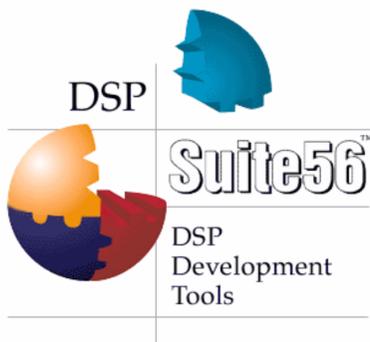


Motorola DSP56000 Family Optimizing C Compiler

User's Manual, Release 6.3

DSP56KCCUM/D
Document Rev. 2.0, 07/1999



Suite56, OnCe, and MFAX are trademarks of Motorola, Inc.

Motorola reserves the right to make changes without further notice to any products herein. Motorola makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Motorola 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 consequential or incidental damages. "Typical" parameters which may be provided in Motorola 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. Motorola does not convey any license under its patent rights nor the rights of others. Motorola 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 Motorola product could create a situation where personal injury or death may occur. Should Buyer purchase or use Motorola products for any such unintended or unauthorized application, Buyer shall indemnify and hold Motorola 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 Motorola 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.

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

Table of Contents

Chapter 1 Introduction

1.1	Overview	1-1
1.2	Error Codes	1-4
1.3	Notation	1-4
1.4	Manual Organization	1-5

Chapter 2 Installation Guide

2.1	Introduction	2-1
2.2	Installation On An MS-DOS Machine (80386 or 80486)	2-1
2.3	Standard Installation On A SUN	2-3
2.4	Alternate Installation On A SUN	2-4
2.5	Test Program	2-4

Chapter 3 Control Program Options

3.1	Overview	3-1
3.2	G56k Command Line Options	3-5
3.2.1	Preprocessor Phase Options	3-7
3.2.2	Compile Phase Options	3-20
3.3	Assemble Phase Options	3-30
3.4	Link Phase Options	3-31

Chapter 4 About g56k

4.1	Introduction	4-1
4.2	Identifiers	4-1
4.3	Predefined Preprocessor Macro Names	4-1
4.4	Data Types and Sizes	4-1
4.4.1	Integral Data Types	4-2
4.4.2	Floating-point Types	4-3
4.4.2.1	Floating-point Format Description	4-3
4.4.2.2	Comparison with IEEE STD 754-1985 Standard	4-5

4.4.3	Pointer Types	4-7
4.5	Register Usage	4-8
4.6	Memory Usage	4-9
4.6.1	Activation Record	4-10
4.6.2	Global/Static Data	4-13
4.7	Compiler Naming Conventions	4-13
4.8	Subroutine Call Sequence	4-14
4.8.1	Caller Sequence	4-14
4.8.2	Callee Sequence	4-15
4.8.3	Return Sequence	4-15
4.9	Software Support for Arithmetic Routines	4-15
4.10	Run-time Safety	4-16
4.10.1	Memory Allocation Checks	4-16
4.10.2	Run-time Stack Checks	4-17
4.11	Optimization Techniques Implemented	4-17
4.11.1	Register Promotion and Lifetime Analysis	4-17
4.11.2	Common Sub-expression Elimination	4-17
4.11.3	Constant Propagation/Folding	4-18
4.11.4	Dead Code Elimination	4-18
4.11.5	Tail Merging	4-18
4.11.6	Strength Reduction	4-19
4.11.7	Loop Invariant Code Motion	4-19
4.11.8	Hardware DO Loop Instruction	4-19
4.11.9	Loop Rotation	4-20
4.11.10	Jump Optimizations	4-20
4.11.11	Instruction Combination	4-20
4.11.12	Leaf Routine Detection	4-21
4.11.13	Function In-lining	4-21
4.11.14	Instruction Scheduling / Microcode Compaction	4-21

Chapter 5 Mixing C and Assembly Language

5.1	Overview	5-1
5.2	In-line Assembly Code	5-1
5.2.1	Instruction Template	5-3
5.2.2	Output/Input Operands	5-5
5.2.3	Explicit Register Saving	5-7
5.2.4	In-line Assembly Code Examples	5-8
5.2.5	Controlling Labels Generated by the Compiler	5-17
5.2.5.1	Calling Assembly Subroutines	5-17

5.2.5.2	Calling C Subroutines from Assembly Code	5-18
5.2.5.3	Referencing Assembly Global Variables from C	5-20
5.2.5.4	Referencing Global C Variables from Assembly Language	5-21
5.2.6	Specifying Registers for Variables.	5-22
5.2.7	Optimizer Effects on Code.	5-22
5.3	#pragma Directive	5-22
5.4	Out-of-line Assembly Code.	5-25
5.4.1	General Template	5-26
5.4.1.1	Prologue	5-27
5.4.1.2	Save all registers.	5-27
5.4.1.3	Main Program	5-28
5.4.1.4	Restore all registers	5-28
5.4.1.5	Epilogue	5-29
5.4.1.6	Out-of-line Assembly Code Example.	5-29
5.4.2	Global C and Static Variables in C	5-33
5.4.3	Using Run-time Stack for Local Data	5-34
5.4.4	Calling C Routines	5-35
5.4.5	Optimization Techniques	5-36

Chapter 6 Software-Hardware Integration

6.1	Overview	6-1
6.2	Run-Time Environment Specification Files	6-1
6.3	crt0 File	6-2
6.4	Bootstrapping the C program.	6-2
6.5	Memory Configuration and Management	6-3
6.6	Interrupt Vectors	6-6
6.7	Miscellaneous Code.	6-8
6.8	Signal File	6-9
6.9	Signal()	6-10
6.9.1	Raise()	6-11
6.10	Setjmp File.	6-12
6.11	Host-Supported I/O (printf (), et al).	6-12
6.11.1	DSP functions __send () and __receive ()	6-13
6.11.2	The Host-Side I/O Driver.	6-13
6.11.3	Communication between the Host and DSP	6-13

Appendix A Programming Support

A.1	Standard ANSI Header Files	A-1
A.2	ANSI C Library Sub-routines	A-2
A.2.1	Hosted vs. Non-Hosted Library Routines	A-2
A.3	Forcing Library Routines Out-of-line	A-2
A.4	Function Descriptions	A-4
A.4.1	abort—Force abnormal program termination	A-9
A.4.2	abs—Absolute value of integer	A-10
A.4.3	acos—Arc cosine	A-11
A.4.4	asin—Arc sine	A-12
A.4.5	atan—Arc tangent	A-13
A.4.6	atan2—Arc tangent of angle defined by point y/x	A-14
A.4.7	atexit—Register a function for execution at normal program termination	A-15
A.4.8	atof—String to floating point	A-16
A.4.9	atoi—String to integer	A-17
A.4.10	atol—String to long integer	A-18
A.4.11	bsearch—Perform binary search	A-19
A.4.12	calloc—Dynamically allocate zero-initialized storage for objects	A-20
A.4.13	ceil—Ceiling function	A-22
A.4.14	clearerr—Clear any error indicators associated with a specified stream	A-22
A.4.15	cos—Cosine	A-23
A.4.16	cosh—Hyperbolic cosine	A-23
A.4.17	div—integer division with remainder	A-24
A.4.18	exit—Terminate program normally	A-25
A.4.19	exp—Exponential, ex.	A-26
A.4.20	fabs—Absolute value of a double	A-26
A.4.21	fclose—Close a stream.	A-27
A.4.22	feof—Test the end-of-file indicator of a specified stream	A-28
A.4.23	ferror—Test the error indicator of a stream	A-29
A.4.24	fflush—Flush all pending output associated with a stream	A-29
A.4.25	fgetc—Read a character from the specified stream	A-30
A.4.26	fgetpos—Get value of file position indicator associated with a stream	A-31
A.4.27	fgets—Read a string from the specified stream	A-32
A.4.28	floor—Floor function	A-33
A.4.29	fmod—Floating point remainder	A-33
A.4.30	fopen—Open a named file on the host's disk	A-34
A.4.31	fprintf—Write formatted output to a stream	A-36
A.4.32	fputc—Write a single character to a stream	A-36
A.4.33	fputs—Write a string to a stream	A-37

A.4.34	fread—Read data directly from a stream	A-38
A.4.35	free—Free storage allocated by calloc, malloc, and realloc	A-39
A.4.36	freopen—Open named file on disk, to be accessed via specified stream . .	A-40
A.4.37	frexp—Break a floating point number into mantissa and exponent	A-41
A.4.38	fscanf—Read formatted input from a stream.	A-42
A.4.39	fseek—Set the file position indicator associated with a stream.	A-44
A.4.40	fsetpos—Set the file position indicator associated with a stream	A-45
A.4.41	ftell—Get the file position indicator associated with a stream	A-46
A.4.42	fwrite—Write data directly to a stream	A-47
A.4.43	getc—Read a character from the specified stream.	A-48
A.4.44	getchar—Read a character from standard input.	A-48
A.4.45	gets—Read a string from standard input	A-49
A.4.46	isalnum—Test for alphanumeric character	A-50
A.4.47	isalpha—Test for alphabetic character.	A-51
A.4.48	iscntrl—Test for control character	A-52
A.4.49	isdigit—Test for numeric character	A-52
A.4.50	isgraph—Test for printing character, excluding space and tab	A-54
A.4.51	islower—Test for lower-case alphabetic characters.	A-54
A.4.52	isprint—Test for printing character, excluding '\t'	A-55
A.4.53	ispunct—Test for punctuation character	A-56
A.4.54	isspace—Test for white-space character	A-56
A.4.55	isupper—Test for upper-case alphabetic character	A-57
A.4.56	isxdigit—Test for hexadecimal numeric character	A-58
A.4.57	labs—Absolute value of a long integer	A-58
A.4.58	ldexp—Multiply floating point number by 2.	A-59
A.4.59	ldiv—Long integer division with remainder	A-60
A.4.60	log—Natural logarithm, base e	A-61
A.4.61	log10—Base ten logarithm	A-62
A.4.62	longjmp—Execute a non-local jump	A-63
A.4.63	malloc—Dynamically allocate uninitialized storage.	A-64
A.4.64	mblen—Length of a multibyte character	A-64
A.4.65	mbstowcs—Convert multibyte string to wide character string	A-66
A.4.66	mbtowc—Convert a multibyte character to a wide character	A-68
A.4.67	memchr—Find a character in a memory area	A-70
A.4.68	memcmp—Compare portion of two memory areas.	A-71
A.4.69	memcpy—Copy from one area to another.	A-72
A.4.70	memmove—Copy storage	A-73
A.4.71	memset—Initialize memory area	A-74
A.4.72	modf—Break a double into it's integral and fractional parts.	A-75
A.4.73	perror—Print error message.	A-76

A.4.74	pow—Raise a double to a power	A-77
A.4.75	printf—Print to standard output	A-78
A.4.76	putc—Write a single character to a stream	A-83
A.4.77	putchar—Write a character to standard output	A-84
A.4.78	puts—Write a string to standard output	A-85
A.4.79	qsort—Quick sort	A-86
A.4.80	raise—Raise a signal	A-87
A.4.81	rand—Pseudo- random number generator	A-88
A.4.82	realloc—change size of dynamically allocated storage area	A-88
A.4.83	remove—Remove a file from the disk	A-90
A.4.84	rename —Rename a file on the disk	A-90
A.4.85	rewind—Reset the file position indicator to the beginning of the file.	A-90
A.4.86	scanf—Read formatted input from standard input.	A-91
A.4.87	setjmp—Save reference of current calling environment for later use by longjmp	A-92
A.4.88	setbuf—Read formatted input from standard input	A-93
A.4.89	setvbuf—Alter stream buffering	A-94
A.4.90	signal—Set up signal handler.	A-94
A.4.91	sin—Sine	A-95
A.4.92	sinh—Hyperbolic Sine	A-96
A.4.93	sprintf—Print to a string.	A-96
A.4.94	sqrt—Square root	A-97
A.4.95	srand—Seed the pseudo-random number generator	A-98
A.4.96	sscanf—Read formatted input from a string	A-98
A.4.97	strcat—Concatenate two strings.	A-99
A.4.98	strchr—Find first occurrence of a character in a string	A-100
A.4.99	strcmp—Compare two strings	A-100
A.4.100	strcoll—Compare two strings based on current locale	A-102
A.4.101	strcpy—Copy one string into another	A-102
A.4.102	strcspn—Compute length of prefix of one string consisting entirely of characters not from another	A-104
A.4.103	strerror—Map error code into an error message string	A-105
A.4.104	strlen—Determine length of a string	A-106
A.4.105	strncat—Concatenate a portion of one string to another	A-106
A.4.106	strncmp—Compare a portions of two strings	A-108
A.4.107	strncpy—Copy a portion of one string into another.	A-109
A.4.108	strpbrk—Find first occurrence of a character from one string in another .	A-110
A.4.109	strrchr—Find last occurrence of a character from one string in another . .	A-111
A.4.110	strspn—Compute length of prefix of one string consisting entirely of characters from another	A-112

A.4.111	strstr—Find the first occurrence of one string in another	A-112
A.4.112	strtod—String to double	A-114
A.4.113	strtok—Break string into tokens	A-115
A.4.114	strtol—String to long integer	A-116
A.4.115	strtoul—String to unsigned long integer	A-118
A.4.116	strxfrm—Transform a string into locale-independent form.	A-120
A.4.117	tan—Tangent	A-120
A.4.118	tanh—Hyperbolic tangent	A-121
A.4.119	tmpfile—Create a temporary binary file	A-121
A.4.120	tmpnam—Create a temporary file name	A-122
A.4.121	tolower—Convert uppercase character to lowercase.	A-122
A.4.122	toupper—Convert lowercase character to uppercase.	A-123
A.4.123	ungetc—Push a character back onto an input stream.	A-124
A.4.124	vfprintf—Write formatted output to a stream using a va_list	A-124
A.4.125	vprintf—Write formatted output to standard output using a va_list	A-125
A.4.126	vsprintf—Write formatted output to a string using a va_list	A-126
A.4.127	wcstombs—Convert wchar_t array to multibyte string	A-126
A.4.128	wctomb—Convert wchar_t character to multibyte character	A-128

Appendix B DSP56000/DSP56001 Instruction Set and Assembler Directive Summary

B.1	Overview	B-1
B.2	Instruction Set	B-1
B.2.1	Arithmetic Instructions.	B-1
B.2.2	Logical Instructions	B-2
B.2.3	Bit Manipulation Instructions	B-3
B.2.4	Loop Instructions	B-3
B.2.5	Move Instructions.	B-4
B.2.6	Program Control Instructions.	B-4
B.3	Directive Summary	B-5
B.3.1	Assembly Control.	B-5
B.3.2	Symbol Definition	B-6
B.3.3	Data Definition/Storage Allocation	B-6
B.3.4	Listing Control and Options.	B-6
B.3.5	Object File Control.	B-7
B.3.6	Macros and Conditional Assembly	B-7
B.3.7	Structured Programming	B-8

Appendix C Utilities

C.1	Utility Programs	C-1
C.1.1	asm56000—Motorola DSP56000/DSP56001 COFF Assembler	C-1
C.1.1.1	asm5600 Options	C-2
C.1.2	cldinfo—Memory size information from Motorola DSP COFF object file	C-4
C.1.3	cldlod—Motorola COFF to LOD Format converter	C-5
C.1.4	cofdmp—Motorola DSP COFF File Dump Utility	C-5
C.1.4.1	cofdmp Options	C-5
C.1.5	dsplib—Motorola DSP COFF Librarian	C-6
C.1.5.1	dsplib Options	C-6
C.1.6	dsplnk—Motorola DSP COFF Linker	C-7
C.1.6.1	dsplnk Options	C-8
C.1.7	gdb56—GNU Source-level Debugger for DSP56000/DSP56001	C-11
C.1.7.1	gdb56 Options	C-11
C.1.8	run56—Motorola DSP56000/DSP56001 Simulator Based Execution Device	C-12
C.1.8.1	run56 Options	C-12
C.1.9	srec—Motorola DSP S-Record Conversion Utility	C-13
C.1.9.1	srec Options	C-13

Appendix D GNU Debugger (GDB)

D.1	Introduction	D-1
D.1.1	Commands Not Implemented	D-1
D.1.2	DSP56000 Family Differences / Specific Requirements	D-1
D.1.3	Summary of GDB	D-3
D.1.4	GNU General Public License	D-3
D.1.5	Preamble	D-3
D.1.6	Terms and Conditions	D-4
D.1.7	No Warranty	D-6
D.1.8	End of Terms and Conditions	D-7
D.2	Input and Output Conventions	D-8
D.3	Specifying GDB's Files	D-10
D.3.1	Specifying Files with Arguments	D-10
D.3.2	Specifying Files with Commands	D-10
D.4	Compiling Your Program for Debugging	D-12
D.5	Running Your Program Under GDB	D-12
D.5.1	Your Program's Arguments	D-13
D.5.2	Your Program's Environment	D-13

D.5.3	Your Program's Working Directory	D-14
D.5.4	Your Program's Input and Output	D-14
D.5.5	Debugging an Already-Running Process	D-15
D.5.6	Killing the Child Process	D-16
D.6	Stopping and Continuing	D-16
D.6.1	Signals	D-16
D.6.2	Breakpoints	D-18
D.6.2.1	Setting Breakpoints	D-18
D.6.2.2	Deleting Breakpoints	D-19
D.6.2.3	Disabling Breakpoints	D-20
D.6.2.4	Break Conditions	D-21
D.6.2.5	Commands Executed on Breaking	D-22
D.6.2.6	"Cannot Insert Breakpoints" Error	D-24
D.6.3	Continuing	D-25
D.6.4	Stepping	D-25
D.7	Examining the Stack	D-27
D.7.1	Stack Frames	D-27
D.7.2	Backtraces	D-28
D.7.3	Selecting a Frame	D-29
D.7.4	Information on a Frame	D-30
D.8	Examining Source Files	D-30
D.8.1	Printing Source Lines	D-30
D.8.2	Searching Source Files	D-32
D.8.3	Specifying Source Directories	D-32
D.9	Examining Data	D-33
D.9.1	Expressions	D-33
D.9.2	Program Variables	D-34
D.9.3	Artificial Arrays	D-35
D.9.4	Format options	D-35
D.9.5	Output formats	D-37
D.9.5.1	Examining Memory	D-38
D.9.6	Automatic Display	D-40
D.9.7	Value History	D-41
D.9.8	Convenience Variables	D-42
D.9.9	Registers	D-43
D.9.9.1	Examples	D-44
D.10	Examining the Symbol Table	D-44
D.11	Altering Execution	D-45
D.11.1	Assignment to Variables	D-46
D.11.2	Continuing at a Different Address	D-46

D.11.3	Giving the Program a Signal	D-47
D.11.4	Returning from a Function	D-48
D.12	Canned Sequences of Commands	D-48
D.13	User-Defined Commands	D-48
D.13.1	Command Files	D-49
D.13.2	Commands for Controlled Output	D-49
D.14	Options and Arguments for GDB	D-50
D.14.1	Mode Options	D-50
D.14.2	File-specifying Options	D-51
D.14.3	Other Arguments	D-51
D.15	Using GDB under GNU Emacs	D-52

Appendix E Additional Support

E.1	Motorola DSP Product Support	E-2
E.2	DSP56000CLASx Assembler/Simulator	E-2
E.2.1	Macro Cross Assembler Features:	E-3
E.2.2	Simulator Features:	E-3
E.3	DSP320to56001 Translator	E-4
E.3.1	DSP320to56001 Translator Features:	E-4
E.4	DSP56000ADSx Application Development Systems (ADS)	E-4
E.4.1	DSP56000ADS ADS Hardware Features:	E-4
E.4.2	DSP56000ADSx ADS Software Features:	E-5
E.4.2.1	Support Integrated Circuits:	E-5
E.5	Dr. BuB Electronic Bulletin Board	E-6
E.5.1	Motorola DSP News	E-14
E.5.2	Motorola Field Application Engineers	E-14
E.5.3	Design Hotline – 1-800-521-6274	E-14
E.5.4	DSP Applications Helpline– (512) 891-3230	E-14
E.5.5	DSP Marketing Information – (512) 891-2030	E-15
E.5.6	DSP Third-Party Support Information – (512) 891-3098	E-15
E.5.7	DSP University Support – (512) 891-3098	E-15
E.5.8	DSP Training Courses – (602) 994-6900	E-15
E.5.9	Reference Books and Manuals	E-16
E.5.10	General DSP:	E-16
E.5.11	Digital Audio and Filters:	E-17
E.5.12	C Programming Language:	E-18
E.5.13	Controls:	E-19
E.5.14	Graphics:	E-19
E.5.15	Image Processing:	E-21

E.5.16	Motorola DSP Manuals:.....	E-21
E.5.17	Numerical Methods:.....	E-21
E.5.18	Pattern Recognition:.....	E-22
E.5.19	Speech:.....	E-23
E.5.20	Telecommunications:.....	E-23

List of Tables

3-1	DSP56KCC Options	3-4
4-1	Identifier Length Limits	4-1
4-2	Predefined Macro List and Explanation	4-2
4-3	Integral Data Type Sizes and Ranges	4-2
4-4	Floating-point Data Type Sizes and Ranges	4-3
4-5	Floating-point Format Description	4-5
4-6	Comparison of DSP56KCC and IEEE 754-1985	4-6
4-7	Pointer Size and Range	4-8
4-8	DSP56000 Family Processor Registers	4-8
4-9	DSP56KCC Registers and Usage	4-9
4-10	Floating-point Arithmetic	4-16
4-11	48-bit Long Integer Arithmetic	4-16
5-1	Operand Constraints/Modifiers Associations	5-7
5-2	reg_save Names	5-8
5-3	Output and Input Operands for Section 5-3	5-9
A-1	ANSI Standard Header Files	A-1
A-2	In-Line Library Routines	A-3
A-3	Function Descriptions	A-4
A-4	Argument Range/Output Range	A-14
A-5	Mode Arguments	A-34
A-6	Conversion Specifyers	A-42
A-7	ANSsbi Flags	A-79
A-8	Conversion Characters	A-80
B-1	Arithmetic Operations	B-1
B-2	Logical Operations	B-2
B-3	Bit Manipulation Instructions	B-3
B-4	Loop Operation Instructions	B-3

B-5	Move Operation Instructions	B-4
B-6	Program Control Instructions	B-4
B-7	Assembly Control Directives	B-5
B-8	Control Symbol Directives	B-6
B-9	Storage Allocation Directives	B-6
B-10	Output Listing Directives	B-6
B-11	Object file Directives	B-7
B-12	Conditional Assembly Directives	B-7
B-13	Programming Directives	B-8
C-1	asm56000 Command Line Options	C-3
C-2	cofdmp Command Line Options	C-5
C-3	cldinfo Command Line Options	C-7
C-4	dsplnk Command Line Options	C-9
C-5	gdb56 Command Line Options	C-11
C-6	run56 Command Line Options	C-12
C-7	srec Command Line Options	C-13
D-1	GDB File Commands	D-11
D-2	Environment Variables	D-14
D-3	Working Directory Commands	D-14
D-4	Keywords	D-17
D-5	Breakpoints	D-18
D-6	Deleting Breakpoints Commands	D-20
D-7	Disabling Breakpoint Commands	D-20
D-8	Stepping Commands	D-25
D-9	Backtraces	D-28
D-10	Frame Commands	D-29
D-11	Frame Command	D-30
D-12	List Commands	D-31
D-13	Single Source Line Commands	D-31
D-14	Directory Commands	D-33
D-15	GDB Operators	D-34

D-16	Print Output Format Letters	D-37
D-17	Memory Output Format Letters	D-38
D-18	Print Output Format Letters	D-38
D-19	Display Commands.	D-41
D-20	Info Commands	D-42
D-21	Convenience Variables	D-43
D-22	Symbol-File Commands	D-45
D-23	Define Commands	D-48
D-24	Controlled Output Commands	D-50
D-25	Mode Options	D-51
D-26	File Specifying Options	D-51
D-27	Emac Commands	D-52
E-1	User Support	E-1
E-2	Available Software	E-6

List of Figures

1-1	Motorola Software Development System	1-2
4-1	Mantissa and Exponent Data Range of C Floating Point	4-4
4-2	Default Program Memory Configuration	4-10
4-3	Activation Record	4-11
4-4	Default Data Memory Configuration	4-12
4-5	Run-time Stack Growth	4-13
6-1	Environment Memory Configuration	6-4

List of Examples

2-1	Test Program	2-4
3-1	Lib56cy.clb	3-5
3-2	G56_EXEC_PREFIX	3-5
3-3	Test New Version of DSP56000/1 C	3-5
3-4	Test New Version of ANSI C Library	3-6
3-5	Test New Version of DSP56KCC.	3-6
3-6	Generate a Preprocessed File	3-6
3-7	Compile file.c / Generate Executable Output File	3-6
3-8	Preprocess Program Foo.c	3-7
3-9	Compile and Run Program Dsp.c	3-8
3-10	Program Dsp.c.	3-9
3-11	Preprocess C Source Program Foo.c.	3-9
3-12	CPP Output	3-10
3-13	-I Option Used to Include the File “myinclude.h”	3-11
3-14	Test Program “file.c”	3-12
3-15	Program Greeting.c Prints Message	3-14
3-16	Program Big.c	3-15
3-17	Program Big.c	3-15
3-18	Test Program Test.c	3-16
3-19	-v Option	3-18
3-20	DEBUG Flag.	3-19
3-21	-Wcomment Option	3-19
3-22	R4	3-22
3-23	Program Foo.c.	3-23
3-24	-mx-memory Option	3-24
3-25	Generate an optimized assembly language file (test.asm) of the C source program (test.c)..	3-26
3-26	Generate an optimized assembly language file test.asm.	3-26
3-27	-W	3-27

3-28	-Wimplicit.	3-28
3-29	-Wreturn-type	3-28
3-30	-Wunused	3-29
3-31	Pass a single option to the assembler.	3-31
3-32	Pass multiple options to the assembler.	3-31
3-33	-crt file	3-32
3-34	Pass a single option to the linker.	3-32
3-35	Pass multiple options to the linker.	3-32
3-36	-ILIBRARY	3-32
3-37	-r CTLFILE.	3-33
5-1	reg_save	5-8
5-2	instruction_template	5-9
5-3	operand_id	5-9
5-4	OES Modifier j	5-10
5-5	OES Modifier e.	5-10
5-6	OES Modifier h	5-10
5-7	OES Modifier k	5-11
5-8	OES Modifier g	5-11
5-9	OES Modifier g	5-12
5-10	OES Modifier i	5-12
5-11	OES Modifier f	5-13
5-12	OES Modifier p and q.	5-13
5-13	Input Expression / Output Expression	5-14
5-14	Read-Only Operand	5-14
5-15	Write-Only Operand.	5-14
5-16	Read-Write Operand.	5-14
5-17	Read-Write Operand.	5-15
5-18	Multiple Instruction — Single-Line	5-15
5-19	Multiple Instruction — Multiple-Line	5-15
5-20	Multiple Use of __asm()	5-15
5-21	Line Separation.	5-16
5-22	Instruction Template Label.	5-16

