 3-50, 7-25
X_mult intrinsic function 3-50, 7-25
X_norm intrinsic function 3-50, 7-25
X_or intrinsic function 3-50, 7-25
X_rol intrinsic function 3-50, 7-25
X_ror intrinsic function 3-50, 7-25
X_round intrinsic function 3-50, 7-25
X_sat intrinsic function 3-50, 7-25
X_set intrinsic function 3-50, 7-25
X_shl intrinsic function 3-50, 7-25
X_shr intrinsic function 3-50, 7-25
X_sub intrinsic function 3-50, 7-25
X_trunc intrinsic function 3-50, 7-25
-Xasm shell option 3-13, 3-22
-xasm shell option 3-12, 3-19
-xc shell option 3-12, 3-19
-Xlnk shell option 3-13, 3-22
-xobj shell option 3-12, 3-19



How to reach us:

Motorola Literature Distribution

P.O. Box 5405
Denver, Colorado 80217
1 (800) 441-2447

Asia/Pacific

Motorola Semiconductors H.K. Ltd., Hong Kong
852-26629298

Japan

Motorola Japan, Ltd., Shinagawa-ku, Japan
81-3-5487-8488

Motorola Fax Back System (Mfax™)

1 (800) 774-1848; RMFAX0@email.sps.mot.com

DSP Helpline

dsphelp@dsp.sps.mot.com

Technical Resource Center

1 (800) 521-6274

Internet

<http://www.motorola-dsp.com>

Lucent Technologies Microelectronics Group

Internet

<http://www.lucent.com/micro/dsp>

Email

docmaster@micro.lucent.com

U.S./Canada

Lucent Technologies Microelectronics Group

1-800-372-2447, FAX 610-712-4106

In CANADA: 1-800-553-2448, FAX 610-712-4106

Asia/Pacific Microelectronics Group

Lucent Technologies Singapore Pte. Ltd., Singapore

Tel. (65) 778 8833, FAX (65) 777 7495

China Microelectronics Group

Lucent Technologies (China) Co., Ltd., Shanghai

Tel. (86) 21 6440 0468, ext. 316, Fax (86)21 6440 0652

Japan Microelectronics Group

Lucent Technologies Japan, Ltd., Shinagawa-ku, Japan

Tel. (81) 3 5421 1600, FAX (81) 3 5421 1700

Europe Microelectronics Group Dataline

Tel. (44) 1189 324 299, FAX (44) 1189 328 148



MOTOROLA

microelectronics group

Lucent Technologies