5-23	Program DSP56001 SCI Timer	5-17
5-24	Calling Assembly from C	5-18
5-25	Calling C from Assembly	5-19
5-26	Calling C from Assembly	5-19
5-27	Generate Data with Assembly Language	5-20
5-28	Access Data with C	5-20
5-29	Generate Data with C	5-21
5-30	4-tap FIR Filter	5-21
5-31	General Template for Out-of-line Assembly Code	5-30
5-32	Global Label in Assembly Language	5-33
5-33	Global Variable Declaration	5-33
5-34	Changing a Global Label	5-34
5-35	Run-time Stack Allocation	5-34
5-36	Run-time Stack Deallocation	5-35
5-37	Calling C Routines	5-35
5-38	ReadX / ReadY Routines	5-37
6-1	DSP56000/DSP56001 Operation Mode Change	6-3
6-2	C Bootstrap Code Location Change	6-3
6-3	Fast Stack	6-5
6-4	Fast Heap	6-6
6-5	User-defined Interrupt Vector Table	6-7
6-6	Interrupt Service Routine	6-8
6-7	Sample Host-Side Glue Code	6-15
A-1	abort	A-9
A-2	abs	A-10
A-3	acos	A-11
A-4	asin	A-12
A-5	attan	A-13
A-6	attan2	A-14
A-7	atexit	A-15
A-8	atof	A-16
A-9	atoi	A-17

A-10	atol	A-18
A-11	bsearch	A-19
A-12	calloc.	A-21
A-13	ceil	A-22
A-14	clearerr	A-22
A-15	cos.	A-23
A-16	cosh.	A-23
A-17	div.	A-24
A-18	exit	A-25
A-19	exp	A-26
A-20	fabs	A-27
A-21	fclose.	A-27
A-22	feof	A-28
A-23	fflush.	A-29
A-24	fgetc	A-30
A-25	fgetpos	A-31
A-26	fgets	A-32
A-27	floor	A-33
A-28	fmod	A-33
A-29	fopen.	A-35
A-30	fprintf	A-36
A-31	fputc	A-36
A-32	fputs	A-37
A-33	fread	A-38
A-34	free	A-39
A-35	freopen	A-40
A-36	frexp	A-41
A-37	fscanf	A-43
A-38	fseek	A-44
A-39	fsetpos.	A-45
A-40	ftell	A-46
A-41	fwrite.	A-47

A-42	getc	A-48
A-43	getchar	A-49
A-44	gets	A-49
A-45	isalnum	A-50
A-46	isalpha	A-51
A-47	iscntrl	A-52
A-48	isdigit	A-53
A-49	isgraph	A-54
A-50	islower	A-55
A-51	isprint	A-55
A-52	ispunct	A-56
A-53	isspace	A-57
A-54	isupper	A-57
A-55	isxdigit	A-58
A-56	labs	A-59
A-57	ldexp	A-59
A-58	ldiv	A-60
A-59	log	A-61
A-60	log10	A-62
A-61	longjmp	A-63
A-62	malloc	A-64
A-63	mblem	A-65
A-64	mbstowcs	A-67
A-65	mbtowc	A-69
A-66	memchr	A-70
A-67	memcmp	A-71
A-68	memcpy	A-72
A-69	memmove	A-73
A-70	memset	A-74
A-71	modf	A-75
A-72	perror	A-76
A-73	pow	A-77

A-74	printf	A-82
A-75	putc	A-83
A-76	putchar	A-84
A-77	puts	A-85
A-78	qsort	A-86
A-79	raise.	A-87
A-80	rand.	A-88
A-81	realloc	A-89
A-82	remove	A-90
A-83	rename	A-90
A-84	rewind.	A-91
A-85	scanf	A-91
A-86	setjmp	A-92
A-87	signal.	A-95
A-88	sin	A-95
A-89	sinh	A-96
A-90	sprintf	A-97
A-91	sqrt	A-97
A-92	srand	A-98
A-93	strcat	A-99
A-94	strchr	A-100
A-95	strcmp	A-101
A-96	strcpy	A-103
A-97	strcspn.	A-104
A-98	strerror	A-105
A-99	strlen	A-106
A-100	strncat	A-107
A-101	strncmp	A-108
A-102	strncpy	A-109
A-103	strbrk.	A-110
A-104	strchr	A-111
A-105	strspn.	A-112

A-106	strstr	A-113
A-107	strtod	A-114
A-108	strtok	A-115
A-109	strtol	A-117
A-110	strtoul	A-119
A-111	tan	A-120
A-112	tanh	A-121
A-113	tmpnam	A-122
A-114	tolower	A-123
A-115	toupper	A-123
A-116	ungetc	A-124
A-117	vfprintf	A-125
A-118	vprintf	A-125
A-119	vsprintf	A-126
A-120	wcstombs	A-127
A-121	wctomb	A-129

Chapter 1

Introduction

1.1 Overview

The DSP56KCC GNU based C cross-compiler is the latest high-level language development system for the Motorola DSP56000 family of digital signal processors (DSP). It includes:

- an integrated control program — g56k
- an ANSI compliant C language preprocessor — mcpp
- an ANSI optimizing C compiler — g56-cc1
- a DSP56000 common object file format (COFF) assembler — asm56000
- a COFF linker — dsplnk
- a COFF librarian — dsplib
- a source level debugger for C — gdb56
- a simulator based execution program — run56
- various object file manipulation tools — srec, cldlod, cofdmp

This integrated software system runs on a variety of machines and operating systems, including the IBM PC™ (80386 family and above — 386sx, 486, etc.), and Sun SPARC™ workstations.

The compiler implements the full C language as defined in *American National Standard for Information Systems - Programming Language - C, ANSI X3.159-1989*. It accepts one or more C language source files as input and generates a corresponding number of assembly language source files which are suitable as inputs to the DSP56000 COFF assembler. The compiler automatically implements numerous optimizations which simplifies implementing fast and efficient DSP algorithms.

The C language preprocessor is an implementation of the ANSI standard which includes support for arbitrary text file inclusion, macro definition and expansion, and conditional compilation. The preprocessor exists as a separate program and may be used as a general-purpose macro preprocessor.

The compiler control program, **g56k**, is the standard compiler interface used to control the sequence of operations required to compile a program. This control program allows the user to select a wide variety of control options which affect the four compilation phases —

preprocessing, compiling, assembling, and linking. The control program parses the command line options, and invokes the required sub-programs on the specified files.

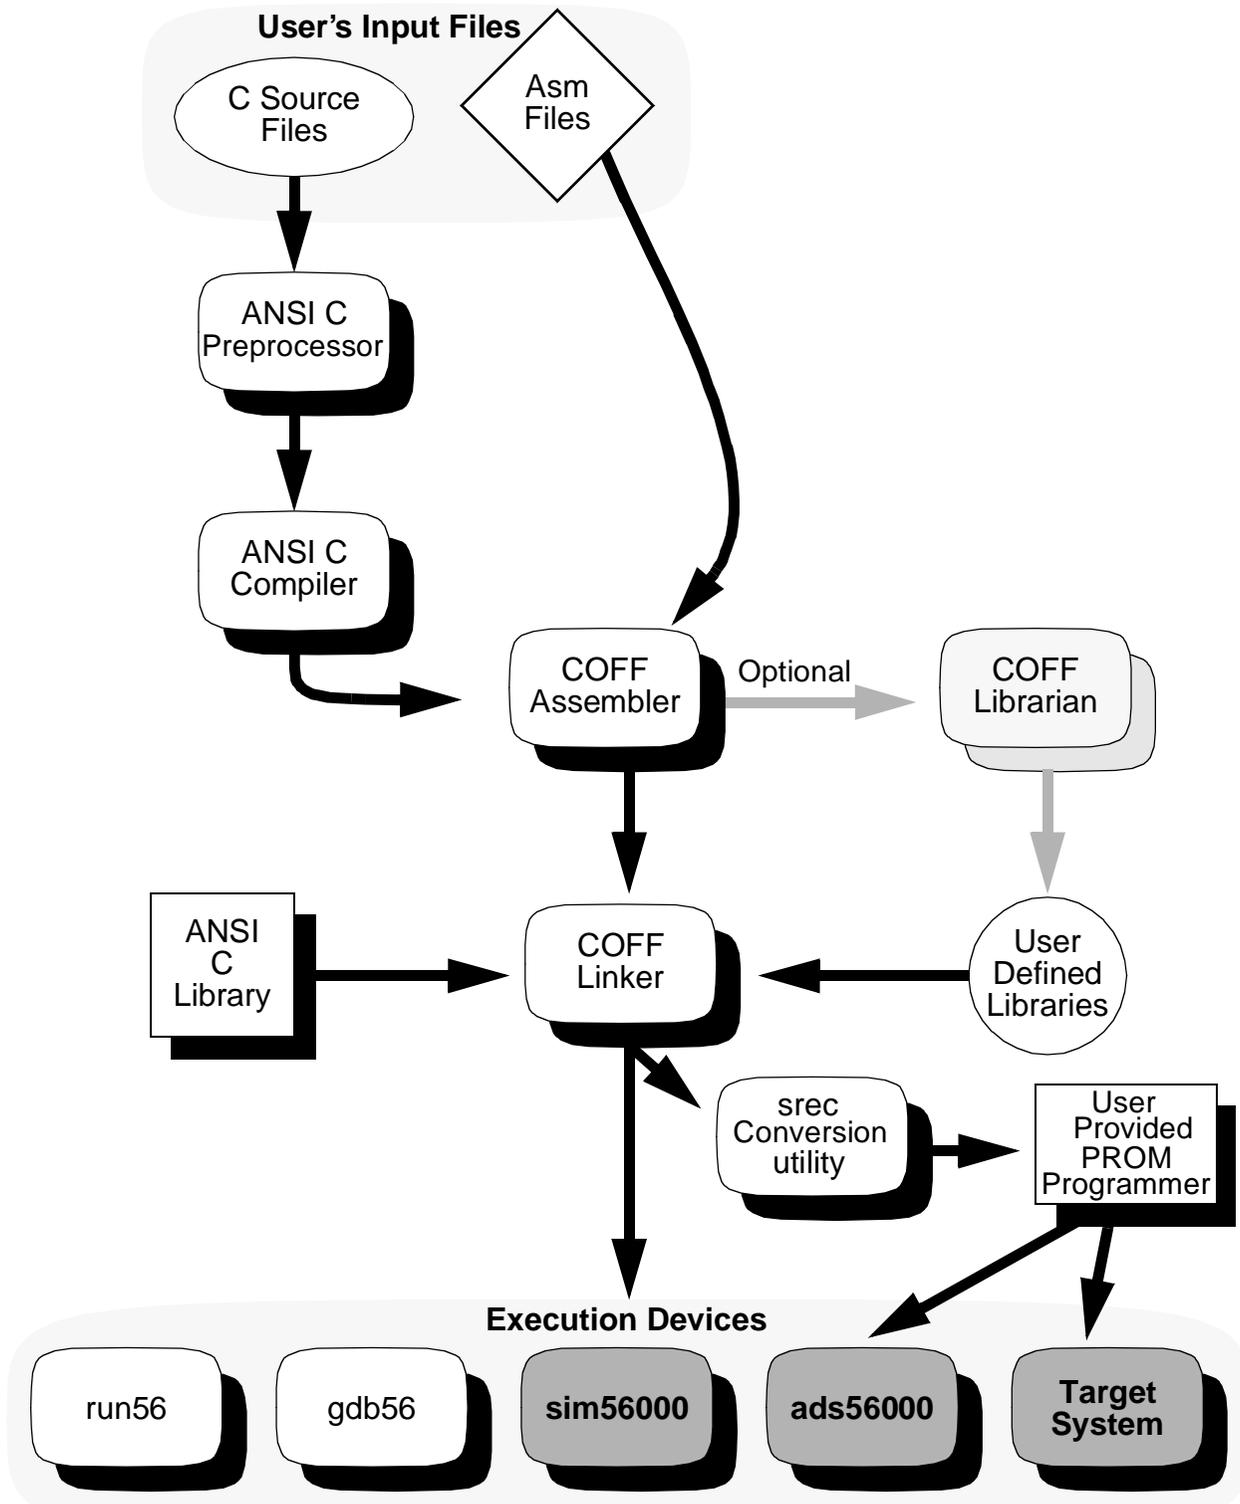


Figure 1-1. Motorola Software Development System

Object files are stored using the COFF format. COFF stands for *Common Object File Format*. Utilities such as `Cldinfo` and `Cldlod` may be used to gain visibility into object files.

1. Given a list of C source files from the user (see Figure 1-1) and options to apply to those files, the control program runs each file through the C preprocessor and the C compiler. The compiler creates individual assembly language source files for each C source file provided on the command line.
2. The control program then sends the compiler output from these files to the assembler, in addition to any assembly language files specified by the user on the G56k command line.
3. The assembler output is routed to the linker for final processing. The linker tries to resolve all unresolved link-time symbols with the standard (and any explicitly requested) C libraries. The COFF linker object file output may then be directed to any of several execution devices. Notice that the assembler can also be used to create library files which can be included in a user defined library.
4. The execution devices shown in Figure 1-1 are:
 - a. RUN56 which allows the DSP56000/1 code (in COFF format) to be executed on the host computer's CPU
 - b. GDB56 which is a debugging tool for trouble-shooting the compiled application
 - c. SIM56000 which is a complete DSP56000/1 simulator that can be used to execute the compiled application (in either COFF format or .lod file format) and allow examination of registers and memory
 - d. ADS56000 is the development system hardware that can then be used to load and execute the compiled application (in either COFF format or .lod file format) on the ADS development system
 - e. The target system shown is the user's custom DSP system.

Note: The three execution devices in the shaded boxes are not part of the C compiler software. The COFF linker output can be used by these execution devices directly. The conversion utility `Srec` (see Figure 1-1) can be used to convert the executable file from the COFF Linker to a suitable format for PROM burning. These PROMs can then be used on the ADS development system or the user's target system. The PROM programmer, ADS development system and user's target system are not part of the DSP56KCC compiler.

The DSP56000 family represents a departure from more conventional architectures on which other implementations of the C language are based. Also, the nature of DSP applications dictates that a greater measure of control be provided to the programmer in specifying the constraints of the run-time environment. For these reasons, the components

of the development system include options for handling separate memory spaces, stack initialization, chip operating modes and other issues.

The purpose of this manual is to:

1. Provide detailed installation procedures for both UNIX based systems and MS-DOS based systems. This manual explains how to install and operate the DSP56KCC ANSI C compiler development system in a hosted environment.
2. Provide an overview of the compiler operation. It also includes information on combining C modules with assembly language programs and advanced topics pertaining to compiler run-time characteristics.
3. Provide reference information on compiler options, ANSI library routines, and utilities.

This manual assumes a familiarity with the C programming language, with the host machine and operating system environment. It also assumes that the programmer understands how to create and edit C language source files on this host system. For more information on the C language and other DSP56000/1 development tools, see the references listed in Appendix D.

1.2 Error Codes

The error messages generated by the compiler are intended to be complete without additional explanation. Since the compiler produces a detailed description of the problem rather than an error code, these error messages have not been reproduced in this manual.

1.3 Notation

The following notation will be used in this text.

1. A prompt is indicated in this manual by:
`C:\>`
2. An example of an MS-DOS directory name is:
`\USR\DIRECTORY`
3. The contents of an MS-DOS directory are examined by entering:
`C:\> DIR`
4. The contents of an MS-DOS file are displayed by entering:
`C:\> TYPE FILE`
5. The program "HELLO.EXE" would be executed by the command line:
`C:\> HELLO`

1.4 Manual Organization

Installation details are provided in Chapter 2, the compiler operation is described in Chapters 3-6 and reference information is in Chapter 3 and Appendices A-E. The contents of each chapter and each appendix are described below.

Chapter 1, "Introduction," describes the overall organization of the DSP56KCC compiler system. It also details the structure of this document, and conventions followed herein.

Chapter 2, "Installation Guide," describes the installation and organization of DSP56KCC. It details how to set up an operating environment on the host system by defining global environment variables and includes a step-by-step installation procedure.

Chapter 3, "Control Program Options," discusses the four passes of the compilation process with particular attention to the functions of the compiler control program G56k. This chapter includes a list of the compiler invocation options along with example command lines for different memory and program configurations.

Chapter 4, "About g56k," provides information on the compiler run-time environment, including explanations of compiler register and memory usage, stack frame architecture, stack overflow checking, and defining/referencing of absolute memory locations. Additionally, this chapter covers implementation issues such as data type sizes.

Chapter 5, "Mixing C and Assembly Language," discusses the methods for using assembly language in conjunction with C language programs. It covers the inclusion of assembly language within a C source file and also describes linking assembly language modules with C modules and linking C modules with assembly language modules.

Chapter 6, "Software-Hardware Integration," describes how to modify a program's run-time environment, how to write software to handle interrupts, and the `set jmp/long jmp` ANSI library routines.

Appendix A, "Programming Support," provides a complete description and brief example for every ANSI library subroutine distributed with the C compiler.

Appendix B, "DSP56000/DSP56001 Instruction Set and Assembler Directive Summary," provides a brief overview of the assembly language instructions and assembler directives.

Appendix C, "Utilities," provides DSP56000/1 Assembler manual pages for each of the supporting programming utilities provided with the compiler.

Appendix D, "GNU Debugger (GDB)," provides a description of software, hardware, support telephone numbers, and suggested reading for the C language and various areas of DSP.

Appendix E, "Additional Support," , provides information on how you can obtain further support from Motorola and third party companies.

Chapter 2

Installation Guide

2.1 Introduction

This chapter describes installation on MS-DOS and Sun computers. Two installation procedures are detailed for the Sun. The first procedure uses the default location for the files. The second procedure allows the user to select the directory where the compiler's files will be located. Only one procedure is needed for MS-DOS machines.

The various parts of the compiler reside in a directory tree named DSP. The default location for the DSP directory tree is /usr/local on UNIX systems. If this default location is acceptable, then perform the standard installation; if it is not acceptable, then perform the alternate installation. The alternate installation procedure allows the user to install the DSP directory tree anywhere.

2.2 Installation On An MS-DOS Machine (80386 or 80486)

1. Insert the supplied floppy labeled "Disk 1" into floppy drive "A:".
2. Change to floppy drive "A", with the command "A:".
3. Run the install program, Install.exe. This installation program will ask questions about the computer being used and about where the compiler's directory tree, DSP, is to reside.
4. Add all new lines of code specified by Install into the `autoexec.bat` file. The only difference between the standard and alternate installation procedure on the PC is whether or not the default output drive or default location is selected. If the defaults are not selected, an environment variable named `DSPLOC` must be set in the `autoexec.bat` file. The Install program will provide directions. The lines to be added to the `autoexec.bat` file follow:

- f. `TERMCAP` must be set to the current location of <compiler's dsp directory tree>\dsp\etc\termcap:

```
SET TERMCAP=<DSP location>\DSP\ETC\TERMCAP
```

for example, if the directory <compiler's dsp directory tree> is `d:\usr\mydir`, then the line to be added is

```
SET TERMCAP=D:\USR\MYDIR\DSP\ETC\TERMCAP
```

- g. `TERM` must be set to indicate the type of display to be used (probably `nansi-mono`):

```
SET TERM=NANSI-MONO
```

- h. DSPLOC need only be set if the default output drive or the default location is not chosen. If DSPLOC is set, it must be set to the location of the dsp directory tree. For example, if the user installed the compiler's directory tree DSP in the directory D:\usr\mydir, then DSPLOC would need to be set as D:\usr\mydir:

```
SET DSPLOC=<compiler's dsp directory tree>
```

if the directory <compiler's dsp directory tree> is d:\usr\mydir, then

```
SET DSPLOC=D:\USR\MYDIR
```

5. Make sure that <compiler's dsp directory tree>\dsp\bin is included in the path instruction. This is needed by `command.com` if it is to find G56k. If the default drive and path were chosen, then the path C:\dsp\bin would need to be added as follows:

```
PATH ...;C:\DSP\BIN;...
```

If, for example the compiler's dsp directory tree was installed in C:\usr\mydir, then C:\usr\mydir\dsp\bin would need to be added as follows:

```
PATH ...;C:\USR\MYDIR\DSP\BIN;...
```

6. Make sure that other DOS memory managers do not interfere with the DOS extended memory manager for G56k. The compiler uses its own DOS extended memory manager called `dos4gw.exe`, and this memory manager may not work if a different memory manager is already installed. Although this DOS extender is DPMI 0.9 compliant, It is recommended that initially all other DOS extended memory managers be removed, in order to test the installation. The DOS extended memory manager `dos4gw.exe` is called during the compiler's execution, and requires at least 2M bytes of RAM. This memory manager uses hard drives for the swap space for the memory management. By default, the swap space location is the C drive and the size of the swap space is 16 Mbytes.
7. The DOS environment DOS4GVM controls the configuration of the DOS extended memory management, and the environment DOS4GVM has the following format.

```
[option[#value]] [option[#value]]
```

The possible parameters for the option are:

MINMEM	The minimum amount of RAM managed by the memory manager. Default value is 512KB.
MAXMEM	The maximum amount of RAM managed by the memory manager. Default value is 4MB.

MINMEM	The minimum amount of RAM managed by the memory manager. Default value is 512KB.
SWAPNAME	The swap file name for the memory manager uses for the swap space. Default is DOS4GVM.SWP on the current drive.
DELETESWAP	Specifies that the swap file should be deleted after memory management.
VIRTUALSIZE	The size of the virtual memory space. Default is 16MB. The value should be entered as numeric value of KB.

As an example, the following line in the autoexec.bat file will enable an 8MB swap file with automatic deletion of the swap file:

```
SET DOS4GVM=DELETESWAP VIRTUALSIZE#8192
```

The following line will use F drive for the swap space instead of the current drive.

```
SET DOS4GVM=DELETESWAP SWAPNAME#F:\BIG.SWP
```

8. There is a READ.ME file included in the package and it should be read for any recent changes in the installation on compiler itself.

2.3 Standard Installation On A SUN

1. Insert the supplied first floppy diskette into the tape drive.
2. Login as `root`.
3. Enter the command: `cd /usr/local`.
4. Enter the command: `bar xZvf /dev/rfd0`. If the floppy drive must be accessed via a different device file than `rfd0`, then use the appropriate device for your system.
5. Logout.
6. Make sure that all users add `/usr/local/dsp/bin` to their path. This enables the shell to find the control program G56k and other programs in the DSP56KCC distribution package.

2.4 Alternate Installation On A SUN

1. Insert the supplied first floppy diskette into the floppy drive.
2. Login as `root`, or as yourself, if access permissions allow.
3. Inside the shell, use the command `cd` to go to the directory where the compiler's DSP directory tree is to reside. For this example, assume that the compiler is to be installed in `/usr/mydir` (referred to by <compiler's DSP directory tree> here).
4. Make sure that you have write permission in the directory.
5. Enter the command: `bar xZvf /dev/rfd0`. If the floppy drive must be accessed via a different device file than `rfd0`, then use the appropriate device for your system.
6. Make sure that every user adds <compiler's DSP directory tree>/dsp/bin to their path. In this example, the path `/usr/mydir/dsp/bin` would be added to everyone's path.
7. Make sure that every user sets the environment variable `DSPLOC` to the path leading to the **DSP** directory tree which is the directory <compiler's dsp directory tree>. In this example, `DSPLOC` would be set to `/usr/mydir`. Note that `DSPLOC` would not be set to `/usr/mydir/dsp`.

2.5 Test Program

The following test program is intended to be a very simple check to see if the installation has been completed correctly. The program should be put in a file named "hello.c". The control program, G56k, compiles the program in the file "hello.c" and generates the output file "a.cld". Do not enter the `C:>` as it is simply a prompt indicating that this line should be entered from the keyboard. The command `run56` executes the program in the file "a.cld" and the result is to print "hello world." on the computer screen.

Example 2-1. Test Program

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

C:\> g56k hello.c

C:\> run56 a.cld

hello
```

Chapter 3

Control Program Options

3.1 Overview

Program G56k is the *control program* for Motorola's optimizing C compiler for the DSP56000/DSP56001 family of digital signal processors. The program G56k automates control of the four C compiler phases – preprocessing, compiling, assembling, and linking. The program G56k utilizes a command syntax similar to those adopted by UNIX utilities. The G56k syntax is:

```
g56k [options] files
```

where:

1. [options] is one or more of the options found in this chapter. One difference between G56k and UNIX-style utilities is that the combination of multiple single character options is not allowed. For example, “-O -g” instructs the compiler to generate an optimized executable with source level debug information, whereas “-Og”, which is acceptable to UNIX-style compilers is not acceptable to G56k.
2. “file ...” is the file(s) to be processed. Program G56k accepts input filenames suffixed with “.c” (ANSI C source files), “.i” (preprocessed C source files), “.asm” (DSP56000/DSP56001 assembly code), and “.cln” (COFF link files). The control program processes each file according to this suffix. The G56k output is controlled by the specific set of command line options provided. For instance, if no command line arguments are provided, the compiler will generate a COFF load file “a.cld”. If the **-c** option is invoked, the compiler will generate a COFF link file suffixed with “.cln”. A complete description of the command line options, with examples, is provided in Section 3.2.

Note: It is strongly recommended that g56k always be used to invoke the C compiler utilities rather than individually executing them.

A standard directory search list is consulted by G56k for:

1. Each of the four executables,
 - i. MCPP– the C compiler preprocessor,
 - j. G56-CC1– the C compiler/optimizer,
 - k. ALO56 –the assembly language optimizer,
 - l. ASM56000 – the DSP56000/DSP56001 assembler,
 - m. DSPLNK – the DSP56000/DSP56001 linker.
2. Start-up file, `crt056*.cln`. Where * is a single character (x, y or l; i.e., x data memory, y data memory or long data memory) to denote the memory model.
3. ANSI C library, `lib56c*.clb`. Where * is a single character (x, y or l) to denote the memory model.

This standard directory search list for UNIX systems is:

1. `/usr/local/dsp/bin/`
2. `/usr/local/dsp/lib/`
3. `/lib/`
4. `/usr/lib/`
5. `./`

The standard MS-DOS directory search list for the path set up in Section 2.2 is:

1. `C:\dsp\bin`
2. `C:\dsp\lib`
3. `C:`
4. `C:\dos`
5. other directories in the path name

Note that if the environment variable DSPLOC is set, the value of DSPLOC will be substituted for 1 and 2 above.

Table 3-1 lists all the user selectable options used by **g56k**. They are grouped to show what program uses each option. All of these options are described in detail later in this chapter; however, these lists provide an overview of what options are available. Notice that there is a `-v` option listed under both **g56k** Command Options and Preprocessor Phase Options. This is actually the same option but it is used by these two programs in different ways (see Section 3.2 and Section 3.2.1).

Under compile phase options, there is a group of `-f` options; these are the machine independent optimization options whereas the `-m` options below are the optimization

options specific to the DSP56000/1. Although these various methods of optimization are all effective, they may have side effects which are undesirable in specific cases, e.g. an optimization option may increase code speed at the cost of increased memory usage. It is often preferable to trade memory space for speed, but in cases where the extra memory space is not available, this particular optimization could not be used.

The various compiler phases will report errors; however, the user has the option to turn off all warnings using **-w** and can enable additional warnings individually or as a group using **-wall**. The warnings which are not enabled by **-wall** are those listed below **-wall** in Table 3-1, “DSP56KCC Options,” on page 3-4.

Table 3-1. DSP56KCC Options

g56k Command Line Options	Preprocesso r Phase Options	Assemble Phase Options	Link Phase Options	Compile Phase Options
-Bdirectory -bPREFIX -o FILE -v	-C -DMACRO -DMACRO=D EFN -E -IDIR -I- -i FILE -M -MM -nostdinc -pedantic -P -v -UMACRO -Wcomment -Wtrigraphs	-asm string -c	-crt file -j string -LIBRARY -r MAPFILE	-alo -fno-opt -fno-peephole -fno-strength-reduce -fno-defer-pop -fforce-addr -finline-functions -fcaller-saves -fkeep-inline-functio ns -fwritable-strings -fcond-mismatch -fvolatile -ffixed-REG -g -O -m56002 -mno-dsp -mno-do-loop-gener ation -mno-linv-plus-biv-pr omotion -mp-mem-switchtabl e -mstack-check -mx-memory -my-memory -ml-memory -pedantic -Q -S -w -W -Wimplicit -Wreturn-type -Wunused -Wswitch -Wall -Wshadow -Wid-clash-LEN -Wpointer-arith -Wcast-qual -Wwrite-strings

3.2 G56k Command Line Options

The default options are:

1. Use strict ANSI syntax.
2. Perform all machine dependent and independent optimizations.
3. Use trigraphs.
4. Locate data in the Y data memory space.

-Bdirectory

Add *directory* to the standard search list and have it searched first. This can also be accomplished by defining the environment variable `G56_EXEC_PREFIX`. Note that only one additional directory can be specified and that the **-B** option will override the environment variable.

Example 3-1. Lib56cy.clb

To test a new version of the ANSI C library, `lib56cy.clb`, which is installed as `\dsp\new\lib56cy.clb` use:

```
C:\> g56k -B\dsp\new\ file.c -o file.cld
```

Example 3-2. G56_EXEC_PREFIX

Using the `G56_EXEC_PREFIX` environment variable to have the same effect as Example 3-1, include in the `.cshrc` or `.login` files:

```
set G56_EXEC_PREFIX=C:\DSP\NEW\
and then execute:
```

```
C:\> g56k file.c -o file.cld
```

Example 3-3. Test New Version of DSP56000/1 C

To test a new version of the DSP56000/1 C preprocessor before permanent installation, install a new MCPP program as `c:\tmp\new\mcpp` and then execute:

```
C:\> g56k -Bc:\tmp\new\ testfile.c
```

-bPREFIX

Direct `G56k` to search for compilation phases, start-up files and libraries whose names are prefixed with the word *PREFIX*. Note that only one additional prefix can be specified. This is very similar to the **-B** option.

Example 3-4. Test New Version of ANSI C Library

Test a new version of the ANSI C library, lib56cy.clb, installed as c:\dsp\lib\new-lib56cy.clb.

```
C:\> g56k -bnew- file.c -o file.cld
```

Example 3-5. Test New Version of DSP56KCC

Test a new version of the DSP56KCC preprocessor before permanent installation.

```
install the new MCPP program as c:\dsp\bin\new-mcpp and  
C:\> g56k -bnew- testfile.c
```

-o FILE

Select *FILE* as the output file. This applies to all phases of the compiler. When the **-o** flag is not in use, the following file names are used by the compiler as the default output file names depending upon the compile options as follows:

-E (preprocess only)	stdout
-S (compile only)	foo.asm
-c (no linkage)	foo.cln
complete process	a.cld

where stdout is “standard output” and prints to the console.

Example 3-6. Generate a Preprocessed File

Only generate a preprocessed file (do not invoke the compiler, assembler or linker) and put the results in file.i.

```
C:\> g56k -E file.c -o file.i
```

Example 3-7. Compile file.c / Generate Executable Output File

Compile file.c and generate the executable output file, fft.cld. If an output name is not given, the default file name is a.cld.

```
C:\> g56k file.c -o fft.cld
```

-v

Verbose mode. The compiler control program announces to Stderr all commands that it attempts to execute for each phase of the compilation process. This command is also used by the preprocessor to print the software version information. If the **-E** option is selected, **-v** will only enable the verbose mode, otherwise it will enable the verbose mode and print the version information.

3.2.1 Preprocessor Phase Options

The options listed below control the C preprocessor, which is run on each C source file before actual compilation. Some options described only make sense when used in combination with the **-E** option (preprocess only), as the requested preprocessor output may be unsuitable for actual compilation. The default option is to use ANSI C syntax. For example, if the **-IDIR** option is not specified then ANSI specifies that the current working directory will be searched first for user defined “include” files.

-C

Tell the preprocessor *not* to discard comments. This option is only valid when used in conjunction with the **-E** option.

Example 3-8. Preprocess Program Foo.c

This example preprocesses a simple program, foo.c, without discarding comments.

```
C:\> type foo.c
/*
 * This COMMENT won't be deleted.
 */
main()
{
    printf("Hello, DSP56000\n");
}

C:\> g56k -E -C foo.c
# 1 "foo.c"
/*
 * This COMMENT won't be deleted.
 */
main()
{
    printf("Hello, DSP56000\n");
}
```

-DMACRO

Define the preprocessor macro **MACRO** with a constant value of "1". This is equivalent to making **MACRO** a constant set to one.

Example 3-9. Compile and Run Program Dsp.c

Compile and run a simple program, **Dsp.c**, and enable or disable a printed message depending on the macro definition given at the command line.

```
C:\> type dsp.c
#include <stdio.h>
main()
{
#ifdef DSP56000
    printf("message: DSP56000.\n");
#else
    printf("message: DSP56001.\n");
#endif
}
```

```
C:\> g56k -DDSP56000 dsp.c
```

```
C:\> dir
a.cld dsp.c
```

```
C:\> run56 a.cld
message: DSP56000.
```

```
C:\> g56k dsp.c
C:\> ls
a.cld dsp.c
C:\> run56 a.cld
message: DSP56001.
```

-DMACRO=DEFN

Define preprocessor macro **MACRO** as **DEFN**.

Example 3-10. Program Dsp.c

The program Dsp.c uses the macro FROM_COMMAND_LINE which prints a message to the standard output using a message code given on the command line.

```
C:\> type dsp.c
#include <stdio.h>
main()
{
    printf("message code: %d.\n", FROM_COMMAND_LINE);
}

C:\> g56k -DFROM_COMMAND_LINE=56000 dsp.c
C:\> dir
a.cld dsp.c

C:\> run56 a.cld
message code: 56000.
```

-E

The input source file will only be preprocessed through CPP and the output results will be sent to the standard output. See the **-O** option to save the output into a named file.

Example 3-11. Preprocess C Source Program Foo.c

This example shows how to preprocess the C source program Foo.c and send the results to the standard output.

```
C:\> type foo.c
#define DELAY 1000
main()
{
    int cnt = DELAY;
    while(cnt--);
}

C:\> g56k -E foo.c
# 1 "foo.c"

main()
{
    int cnt = 1000 ;
    while(cnt--);
}
```

Example 3-12. CPP Output

The CPP output can be saved into file "foo.i" by using the **-o** option.

```
C:\> type foo.c
#define DELAY 1000
main()
{
    int cnt = DELAY;
    while(cnt--);
}

C:\> g56k -E foo.c -o foo.i
C:\> dir
foo.c  foo.i

C:\> type foo.i
# 1 "foo.c"

main()
{
    int cnt = 1000 ;
    while(cnt--);
}
```

-IDIR

The control line of the C source program of the form

```
#include <filename>
```

will cause the replacement of that line by the entire contents of the file *filename*. This is normally referred to as "file inclusion". The named file is searched for in a sequence of implementation-dependent directories. The standard include directory for this compiler is /usr/local/dsp/include. Similarly, a control line of the form

```
#include "filename"
```

searches first in the current working directory, and if that fails, then searches as if the control line were `#include <filename>`.

The option **-IDIR** directs the C compiler to include the directory *DIR* in addition to the standard include directory. For the file inclusion `<filename>`, the compiler searches first in the *DIR* directory and if that fails, then searches /usr/local/dsp/include. For the file inclusion "filename", the compiler searches first in the *DIR* directory and if that fails, then searches the current working directory, and if that fails also, then searches /usr/local/dsp/include.

A delay program Foo.c uses delay constant **DELAY** which is defined in the include file, myinclude.h. The program uses the control line #include "myinclude.h" to include the definition of the constant **DELAY**. Without any option the include file should be located in the currently working directory since it is not in the standard include directory. Assuming that the include file "myinclude.h" is desired to be in the directory ./inc, the following sequence of the commands explains how the **-I** option is used to include the file myinclude.h in the ./inc directory with the control line #include "myinclude.h" in the Foo.c program.

Example 3-13. -I Option Used to Include the File "myinclude.h"

```
C:\> dir
foo.c  inc/

C:\> dir inc
myinclude.h

C:\> type foo.c
#include "myinc.h" /* this is the control line to
include it */
main()
{
    int cnt;
    cnt = DELAY;
    while(cnt--); }

C:\> type inc\myinc.h
#define DELAY 100

C:\> g56k -I.\inc foo.c
C:\> dir
a.cld  foo.c  inc/
```

-I-

This option is always used in conjunction with the **-IDIR** option and limits the file search to look for file inclusions #include "filename", whereas **-IDIR** alone directs C compiler to search the directory DIR for both file inclusion <filename> and "filename". Any directories specified with **-I** options before the **-I-** option are searched only for the case of #include "filename"; they are not searched for #include <filename>.

If additional directories are specified with **-I** options after the **-I-** option, these directories are searched for both #include "filename" and #include <filename> directives.

As an example, the sequence of the options

-IDIRA -I- -IDIRB

directs C compiler to use both the directories *DIRA* and *DIRB* for the file inclusion “filename” and *DIRB* only for file inclusion <filename>.

Note: The **-I-** option inhibits the use of the current directory as the first search directory for `#include "filename"`. There is no way to override this effect of **-I-**. However, the directory which is current when the compiler is invoked can be searched by using **-I**. This is different from the preprocessor’s default search list, but it is often satisfactory. **-I-** does not inhibit the use of the standard system directories for header files. Thus, **-I-** and **-nostdinc** are independent.

A test program `File.c` is used to test a file operation **fopen()** which is, in this example, desired to be developed for a DSP56000/1 system. The file include <stdio.h> is used as if it is in the standard include directory. The file is desired to be developed or debugged, and it is located in the user working directory `./mysys`. This example shows how to use **-IDIR** and **-I-** combination to test file inclusion <filename>. Notice that the **-I./inc -I- -I./mysys** option specifies the `Inc` directory only for the file inclusion “`cnt.h`” and `mysys` directory for the file inclusion <stdio.h>.

Example 3-14. Test Program “file.c”

```
C:\> dir
file.c  inc/                               mysys/

C:\> dir inc
cnt.h

C:\> dir mysys
stdio.h

C:\> type file.c
#include <stdio.h>
#include "cnt.h"
main()
{
    int delay = COUNT;
    FILE *fp;
    fp = fopen("myfile", "w");
    while(--delay);
}

C:\> type inc\cnt.h
#define COUNT 25
```

```

C:\> type msys\stdio.h
typedef struct FILE { /* FILE data structure to develop
*/
    char    name[10];
    char    buffer[1024];
} FILE;
FILE *fopen(char *, char *); /* new function to develop
*/

C:\> g56k -I.\inc -I- -I.\mysys -E file.c
# 1      "file.c"
# 1      ".\mysys\stdio.h"1
typedef      struct  FILE {
    char    name[10];
    char    buffer[1024];
} FILE;
FILE*fopen(char*,char*);
# 1      "file.c"2
# 1      ".\inc\cnt.h"1
# 2      "file.c"2

main()
{
    int delay = 25;
    FILE      *fp;
    fp = fopen ("myfile", "w");
    while (--delay);
}

```

Notice that the file inclusion “cnt.h” is from the directory `./inc` as shown in the line # 1 “.\inc\cnt.h” 1 and the file inclusion `<stdio.h>` is from the directory `.\myinc` as shown in the line # 1 “.\myinc\stdio.h” 1.

-i FILE

Process FILE as an input, discarding the resulting output, before processing the regular input file. Because the output generated from FILE is discarded, the only effect of **-i FILE** is to make the macros defined in FILE available for use in the main input.

Example 3-15. Program Greeting.c Prints Message

The program Greeting.c prints a simple message using the macro MESSAGE. The file macros.c contains the macro definition, i.e. the actual message. The only role of the file macros.c is to provide the macro definitions and will not affect any other code or data segments.

```
C:\> dir
  macros.c  greeting.c

C:\> type macros.c
#define MESSAGE "Hello, world."

C:\> type greeting.c
#include <stdio.h>
main()
{
    printf("Greeting: %s\n", MESSAGE);
}

C:\> g56k -i macros.c greeting.c
C:\> run56 a.cld
Greeting: Hello, world.
```

-M

Cause the preprocessor to output the makefile rules to describe the dependencies of each source file. For each source file, the preprocessor outputs one make-rule whose target is the object file name for that source file and whose dependencies are all the files needed to generate the object target file. This rule may be a single line or may be continued with ‘\’-newline if it is long. **-M** implies **-E** with makefile rules.

Example 3-16. Program Big.c

The program `Big.c`, which prints the larger of two integers, uses the macro `greater(x,y)` which is defined in the file `greater.h`. A command line output using the `-M` option can be used for makefile utilities. For more information on how to use this dependency check the make utility information in any UNIX utility manual.

```
C:\> dir
big.c      greater.h

C:\> type big.c
#include <stdio.h>
#include "greater.h"
main()
{
    printf("big:%d\n", greater(10,20));
}

C:\> type greater.h
#define greater(x,y) ((x)>(y)?(x):(y))
C:\> g56k -M big.c
big.o : big.c \dsp\include\stdio.h
\dsp\include\ioprim.h \dsp\include\stdarg.h greater.h
```

-MM

Like `-M` but the output mentions only the header files described in the `#include "FILE"` directive. System header files included with `#include <FILE>` are omitted. `-MM` implies `-E` with makefile rules.

Example 3-17. Program Big.c

The program `Big.c`, which prints the larger of two integers, uses the macro `greater(x,y)` defined in the file `greater.h`. The `-MM` option is used to generate a makefile rule. Notice that the rule that generates an output file appended by `".o"` can be modified to generate `".cld"` which is required for the Motorola Cross C Compiler.

```
C:\> dir
big.c      greater.h

C:\> type big.c
#include <stdio.h>
#include "greater.h"
main()
{
    printf("big:%d\n", greater(10,20));
}
```

```
C:\> type greater.h
#define    greater(x,y)    ((x)>(y)?(x):(y))

C:\> g56k -MM big.c
big.o : big.c greater.h

C:\> dir
big.c      greater.h  makefile   text

C:\> type makefile
a.cld : big.o
        g56k big.o
big.o : big.c greater.h
        g56k -c -o big.o big.c

C:\> make
g56k -c -o big.o big.c
g56k big.o

C:\> run56 a.cld
big:20
```

-nostdinc

Do not search the standard system directories for file inclusions. Only the directories specified with **-I** options (and the current directory, if appropriate) are searched.

Using both **-nostdinc** and **-I** options, all directories from the search path except those specified can be eliminated.

Example 3-18. Test Program Test.c

A test program, Test.c, is used to test a new version of the function **printf()** which is declared in a new header file `inc/stdio.h`. The directive `#include <stdio.h>` causes the program to use `stdio.h`; however, it would normally find it in the standard search directory, `/usr/local/dsp/include/`. Using the **-nostdinc** option prevents the standard search directory from being searched and allows the **-I** option to point to the correct directory.

```
C:\> dir
inc/      test.c

C:\> dir inc
stdio.h

C:\> type test.c
```

```
#include <stdio.h>
main()
{
    printf("Hello, there.\n");
}

C:\> type inc\stdio.h
void printf(char *);

C:\> g56k -nostdinc -I.\inc -E test.c
# 1 "test.c"
# 1 ".\inc\stdio.h" 1
void printf(char *);
# 1 "test.c" 2

main()
{
    printf("Hello, there.\n");
}
```

-pedantic

The **-pedantic** option is used by both the preprocessor and the compiler (see **-pedantic** in the section 3.2.2, “Compile Phase Options.” for an explanation of this option).

-P

Only run the source file through the C preprocessor. This option sends the output to a file with a .i suffix. This output file does not include # line information in the output file.

-v

Output preprocessor version information. The primary purpose of this command is to display the software version. This information is needed when calling the Motorola DSP helpline for assistance (see Appendix D — Additional Support for the helpline phone number). Although information pertaining to the internal flags and switch settings is included in this information, it is not intended for use by the programmer and may be misleading. This command is also used by the command program to initiate the verbose mode of operation.

Example 3-19. -v Option

The **-v** option is selected using the control program G56k. The version numbers for G56k, MCPP and G56-cc1 are printed. This information is showing the commands that the control program invokes along with the selected options. In this case it is showing the default options plus the **-v** option. However, the user should not invoke these programs independently but should always use the control program to invoke them.

```
C:\> dir
foo.c
C:\> g56k -v foo.c
g56k version Motorola Version: g1.11 -- GNU 1.37.1
c:\dsp\bin\mcpp -v -undef -D__Y_MEMORY -trigraphs
-$ -D__STRICT_ANSI__ -D__DSP56K__ -D__OPTIMIZE__
foo.c cca00527.cpp
GNU CPP version 1.37.1
c:\dsp\bin\g56-cc1 cca00527.cpp -ansi
-fstrength-reduce -quiet -dumpbase foo.c -O -version
-o cca00527.asm
GNU C version 1.37.1 Motorola DSP56000/1 Motorola
Version: g1.11 compiled by GNU C version 1.37.1.
default target switches: -mdsp
-mlinv-plus-biv-promotion -mdo-loop-generation
-my-memory -mcall-overhead-reduction -mrep
-mreload-cleanup -mnormalization-reduction
c:\dsp\bin\asm56000 -c -B foo.cln -- cca00527.asm
c:\dsp\bin\dsplnk -c -B a.cld
--c:\dsp\lib\crt056y.cln foo.cln -L
c:\dsp\lib\lib56cy.clb
C:\> dir
a.cld foo.c
```

-UMACRO

Undefine macro MACRO.

Example 3-20. DEBUG Flag

An application program, Test.c, is being tested and some portions of the code need to be debugged. The flag DEBUG may be turned on or off through the command line with the **-D** and **-U** options respectively. This flag can then be used inside the program to enable/disable debugging features within the program.

```
C:\> dir
  debug.c

C:\> type debug.c
#include <stdio.h>
main()
{
#ifdef DEBUG
    printf("debug: a message.\n");
#endif
    printf("normal operation.\n");
}

C:\> g56k -UDEBUG debug.c
C:\> run56 a.cld
normal operation.
```

-Wcomment

Warn the user whenever the comment start sequence `/*` appears within a comment.

Example 3-21. -Wcomment Option

A comment is enclosed with `/*` and `*/` and therefore is ignored by the preprocessor. Any number of leading `/*`'s are permitted within the comment and will not be reported; however, a warning message can be enabled by using the **-Wcomment** option.

```
C:\> dir
  foo.c

C:\> type foo.c
/* foo.c */
main() { /* begin */
    int d = 1000; /* /* delay */
    while(d--); /* /* main /* loop */
} /* end */

C:\> g56k -Wcomment foo.c
foo.c:3: warning: `/*' within comment
```

```
foo.c:4: warning: `/*' within comment
foo.c:4: warning: `/*' within comment
C:\> g56k foo.c
```

-Wtrigraphs

Warn if any trigraphs are encountered (Trigraphs are sequences of three characters which are replaced by a single character. These trigraph sequences enable the input of characters that are not defined in the Invariant Code Set as described in ISO 646:1983, which is a subset of the seven-bit ASCII code set.).

3.2.2 Compile Phase Options

The default options are:

1. Perform all machine dependent and independent optimizations.
2. Do *not* run the assembly language optimizer (alo56).
3. Do not generate debugging information.
4. Locate data only in the Y data memory space.

-alo

Run the assembly language optimizer on the assembly language output of **G56-cc1**. This improves the utilization of parallel moves.

-fno-opt

Disable all optimizations.

-fno-peephole

Disable the peephole portion of optimization.

-fno-strength-reduce

Disable the optimizations of loop strength reduction and elimination of iteration variables as well as DSP56000/1 specific looping optimizations (**DO** instruction usage, etc.).

-fno-defer-pop

By default, the compiler will try to defer (delay) restoring the stack pointer upon the return of a function call. The purpose of deferring restoration of the stack pointer is to reduce code size and decrease execution time; however, the stack penetration may increase (see the *DSP56000 Family Manual* for information on stack overflow).

Examples of function calls that will not incur deferred pops whether or not the **-fno-defer-pop** option is specified are:

- calls as function arguments
- calls in conditional expressions
- calls inside a statement expression

-fforce-addr

Force memory address constants to be copied into registers before doing arithmetic on them. The code generated with this option may be better or it may be worse depending on the source code. This option forces memory addresses into registers which, in turn, may be handled as common sub-expressions.

-finline-functions

Attempt to insert all simple functions in-line into their callers. The compiler heuristically decides which functions are simple enough to merit this form of integration.

If all calls to a given function are inserted, and the function is declared static, then the function is no longer needed as a separate function and normally is not output as a separate function in assembly code.

-fcaller-saves

Enable values to be allocated in registers that will be overwritten by function calls by emitting extra instructions to save and restore the registers around such calls. Such allocation is done only when it seems to result in better code than would otherwise be produced.

-fkeep-inline-functions

Output a separate run-time callable version of the function even if all calls to a given function are integrated and the function is declared static.

-fwritable-strings

Store string constants in the writable data segment without making them unique. This is for compatibility with old programs which assume they can write into string constants. Writing into string constants is poor technique; *constants* should be constant.

-fcond-mismatch

Allow conditional expressions with mismatched types in the second and third arguments. The value of such an expression is void.

-fvolatile

Consider all memory references through pointers to be volatile.

-ffixed-REG

Treat the register named REG as a fixed register; generated code should never refer to it (except perhaps as a stack or frame pointer). Legal values for REG are:

r1, r2, r3, r4, r5, r7

This flag should be used *sparingly* as it can have devastating results on the code generated.

Example 3-22. R4

Reserve r4 for later special purpose.

```
C:\> g56k -O -ffixed-r4 file.c -o file.cld
```

Warning

C code that utilizes library code can produce non-deterministic results, as the libraries have been written to utilize the complete set of registers.

-g

Produce COFF debugging information for Gdb56, the GDB cross debugger for the DSP56000/DSP56001.

A key feature afforded by the use of the GNU C compiler (G56k) teamed with the GNU source level debugger (Gdb56) is that the programmer is allowed to generate optimized code with debug information (select options **-g -O**) making it possible for the programmer to debug optimized code directly. Due to the optimizations performed, it is possible that variables will not be defined (unused variable

elimination), statements may not be executed (dead code elimination), and code may be executed early (code motion). This is a partial list of the oddities that may be encountered when debugging optimized code. However, the improved code performance due to optimization normally outweighs the problems encountered.

Example 3-23. Program Foo.c

The program Foo.c has a bug. In this application the line `i *= 2` should be `i += 2`. In order to use the symbolic debugger, Gdb56, to locate this problem, the `-g` option is used when Foo.c is compiled. This causes the compiler to generate additional information used by the debugger.

```
C:\> ls
foo.c

C:\> type foo.c
main()
{
    int i = 100;
    i *= 2;
}

C:\> g56k -g foo.c
C:\> dir
a.cld  foo.c
C:\> gdb56 a.cld
```

-O

Perform machine dependent and independent optimizations. This is the default mode of the compiler.

Invoking the compiler with the optimizer may cause compile times to increase and require more system memory.

Invoking the compiler without the optimizer should be done only when the programmer requires additional flexibility while debugging code. An example of such flexibility includes the ability to assign new values to declared C variables. Additionally, non-optimized code takes register usage clues from the storage class specifier register, something not done with the optimizer invoked.

Disabling the optimizer is done via `-f` options listed above.

-m56002

Enables the generation of DSP56002 instructions.

-mno-dsp-optimization

Disables all Motorola optimizer enhancements.

-mno-do-loop-generation

Disable **DO** instruction usage by optimizer.

-mno-biv-plus-linv-promotion

Disable the promotion of address expressions to address registers within loops. This optimization transforms array base address plus induction variable expressions into auto-increment/decrement style memory references.

-mp-mem-switchtable

Forces the compiler to locate all switch tables in P memory.

-mstack-check

Generate *extra* run-time code to check for run-time stack collision with the heap. This option causes run-time execution times to increase dramatically.

-mx-memory

Direct the compiler to locate data in the X data memory space. Memory modes cannot be mixed, i.e. only one of **-mx-memory**, **-my-memory** or **-ml-memory** may be selected.

Example 3-24. -mx-memory Option

An application is programmed to utilize only the DSP56001 X data memory space and therefore must be compiled using the **-mx-memory** option.

```
C:\> ls
x.c

C:\> type x.c
void function(int a, int b);
int X;
main()
{
    int arg1,arg2;
    function(arg1, arg2);
}
void function(int a, int b)
{
    X = a + b;
```

```
}  
  
C:\> g56k -S -mx-memory x.c  
C:\> dir  
x.asm x.c
```

-my-memory

Direct the compiler to locate data in the Y data memory space. This is the default memory mode. Memory modes cannot be mixed, i.e. only one of **-mx-memory**, **-my-memory** or **-ml-memory** may be selected.

-ml-memory

Direct the compiler to locate data in the L data memory space. This has 2 side effects.

1. A performance increase for 48-bit data code (long, float, and double).
2. This requires that the X and Y memory spaces be evenly populated.

Memory modes cannot be mixed, i.e. only one of **-mx-memory**, **-my-memory** or **-ml-memory** may be selected.

-pedantic

Issue all the warnings demanded by strict ANSI standard C; reject all programs that use forbidden extensions.

Without this option, certain GNU extensions and traditional C features are supported. With this option, they are rejected. Valid ANSI standard C programs will compile properly with or without this option.

-pedantic does not cause warning messages for use of the alternate keywords whose names begin and end with “_”.

-Q

Direct the compiler to execute in verbose mode.

-S

Compile to DSP56000/1 assembly code with the original C source lines as comments but do not assemble. The assembly language output is placed into a file suffixed .asm.

Example 3-25. Generate an optimized assembly language file (test.asm) of the C source program (test.c).

```
C:\> dir
test.c

C:\> type test.c
#include <stdio.h>
main()
{
    int i = 100;
    printf("value:%d\n", i++);
}

C:\> g56k -S test.c
C:\> dir
test.asm test.c
```

Example 3-26. Generate an optimized assembly language file test.asm.

```
C:\> g56k -O -S test.c
```

-w

Inhibit all warning messages.

-W

Print extra warning messages for the following events:

- An automatic variable is used without first being initialized.
This warning is possible only when the optimizer is invoked during compilation (default). The optimizer generates the *data flow* information required for reporting. This warning will only occur for variables that are candidates for register promotion. Therefore, they do not occur for a variable that is declared volatile, whose address is taken, or whose size is other than 1 or 2 words (integral and float data types). Warnings will not occur for structures, unions or arrays, even when they are in registers.
There may be no warning about a variable that is used only to compute a value that is never used because such computations may be deleted by data flow analysis before the warnings are printed.
Spurious warnings may be avoided by declaring functions that do not return as volatile.
- A nonvolatile automatic variable may be changed by a call to longjmp.
This warning also requires that the optimizer be invoked.

The compiler sees only the calls to `setjmp`. It cannot know where `longjmp` will be called; in fact, a signal handler could call it at any point in the code. As a result, a warning may be issued even when there is no problem because `longjmp` cannot be called at the place which would cause a problem.

A function can return either with or without a value. (Falling off the end of the function body is considered returning without a value.) For example, this function would evoke such a warning:

```
foo (a)
{
    if (a > 0)
        return a;
}
```

Spurious warnings can occur because GNU CC does not realize that certain functions (including 'abort' and 'longjmp') will never return.

- An expression-statement contains no side effects.

Example 3-27. -W

Extra warning messages are wanted to help find potential problems in a test function, `foo()`, which is programmed to return a value only if `a > 0`.

```
C:\> dir
foo.c

C:\> type foo.c
int foo(int);
main()
{
    int i;
    foo(i);
}
int foo(a)
{
    if(a > 0)
        return a;
}

C:\> g56k -W foo.c
foo.c: In function main:
foo.c:4: warning: 'i' may be used uninitialized in this
function
foo.c: In function foo:
foo.c:11: warning: this function may return with or
without a value
```

-Wimplicit

Warn whenever a function is implicitly declared.

Example 3-28. -Wimplicit

The function `foo()` is declared implicitly in the program `foo.c`, the `-Wimplicit` option will generate a warning message for this situation.

```
C:\> dir
foo.c

C:\> type foo.c
main()
{
    foo();
}
int foo(){}
C:\> g56k -Wimplicit foo.c
foo.c: In function main:
foo.c:3: warning: implicit declaration of function
'foo'
C:\> dir
a.cld foo.c
```

-Wreturn-type

Warn whenever a function is defined with a return-type that defaults to `int`. Also warn about any return statement with no return-value in a function whose return-type is not `void`.

Example 3-29. -Wreturn-type

The function `foo()` is declared as a function that should return an integer but in this case does not return an integer. The `-Wreturn-type` option generates a warning message in this situation.

```
C:\> dir
foo.c

C:\> type foo.c
int foo(), main();
int main()
{
    return foo();
}
```

```
int foo(){}

C:\> g56k -Wreturn-type foo.c
foo.c: In function foo:
foo.c:6: warning: control reaches end of non-void
function

C:\> dir
a.cld  foo.c
```

-Wunused

Warn whenever a local variable is unused aside from its declaration, whenever a function is declared static but never defined and whenever a statement computes a result that is explicitly not used.

Example 3-30. -Wunused

The file `foo.c` contains an undefined static function, unused local variable, and a dead statement. The **-Wunused** option will issue warnings to indicate these situations.

```
C:\> dir
foo.c

C:\> type foo.c
static int foo();
main()
{
    int x;
    2+3;
}

C:\> g56k -Wunused foo.c
foo.c: In function main:
foo.c:5: warning: statement with no effect
foo.c:4: warning: unused variable 'x'
foo.c: At top level:
foo.c:1: warning: 'foo' declared but never defined

C:\> dir
a.cld  foo.c
```

-Wswitch

Warn whenever a switch statement has an enumeration type of index and lacks a case for one or more of the named codes of that enumeration. (The presence of a default label prevents this warning.) case labels outside the enumeration range also provoke warnings when this option is used.

-Wall

All of the above **-W** options combined. The remaining **-W** options described below are not implied by **-Wall** because certain kinds of useful macros are almost impossible to write without causing those warnings.

-Wshadow

Warn whenever a local variable shadows another local variable.

-Wid-clash-LEN

Warn whenever two distinct identifiers match in the first LEN characters. This may help prepare a program that will compile with certain obsolete compilers.

-Wpointer-arith

Warn about anything that depends on the sizeof a function type or of void. GNU C assigns these types a size of 1, for convenience in calculations with void * pointers and pointers to functions.

-Wcast-qual

Warn whenever a pointer is cast so as to remove a type qualifier from the target type. For example, warn if a const char * is cast to an ordinary char *.

-Wwrite-strings

Give string constants the type const char[LENGTH] so that copying the address of one into a non-const char * pointer will generate a warning. These warnings help at compile time to find code that can try to write into a string constant, but only if const in declarations and prototypes have been used carefully.

3.3 Assemble Phase Options

This group of assemble phase options is the sub-set of the available assembler options that are compiler oriented (see the *Motorola DSP56000 Macro Assembler Reference Manual* for a complete option list). The default option is to add to the standard search list the directory that the C compiler writes its output into and then search that directory first.

-asm string

Pass the argument *string* directly to Asm56000, the DSP56000/1 assembler.

Example 3-31. Pass a single option to the assembler.

```
C:\> g56k -asm -v file.c
```

Example 3-32. Pass multiple options to the assembler.

```
C:\> g56k -asm "-v -OS,CRE" file.c
```

-c

Compile and/or assemble the source files, suppressing the link phase. This option generates corresponding output files suffixed “.cln”. Affected input files are suffixed with “.c” and “.asm”.

3.4 Link Phase Options

The options listed below control the link phase of the compilation process. This group of link phase options is the sub-set of the available linker options that are compiler oriented (see the *Motorola DSP56000 Linker/Librarian Reference Manual* for a complete option list). The **-crt** and **-l** options locate the file provided as an argument by searching a standard list of directories, see section 3.1, “Overview.” for this directory list. The default option is to add the C compiler output directory into the standard search list and search that directory first.

-crt file

Replace the default start-up file (`crt056.cln`) with *file*. G56k searches the standard list of directories to find the start-up file. In addition, any directory defined using the **-B** option or the `G56_EXEC_PREFIX` environment variable will be searched. For additional information, see Chapter 6.

Example 3-33. -crt file

Compile the C program `foo.c` with the crt0 file `crt.asm`. Notice that the crt0 file `crt.asm` should be assembled before use since the option `-crt` takes `.cln` file not `.asm` file.

```
C:\> dir
crt.asm foo.c
C:\> g56k -c crt.asm
C:\> dir
crt.cln crt.asm foo.c
C:\> g56k -crt crt.cln foo.c
```

-j string

Pass the argument string directly to `dsplnk`, the DSP56000/1 linker.

Example 3-34. Pass a single option to the linker.

```
C:\> g56k -j -v file.c
```

Example 3-35. Pass multiple options to the linker.

```
C:\> g56k -j "-v -i" file.c
```

-ILIBRARY

Search the standard list of directories for a library file named `libLIBRARY.clb`. The linker automatically expands `LIBRARY` from the option command into `libLIBRARY.clb` and uses this file as if it had been specified precisely by name.

Example 3-36. -ILIBRARY

Compile the source code using the special dsp application library. Searching the standard list of directories for a library named `libdspaps.clb`.

```
C:\> g56k -O file.c -ldspaps
```

-r CTLFILE

Search the standard list of directories for the memory control file `CTLFILE` to be passed as an argument to the DSP56000/1 relocatable linker. This control file will be used as a table to locate object files sections to be linked. For more detailed information, see the **-R** options and the section on “Memory Control File” in the *Motorola Linker/Librarian Reference Manual*.

Compile the source code `main.c` and `data.c` with the memory configuration described in the control file `map.ct1`. Notice that the section `main_c` of the program `Main.c` is located at the memory address `p:$3000` and the section of `data_c` of the data `data.c` is located at the memory address `y:$5000`. See chapter 5 for detailed information on the in-line assembly code (`__asm(...)`).

Example 3-37. -r CTLFILE

```
C:\> type map.ct1
sectionmain_c p:$3000
section data_c y:$5000

C:\> type data.c
int data = 0x1; /* test value */

C:\> type main.c
extern int data;
main()
{
    int i;
    for(i = 0; i < 10; i++)
        __asm("rol %0" : "=D" (data) : "0"
(data));
}

C:\> g56k -j "-mmap.map" -r map.ct1 main.c data.c

C:\> type map.map
...
Y Memory (default)
Start   End     Length Section
5000    5000    1      data_c
...
P Memory (default)
Start   End     Length Section
3000    3008    9      main_c
...
```

Chapter 4

About g56k

4.1 Introduction

The DSP56000 digital signal processors are designed to execute DSP oriented calculations as fast as possible. As a by-product, they have an architecture that is somewhat unconventional for C programming. Because of this architecture, there are characteristics of the compiler, and the code generated by the compiler, that the programmer must understand in order to take full advantage of the DSP56KCC programming environment. All programmers, whether they are familiar with DSP or not, should understand the DSP56000 family architecture before attempting to program it in C. The following sections provide important information on data types, storage classes, memory and register usage, and other topics which will be useful to the DSP56000 application developer programming in C.

4.2 Identifiers

An identifier is defined as a sequence of letters, digits and underscore characters ('_'). The first character must be a letter or underscore. DSP56KCC identifier length limits are listed in Table 4-1.

Table 4-1. Identifier Length Limits

Identifier Storage Class	Length
Global/Static (External Linkage)	255
Auto	unlimited

4.3 Predefined Preprocessor Macro Names

DSP56KCC supports and expands all ANSI defined macros and three additional non-ANSI predefined macro names. Table 4-2 lists the macros and their explanation.

4.4 Data Types and Sizes

Due to the word orientation of the DSP56000 family (24-bit words), all data types are aligned on word boundaries. This has several side effects, one of which is that `sizeof(char)` is equal to `sizeof(int)`.

Table 4-2. Predefined Macro List and Explanation

MACRO	ANSI	Explanation
<code>__LINE__</code>	YES	The line number of the current source line (a decimal constant).
<code>__FILE__</code>	YES	The name of the source file (a character string).
<code>__DATE__</code>	YES	The compilation date (a character string of the form "Mmm dd yyyy" e.g., Jul 22 1991).
<code>__TIME__</code>	YES	The compilation time (a character string of the form "hh:mm:ss").
<code>__STDC__</code>	YES	Decimal constant 1, indicates ANSI conformation.
<code>__DSP56K__</code>	NO	Decimal constant 1, indicates that code is being generated for the DSP56000/DSP56001 Digital Signal Processor.
<code>__VERSION__</code>	NO	The version number of the compiler (a character string of the form "d.dd.d").
<code>__INCLUDE_LEVEL__</code>	NO	Decimal constant, indicates the current depth of file inclusion.

4.4.1 Integral Data Types

The type char, unsigned char, signed char, short int, int, long int and the enumerated types comprise the integral data types. The type sizes and ranges are defined in Table 4-3.

Table 4-3. Integral Data Type Sizes and Ranges

Data Type	Size (words)	Min value	Max value
char	1	-8388608	8388607
unsigned char	1	0	0xFFFFFFFF
short	1	-8388608	8388607
unsigned short	1	0	0xFFFFFFFF
int	1	-8388608	8388607
unsigned int	1	0	0xFFFFFFFF
long	2	-140737488355328	140737488355327

4.4.2 Floating-point Types

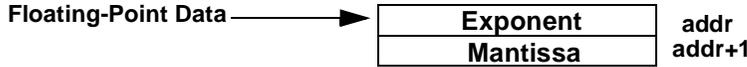
In DSP56KCC, the C data types float and double are both implemented as single precision (see Table 4-5). DSP56KCC does not implement the IEEE STD 754-1985 standard format for binary floating-point arithmetic. A description of the format and a comparison with the IEEE standard follow.

Table 4-4. Floating-point Data Type Sizes and Ranges

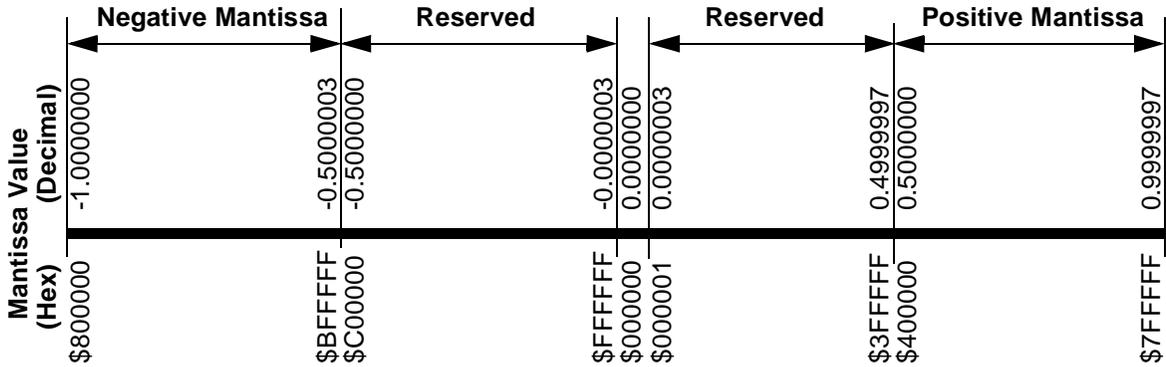
Data Type	Size (words)	Min value	Max value
float	2	1.175494351e-38	3.402823466e+38

4.4.2.1 Floating-point Format Description

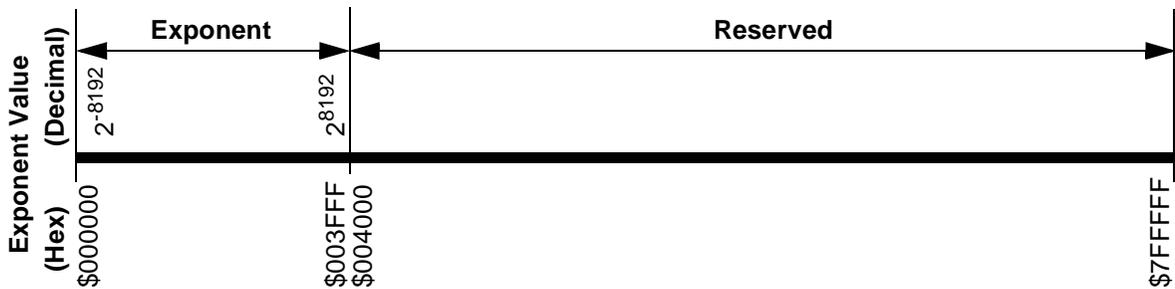
Figure 4-1 illustrates the floating-point format used in DSP56KCC. Figure 4-1a shows that the exponent and mantissa occupy consecutive memory locations. Figure 4-1b is in number line format and shows the fractional nature of the mantissa and the fact that, due to the nature of a fractional arithmetic mantissa, the numbers between -0.5 and +0.5 (except for zero) are not needed and are therefore reserved. Figure 4-1c shows the range used by the exponent in this implementation. Notice how this compares with the IEEE implementation shown in Table 4-6. Figure 4-1d is a combined number line showing the range of numbers which can be represented in DSP56KCC. The mantissa \$C00000 (-0.5) is not included as the smallest negative floating-point number because the normalization routine automatically detects the two leading ones and decrements the exponent which, if at \$003FFF, will result in an underflow. Therefore, the smallest negative mantissa has been set to \$BFFFFFF (-0.916). Table 4-5 lists the specific floating-point format information for DSP56KCC and is a tabular version of the information in Figure 4-1.



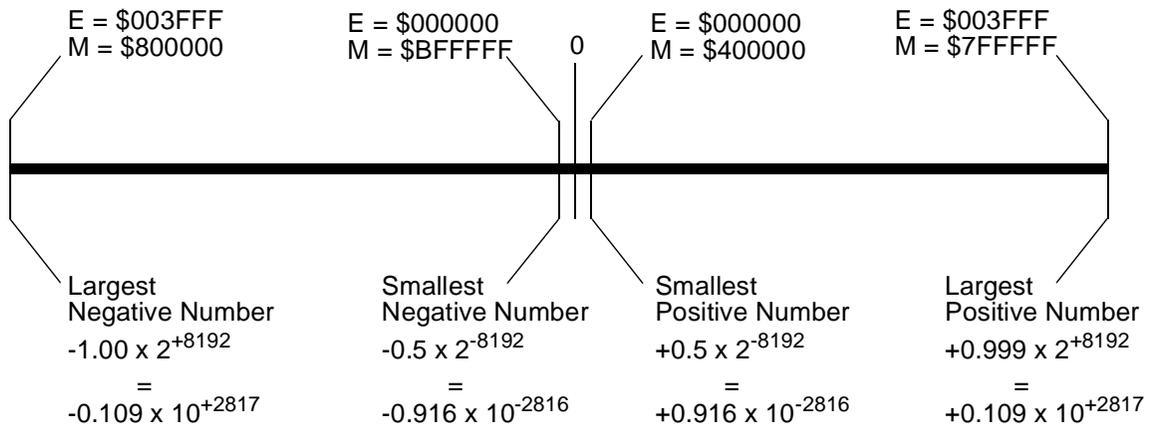
(a) Floating-point Data Arrangement in Memory



(b) Mantissa Data Range



(c) Exponent Data Range



Note: E = Exponent and M = Mantissa

(d) Mantissa and Exponent Data Range

Figure 4-1. Mantissa and Exponent Data Range of C Floating Point

4.4.2.2 Comparison with IEEE STD 754-1985 Standard

Since the IEEE floating-point arithmetic standard is well publicized, it is useful to compare these two floating-point formats. The DSP56KCC floating-point format differs from the IEEE standard primarily in its handling of floating-point exceptions. Other differences are noted in Table 4-6. Conversion between the IEEE standard format and DSP56KCC format is straightforward.

As shown in Table 4-6, the DSP56KCC mantissa precision is one bit less than the IEEE single precision format. This is the result of using two's complement arithmetic in the DSP56000/DSP56001. The DSP56KCC exponent width is three bits more than the IEEE double precision format. This provides an extremely large (approximately 100,000 dB) dynamic range which eliminates exponent overflow for most applications. If exponent overflow occurs, the result is limited to the maximum representable floating-point number of the correct sign. If exponent underflow occurs, the result is limited to the minimum representable floating-point number, which is zero. Although the DSP56KCC format does not provide the arithmetic safety offered by the IEEE standard, it avoids extensive error checking and exceptions in favor of real-time execution speed and efficient implementation on the DSP56000/DSP56001. All exception conditions are handled *in-line* (i.e., an exception routine is not invoked) according to predefined rules. The reason for this is the fact that real-time systems must continue to process data non-stop. It is not possible to stop execution until the application program determines a solution to the problem and corrects it and so there is no choice but to provide an output with some amount of error should an exception occur.

Table 4-5. Floating-point Format Description

IEEE Format Characteristic	DSP56KCC Value
Decimal Value	$m * 2^{(e - ebias)}$
Mantissa	24-bit two's complement, normalized fractional mantissa. This gives a precision of approximately 7 decimal digits. A hidden leading 1 is not implemented in this format (see Figure 4-1).
Largest Positive Mantissa	\$7FFFFFF
Smallest Positive Mantissa	\$400000
Floating-point Zero Mantissa	\$000000
Smallest Negative Mantissa	\$BFFFFFF
Largest Negative Mantissa	\$800000
Reserved Mantissas	\$000001 through \$3FFFFFF and \$C00000 through \$FFFFFF

Table 4-5. Floating-point Format Description

IEEE Format Characteristic	DSP56KCC Value
exponent	14-bit exponent (unsigned integer, biased by ebias = \$1FFF). Stored as a 24-bit unsigned integer with 10 leading zeros. The 14-bit exponent used by DSP56KCC provides a larger dynamic range than IEEE double precision format.
Largest exponent (biased)	\$003FFF = 2^{+8192}
Smallest exponent (biased)	\$000000 = 2^{-8192}
Reserved exponents	\$004000 through \$FFFFFF

Note:

1. No distinct exponents are reserved for plus infinity, minus infinity, Not-a-Number (IEEE NaN), minus zero or denormalized numbers as is done in IEEE format.
2. All reserved mantissas are illegal since they represent denormalized mantissas.
3. If the 15th bit is set, exponent overflow has occurred.
4. If the 16th bit is set, exponent underflow has occurred.

Table 4-6. Comparison of DSP56KCC and IEEE 754-1985

CHARACTERISTIC	DSP56KCC FORMAT	IEEE FORMAT
Mantissa Precision	23 bits	24 bits
Hidden Leading One	No	Yes
Mantissa Format	24-bit Two's Complement Fraction	23-bit Unsigned Magnitude Fraction
Exponent Width	16 bits (14 bits used)	8 bits 11 bits
Maximum Exponent	+8192	+127 (8 bit case) +1023 (11 bit case)
Minimum Exponent	-8191	-127 (8 bit case) -1022 (11 bit case)
Exponent Bias	+8191	+127 (8 bit case) +1023 (11 bit case)
Format Width	48 bits	32 bits (8 bit case) 64 bits (11 bit case)

Table 4-6. Comparison of DSP56KCC and IEEE 754-1985

CHARACTERISTIC	DSP56KCC FORMAT	IEEE FORMAT
Rounding	Round to Nearest	Round to Nearest Round to +/- Infinity Round to Zero
Infinity Arithmetic	Saturation Limiting	Affine Operations
Denormalized Numbers	No (Forced to Zero)	Yes (with Min Exp)
Exceptions	Divide by Zero Overflow Negative Square Root	Invalid Operation Divide by Zero Overflow Underflow Inexact Arithmetic

One major difference is the use of affine arithmetic in the IEEE standard versus the use of saturation arithmetic in the DSP56KCC format. Affine arithmetic gives separate identity to plus infinity, minus infinity, plus zero and minus zero. In operations involving these values, finite quantities remain finite and infinite quantities remain infinite. In contrast, DSP56KCC format gives special identity only to unsigned zero. This format performs saturation arithmetic such that any result out of the representable floating-point range is replaced with the closest floating-point representation. Since the dynamic range of this format is quite large, it is adequate for most applications.

The IEEE floating-point standard provides extensive error handling required by affine arithmetic, denormalized numbers, signaling Not-a-Numbers (NaNs) and quiet NaNs. It postpones the introduction of computational errors by using internal signaling and user traps to process each exception condition. Computational errors will be introduced by the application program if the calculation is completed instead of aborting the program. The DSP56KCC format introduces computation errors when an exception occurs in order to maintain real-time execution. An error flag (L bit in CCR) is set to inform the application program that an exception has occurred. This bit will remain set until reset by the application program. The user can then eliminate the exception by algorithm modifications.

4.4.3 Pointer Types

With DSP56KCC, all pointers are 16-bits (see Table 4-7). When computing addresses with integer arithmetic, only the least significant 16-bits are relevant.

Table 4-7. Pointer Size and Range

Data Type	Size (words)	Min value	Max value
pointers	1	0	0xFFFF

4.5 Register Usage

The DSP56000 family digital signal processor register set is shown in Table 4-8. DSP56KCC uses all of the registers listed in Table 4-8 with the exception of the M_n address modifier registers.

Table 4-8. DSP56000 Family Processor Registers

Data ALU	
x_n	Input Registers x1, x0 (24-bits)
y_n	Input Registers y1, y0 (24-bits)
a_n	Accumulator Registers a2 (8-bits), a1, a0 (24-bits)
b_n	Accumulator Registers b2 (8-bits), b1, b0 (24-bits)
x	Input Register x (x1:x0, 48-bits)
y	Input Register y (y1:y0, 48-bits)
a10	Input Register a10 (a1:a0, 48-bits)
b10	Input Register b10 (b1:b0, 48-bits)
a	Accumulator a (a2:a1:a0, 56-bits)
b	Accumulator b (b2:b1:b0, 56-bits)
Address ALU	
r_n	Address Registers R0-R7 (16-bits)
n_n	Address Offset Registers N0-N7 (16-bits)
m_n	Address Modifier Registers M0-M7 (16-bits)

Warning

The mn address modifier registers are not used directly by DSP56KCC. Some of these registers are implied whenever any address registers (r0 - r7) are referenced either in C library or in C. While assembly code can access and use these registers, the programmer must restore them to their previous state (\$FFFF) before returning control to DSP56KCC. Failing to do so will cause unpredictable errors when DSP56KCC accesses these registers.

The programmer is required to preserve any registers that are directly used in *in-line* and in *out-of-line* assembly language code (see Chapter 5, *Mixing C and Assembly Language*). Table 4-9 outlines the compiler's usage of each register.

Table 4-9. DSP56KCC Registers and Usage

Register	Usage
r0	Frame Pointer
r6	Stack Pointer
r1 - r5, r7	Register promotion by the optimizer
n0 - n7	Code generator temporary
m0 - m7	Used by compiler; keep this as \$FFFF
a	48-bit function return value. float, double, or long
a1	24-bit and 16-bit return value. Integer or pointer
b, x, y	48-bit register promotion by optimizer
x1, x0, y1, y0	24-bit register promotion by optimizer

4.6 Memory Usage

Memory can be partitioned in several ways to provide high-speed operation and additional off-chip memory expansion. Program and data memory are separate and the data memory is divided into two separate 24-bit wide memory spaces, X and Y which can be combined to form a third 48-bit wide memory space, L.

DSP56KCC generates code that can utilize X or Y or L data memory. By default, the Y memory model is selected. Each separate memory model is supported by separate run-time libraries that have been written and compiled/assembled to handle each memory space. See Chapter 3, *Control Program Options*, for examples of the **-mx-memory**,

-my-memory and **-ml-memory** command-line options for changing target memory spaces.

By default, the compiler expects that each memory space is fully populated and several global C variables are defined (see Chapter 6 — *Software-Hardware Integration* for information about customizing the memory configuration). Figure 4-2 and Figure 4-4 illustrate the default program and data memory configuration.

The L memory model is recommended for best performance when long, float or double data types are extensively used.

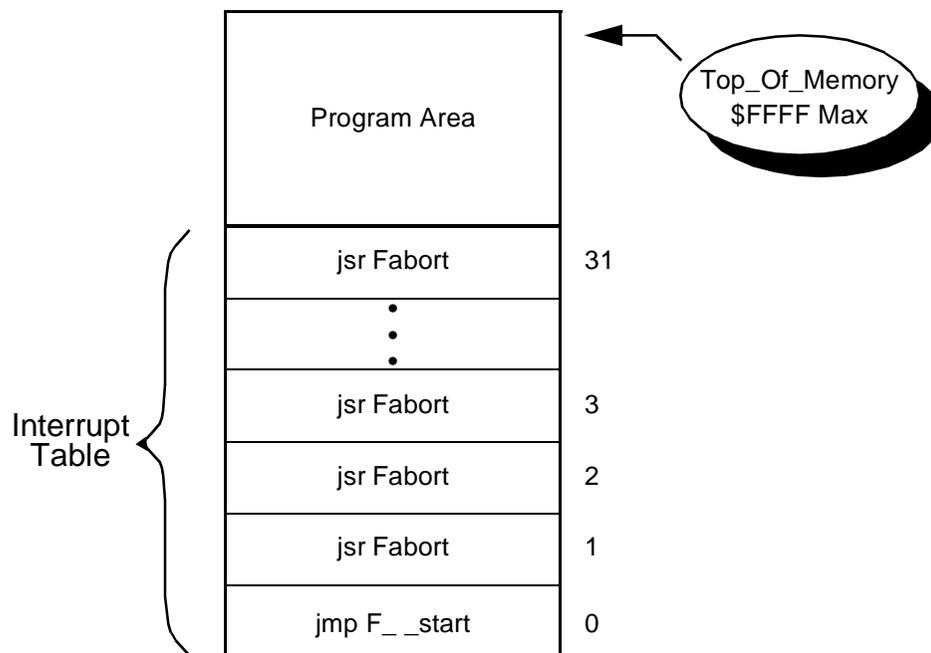


Figure 4-2. Default Program Memory Configuration

4.6.1 Activation Record

An activation record is the run-time representation of a C subroutine. A typical DSP56KCC activation record consists of the following elements and is illustrated in Figure 4-3:

1. Parameter data space. Information passed to C subroutines is stored in a parameter data space which is similar to the local data space (see Figure 4-3). However, the data is in reverse order and each parameter is referenced via a negative offset from the frame pointer. Actual parameters are pushed onto the activation record in reverse order by the calling subroutine.
2. Old frame pointer. Once the called subroutine has completed execution, the frame pointer will be updated with this value.

3. Return address — which is pushed on the DSP's system stack high (ssh) register. This is the return address to the calling subroutine. The return address is *not* saved for subroutines that have been determined to be a leaf. A leaf subroutine is one that makes no subroutine calls.
4. Local data space. The location of C variables that have a lifetime that extends only as long as the subroutine is active and that could *not* be explicitly promoted to register storage class by the optimizer.

Note: The frame pointer (r0) points to the first element in the local data space.

5. Register spill and compiler temporary space. This area is utilized by the compiler to store intermediate results and preserve registers.

Note: The Stack Pointer (r6) points to the next available data memory location.

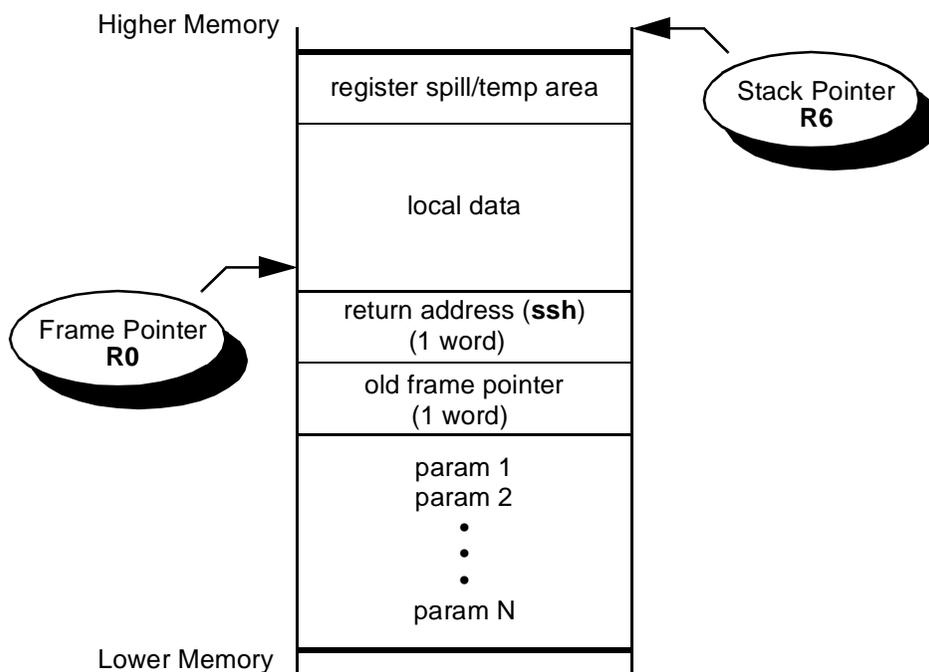


Figure 4-3. Activation Record

Each subroutine called puts a new copy of the subroutine activation record in the run-time stack and returning from the subroutine removes the activation record. The run-time stack is described in Figure 4-4, “Default Data Memory Configuration,” on page 4-12. The variables shown in the bottom of the “X or Y memory option selected” memory are controlled by the crt0 file. For example, the `F__fp_shift` variable is typically 23 words but can be changed by the user or may vary with later releases of this compiler. When the L memory option is selected, the heap, run-time stack, global/static data and data that is more than 24 bits in length will occupy one word in L memory i.e., 48-bit memory. 16-bit

and 24-bit data memory will occupy only one word in Y memory. DSIZE is set by the linker and points to the top address of the global and static data. DSIZE is used by the crt0 file as the default initial stack pointer.

The dynamic run-time stack growth is illustrated in Figure 4-5. In this example, there is one activation record as execution of the sample C code begins. This activation record is pushed onto the stack and a new activation record is built with a dynamic link to the old frame pointer. When the function returns, the function's activation record is popped and the original activation record is restored to its original place on the stack.

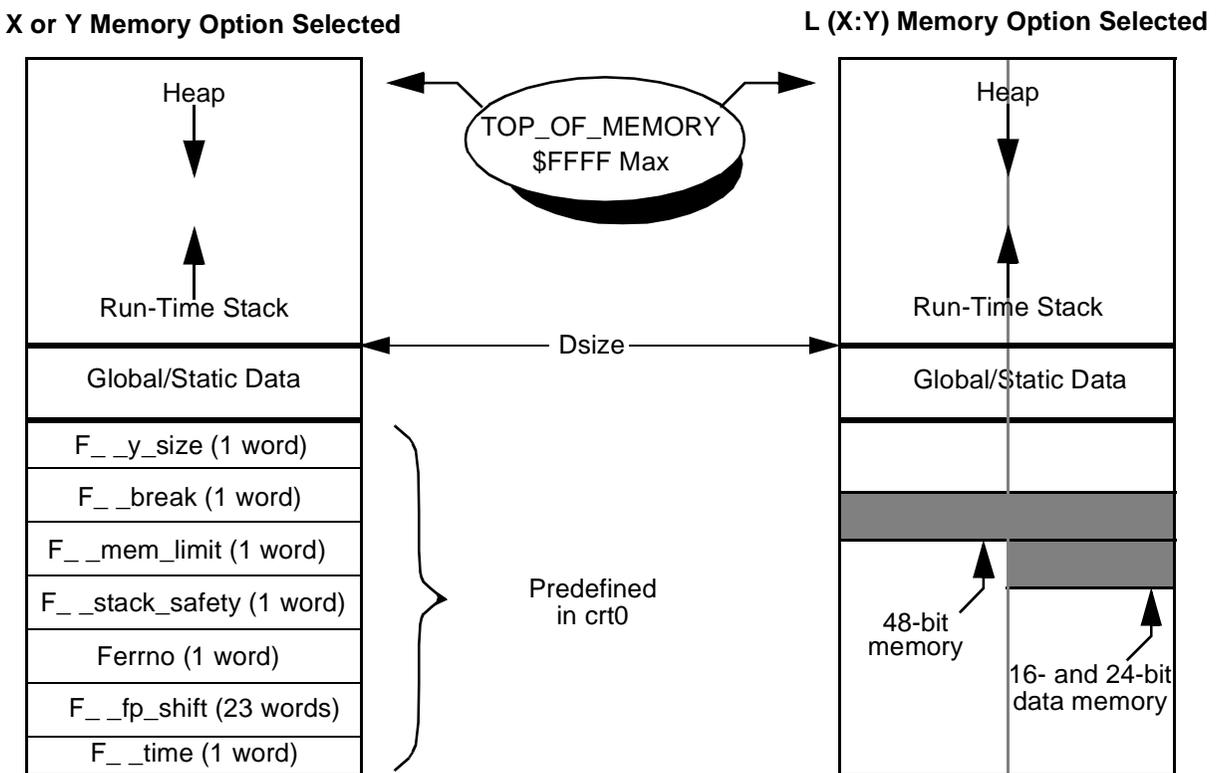


Figure 4-4. Default Data Memory Configuration

```
main()
{
    func_1();
}
```

Sample C code

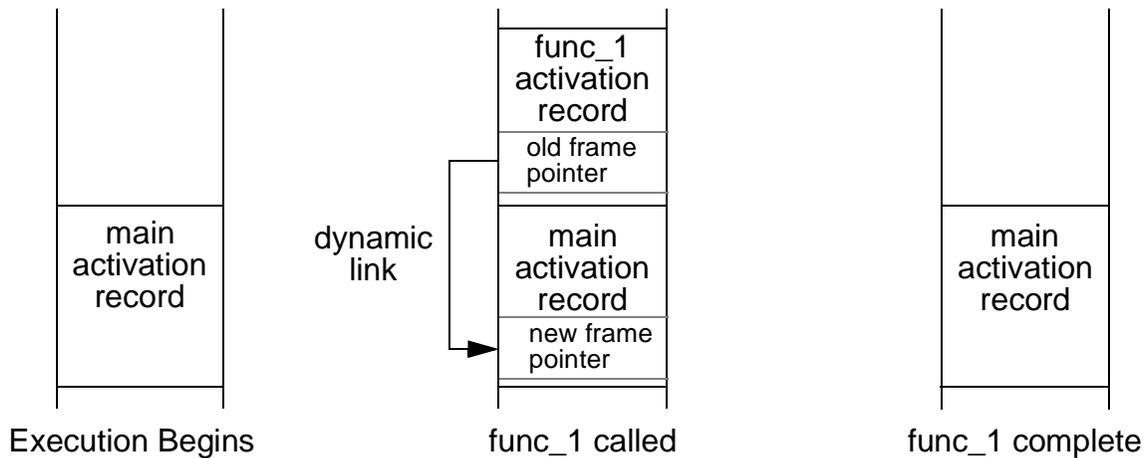


Figure 4-5. Run-time Stack Growth

4.6.2 Global/Static Data

By default, global and static data elements are located below the run-time stack and each element is referenced by a unique label that is known at compile-time (see Chapter 6, *Software-Hardware Integration* for additional information).

4.7 Compiler Naming Conventions

The compiler uses five different internal label formats and a special section naming format. These six separate formats simplify the procedures to combine hand written assembly language and C language statements. Use of these formats also makes compiler generated assembly language listings easier to read. It is strongly recommended that the programmer avoid using labels with these formats.

- L# Local labels. Generally the targets of conditional and unconditional jumps. Where # is a decimal digit.

- LC# String Constant labels. The data memory location for C string constants, such as “string constant”.

L#	Local labels. Generally the targets of conditional and unconditional jumps. Where # is a decimal digit.
F<identifier>	Global C variables, global subroutines, static C variables and static subroutines. A static C variable or subroutine is one which is not visible to any C code outside the file in which it has been declared, thus making it possible to reuse variable names across file boundaries. Where <i>identifier</i> is the variable or subroutine name.
F_ _<identifier>#	Variables static to a function.
ASM_APP_#	In-line assembly code delimiters. Required to allow the programmer to define and use local labels (labels beginning with an underscore character ‘_’).
<filename_c>	Section names. The contents of each assembly language file generated by the compiler are contained in a unique section . Where <i>filename</i> is the file name minus any ‘.’ extensions.

4.8 Subroutine Call Sequence

Each time a C language subroutine is called, a strict calling convention is followed. The subroutine calling sequence is broken down into three sub-sequences that are strictly defined. The three sub-sequences are caller, callee and return sequence.

Note: This calling convention must be followed when writing in-line or out-of-line assembly language subroutines that call subroutines written in C.

4.8.1 Caller Sequence

The *caller* portion of the subroutine calling sequence is responsible for:

1. Pushing arguments onto the activation record (in reverse order).
2. Actual subroutine call (jsr).
3. Stack pointer adjustment.

Additional *caller* sequence when the subroutine called will return a structure:

4. Allocate space in the caller’s activation record to store the return structure.

5. Pass the return value address in accumulator a.

4.8.2 Callee Sequence

During the initial portion of the subroutine calling sequence, the *callee* is responsible for:

1. Saving return address (ssh) and the old frame pointer (R0).
2. Updating frame and stack pointers.
3. Saving the following registers, as required, B1, B0, X1, X0, Y1, Y0, R1 - R5 and R7.

4.8.3 Return Sequence

During the final portion of the subroutine calling sequence, the *callee* is responsible for:

1. Placing the return value in accumulator a.
2. Testing the return value. This optimizes the case where function calls are arguments to conditional operators.

Additional *callee* sequence when the subroutine called will return a structure:

3. The return value is not passed in accumulator a. A copy of the return structure is placed into the space allocated in the caller's activation record and pointed to by accumulator A.

4.9 Software Support for Arithmetic Routines

The DSP56000 family architecture provides full hardware support for all 24-bit arithmetic operations, and partial support for 48-bit *integer* operations. Support for all float/double and a portion of the 48-bit long is provided via special software library routines. These special library routines *do not* pass arguments to the routines according to the normal subroutine calling convention, instead, each routine has its arguments passed in the A and B accumulators and the result returned in the A accumulator, limiting the overhead involved in the call/return sequence. See Table 4-10 and Table 4-11 for a list of the routines supported in software. All but two of these routines restore all registers used except for the A accumulator which contains the result. The two routines that do not restore the registers are **F__uldiv** and **F__ulmod** which use the registers according the C standard shown in Table 4-9.

Table 4-10. Floating-point Arithmetic

Routine	Source Module
fadd_ba	fadd56.asm
fsub_ba	fsub56.asm
fmpy_ba	fmpy56.asm
fdiv_ba	fdiv56.asm
fneg_a	fneg56.asm
fcmp_a	fcmp56.asm

Table 4-11. 48-bit Long Integer Arithmetic

Routine	Source Module
div_ba	idiv56.asm
mod_ba	imod56.asm
udiv_ba	uidiv56.asm
umod_ba	uimod56.asm
lmpy_ba	l_mpy56.asm
ldiv_ba	l_div56.asm
lmod_ba	l_mod56.asm
F__uldiv	ulongdivmod.c
F__ulmod	ulongdivmod.c

4.10 Run-time Safety

DSP56KCC provides two methods for providing run-time memory utilization checks. The first method, memory allocation check, is automatic. The second method, run-time stack probes, is provided by selecting the command-line argument **-mstack-check**.

4.10.1 Memory Allocation Checks

Memory allocation checks are provided during every call to the run-time library routines malloc, calloc and realloc. This automatic run-time check determines when the heap is about to collide with the run-time stack. When this occurs, the library routine returns a NULL pointer and sets the global variable **errno** to **ENOMEM**.

4.10.2 Run-time Stack Checks

If the programmer does not use `malloc`, `calloc` or `realloc` there is no automatic run-time collision check. A tool is provided to allow the programmer to check the code during debugging to see if a stack collision is likely to occur when the code is executing normally. By selecting the **-mstack-check** option on the command-line, the compiler is instructed to generate extra code to watch the stack and heap and detect when the run-time stack is about to collide with the heap. This may be important when writing code for embedded applications. When/if a stack collision occurs, the execution device, such as RUN56, will report that the collision occurred.

Note: All run-time libraries provided have been compiled/assembled without the stack checking option. Thus it is possible to have a run-time stack/heap collision during execution of library routines.

4.11 Optimization Techniques Implemented

This section provides a brief overview of the optimization techniques currently implemented by DSP56KCC. DSP56KCC implements most machine-independent C language optimization techniques as well as machine-specific optimizations. By default, the control program G56k enables all levels of optimization (see Chapter 3, “Control Program Options.”)

4.11.1 Register Promotion and Lifetime Analysis

The compiler automatically identifies all variables that are candidates for promotion to the register storage class. Using standard data flow techniques, lifetime analysis is performed to determine the best variables for promotion. When variable lifetimes do not overlap, more than one variable may be promoted to a single register.

4.11.2 Common Sub-expression Elimination

A Common Sub-Expression, or CSE, is created when two or more statements compute the same value. When CSEs are detected during data flow analysis, the compiler eliminates all but one of the computations and propagates the result. A classic example of a CSE is the array element assignment from another array,

```
array_1[index + 1] = array_2[index + 1];
```

where the expression “`index + 1`” is the CSE.

This optimization is especially effective as CSEs become candidates for register promotion.

4.11.3 Constant Propagation/Folding

Propagation of constants is detected during data flow analysis and is simply the replacement of variable references with constant references when possible. For example,

```
a = 3;
/* block of C code with no references to a */
func_call( a + 709 );
```

becomes:

```
/* block of C code */
func_call( 3 + 709 );
```

Constant folding is the replacement of run-time computations with compile-time computations. This optimization works well with constant propagation, as it exposes many new opportunities for folding.

The example above would be further transformed to:

```
/* block of C code */
func_call( 712 );
```

4.11.4 Dead Code Elimination

During data flow analysis, the compiler detects when the results of C expressions are never used. When this is detected, the offending C statements are eliminated from code generation. This includes (1) the initialization of variables that are not referenced in the subroutine and (2) the situation where the variables next reference is not an assignment.

To guarantee code generation, a volatile type specifier can be used to declare the left side of the expression.

```
main()
{
  int volatile i = 0, j = 1;
}
```

The example above generates code to initialize variables *i* and *j* even though they are not used anywhere else. Without the key word *volatile*, the optimizing C compiler will eliminate the two local variables because they are not referenced anywhere in the main function.

4.11.5 Tail Merging

When two or more conditional blocks of code have similar ending sequences, the sections of code are rewritten to generate similar code only once.

This is a space saving optimization technique. For example:

<pre>if (a>b) { b = a; func (a); } else func (a);</pre>	becomes	<pre>if (a>b) { b = a; } func (a);</pre>
--	----------------	---

4.11.6 Strength Reduction

Strength reduction replaces expensive operators with less expensive operators. This optimization method is very machine specific. For instance, a popular strength reduction for many compilers is to replace a multiplication by a constant with a series of shifts, additions and subtractions. The exact opposite is the case on the DSP56000/DSP56001, however since a series of left shifts is replaced with a single multiply by a constant power of 2.

4.11.7 Loop Invariant Code Motion

Loop Invariant Code Motion is a method in which all C expressions that yield the same result each time through the loop are moved outside of the loop and are executed once prior to entering the loop.

4.11.8 Hardware DO Loop Instruction

The DSP56000 family architecture provides a method in hardware to perform zero overhead looping via the **do** instruction. DSP56KCC may exchange the standard increment/compare/conditional jump sequence with a single **do** instruction (this is called do loop promotion) when the following conditions are met:

1. The body of the loop contains *no* subroutine calls,
2. The loop is entered from the top, i.e., no **goto label** entries.
3. No conditional exits from the loop are allowed.
4. The loop's induction variable is only altered in the body of the loop once per iteration. Please note that this includes any modifications to the induction variable within the actual **for** or **while** statement.

4.11.9 Loop Rotation

Loop rotation is the elimination of the code generated to check the loop's entry conditions. When a loop fails to qualify for *do loop* promotion i.e., it does not meet the four conditions listed above, it will qualify for loop rotation if the length of the loop is known at compile-time, for example:

```
for ( i = 0 ; i < 10 ; i ++ )
```

The loop created with this *for* statement will always be executed at least one time. Therefore, the "is $i < 10$?" test does not have to be run the first time through the loop and as a result, can be eliminated during the first pass through the loop only. If the result of the first test cannot be predetermined then it cannot be eliminated. In the example below, the number of loops is a variable (and therefore cannot be predetermined) that may equal zero.

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

4.11.10 Jump Optimizations

All cases of jumps (conditional and unconditional) to unconditional jumps are eliminated to create a single jump and target label.

4.11.11 Instruction Combination

Instruction combination replaces several operators with a single, less expensive operator. This optimization method is very machine specific.

Sequences that are commonly combined by the optimizer include:

1. Integer add/multiply becomes a **mac** instruction.
2. Integer subtract/multiply becomes a **mac** instruction.
3. A memory reference combined with a pointer adjustment becomes an **autoincrement** or **autodecrement** addressing mode. This is very powerful when combined with register promotion and do loop promotion. For example,

```
for ( i = 0 ; i < 10 ; i ++ )  
{  
    array_1[ i ] = array_2[ i ];  
}
```

the **for** loop becomes a **do** instruction, the array references are promoted to address registers and the induction variable is eliminated with array pointer advancement done via the **autoincrement** addressing mode.

4.11.12 Leaf Routine Detection

A leaf routine is a subroutine that makes no further subroutine calls. When the compiler identifies such routines, the prologue and epilogue code are optimized (no save and restore of the **ssh**).

4.11.13 Function In-lining

When explicitly requested via the command-line option **-finline-function**, the compiler will replace calls to subroutines with an in-line copy of the subroutine if the subroutine meets these requirements:

1. The subroutine must be a non-volatile leaf function.
2. The subroutine must be in the same module.
3. The definition must precede use of the subroutine.

Function in-lining eliminates the overhead associated with subroutine calls and provides additional opportunities for the optimizer and register allocator to further increase code quality. Function in-lining can also be performed explicitly by the programmer by utilizing the additional non-ANSI function type specifier **__inline**. By default, many run-time libraries are in-lined by the compiler.

Note: The function in-lining method can cause program memory requirements to grow significantly. See Appendix A, *Programming Support*, for instructions on disabling library routine in-lining.

4.11.14 Instruction Scheduling / Microcode Compaction

The command line switch **-alo** causes an assembly language optimizer (**alo56**) to be run, using the assembly code emitted by the compiler as input. This optimizer attempts to compact multiple operations into a single instruction word, while simultaneously avoiding the pipeline hazards exposed by the address generation unit. Because this optimizer mixes together instructions from different C language statements, debugging code compiled with **-alo** may be more difficult.

Chapter 5

Mixing C and Assembly Language

5.1 Overview

In cases where the DSP56000/DSP56001 programmer requires direct access to the hardware or greater performance at the inner-loop of an algorithm, C programs can be mixed with assembly programs. This chapter describes two methods for combining C and assembly language source code. The first is the in-line method which is a convenient way to put assembly code into the compiler output via the non-ANSI C directive `__asm()`. The second is the out-of-line method, a generic method of combining C and assembly language source files.

Warning

Before mixing C and assembly, read and understand Chapter 4, About G56k, and the DSP56000 Family Manual.

Attempting to write programs for this DSP without knowledge of the chip and how the compiler utilizes registers, memory, defines labels, etc. may generate unsatisfactory results.

However, with an understanding of the DSP architecture and how this implementation of C uses it, programming should be straightforward.

Labels which begin with a double underline (e.g., `__asm()`) in this manual have a space between the double underlines to visually separate them. Do not separate the leading double underlines with a space when coding them (i.e., `code __asm() as __asm()`).

5.2 In-line Assembly Code

In-line assembly code is assembly code that is inserted inside a C statement in a C source file. Since assembly code is generated from this C statement directly, the C statement looks like assembly code in the C source and is referred to as *in-line* assembly code. All of the assembly code to be generated is visible at the C source-level and it is often convenient to intermix assembly code with a C source program in this fashion.

Typically, in-line assembly code is used when:

1. Inserting a small amount of assembly code directly into the compiler output i.e., inner loop kernels.
2. Writing fast assembly language routines to be called by C subroutines. This eliminates the need to manage data referencing, register saving and allocation, and function call/return overhead.

The key word `__asm` is introduced as an extension to the ANSI C language standard and this key word is used like a function call to specify the in-line assembly code generation rule discussed below.

The in-line assembly statement syntax is:

```
__asm ("instruction_template": output_operands: input_operands:  
reg_save);
```

where:

1. **instruction_template** is a string used to describe a sequence of assembly code instructions that are to be inserted into the compiler output stream. It may specify arguments, which are listed in `output_operands` and `input_operands`. It does this via a substring called an operand expansion string (OES). An OES starts with a '%'. OES and `instruction_template` interpretation is described in Section 5.2.1.
2. **output_operands** are optional operands used to provide specific output information to the compiler. Each `output_operand` string is separated by a comma and should begin with the character '='. As an example, the `output_operand` "=A" (cptr) means "the C variable cptr will get its value from this output operand, which will be in an address register". See Section 5.2.2 for more details.
3. **input_operands** are optional operands to provide specific input information to the compiler. Each `input_operand` is separated by a comma and may include a C variable. As an example, the `input_operand` "A" (cptr) means "the value of this input operand will be taken from the C variable cptr, and placed in an address register". Again, full descriptions of the input and output operands can be found in Section 5.2.2.
4. **reg_save** specifies registers that are to be explicitly reserved for the `__asm ()` statement. The registers to be saved must be named as strings such as "a" and "b". Each register is separated by a comma (see Section 5.2.3 for additional information) The compiler assumes that all data residing in the `reg_save` registers will be invalidated by the `__asm ()` statement.

5.2.1 Instruction Template

The first argument of the `__asm()` extension is the instruction template or assembler instruction template. This instruction template is mandatory, and describes the line or lines of assembly language to be placed into the compiler's assembly language output (see Section in Section 5.2.4). This template is not parsed for assembly language syntax violations and is simply written to the compiler output. As a result, the compiler will not detect assembly-time errors. These errors will be caught by the assembler.

More than one assembly instruction can be put into a single instruction template by using the line separator “\n”. The line separator, or newline, can be utilized as in a normal C statement. The line continuation sequence, a “\” followed by an immediate newline, is particularly useful when an instruction template contains an assembly instruction that is too long to fit in one source line (see Section 5-18 and Section in Section 5.2.4). Other C language character constants such as “\t”, “\f”, etc. can also be used in the instruction template.

In many situations, it is desirable to use the values of C variables and/or expressions directly in the instruction template. Since all memory and register accesses are accomplished through variables, manipulating memory and registers directly using assembly code requires knowledge of their locations. Without optimizations, the current value of a variable will be maintained in memory at a specific address. However, an optimizing C compiler such as DSP56KCC may retain a variable in a register and perform operations on that variable without updating the memory location corresponding to the variable between operations. This enhances the performance of the compiled code but makes accessing variables in memory risky. In order to guarantee that the correct value of a variable is returned when it is referenced, a mechanism called operand expansion string (OES) is provided.

The OES allows a variable to be securely accessed even though its current location is unknown. The operand expansion string is a substring of the instruction template and begins with the character “%”. This string is usually two or three characters long and provides the compiler with special information about an operand, and how its reference should be printed. An OES must reference only one C variable or expression, which in turn must be listed in either one or both operand lists (see Section 5.2.2). The OES is parsed by the compiler and gives sufficient information to allow the variable to be correctly referenced by the assembly language instruction in the instruction template. Most examples in Section 5.2.4 include an OES.

The OES syntax is:

```
% [modifier] operand_id_number
```

where:

1. **modifier** is a single optional character which specifies a particular class of operand. The available modifiers are 'j', 'e', 'h', 'k', 'g', 'i', 'f', 'p', and 'q'.
 - j**—an offset register (**nx**) associated with the corresponding address registers (**rx**). Since the offset register is paired with the address register, the allocated offset register has the same index as the address register (see Section in Section 5.2.4).
 - e**—**a1** or **b1**, upper word of the destination registers **a** or **b** (see Section in Section 5.2.4).
 - h**—**a0** or **b0**, lower word of the destination registers **a** or **b** (see Section in Section 5.2.4).
 - k**—**a2** or **b2**, extension register of the destination register, **a** or **b** (see Section in Section 5.2.4).
 - g**—Select the 24-bit portion of the 48-bit ALU register (**x** or **y**) that is not occupied by data pointed to by the operand id — e.g., if the operand id points to **x0** then **x1** is selected and similarly **x1**→**x0**, **y0**→**y1**, **y1**→**y0** (see Section and Section 5-9 in Section 5.2.4).
 - i**—strip the **0** or **1** from the allocated register name i.e., **a1**→**a**, **a0**→**a**, **b1**→**b**, **b0**→**b** (see Section in Section 5.2.4).
 - f**—insert the memory space identifier (**x** or **y**) for the memory location (see Section in Section 5.2.4).
 - p**—generate an immediate 16-bit constant without # sign (see Example in Section 5.2.4).
 - q**—generate an immediate 24-bit constant without # sign (see Section in Section 5.2.4).
2. **operand_id_number** specifies the operand location in the operand descriptor list (see Section in Section 5.2.4). The operand descriptor list is a concatenation of the output operands and input operands (see Section 5.2.2). The first operand is labeled zero and there can be up to 31 more operands in the list. More than one instruction template can be used if more than 32 operands are needed.

In-line assembly code can also be used to insert a label directly into the compiler output. A label without any white spaces in the in-line assembly code instruction template will guarantee that the same template label will be in the compiler output (see Section). Care should be taken not to use labels that the C compiler generates. Using the same labels as the C compiler will cause a duplicate label error (see Section 4.7, "Compiler Naming Conventions," on page 4-13.)

5.2.2 Output/Input Operands

The operand list is a concatenation of output and input operands which the OES can access via the `operand_id_number` (see Section 5.2.1). Output or input operands consist of operands separated by a comma (`,`). Each operand should be associated with a C expression and its operand constraint described below.

A colon, `:`, is used to separate the assembler instruction template from the first output operand. A second colon separates the final output operand from the first input operand. A third colon can be used to separate the input operands from the optional field `reg_save`. Two consecutive colons are required when only input operands are needed in the operand list, leaving us with the empty list of output operands.

The operand syntax is:

```
" [= ]operand_constraint" (C_expression)
```

where:

1. `=` differentiates input and output operands. All output operands must use this character first.
2. **operand constraint** is a single character that describes the type of resource (memory or register) that an operand is to be loaded into or read from. Each operand constraint has an optional set of modifiers that may be applied in the instruction template.
3. `C_expression` is any valid C expression defined by the ANSI C standard. The C expression can be either l-value or r-value. Any output operand should use the l-value to specify the memory location to store the data.

The available operand constraints are “**A**”, “**D**”, “**R**”, “**S**”, “**i**” and “**m**”. All of these constraints originate from the DSP56000 family architecture; therefore, a full understanding of these constraints requires that the programmer understand this architecture. The constraints are:

A —One of the Address Registers (**rx**, where $x = 0$ through 7; see Section 5 of the *DSP56000 Family User’s Manual*) will be allocated to the C expression (see “OES modifier j: The following in-line assembly code is used to generate executable assembly code. Notice that the actual register selection is totally dependent on the C compiler but the register selected (r3 in this example) is guaranteed to be related to the C expression used (in this case `cptr`, see Section 5.2.)” on page 10). The C expression will be promoted to this register. Typically the C expression should be a pointer to be assigned to an address register. The OES modifier, **j**, can only be associated with operand constraint **A** (see Section 5.2.1).

D —One of the 56-bit accumulators (**a** or **b** which are referred to as Destination Registers; see Section 4 of the *DSP56000/DSP56001 User's Manual*) will be allocated to the C expression (see “OES modifier e: The modifier e can be used to generate the assembly code below because a1 is the upper part of register A.” on page 10 through “OES modifier k: The k modifier can be used to generate the following assembly code because a2 is the extension portion of register A.” on page 11). The C expression will be promoted to this register. Care must be exercised to make sure the size of the C expression is consistent with the accumulator length. The OES modifiers ‘**e**’, ‘**h**’ and ‘**k**’ can be associated with operand constraint D (see Section 5.2.1).

R —One of the Input Registers (**x0** or **y0** which are also called Multiplier Registers; see Section 4 of the *DSP56000/DSP56001 User's Manual*) will be allocated to the C expression (see “OES modifier i: The modifier can be used to generate the following assembly code because x is a register without a 0 or 1 portion.” on page 12). The C expression will be promoted to this register. The OES modifiers ‘**g**’ and ‘**i**’ can only be associated with operand constraint R (see Section 5.2.1).

S —One of the Input Registers (**x0**, **x1**, **y0** or **y1** which are also called Source Registers; see Section 4 of the *DSP56000/DSP56001 User's Manual*) will be allocated to the C expression (see Example 5-13 on page 5-14 through Example 5-17 on page 5-15). The C expression will be promoted to this register. The OES modifiers “**g**” and “**i**” can only be associated with operand constraint **S** (see Section 5.2.1).

i —An immediate constant; a constant is generated in the form of #constant if no modifier is specified. With **p** or **q** modifier, the value is generated without the # sign. The value is the 16-bit constant with **p** modifier, and 24-bit constant with **q** modifier (see Example 5-12).

m —A memory location will be allocated to the C expression (see Example 5-11). The DSP56000/DSP56001 has four memory spaces: X, Y, L and P but the C compiler may only use the X, Y or L memory spaces for this expression. The OES modifier ‘**f**’ can only be associated with operand constraint m (see Section 5.2.1).

number —Inherit all memory or register characteristics of the operand indicated by the **operand id number** (see Example 5-3). This constraint is usually used for read/write operands which are declared by using the same C variable as both the input and output operand.

The operand is sometimes referred to as a read-only operand if it is used only as input (see Example 5-14) and it is called a write-only operand if it is used only as an output (see

Example 5-15). In most of cases, the operand is used as both an input and an output operand (see Example 5-16 and Example 5-17). In these cases the operand should be described as both. Since output operands should be listed first, the operand id number is determined when the input operand is declared. The id number will be used as the operand constraint of the input operand.

Section lists the operand constraints and their related modifiers.

Table 5-1. Operand Constraints/Modifiers Associations

Operand Constraint		Modifier	
A	-Address Register (rx)	%j	-Rn's paired offset register
D	-Destination Register (a,b)	%e %h %k	-Upper word (a1, b1) -Lower word (a0, b0) -Extension (a2, b2)
R	-Multiplier Register (x0,y0)	%g %i	-Select other register portion -Strip 0/1 from register name
S	-Source Register (x0,x1,y0,y1)	%g %i	-Select other register portion -Strip 0/1 from register name
i	-Constant (#number)	%p %q	-16 bit immediate value - 24 bit immediate value
m	-Memory Location	%f	-Memory space identifier
number	-Any one of above	%(corresponding modifier	

5.2.3 Explicit Register Saving

It is possible to manually perform register allocation, in fact this may simplify the process of converting an existing body of DSP56000/DSP56001 assembly language subroutines to in-line assembly code. The programmer need only identify each register explicitly referenced in the assembly code and list them in the **reg_save** argument region (see Section 5.2). This guarantees that the content of each register is preserved across **__asm()** calls. Use of registers **r0** and **r6** is prohibited in the assembly code because they are reserved for the C compiler during run-time where **r0** is the frame pointer and **r6** is the stack pointer. Registers **Mn** are used by the compiler as temporary registers and **Nn** are not used by the C compiler at all. As a result, these registers do not need to be saved unless the programmer uses them in assembly code. If they are used in assembly code, they should only be used as local variables.

Explicit register saving is done through specifying the registers to be saved. A string is used to specify each register. The valid register names are listed in Table 5-2.

Table 5-2. reg_save Names

Register Type	Valid Names
Accumulator	"a","b"
Input	"x0", "x1", "y0", "y1"

An example for **reg_save** is a program performing an IIR filter operation which is based on the Dr. BuB IIR filter program "iir1.asm". The program data structure is the same as the Dr. BuB program and data is passed through the C variable "data". Pointers "sptr" and "cptr" are used to point to static variables.

Example 5-1. reg_save

reg_save: This in-line assembly code is converted from the Dr. BuB IIR filter program using **reg_save**. This program is based on Y Memory Model, and there are other in-line assembly features used here for the demonstration of the reg save. Section 5.2.4 may need to be referenced for the other feature of the in-line assembly.

```

iir1( int data, int* sptr, int* cptr )
{
    __asm volatile ( "move %0,a\n\
move %1,r1\n\
move %2,r4\n\
move y:(r1)+,x0\n\
move y:(r4)+,y0\n\
mac x0,y0,a          y:(r4)-,y0\n\
move y:(r1),x1\n\
macr x1,y0,a         x0,y:(r1)-\n\
move a,y:(r1)" : /* no output */ : "D" (data), "A"
(sptr), "A" (cptr) :
    "a", "x1", "x0", "y0", "r1", "r4" );
}

```

5.2.4 In-line Assembly Code Examples

The examples in this section illustrate the practical application of the **__asm()** extension. The main purpose of this section is to show how to write in-line assembly code. Since these examples are intended to illustrate the information presented earlier in this chapter, references to the appropriate subjects have been included.

Example 5-2 illustrates the use of the in-line assembly code "instruction_template". Since this in-line assembly code directly clears register A, the programmer should check to be sure that the contents of A are not needed.

Example 5-2. instruction_template

Instruction_template: The following are a few examples of how to utilize the instruction template in in-line assembly code. This feature allows the generation of any valid assembly instruction and it is probably the most frequently used feature in in-line assembly coding.

```

_ _asm("clr    a"); /* clears the register A */
_ _asm("move  #$10, a2"); /* load the register A2 with
the hex value 10 */
_ _asm("HCR   EQU  $FFE8"); /* equate the Host Control
Register to
$FFE8 */

```

A pseudo operand will be used to illustrate use of the OES operand id number. The pseudo operand functions as an input or output operand. Example 5-3 uses five pseudo operands: V, W, X, Y and Z each of which is referenced by operand ids **0**, **1**, **2**, **3** and **4**, respectively. The pseudo operands are used as in the OES “%0”, “%1”, “%2”, “%3” and “%4”. Table 5-3 shows which operands in Example 5-3 are input or output operands.

Example 5-3. operand_id

Instruction template with **operand_id**: In order to illustrate how to use output or input operands, pseudo operands V, W, X, Y, and Z are used. The **operand_id** listed in this example can be used as part of an **instruction_template**.

```

_ _asm("instruction_template" : V, W, X : Y, Z );

```

Example 5-3 through Example 5-12 illustrate the use of OES modifier (see Section 5.2.1).

Table 5-3. Output and Input Operands for Section 5-3

operand id	pseudo operand	operand type	OES
0	V	output	%0
1	W	output	%1
2	X	output	%2
3	Y	input	%3
4	Z	input	%4

Example 5-4. OES Modifier j

OES modifier **j**: The following in-line assembly code is used to generate executable assembly code. Notice that the actual register selection is totally dependent on the C compiler but the register selected (r3 in this example) is guaranteed to be related to the C expression used (in this case `cptr`, see Section 5.2).

In-line Assembly code:

```
char *cptr;  
_ _asm("move (%0)+%j0"::"A"(cptr));
```

Assembly Code Generated:

```
move (r3)+n3
```

Example 5-5. OES Modifier e

OES modifier **e**: The modifier **e** can be used to generate the assembly code below because `a1` is the upper part of register A.

In-line Assembly code:

```
int foo;  
_ _asm("move #ffffff,%e0"::"=D"(foo));
```

Assembly Code Generated:

```
move #ffffff,a1
```

Example 5-6. OES Modifier h

OES modifier **h**: The **h** modifier can be used to generate the following assembly code because `a0` is the lower part of register A.

In-line Assembly code:

```
int foo;  
_ _asm("move #ffffff,%h0"::"=D"(foo))
```

Assembly Code Generated:

```
move #ffffff,a0
```

Example 5-7. OES Modifier k

OES modifier **k**: The **k** modifier can be used to generate the following assembly code because a2 is the extension portion of register A.

In-line Assembly code:

```
int foo;
__asm("move #fff,%k0" : "=D" (foo));
```

Assembly Code Generated:

```
move #fff,a2
```

Example 5-8. OES Modifier g

OES modifier **g**: Swap the most significant 24-bit portion and the least significant 24-bit portion of 48-bit registers x and y to allow the OR instruction to operate on an entire 48-bit register.

```
/*
 * The following assembly code could be generated
 (note that the
 * optimizer may vary the code actually generated).
     move x1,a1
     move x0,x1
     move a1,x0
 *
 * The variable foo can be allocated to either x0,
 x1, y0, or y1
 * by using the operand constraint S. The swap
 operation can
 * be applied to the register allocated to the
 variable foo by
 * using the following in-line assembly code.
 *
 */
main()
{
    int foo;
    __asm volatile ("move %g0,a1" : : "S" (foo));
    __asm volatile ("move %0,%g0" : "=S" (foo) : "0"
(foo));
    __asm volatile ("move a1,%0" : "=S" (foo));
}
```

Example 5-9. OES Modifier g

OES modifier **g**: Program bit checker looks to see if any bit in the 48-bit registers **x** or **y** is a one. In this case, Bit Checker looks to see whether the variable **foo** which is placed in either the **x** register is all zeros or contains a one. The result is stored in register **a1**. If register **a1** is not **0x000000**, then **foo** has one or more bits set to one.

```

/*
 * The variable foo can be allocated to either the x or
 y register by using
 * the operand constraint S. The OR instruction only
 operates on 24-bit
 * registers so that to OR the x register with another
 register, x1 must
 * be ORed separately from x0. The same applies for
 the y register.
 */
main()
{
    long volatile foo;

    _asm volatile ("clr a");
    _asm volatile ("or %0,a" :: "S" (foo));
    _asm volatile ("or %g0,a" :: "S" (foo));
}

```

Example 5-10. OES Modifier i

OES modifier **i**: The modifier can be used to generate the following assembly code because **x** is a register without a 0 or 1 portion.

In-line Assembly code:

```

int foo;
_ _asm("move l:<$0,%i0" : "=R"(foo));

```

Assembly Code Generated:

```

move l:<$0, x

```

Example 5-11. OES Modifier f

OES modifier **f**: The **f** modifier can be used to generate the following assembly code. Assuming that Y memory space is used and the memory location of the variable “foo” is 233, then the desired memory space indicator “y” will be automatically generated by the **f** modifier.

In-line Assembly code:

```
int foo;
__asm("move #ffffff,%f0" : "=m" (foo));
```

Assembly Code Generated:

```
move #ffffff,Y:233
```

Example 5-12. OES Modifier p and q

OES modifier **p**, and **q**: This in-line assembly code programs the SCI of DSP56001 by setting up the SCI registers located at X:\$FFF0 and X:\$FFF2. You may use modifier **p** for any 16-bit value.

```
#define SCR      0xFFF0
#define SCCR     0xFFF2
#define V_SCR    0x2000
#define V_SCCR   0x013F
main()
{
  __asm("movep %0,x:%q1" : : "i" ( V_SCR), "i" (SCR));
  __asm("movep %0,x:%q1" : : "i" ( V_SCCR), "i" (SCCR));
}
```

Assembly Code Generated:

```
movep    #>$002000,x:65520
movep    #>$00013f,x:65522
```

Example 5-13. Input Expression / Output Expression

Input Expression / Output Expression: This in-line assembly code uses the pseudo assembly mnemonic “**asm_instruction**” and refers to two C expressions: `output_expression` and `input_expression`. This example illustrates how to interpret the operand constraint (see Section 5.2.2) and operand id (see Section 5.2.1 and Example 5-3). The example implies that the C expression `output_expression` is expanded with constraint D and is an output of the assembly code instruction `asm_instruction`. Similarly, the C expression `input_expression` is expanded with constraint S and used as an input to the assembly code instruction `asm_instruction`.

```
_ _asm("asm_instruction %1,%0" : "=D"  
      (output_expression) : "S" (input_expression));
```

Example 5-14. Read-Only Operand

Read-Only Operand: This in-line assembly code uses the pseudo assembly mnemonic “**asm_instruction**” and uses `input_expression` as a read-only operand.

```
_ _asm("asm_instruction %0" :: "S" (input_expression));
```

Example 5-15. Write-Only Operand

Write-Only Operand: This in-line assembly code uses the pseudo assembly mnemonic “**asm_instruction**” and uses `output_expression` as a write-only operand.

```
_ _asm("asm_instruction %0" : "=D" (output_expression));
```

Example 5-16. Read-Write Operand

Read-Write Operand: An addition is programmed using in-line assembly code and the C expression result is used as a read-write operand. The variable, `foo`, is used as a read only operand. Notice that operand constraint ‘0’ was used to reference the add instructions second source operand which is also the destination operand (see the *DSP56000/DSP56001 User’s Manual* — Appendix A for the syntax of the add instruction).

```
int foo, result;  
_ _asm("add %1,%0" : "=D" (result) : "S" (foo), "0"  
      (result));
```

Example 5-17. Read-Write Operand

Read-Write Operand: The same result will be obtained as in Example 5-16. Notice how the operand id is changed according to the placement of the C variables.

```
int foo, result;
__asm("add %2,%0" : "=D" (result) : "0" (result), "S"
(foo));
```

Example 5-18. Multiple Instruction — Single-Line

Multiple Instruction — Single-Line: An in-line assembly program which places a value (e.g. \$709) in register a and negates the result is written in one line. This one line will generate two lines of assembly code in the C compiler output.

```
__asm("move #$709,A\n neg A");
```

Example 5-19. Multiple Instruction — Multiple-Line

Multiple Instruction — Multiple-Line: The two lines of in-line assembly code in this example have the same effect as the one line in Section 5-18. Notice that using two lines increases the in-line assembly code readability. The line continuation character ‘\’ used at the end of the in-line assembly codes first line makes this possible.

```
__asm("move #$709,A\n\
neg A");
```

Example 5-20. Multiple Use of __asm()

Multiple use of `__asm()`. This example and Example 5-21 are done in-line with the compiler performing all register allocation and all operands are referenced via C expressions. This in-line program is a single memory space version of the Dr. BuB IIR filter program `iir1.asm`. The method used to write this in-line assembly program is to use an `__asm()` statement for each assembly language instruction.

```
int iir1( int data, int* sptr, int* cptr )
{
    int state1, state2, coef;
    __asm ( "move y:(%1)+,%0" : "=R" (state1) : "A" (sptr)
);
    __asm ( "move y:(%1)+,%0" : "=R" (coef) : "A" (cptr)
);
    __asm ( "mac %2,%3,%0\ty:(%4)-,%1"
: "=D" (data), "=R" (coef)
: "R" (state1), "1" (coef), "A" (cptr), "0" (data) );
    __asm ( "move y:(%1),%0" : "=S" (state2) : "A" (sptr)
);
    __asm ( "macr %1,%2,%0\t%3,y:(%4)-"
```

```

: "=D" (data)
: "S" (state2), "R" (coef), "R" (state1), "A" (sptr) );
- _asm ( "move %0,y:(%1)" : : "D" (data), "A" (sptr) );
return data;
}

```

Example 5-21. Line Separation

Line Separation. This in-line program is functionally identical to Section except that line separation is used to insert the entire assembly language program (Dr. BuB IIR filter program, iir1.asm) into a single `__asm()` statement. Notice how much easier it is to read the program.

```

int iir1( int data, int* sptr, int* cptr )
{
    int state1, state2, coef;
    __asm ( "\n\
move y:(%4)+,%2\n\
move y:(%5)+,%1\n\
mac %8,%7,%0\ty:(%5)-,%1\n\
move y:(%4),%3\n\
macr %9,%7,%0\t%8,y:(%4)-\n\
move %6,y:(%4)":
state2):
"A" (sptr), "A" (cptr), "0" (data), "1"
(coef),"2" (state1),
"3" (state2) );
return data;
}

```

Example 5-22. Instruction Template Label

Instruction Template Label: The following in-line assembly code which generates the label "foo" uses a return character "\n" to insure that there is no white space in front of the label.

```

__asm("\nfoo");

```

Example 5-23. Program DSP56001 SCI Timer

Program DSP56001 SCI Timer: The SCI timer can be programmed so that a SCI interrupt service routine is accessed periodically. The following in-line assembly program is based on the SCI timer described in section 11.2 of the DSP56000/DSP56001 User Manual.

```

_ _asm( "\nSCR      EQU      $FFF0" );
_ _asm( "\nSCCR    EQU      $FFF2" );
_ _asm( "\nIPR     EQU      $FFFF" );
main()
{
_ _asm( "move      #0,r7\n\
      movep     #$2000,x:SCR\n\
      movep     #$013f,x:SCCR\n\
      movep     #$C000,x:IPR\n\
      andi      #$fc,mr\n\
\nEND      jmp      END" );
_ _asm( "org      p:$001c\n\
      move      (r7)+\n\
      nop" );
}

```

5.2.5 Controlling Labels Generated by the Compiler

Using the `__asm()` keyword, it is possible for the programmer to override the compilers label generation conventions for subroutines and global variables. This may be useful for:

1. Calling assembly language subroutines.
2. Calling C subroutines from assembly code.
3. Referencing assembly language global variables from C.
4. Referencing global C variables from assembly code.

5.2.5.1 Calling Assembly Subroutines

Calling a subroutine or function requires using a label that points to the subroutine or function. The C compiler uses a predetermined labeling convention (see Section 4.7, "Compiler Naming Conventions," on page 4-13), and does not generate arbitrary labels like most assembly programs. In order to call assembly subroutines labeled in an arbitrary fashion, `__asm()` can be used to overwrite the C convention label with an arbitrary label.

To illustrate how to use the `__asm` directive for this purpose, Section 4-7 reads the data at x memory location \$100 and y memory location "y+2". For test purposes, the y memory space is filled with the integer sequence 0 through 9 and the x memory space is left uninitialized. The `printf()` statement prints the data returned from the functions `ValOfX(100)` and `ValOfY(y+2)`. These two functions are constructed in assembly code and reside in another file. The two assembly subroutines, `ReadX` and `ReadY`, are written

using the out-of-line assembly technique described in Section 5.4. The listings for ReadX and ReadY are shown in Example 5-38 in Section 5.4.5.

By using the statement

```
extern int ValOfX() _ _asm("ReadX"), ValOfY() _
_asm("ReadY");
```

all C compiler generated function labels for **ValOfX()** and **ValOfY()** are replaced by the labels ReadX and ReadY, respectively.

Example 5-24. Calling Assembly from C

Calling assembly from C. This C program (called test.c) can be used to examine the data in x or y memory by calling assembly routines "ReadX" or "ReadY". Notice that the assembly code for ReadX and ReadY is listed in Section 5-38 of Section 5.4.5.

```
C:\> type test.c

#include <stdio.h>
extern int ValOfX() _ _asm("ReadX"), ValOfY() _
_asm("ReadY");
unsigned Y[] = {0x1,0x2,0x3,0x4,0x5,0x6,0x7,0x8,0x9};
main()
{
    printf("<%x><%x>\n", ValOfX(100), ValOfY(Y+2));
}
```

The following two command lines test Section .

```
C:\> g56k test.c memread.asm
C:\> run56 a.cld
```

5.2.5.2 Calling C Subroutines from Assembly Code

Any C program can be called from an assembly program to test the assembly program data or utilize built-in standard C libraries such as floating-point operations. Calling a C subroutine from assembly code requires using the C subroutine calling convention (see Section 4.8, "Subroutine Call Sequence," on page 4-14 and Section 5.4.4, "Calling C Routines," on page 5-35) and matching the C function labels. The in-line-assembly directive, `_ _asm()`, can be used as shown in Section to change the C program labels.

Example 5-25. Calling C from Assembly

Calling C from assembly. This C subroutine (called `C_print.c`) uses the standard C library routine, `printf()`, to print the input argument as a string.

```
C:\> type c_print.c
#include <stdio.h>
int C_printf() __asm("print");
C_printf(char *msg)
{
printf("%s\n", msg);
}
```

Example 5-26. Calling C from Assembly

Calling C from assembly. This assembly program (called `greeting.asm`) prints the message “greeting: hello, there” on the screen. It uses the C subroutine, `printf()`, to print this message. Notice that the assembly program name is `Fmain` because the control program, `G56k`, uses the default start-up file `crt056y.cln` and `crt056y.cln` uses `Fmain` to start up the main program (see Chapter 6 in this manual).

```
C:\> type greeting.asm

        section greeting
        org      y:
LC0    dc      "greeting: hello,there.", $00
        org      p:
        global  Fmain
Fmain  move     ssh,y:(r6)+
        move     r1,y:(r6)+
        move     #LC0,r1
        move     r1,y:(r6)+
        jsr      print
        move     (r6)-
        move     y:-(r6),r1
        move     y:-(r6),ssh
        rts
        endsec
```

The following two command lines are used to test two programs: `c_print.c` and `greeting.asm`.

```
C:\> g56k greeting.asm c_print.c
C:\> run56 a.cld
```

5.2.5.3 Referencing Assembly Global Variables from C

The data in assembly language programs must be accessible to C programs to take full advantage of the DSP56000 family architecture since the C language cannot access all of the DSP56000 features directly. One way to access this data is through global data which can be defined in assembly language and accessed in the C program environment. This feature is particularly useful to allocate modulo buffers for C variables. Detailed information on modulo buffers can be found in the *DSP56000/DSP56001 User's Manual*, Section 5 — Address Generation Unit and Address Modes.

Example 5-27. Generate Data with Assembly Language

Generate data with assembly language. The data file, `sqtbl.asm`, is generated in assembly language and consists of the squares of the numbers 0 - 8. Notice that directives BSC, DC, DS, DSM and DSR (see Section 6.3 of the *Motorola DSP56000 Macro Assembler Reference Manual*) can be used depending on the application.

```
C:\> type sqtbl.asm

    section data
    global  table
    org    y:
table  dc 0,1,4,9,16,25,36,49,64
    endsec
```

Example 5-28. Access Data with C

Access data with C. This test program (called `test.c`) prints the value of 5^2 on the screen.

```
C:\> type test.c
#include <stdio.h>
extern int SQUARE[] _ _asm("table");
main()
{
printf("square of %d is %d\n", 5, SQUARE[5]);
}
```

The following two command lines for Example 5-28 test the two programs `sqtbl.asm` and `test.c`.

```
C:\> g56k test.c sqtbl.asm
C:\> run56 a.cld
```

5.2.5.4 Referencing Global C Variables from Assembly Language

One DSP56KCC feature is that global data in a C program is available to assembly language programs. This feature is particularly useful when the data to be processed by an assembly language program is generated by the C program. Example 5-29 provides coefficients that are used in Example 5-30.

Example 5-29. Generate Data with C

Generate data with C. Data file, `data.c`, is generated by a C language program and contains the coefficients of an average filter which takes the average of the last four input data.

```
% cat data.c
/*
 * data.c
 */
int Cwaddr[]    __ _asm("waddr");
int Ccaddr[]    __ _asm("caddr");
int NTAP        __ _asm("N_1");

int Cwaddr[4];
int Ccaddr[] = { 0x200000, 0x200000, 0x200000, 0x200000
};
int NTAP = 4;
```

Example 5-30. 4-tap FIR Filter

4-tap FIR filter — Access data with assembly language: This FIR filter reads an input sample from memory location `Y:input` and writes the filtered output to memory location `y:output`. The input data array is stored in `X` memory starting at `waddr` and the FIR coefficients are stored in `Y` memory starting at `caddr`. Notice that the memory space for `waddr` and `caddr` is allocated in the C routine described in Example 5-29.

```
% cat fir.asm
org      p:
move     #waddr,r0
move     #caddr,r4
move     y:N_1,m0
move     m0,m4

movep    y:input,x:(r0)

clr      a          x:(r0)+,x0y:(r4)+,y0
rep      #n-1
```

```
mac      x0,y0,a x:(r0)+,x0y:(r4)+,y0
macr     x0,x0,a (r0)-

movep   a,y:output

end

C:\> g56k -S data.c
C:\> g56k -c fir.asm
C:\> g56k -c data.asm
C:\> g56k fir.cln data.cln
```

5.2.6 Specifying Registers for Variables

DSP56KCC allows the programmer to identify a specific register for local and global variables, but due to the limited number of registers available, this may not have a positive effect on run-time performance. With this in mind, this feature should be used sparingly.

Both global and local variables are candidates for promotion to specific registers and syntactically they look the same:

```
register int *ptr _ _asm("r5");
```

By specifying a specific register for a local or global variable, the programmer is reserving the register for the variable's entire scope (global for the entire program, local for the function in which they are declared). This means that the compiler will not use the register for any other purpose and the register **will not be saved and restored** by the C function call.

5.2.7 Optimizer Effects on Code

All in-line assembly code is visible to the optimizer and as such it is possible that the optimizer will convert it into a new form or eliminate it entirely if it is determined to be unreachable or dead. In order to guarantee that code is not removed by the optimizer, the ANSI keyword `volatile` must be used.

```
_ _asm volatile ( ... );
```

5.3 #pragma Directive

The purpose of this section is to explain the proper techniques for manipulating the assembler's run-time and load-time counters while programming in the C language.

Currently the Motorola DSP assemblers allow the programmer to specify both a run-time location and a load-time location for any object; however, there is no corresponding concept within C. The generic `#pragma` facility is used to add this capability rather than

extending the C language. Users now have complete freedom in specifying both the run-time and load-time counters for any static or global object. These directives may be used with either code or data.

This flexibility is achieved by allowing the user to modify any of eight counter strings maintained by the compiler — two for each memory space: x, y, l and p. When an object is or defined, the current values of those counter strings are bound to that object.

Syntax for the pragma directive is

```
#pragma counter_string argument
```

C function or data storage definition

```
#pragma counter_string
```

where

1. The two **#pragma** statements must encase the entire definition.
2. **counter_string** in the first **#pragma** specifies which phase (run or load time) and memory space is to be affected. It can be **x_run**, **y_run**, **l_run**, **p_run**, **x_load**, **y_load**, **l_load**, or **p_load**. The memory model used in the C compiler should match the memory model specified by **counter_string**. Note that the **p_run** and **p_load counter_strings** will always have an effect.
3. The argument in the first **#pragma** is the string that will be passed as either the runtime or load-time argument to the `org` assembler directive. This address, which is optional, is of the form **x:address_value** where x is the counter associated with the **x**, **y**, **l**, or **p** memory, and **address_value** is the value to be initially assigned to that counter. As an example, `p: $300` might be used for the counter string **p_load**.
4. The C function or data storage definition is a declaration that reserves storage.
5. The second `counter_string` should be the same as the first **counter_string** and will return the memory specification to the default setting.

If and only if *the memory space of the counter string in the #pragma directive matches the memory model of the C compiler*, then the compiler will insert an assembly `org` statement of the form:

```
(1)      org      a:runtime_address,b:loadtime_address
```

or

```
(2)      org      a:runtime_address
```

where “a” is the run time counter and **runtime_address** is the optional initial value for that counter, as specified in the “argument” to **#pragma**.

“b” is the load time counter and **loadtime_address** is the optional initial value for that counter, as specified in the “argument” to **#pragma**.

The following two examples illustrate that the load time counter is optional. See the section on the ORG statement in the *Motorola DSP56000/DSP56001 Assembler Manual* for a complete description and list of options.

Notice that the pragma directive run-time counter string will only affect the run-time address and the pragma directive load-time counter string will only affect the load-time address.

As a simple example, assuming that y-memory model (or default memory model) is used (see Chapter 3 to change the memory model to be used), the following C segment

```
#pragma y_load x:$100
int coeff[5] = {0x19999a, 0x200000, 0x266666, 0x2ccccd,
0x333333};
#pragma y_load
```

produces the following assembly language code:

```
global Fcoeff
orgy: ,x:$100
Fcoeff
dc 1677722
dc 2097152
dc 2516582
dc 2936013
dc 3355443
```

Notice that the second **#pragma** directive will remove the affect of the first memory specification, i.e., **#pragma y_load x:\$100**, and the rest of the code generated by the C compiler will be in the default memory area (see Chapter 3 in order to change this default memory model).

The above example code will be loaded in the X memory location \$100, and it should be copied to the Y memory space upon system start-up. When burning a PROM, only one memory space is desired to be used, as an example, P memory space, so that only one PROM is enough for both data and program. In such case, both the data and the program will be burned in the PROM and the data should be moved to the data memory space (either, X, Y or L memory space depending on the program) upon system start-up.

Let's assume that the coefficients of the above example is desired to be in the program space when burning the PROM. Then the following C segment

```
#pragma y_load p:$100
int coeff[5] = {0x19999a, 0x200000, 0x266666, 0x2cccd,
0x333333};
#pragma y_load
```

produces the following assembly language code:

```
global    Fcoeff

org      y:,p:$100
Fcoeff
dc 1677722
dc 2097152
dc 2516582
dc 2936013
dc 3355443
```

The above assembly code will be loaded into the P memory space at p:\$100 for the PROM burning, and it should be copied to the L memory space before the actual program is executed. Manipulating the assembler's run-time and load-time counters requires a thorough understanding of the underlying assumptions about memory layout, which are made by the compiler (see Chapter 6).

Note: Incorrect use of this feature may cause compile-time, link-time and even run-time errors.

5.4 Out-of-line Assembly Code

Out-of-line assembly code is assembly code written in a separate source file that is called from a C program. Separating the assembly code and C code in this way provides a powerful and flexible interface to the DSP56000 family architecture. This out-of-line method may be used to convert existing assembly subroutines, write new subroutines completely in assembly language or access both data spaces (X and Y). The key advantage of out-of-line assembly code is that it provides a complete assembly programming environment for the DSP56000 family whereas the in-line assembly code must follow the C programming environment rules.

Writing out-of-line assembly code requires a complete understanding of the C Cross Compiler and the DSP56000 family architecture. For out-of-line assembly code to be callable from a C program, the following five basic elements should be included in the assembly source file in sequence.

1. C subroutine entry code (prologue code).
2. Save all registers to be used.
3. Main program.
4. Restore all registers used.
5. C subroutine exit code (epilogue code).

Each of these steps can be programmed as a macro; however, using generic macros can unnecessarily increase the size of the resulting assembly code program. The DSP's ability to move data in parallel with arithmetic operations provides opportunities to significantly optimize the assembly code by combining steps within a macro and by combining the code of two or more macros.

In order to illustrate the steps listed above, the out-of-line assembly code template is described first and each element of the template is then explained in detail. After reviewing the five elements, some optimization techniques are discussed.

5.4.1 General Template

The following template is a generic form used to make the C function “foo”. There are five distinct elements in the template which are numbered as above. The actual code for the prologue and epilogue is shown but the “Save all registers to be used”, “Main Program”, and “Restore all registers used” are listed as comments because their actual code depends on the main program.

```
    global  Ffoo                ; prologue:
Ffoo                ; sets up entry point (C
function address),
    move   r0,y:(r6)+         ; updates frame pointer (r0)
and saves
    lua    (r6)+,r0           ; return address
    move   ssh,y:(r6)+        ;
; Save all registers to be used
; Main Program
; Restore all registers used
    move   y:-(r6),ssh        ; epilogue:restores return
address, updates
    move   y:-(r6),r0         ; restores return address,
updates
    tst    a                  ; frame pointer (r0), and
updates the SR with
                                ; the contents of register A

    rts
```

5.4.1.1 Prologue

The first two lines of the prologue make the assembly program visible to the C program so that the subroutine or function is callable from the C program. In this case, any one of the following C statements can be used to access out-of-line assembly code.

```
foo();  
x = foo();  
x = foo(arg1, arg2, arg3);
```

The first function call assumes that the C function does not use any arguments and does not return any values. The second only returns a value which is the same data type as the variable `x`. The last call assumes that the C function uses the three arguments: `arg1`, `arg2` and `arg3` and then returns the value `x`.

The rest of the prologue saves the old frame pointer, updates the current frame pointer and saves the return address. The return address is saved when the JSR instruction pushes the hardware stack which saves the program counter in the high 16 bits of the system stack i.e., SSH. This section of the prologue provides bookkeeping for the C compiler activation record. Detailed information on the structure of the activation record can be found in Section 4.6.1, "Activation Record," on page 4-10. After this prologue, the R0 register is the frame pointer and R6 is the stack pointer.

5.4.1.2 Save all registers

All registers used in the main program should be saved before the main program changes them. This step is the second element of the template — "Save all registers to be used".

In order to save the registers, R6 is used as a stack pointer. The stack grows upward and the current stack pointer (R6) points to the next element above the top of stack. The following statement saves one register to the top of stack and sets the stack pointer to the next available stack location.

```
move    x0,y:(r6)+
```

Two registers, `y1` and `x1`, can be saved with the following statements:

```
move    y1,y:(r6)+  
move    x1,y:(r6)+
```

Since saving and restoring the registers are the subroutine's responsibility, the order of saving the registers should be in accordance with their restoration. The restore process should be exactly the reverse order of the register saving sequence.

5.4.1.3 Main Program

The main C function program accesses the parameters passed, executes the operations on the parameters and returns a value. Passing parameters is done through the activation record which has been pushed onto the stack. The activation record frame pointer is pointed to by **r0** and the first parameter is placed 3 locations below the frame pointer (see Section 4.6.1, "Activation Record," on page 4-10). The parameters are arranged in reversed order. For example, the following statement should be used to move the first single word parameter to register **r3**.

```

move    #-3,n0           ; the first parameter location
move    y:(r0+n0),r3     ; load it to r3.

```

Assuming the first three parameters are one word long, the following statements move the second and third parameters to registers **r4** and **r5**, respectively.

```

move    #-4,n0
move    y:(r0+n0),r4
move    #-5,n0
move    y:(r0+n0),r5

```

5.4.1.4 Restore all registers

The stack pointer, **r2**, is needed to restore the registers. The following code will restore one register. At this point in the function's execution, the stack pointer points to the location above the last saved register, hence the pre-decrement.

```

move    y:-(r2),x0      ; restore

```

The restoring procedure can be simplified if more than one register is to be restored. Restoring registers x0, x1, y0 and y1 can be done by the statements below.

```

move    (r2)-
move    y:(r2)-,x0
move    y:(r2)-,x1
move    y:(r2)-,y0
move    y:(r2),y1

```

After the function has finished, a return value can be passed to the caller. Any 32-bit or 16-bit value must be returned through register a. If the return value is larger than 32 bits, then the compiler allocates the proper amount of buffer space and register a1 becomes the pointer to this buffer space upon callee execution. It is the callee's responsibility to copy any return values from the buffer whose address resides in register a1. This is the method used for returning structs.

5.4.1.5 Epilogue

The out-of-line template epilogue is the reverse of the prologue. The epilogue restores the return address and updates the frame pointer, R0. Notice that the stack pointer, R6, should be decremented before each move operation. In addition, register A is tested to update the SR register. This testing is a part of the C compiler code generation feature and should be included (see Section 5.4.5 for optimization).

5.4.1.6 Out-of-line Assembly Code Example

The following example illustrates using the general template as well as the fact that the DSP56000/DSP56001 performs fractional arithmetic. Assume that section `mod1` contains two C callable subroutines, `mac01()` and `mac02()` and section `mod2` contains a single C callable subroutine, `mactwo()`. The two functions `mac01()` and `mac02()` take one argument each and `mactwo()` takes two arguments. These functions perform multiplication according to the following formulas and return their results:

<code>mac01(arg)</code>	calculates	<code>arg * 0.01</code>
<code>mac02(arg)</code>	calculates	<code>arg * 0.02</code>
<code>mactwo(arg1, arg2)</code>	calculates	<code>arg1 * arg2</code>

Note that multiplication on the DSP56000/DSP56001 treats data as fractions; however, the C programming environment does not support a fractional data type. Therefore data will be passed as integers from the C programming environment and will be treated as fractional data in the DSP56000/DSP56001. This can be an advantage if fractional arithmetic is desired since this is normally difficult to accomplish in C. Note that the calling routine must ensure `arg` is in range to prevent overflow.

The following function declarations (ANSI standard) can be used to declare functions which are known to perform fractional arithmetic. Internally, the size of integers is the same as the size of fractions.

```
int mac01(int);
int mac02(int);
int mactwo(int, int);
```

These three functions can be called as follows:

```
int fractvalue;
fractvalue = mac01(0x123456); /* the value 0x123456 is
0.142222166 in fractional */
fractvalue = mactwo(0x123456, 0x0147ae); /* 0x0147ae is
0.001 */
```

The conversion between fractional data and integer data (in hex form) can be performed using the evaluate command in the DSP56000 simulator.

Example 5-31. General Template for Out-of-line Assembly Code

General Template for Out-of-line Assembly Code: The two sections of this out-of-line assembly code are mod1 and mod2. The first section implements:

1. **mac01(arg)** which takes a fractional argument and returns 0.01 times the argument and
2. **mac02(arg)** which takes a fractional argument and returns 0.02 times the argument.

The second section implements **mactwo(arg)** which takes two fractional arguments and returns **arg1*arg2**.

```
; section mod1:
; int mac01(int arg);
;     takes a fractional argument and returns 0.01 *
;     argument.
; int mac02(int arg);
;     takes a fractional argument and returns 0.02 *
;     argument.
    section mod1
    global Fmac01
Fmac01
    move    r0,y:(r6)+      ; prologue
    lua    (r6)+,r0
    move    ssh,y:(r6)+
    move    x1,y:(r6)+      ; save register X1
    move    y0,y:(r6)+      ; save register Y0
    move    n0,y:(r6)+      ; save register n0
    move    #-3,n0
    move    y:(r0+n0),x1    ; argument
    move    #0.01,y0        ; operand 0.01
    clr    a
    macr    x1,y0,a          ; main program: calculates
multiplication of          ; x1 and y0
```

```

    move    y:-(r6),n0      ; restore register n0
    move    y:-(r6),y0      ; restore register Y0
    move    y:-(r6),x1      ; restore register X1
    move    y:-(r6),ssh     ; epilogue
    move    y:-(r6),r0
    tst     a
    rts
    global  Fmac02
Fmac02
    move    r0,y:(r6)+      ; prologue
    lua     (r6)+,r0
    move    ssh,y:(r6)+

    move    x1,y:(r6)+      ; save register X1
    move    y0,y:(r6)+      ; save register Y0
    move    n0,y:(r6)+      ; save register n0
    move    #-3,n0
    move    y:(r0+n0),x1    ; argument
    move    #0.02,y0        ; operand 0.02
    clr     a
    macr    x1,y0,a         ; main program: calculates
multiplication of                                     ; x1 and y0
    move    y:-(r6),n0      ; restore register n0
    move    y:-(r6),y0      ; restore register Y0
    move    y:-(r6),x1      ; restore register X1

    move    y:-(r6),ssh     ; epilogue
    move    y:-(r6),r0
    tst     a
    rts
    endsec

; section mod2
; int mactwo( int arg1, int arg2);
;     takes two fractional arguments and returns
arg1 * arg2.
    section mod2

    global  Fmactwo
Fmactwo
    move    r0,y:(r6)+      ; prologue
    lua     (r6)+,r0

```

```
    move    ssh,y:(r6)+
    move    x1,y:(r6)+      ; save register X1
    move    y0,y:(r6)+      ; save register Y0
    move    n0,y:(r6)+      ; save register n0

    move    #-3,n0
    move    y:(r0+n0),x1    ; the first argument
    move    #-4,n0
    move    y:(r0+n0),y0    ; the second argument
    clr     a
    macr    x1,y0,a         ; main program:
                                ; calculates multiplication of
x1 and y0

    move    y:-(r6),n0      ; restore register n0
    move    y:-(r6),y0      ; restore register Y0
    move    y:-(r6),x1      ; restore register X1

    move    y:-(r6),ssh     ; epilogue
    move    y:-(r6),r0
    tst     a
    rts
    endsec
```

The DSP56KCC control program, G56k, should be used to assemble the out-of-line assembly code. The two sections of this code, mod1 and mod2, can be put in the same file or in separate files. The following command lines and source files provide a test for the case where each program is in a separate file.

```
C:\> type main.c
```

```
#include <stdio.h>
int mac01(int), mac02(int), mactwo(int,int);
main()
{
printf("%x, %x\n", mac01(0x123456), mactwo(0x123456,
0x0147ae));
printf("%x, %x\n", mac02(0x123456), mactwo(0x123456,
0x28f5c));
}
```

```
C:\> g56k main.c mod1.asm mod2.asm
```

```
C:\> run56 a.cld
```

```
 2e9a, 2e9a
5d35, 5d35
```

5.4.2 Global C and Static Variables in C

The global C variables are accessed using labels generated by the C compiler. Any variables that are static to an assembly language subroutine will be accessed the same way. These variables are placed into memory at compile-time and are referenced symbolically according to the labels automatically generated by the compiler. However, it is possible to override the default labels by using the `__asm()` keyword as explained in Section 5.2.5, "Controlling Labels Generated by the Compiler,".

For example, using the default labeling convention, the global integer, **Ginteger** which can be declared within the C statement "extern int Ginteger;" is loaded into the input register x0 in assembly code as follows:

```
move    y:FGinteger,x0
```

When declaring C global variables in an assembly language file, the programmer must be careful to follow the label generating convention or use the `__asm()` keyword to report to the compiler that the labeling convention has been changed. In both cases, the assembler directive `global` is used to export the labels to the C files. DO NOT use the XDEF/XREF pair of directives. NOTE that it is the programmers responsibility to *allocate* space for the global variables declared in this manner. In the example below, this is done with the assembler directive `dc`. Also, ANSI C requires that all global variables be initialized to zero if they are not explicitly initialized.

Example 5-32. Global Label in Assembly Language

Global Label in Assembly Language. This example shows assembly code that defines a global integer (named `FGinteger`) which is normally accessed as `Ginteger` in the C environment and `FGinteger` in the assembly programming environment.

```
FGinteger    dc    $0
globalFGinteger
```

Example 5-33. Global Variable Declaration

Global Variable Declaration. This is the C code equivalent to Section which defines the global integer `Ginteger`.

```
extern    int    Ginteger;
```

Example 5-34. Changing a Global Label

Changing a Global Label. This example shows C code that generates a global integer (Ginteger) which is accessed as Ginteger in both the C environment and the assembly programming environment.

```
extern int Ginteger _ _asm("Ginteger");
    Which will appear in assembly language code as:
Ginteger dc      $0
    global Ginteger
```

5.4.3 Using Run-time Stack for Local Data

The run-time stack may be used when the programmer requires a temporary data space for *automatic* style variables — i.e., local variables in subroutines. Using the run-time stack requires additional steps in the prologue and epilogue sections. It is the subroutine's responsibility to automatically allocate and deallocate the stack at run-time.

In the prologue, an extra step is required to save the run-time stack space. Keeping in mind that the stack pointer **must** always point to the next available stack location, the stack space is *allocated* by advancing the stack pointer by the amount of space required. One way to allocate this space is shown in the Section .

Example 5-35. Run-time Stack Allocation

Run-time stack allocation: This code segment can be inserted in the general template prologue for out-of-line assembly code. Notice that “size” in the move statement below should be replaced with the appropriate constant.

```
    move    #size,n6; the stack size needed
    nop                    ; wait until n6 is available to operate
on r6
    move    (r6)+n6 ; allocate the run-time stack for
locals
```

Referencing the data space can then be accomplished using negative offsets from the stack pointer or via initialized address registers. There are many alternatives to these methods but they are all similar.

In the epilogue, an extra step is required to restore the stack pointer — i.e., deallocate the run-time local stack. This is simply the reverse of the allocation process in the prologue.

Example 5-36. Run-time Stack Deallocation

Run-time stack deallocation: This code segment can be inserted in the general template epilogue for out-of-line assembly code. Notice that “size” in the move statement below should be replaced with the appropriate constant.

```

move    #size,n6; the stack size used before
nop     ; wait until n6 is available to operate on r6
move    (r6)-n6 ; deallocate the run-time stack

```

There are many ways to do this, one simple optimization would be to advance the n6 load instruction in the program to eliminate the nop.

5.4.4 Calling C Routines

C routines are routines that are callable by a C program and may be written in either C or assembly language. When writing assembly language subroutines, it may be necessary to call library routines that have been provided or that have been written by the programmer — e.g., a call to **sin()** or **printf()**. In order to do this, the programmer must follow 3 steps:

1. Push arguments onto the run-time stack in reverse order.
2. Make the subroutine call.
3. Restore the stack pointer.

The following example assumes that the four parameters to be passed to the C function **foo()** are located in registers a1, b1, x0 and x1; respectively. The first four lines of assembly code push the arguments onto the run-time stack in reverse order. The **jsr** statement makes the subroutine call and the last two statements restore the stack pointer.

Example 5-37. Calling C Routines

Calling C Routines. The C function, **foo()**, is called from the following assembly code. Function **foo()** is declared as **int foo(int, int, int, int);**

```

move    x1,y:(r6)+      ; pushing arguments onto
move    x0,y:(r6)+; run-time stack in reverse
move    b1,y:(r6)+; order
move    a1,y:(r6)+
jsr     Ffoo    ; subroutine call
move    #4,n6    ; the stack size restored.
nop
move    (r6)-n6 ; restore it.

```

5.4.5 Optimization Techniques

The general template for out-of-line assembly code provides a clean template to build C callable functions. However, the DSP56000 family microprocessor chips have powerful features such as multiple instruction execution (multiply and accumulate) and parallel data move operations that may allow additional optimization. After constructing the out-of-line assembly code from the general template, some hand-optimization can be performed by combining several assembly statements.

Information about these optimization techniques can be obtained from the *DSP56000/DSP56001 User's Manual*. Some optimization techniques which are related to the C compiler are discussed in this section but additional optimization can be achieved using the architectural features described in the user's manual.

The R0 register (which is updated by the out-of-line assembly code template prologue) is required only when parameters are passed to the C function. If there are no parameters to be passed, then the following two assembly code lines are not required and can be omitted:

```
move    r0,y:(r6)+
lua     (r6)+,r0
```

The return address (SSH) was saved in the out-of-line assembly code prologue but it is only required when a C function calls another C function. A C function is called a leaf function if it does not call any other C function. In leaf functions, the return address is not needed and does not have to be saved.

Similarly the epilogue can be optimized in the same way as the prologue. For any leaf function or non parameter C function, the epilogue size can be reduced by eliminating:

```
move    y:-(r6),ssh; for a leaf C function
```

or

```
move    y:-(r6),r0; for a non-parameter C function
```

The test statement “**tst a**” in the epilogue can be eliminated if the function does not return any values. The test statement may be required due to the C compiler's optimization features since it provides condition flags for an **if** statement in a function call. For example, if the out-of-line assembly function `foo()` is used in the statement `if(foo()) { }`, then the C compiler will not generate code to test the return value when a **jcc** statement is issued. This is primarily because the C compiler uses the condition flags which were executed at the end of the epilogue of `foo()`.

A variety of optimizations can be achieved by combining the move instructions and main program code to utilize parallel moves (see Section 5.4.5 and Section 5-38 which point out possible optimizations in the comments). These and other DSP56000 specific

optimizations can dramatically improve the quality of the application specific library routines. A careful review of the *DSP56000/DSP56001 User's Manual* will be worthwhile for efficient library development.

Example 5-38. ReadX / ReadY Routines

ReadX / ReadY Routines. This out-of-line assembly program (called memread.asm) reads the contents of x or y memory and returns the data to the caller. The ReadX subroutine returns the x memory space contents pointed to by the input argument and the ReadY subroutine returns the y memory contents.

```

section memread
org      p:
global ReadX    ; prologue

ReadX

mover    0,y:(r6)+;
lua      (r6)+,r0 ;
move     r1,y:(r6)+; save register r1
move     #-3,n0 ; the first parameter
move     x:(r0+n0),r1 ;
move     x:(r1),a; main program
                                ; restore register & epilogue
move     y:-(r6),r1 ; Note: r6 decrement can
tst      a          y:-(r6),r0; be a parallel move in the
rts      ; main program.

global ReadY    ; prologue

ReadY

move     r0,y:(r6)+ ;
lua      (r6)+,r0 ;
move     r1,y:(r6)+ ; save register r1
move     #-3,n0 ; the first parameter
move     y:(r0+n0),r1 ;
move     y:(r1),a ; main program
                                ; restore register & epilogue
move     y:-(r6),r1 ; Note: r6 decrement can
tst      a          y:-(r6),r0 ; be a parallel move in the
rts      ; main program.
endsec

```

Chapter 6

Software-Hardware Integration

6.1 Overview

This chapter explains how the run-time environment may be changed and provides examples of some changes and their effects. The run-time environment provided with the compiler assumes, as a default, that the simulator is the target execution device. Several aspects of the default run-time environment must be altered in order to adapt the compiler to work with a custom hardware configuration.

The files which are alterable are discussed and classified according to effect. Aspects of the run-time environment such as: bootstrapping, interrupts and memory management are addressed individually.

6.2 Run-Time Environment Specification Files

The run-time environment is specified by three assembly language files: `crt056[x,y,l].asm`, `signal56[x,y,l].asm` and `setjmp56[x,y,l].asm` where `x`, `y`, or `l` denote the memory model (see chapters 2 and 4). These files may need to be modified if the run-time environment is to be customized.

The `crt0` file contains parameters which specify the C bootstrap code, memory configuration, memory management, interrupt vectors and other miscellaneous code. This file must be modified to match the software and hardware configuration because the memory configuration and interrupt vectors are determined by the hardware. The information in this manual on the `crt0` file applies to DSP56KCC Version g1.11. The `crt0` file may be different for other versions of this compiler.

The `signal` file, which is equivalent to a hardware interrupt, is implemented in the C environment. The `signal` file contains the code and data structures used to implement the **signal()** and **raise()** library functions. Changing this file is not recommended unless necessary since any change to this file requires detailed knowledge of the DSP56000 family interrupt mechanism in addition to the `signal` and `raise` functions. This file is closely tied to the `signal.h` file.

The `setjmp` file contains code which implements the functions **setjmp()** and **longjmp()**. This file will probably never need to be modified unless the `signal` file is changed; however, if either the `setjmp` file or `setjmp.h` are modified, the code in both files must be made consistent. The source code for **setjmp()** and **longjmp()** is provided with DSP56KCC to allow modification should the signal mechanism need to be changed.

The operation of `setjmp()` and `longjmp()` is described in Section 6.10 and detailed implementation information can be obtained from the files provided with DSP56KCC.

6.3 crt0 File

The following subsections describe the various functions of the `crt0` file.

6.4 Bootstrapping the C program

The processor enters a C program through the C bootstrap code in the `crt0` file. The C bootstrap code in `crt0` provides the C environment required to execute a C program. This environment includes a global data area, static data area, stack area, heap area, function calling mechanism, etc. This environment must be established before C programs will execute correctly.

The following steps are normally taken to initialize the processor to execute C code:

1. Jump from the chip reset vector to the C bootstrap code labeled at `F__start` in `crt0`. Remember that the mode select pins on the chip control the chip operating mode when leaving reset which, in turn, controls the reset vector address (see the *Motorola DSP56000/DSP56001 User's Manual* for more details).
2. Configure all hardware registers needed (i.e., ovr, host port, etc.). This is also a proper place to initialize any non-C related data structures or peripheral hardware such as the bcr, Port B and Port C.
3. Load the Stack Pointer, r6, with a pointer to the base of the stack. Remember that the stack grows up (the value in the stack pointer gets greater as data is pushed). The value of `DSIZE` is generated by the linker and is the first address above the statically allocated data (C global and static variables). By default, this value is used as the initial stack pointer.
4. Optional: Initialize the frame pointer, r0. This may be of use to some debuggers when stack back traces are performed. Program execution will not be affected if this initialization is omitted.
5. Call `main()` with instruction `jsr Fmain`. Notice that the label is `Fmain` and there are no parameters passed to the main function. The normal C compiler start-up passes two arguments but `g56k` does not pass any arguments because `DSP56KCC` does not support a particular hosted environment.

The bootstrap code is followed with the label `F__crt0_end`. This label is used by `Gdb56` and `run56` to detect program termination.

Note: Labels which begin with a double underline (e.g., `__ crt0_end`) in this manual have a space between the double underlines to visually separate them. Do not separate the leading double underlines with a space when coding them (i.e., code `__ crt0_end` as `__ crt0_end`).

Example 6-1. DSP56000/DSP56001 Operation Mode Change

Mode 2 has a reset vector of \$E000 which must contain a jmp to the C program bootstrap code. Adding the following code segment to the crt0 file will change the bootstrap mode to Mode2.

```
section mode2_reset
org p:$e000
jmp F_ _start; jump to the C start-up program.
endsec
```

Example 6-2. C Bootstrap Code Location Change

Hardware was designed to have a 256 byte ROM monitor located in the program memory space starting at \$0000 and ending at \$FF. Program RAM starts at location p:\$100. The following changes to the crt0 file will change the beginning location of the C bootstrap code to the first available RAM location (p:\$100). The DS statement allocates program space starting at p:\$0000 and lets the ROM be located at address p:\$0000. The org statement places the C bootstrap code at memory location p:\$100.

Change this portion of the crt0 file:

```
org      p:
F_ _start
to:
  org p:$0000
  ds $100
  org p:$100
F_ _start
```

6.5 Memory Configuration and Management

The DSP56000 family supports three memory spaces: program memory (p memory), x and y data memories. The DSP56000 family design philosophy is to support concurrent access to these memory spaces to accelerate DSP operations. The C programming language does not support separate memory spaces but instead treats memory as a single memory space. As a result, there is a difference between hardware memory organization and C memory utilization.

DSP56KCC provides three different memory models to minimize these differences. They are the x memory model, y memory model and l memory model. Selection of the appropriate memory model is made using the `-m` option of the control program, `G56k` (see Section 3.2.2, "Compile Phase Options," on page 3-20). The memory configuration for each memory model is discussed later in this section. The `crt0` file can be modified to configure memory for any hardware design.

There are four data segments in the C programming environment. These are the program segment, global-static data segment, stack data segment and heap data segment. The program segment is located in program memory. The three data segments can be located in x memory, y memory or l memory space (which is a concatenation of x and y memory) depending on the memory model used.

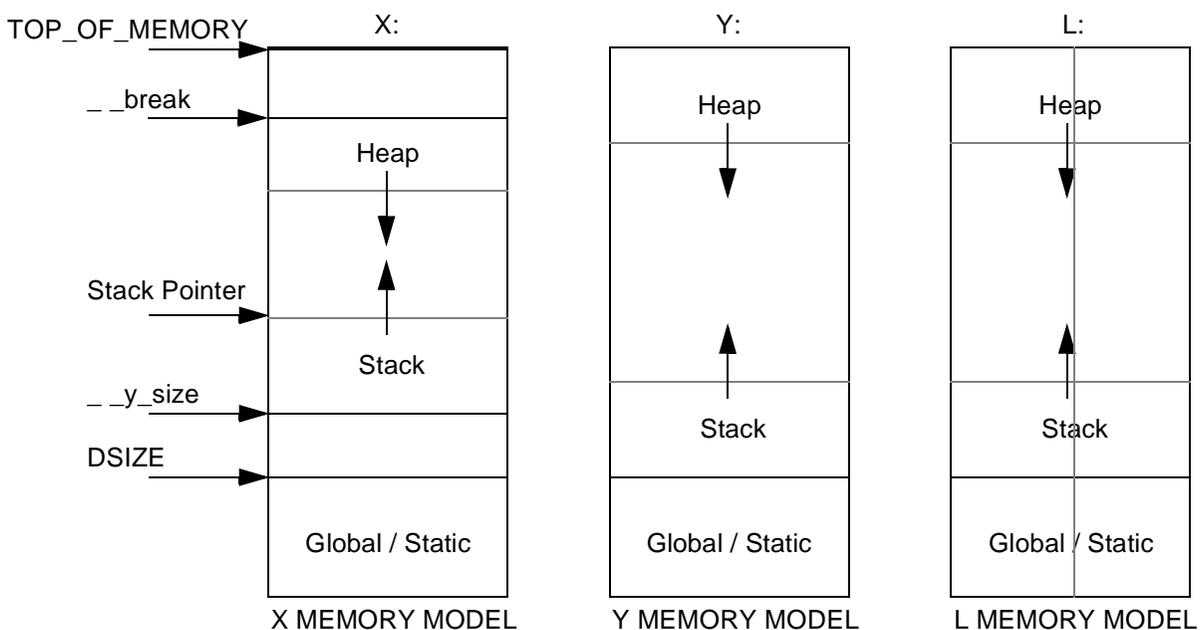


Figure 6-1. Environment Memory Configuration

As indicated in Section 5.1, "Overview," on page 5-1, global and static data reside at the bottom of the available data memory and the top address of the global and static data area, which is called `DSIZE`, is set by the linker. The constant `TOP_OF_MEMORY` is defined to indicate the top of the entire available memory.

The two data segments, heap and stack, are located as shown Section 5.1, "Overview," on page 5-1. The stack is located so that it can grow up and the heap is located so that it can grow down. There are two locations used to indicate the initial values for the heap pointer and stack pointer. These locations are `__y_size` and `__break` and are initialized in the `crt0` file as `DSIZE` and `TOP_OF_MEMORY`, respectively.

In summary, two variables `__y_size` and `__break` and the constant `TOP_OF_MEMORY` are used to configure the data segment. The program segment is configured using the `org` statement in the `crt0` file. The variable, `__y_size`, should be initialized with the initial stack pointer and the variable, `__break`, should be initialized with the initial heap pointer.

Warning

The stack and heap regions must not contain on-chip peripheral memory or the static or global data regions. Also, no region may be reconfigured after the C main function is called. Variables `__y_size` and `__break` should not be altered by an arbitrary function since they are utilized by system level libraries such as `malloc` and `free`.

Example 6-3. Fast Stack

In this example, the stack is required to reside in an 8k SRAM starting at L:\$4000. The following program reserves the stack space using `org` and `ds` statements and sets the initial stack pointer to the SRAM stack area.

Add this section to the `crt0` file:

```
section fast_ram
org      L:$4000
ds       $2000
endsec
```

Change the following line of C bootstrap code in the `crt0` file:

```
move y:F_ __y_size,r6
to:
move # $4000,r6
```

Example 6-4. Fast Heap

The heap is required to reside in an 8k SRAM starting at L:\$4000. The following program reserves the heap space using `org` and `ds` statements and sets the initial heap pointer to the SRAM heap area.

Add this section to the `crt0` file:

```
section fast_ram
org     L:$4000
ds     $2000
endsec
```

Change the following line of C bootstrap code in the `crt0` file:

```
TOP_OF_MEMORY     equ     $ffbe
```

to:

```
TOP_OF_MEMORY     equ     $5fff
```

Sometimes hardware configurations map more than one memory space into a single physical memory. Other implementations partially populate various address spaces leaving holes. Some may have different regions with fast memory and slow memory. All of these special cases can usually be handled by modifying the `crt0` file.

When multiple memory spaces are mapped into a single physical memory, the memory must be partitioned among the memory spaces. A way to restrict the linker from overlapping these memory spaces is needed. For example, suppose that both the Y and P spaces are mapped into the same 64k physical RAM and need to be partitioned with the low 48k for program memory and the high 16k for data memory.

The linker can be restricted from allocating across holes in physical memory by using the **org** and **ds** directives to confiscate those addresses. Note that the linker may not automatically take advantage of memory which is present between holes. It may be required to manually place assembly language data structures to utilize this memory.

6.6 Interrupt Vectors

The interrupt vector locations for the DSP56000 family (a.k.a. “interrupt source” in Section 8 of the *DSP56000/DSP56001 User’s Manual*) contain one or two instructions each to be executed when the interrupt assigned to that location occurs. There are 32 interrupt vectors available, all of which should be initialized with some value to avoid undefined behavior resulting from an unexpected interrupt.

The `crt0` file contains code to initialize these interrupt vectors. By default, all vectors are initialized with the instruction `jsr Fabort`. The first element of the vector table, which is the hardware RESET interrupt, is initialized with the instruction **jmp F_start**. The purpose of the C function **abort()**, which is labeled as `Fabort` in the assembly environment, is to stop program execution and leave the chip in the STOP mode. The **F_start** label is the program address of the C program bootstrap code which calls the C **main()** function after execution.

Interrupt vectors that are to be used must be reprogrammed to point to the interrupt service routines instead of the **abort()** function. Initialization of the interrupt vectors in the `crt0` file reduces the size of the resulting application program and may increase its speed.

The following `crt0` code segment (called section `reset`) is the default interrupt vector table initialization.

```
section reset
org      p:$0
jmp      F_start
org      p:$2
dup      31
jsr      Fabort
endm
endsec
```

The interrupt vector table can be changed to point to user-provided interrupt service routines instead of the **abort()** routine in this portion of `crt0`. Example 6-5 illustrates how to initialize pointers to these user-provided interrupt service routines.

Example 6-5. User-defined Interrupt Vector Table

Assume the hardware supports all interrupts and each interrupt service routine is located at the address labeled `interruptXX` (where `XX` is the value of the interrupt vector). The following code initializes the interrupt vector table. Each service routine starting at `interruptXX` can be programmed in assembly language as shown in Section .

```
section reset
org      p:$0
jmp      F_start
jsr      interrupt02
jsr      interrupt04
jsr      interrupt06
jsr      interrupt08
...
jsr      interrupt3e
endsec
```

Example 6-6. Interrupt Service Routine

Interrupt Service Routine: This service routine updates the global variable `F__time` at each SCI timer interrupt. The SCI timer period can be set by programming the SCCR (see Section 11.2.2.3 of the *DSP56000/DSP56001 User's Manual*). This program is based on the y memory model and the SCI interface control registers (`X:FFF0` and `X:FFF2`) should be initialized for proper SCI timer operation (see Section 11.2 in the *DSP56000/DSP56001 User's Manual* for detailed information on the SCI operation).

```

section interrupt
org      p:
global  interrupt1c
interrupt1c
move    (r6)+          ; secure the stack pointer
                        ; (refer to section 5.4.1.4)
move    r2,y:(r6)      ; save the r2 register
move    y:F__time,r2   ; retrieve the variable __time
move    (r2)+          ; increment the variable
move    r2,y:F__time   ; save the result
move    y:(r6)-,r2     ; restore the r2 register
rti     ; return from interrupt service
endsec

```

Notice that fast interrupts can also be programmed by modifying the `crt0` file in the same way as for the long interrupts (see Chapter 8 of the *DSP56000/DSP56001 User's Manual* for more information on fast and long interrupts.)

6.7 Miscellaneous Code

There are other data structures and code related to the run-time environment in the `crt0` file. They are:

1. Error code variable, `errno`, is a global integer used to record failure codes in C library calls. This error code variable will exist if C library calls are used. This variable can be utilized as a debugging aid in order to check which error code is returned from the C function calls.
2. Heap-stack checking window variable, `__stack_safety`, is a global variable declared in the `crt0` file that is used by the heap allocation routines, `malloc()`, `calloc()` and `realloc()` to avoid heap-stack growth collisions. If the distance between the bottom of the heap and the top of the stack is less than the value contained in `__stack_safety`, then the heap allocation routine will return an error code indicating that no more memory is available. The value may be set as required

by the application since the window `__stack_safety` is declared as a global variable.

3. Memory limit variable, `__mem_limit`, is a global variable declared in the `crt0` file and used by the library routine `brk()` to disallow any meaningless memory requests. This variable should have the same value as `__break` upon entry into `main()`. For information on how to use the function `brk()`, refer to Appendix B in this manual.
4. Dynamic time variable, `__time`, is a global variable declared in the `crt0` file and used as a volatile timer counter by the simulator. This variable is updated by the DSP56000 simulator (`run56`) every clock cycle. Examining this variable allows the programmer to determine program execution time. This variable is only used by the simulator and can be omitted when the program is to be executed by hardware.
5. Host I/O stub functions, `__send()` and `__receive()`, are defined in `crt0` and are called by the standard I/O library functions. The provided `crt0` file only has stub functions, as both `Gdb56` and `run56` watch these addresses and perform all I/O directly.
6. Floating-point shift table, `__fp_shift`, is a pointer used by floating-point library functions. This pointer may be removed if the application does not use floating-point.

All of these variables, constants, functions and pointers are related to the run-time environments that are used by C library functions and must be properly set.

6.8 Signal File

The hardware level interrupt mechanism (see section 8.3 of the *DSP56000/DSP56001 User's Manual*) is more efficient than the `signal()` function. However, in many cases interrupts can be handled in the C environment and it is often more preferable to do so. There are two functions, `signal()` and `raise()`, used to support programming interrupt service routines in C. These functions are not associated with the `crt0` file. Although they are more complicated than the simple hardware interrupt vector table discussed in Section 6.6 (see section 8.2 of the *DSP56000/DSP56001 User's Manual*) they provide very handy tools for the C programmer. A thorough knowledge of the `signal()` function and the C environment is needed in order to modify the `signal()` function. This section describes how the `signal()` function is currently implemented.

6.9 Signal()

The **signal()** function is passed two arguments:

1. A signal number — On the DSP56000 family processor, the signal number corresponds directly to the interrupt vector address. Notice that the signal number is not an integer sequence but is an even number.
2. A function pointer — The function pointer passed is assumed to belong to a C function either generated by this compiler or by assembly code.

Signal() performs the following three steps when binding the specified signal number and function:

1. The instruction **jsr F_ _c_sig_goto_dispatch+<signal number>** is placed into the interrupt table location specified by the signal number.
2. The function pointer passed is entered into the table **_ _c_sig_handlers**, which is used to store pointers to C signal handlers, indexed by the signal number.
3. The old signal handler address is returned.

Once the signal number and specified function are bound, the instruction

```
jsrF_ _c_sig_goto_dispatch+<signal number>
```

is executed upon receiving the interrupt, where the **F_ _c_sig_goto_dispatch** variable is the starting address of a table of

```
jsrF_ _c_sig_dispatch
```

instructions and each **jsr** instruction points to an interrupt service routine. The pseudo-function **_ _c_sig_dispatch()** is used to calculate the actual C interrupt routine.

All registers are saved before the **_ _c_sig_dispatch()** function calls the C signal handler. Pseudo function **_ _c_sig_dispatch()** then calculates the signal number using the return address program counter of the **ssh**. Since the signal number is the same as the interrupt vector address, each entry of the **_ _c_sig_goto_dispatch** table corresponds to an interrupt vector. The pseudo function uses the signal number to fetch the actual C signal handler from the **_ _c_sig_handlers** table which is the C function pointer table.

Once the C signal handler is fetched from the **_ _c_sig_handler** table, its entry is replaced with the default signal handler **SIG_DFL**. This replacement is in compliance with the ANSI standard and forces the next signal service to abort. In most situations, this feature is not needed because any given interrupt will always invoke the same interrupt service routine. Resetting **signal()** after each C service routine or modifying this file so that it does not replace the table entry with **SIG_DFL** will change the interrupt service scheme.

Modification of the signal file is only recommended when optimization of the service time is critical to the application.

Upon return from the C signal handler, all the registers are restored. Finally, the `rti` instruction is executed to return to the code that was executing when the interrupt occurred. Notice two factors in this scheme,

1. all registers are saved and restored before and after the C signal handler and
2. the `rti` instruction is executed by the `__c_sig_dispatch()` function.

Warning

The signal handler must not contain the `rti` instruction at the end of the program regardless which language is used to program the interrupt. The signal handler does not need to save or restore any context or registers. The function `__c_sig_dispatch()` will not act like a normal C function because it never returns to its caller. Instead, it will return to the code that was executing when the interrupt happened by executing the `rti` instruction.

Assembly language interrupt handlers can coexist with C signal handlers. The code in the signal file will not alter any interrupt vector except the one specified by the signal number passed to the `signal()` function (see the first of the three steps above). The C signal interface could be used with an assembly routine but would be unnecessarily slow. To use an assembly language interrupt handler, alter the vector (e.g., interrupt 08) with a `jsr` to it (e.g., `jsr interrupt 08`) or use a fast interrupt routine.

6.9.1 Raise()

The `raise()` function is used to simulate an interrupt. The code in `raise()` simply calls the entry in `__c_sig_goto_dispatch` that is matched to the interrupt vector specified by the signal number passed.

The ANSI standard signal handlers `SIG_DFL`, `SIG_ERR` and `SIG_IGN` are implemented by the hand-coded functions `__sig_dfl()`, `__sig_err()` and `__sig_ign()`, respectively.

1. `SIG_DFL` notes that the interrupt happened by incrementing `__sig_drop_count` and then returns.
2. `SIG_ERR` calls `abort()` and never returns.
3. `SIG_IGN` returns without any effect (i.e., ignore).

The mechanisms used to implement the C signal interface may be altered to fit a particular hardware application. Any series of alterations applied to the signal file must leave an implementation conforming to the ANSI standard X3.159 for C. Alteration of the signal file is done at one's own risk and is not generally advised. Again, the contents of the signal file must remain consistent with the include file `signal.h`.

6.10 Setjmp File

The functions `setjmp()` and `longjmp()` are implemented in the `setjmp` file. The `setjmp()` function stores the current process status (i.e., the current execution context) in a buffer that is passed. The `longjmp()` function is used to restore the process status to the execution context which was saved by `setjmp()`.

Saving the current execution context is done by saving the stack pointer, the return program counter value, the frame pointer and all of the callee-save registers into the buffer. The buffer that is passed to `setjmp()` should have enough space for the saving process. The structure `jmp_buf` defined in `setjmp.h` allocates the buffer space needed for the operation. The function `setjmp()` always returns a zero.

The function `longjmp()` takes two arguments, an environment buffer and a return value. It restores all registers from the buffer passed, including the frame pointer and stack pointer. It then places the return value passed into accumulator A and sets the ccr to reflect the return value just stored in accumulator A. The function `longjmp()` discards the return program counter on the hardware stack and jumps to the address pointed to by the program counter stored in the buffer.

This file must conform to the include files `setjmp.h` and `longjmp.h`. Since these two algorithms are very straightforward, modification of the file may be not needed. If modification is absolutely necessary, then the ANSI standard of the functions `setjmp()` and `longjmp()` should be followed.

6.11 Host-Supported I/O (`printf()`, et al)

The library provided with DSP56KCC includes a full implementation of the ANSI C standard input and output routines. These routines already work with Gdb56 and run56, and can easily be embedded in custom applications. Anywhere that formatted I/O is desired, these library routines can be included to simplify development. The entire suite of routines is based upon a simple communication protocol between the DSP and a host resident I/O driver, so porting the entire system to custom hardware is trivial.

6.11.1 DSP functions `__send ()` and `__receive ()`

All standard I/O functions, no matter how complicated, are built upon two simple communication functions, `__send ()`, and `__receive ()`. `__send ()` sends a message to the I/O driver code residing on the host. `__receive ()` retrieves a message from that same driver. Implementing these two functions is all that need be done on the DSP in order to support standard I/O on custom hardware.

It is assumed that some sort of hardware communication channel exists between the host and the DSP. `__send ()` and `__receive ()` implement a simple message passing mechanism on top of such a channel. `__send ()` accepts two arguments: the address of the buffer to send, and the number of words to draw from that buffer. `__receive ()` accepts the address of a buffer in which to place the received message as its single argument. All of the interactions between the host and DSP are driven by the library code running on the DSP; because the DSP is in control, it knows the size of return messages from the host, rendering a count argument to the function `__receive ()` superfluous. `__send ()` and `__receive ()` are as simple as they seem; the complexity of the standard I/O package is embedded in the host-side driver and the `lib56c[xyl].clb` library routines.

6.11.2 The Host-Side I/O Driver

The application running on the host must have the provided I/O driver embedded in it. The driver is written in C, and uses typically available library routines such as `open ()`, `close ()`, `read ()`, and `write ()` to perform the actions requested by the DSP. The I/O code is written as an event-driven state machine, so that the host side application can perform concurrently with the DSP when the DSP is not requesting I/O activity. In fact, the I/O driver on the host may be interrupt driven.

The host-side driver consists of the code in `dsp/etc/hostio.h` and `dsp/etc/hostio.c`. The meat of the package consists of two functions, `init_host_io_structures ()` and `process_pending_host_io ()`, and two buffers `hio_send` and `hio_receive`. The function `init_host_io_structures ()` is called to initialize host-side driver data structures before DSP execution is commenced. The buffer `hio_send` is used to send messages to the DSP, and `hio_receive` is used to receive messages from the DSP. The function `process_pending_host_io ()` considers the current state of the buffers and its own internal state, and then performs any required buffer modification or host I/O.

6.11.3 Communication between the Host and DSP

The messages passed between the DSP and the host's I/O driver are defined in the file `dsp/include/ioprim.h`. All sequences of communication are initiated by the DSP as a direct result of a call to a standard I/O function. Each standard I/O call may initiate a series of

messages between the DSP and the host, with the host eventually returning a message containing the completion status of the original request. The file `dsp/include/ioprim.h` is included by both the host-side I/O driver, and the standard I/O library code; it defines the constant definitions used in the aforementioned messages. A typical series of events and messages that comprise a standard I/O call might look like this:

1. The application running on the DSP makes a call to **fopen ()**.
2. The library code in `lib56c[xyl].clb` calls **__send ()**, with a buffer that contains the code `DSP_OPEN`, the flags, the mode, and the string length of the path.
3. The host receives the message into the buffer **hio_receive**, sets its valid flag, and calls **process_pending_host_io ()**. The state machine inside **process_pending_host_io ()** notes that it is now in the middle of an open file request, records the values from the first message, and then returns. At this point, code written by the application developer must check the valid flag of the buffer `hio_send`; in this case, the buffer **hio_send** has not been marked valid.
4. The library code in `lib56c[xyl].clb` calls **__send ()** again, this time sending the path.
5. Again, the host receives the message into the buffer **hio_receive**, and calls **process_pending_host_io ()** after setting the buffer's valid flag.
6. **process_pending_host_io ()** uses the information from the two messages to perform the file open. It then builds an operation status message, places it in the buffer **hio_send**, and sets that buffer's valid flag. **process_pending_host_io ()** resets its internal state and returns.
7. The host checks the buffer valid flag on **hio_send**, sees that it is true, and transmits the message to the DSP.
8. The library code running on the DSP finishes the **fopen ()** call and returns.

On the host side of the interface, the application writer must write the code that exchanges data with the DSP, the code that calls **process_pending_host_io ()**, and the code that checks buffer valid flags. On the DSP side of the interface, the application writer must write the routines **__send ()** and **__receive ()**. The communication between the DSP and the host is always initiated by the DSP and always follows a predetermined pattern, depending on the initial message. Because this communication is so simple, the code that calls **process_pending_host_io ()** can also be quite simple.

Example 6-7 is a *hypothetical* non-reentrant interrupt handler written in C. It uses two functions, **peek ()** and **poke ()**, to access some sort of hardware communication device connected to the DSP. The functions **peek ()** and **poke ()** aren't provided; they're simply an abstraction for host-side hardware access. This code assumes that the DSP sends the size of a message directly before sending a message. `CHECK_BUFFER_SIZE` is a macro

defined in dsp/etc/hostio.h. It should always be used to ensure that the buffer **hio_receive** is large enough to handle the incoming message. Finally, this example assumes that the function **signal ()** is available to register interrupt handlers.

This code doesn't have to be implemented in an interrupt driven manner; periodic polling could be used as well. The critical issues are that the communication must be reliable, and that the system must not deadlock; the latter is easy to ensure, given the simple nature of the communication protocol.

Example 6-7. Sample Host-Side Glue Code

```

void interrupt_driven_io ( )
{
    int i;

    /* get the size of the message from the DSP. */
    int size = peek ( IN_PORT );

    /* make sure that hio_receive is large enough. */
    CHECK_BUFFER_SIZE ( & hio_receive, size );

    /* read the message via hardware, into the buffer. */
    for ( i = 0; i < size; ++ i )
    {
        hio_receive.buffer[i] = peek ( IN_PORT );
    }

    /* mark the buffer as valid, perform any requested I/O.
    */
    hio_receive.valid_p = TRUE;
    process_pending_host_io ( );

    /* if the driver wants to send to the DSP, then do so
    now. */
    if ( hio_send.valid_p )
    {
        for ( i = 0; i < hio_send.length; ++ i )
        {
            poke ( OUT_PORT, hio_send.buffer[i] );
        }
        hio_send.valid_p = FALSE;
    }

    /* re-register this handler for future DSP message
    interrupts. */
    signal ( SIG_IN_PORT, interrupt_driven_io );
}

```

Appendix A

Programming Support

A.1 Standard ANSI Header Files

Each function provided in the ANSI C library `lib56c[xy].c1b` is declared with the appropriate prototype in a *header*, whose contents are made available by the `#include` preprocessing directive. In addition to function prototypes, the header also defines all `type` and `MACRO` definitions required for ANSI conformance. Table A-3 lists the ANSI standard header files:

Table A-1. ANSI Standard Header Files

Header Files		
<code><assert.h></code>	<code><locale.h></code>	<code><stddef.h></code>
<code><ctype.h></code>	<code><math.h></code>	<code><stdio.h></code>
<code><errno.h></code>	<code><setjmp.h></code>	<code><stdlib.h></code>
<code><float.h></code>	<code><signal.h></code>	<code><string.h></code>
<code><limits.h></code>	<code><stdarg.h></code>	<code><time.h></code>

In general, header files may be included in any order and as often as desired with no negative side effects. The one exception occurs when including `<assert.h>`. The definition of the `assert` macro depends on the definition of the macro `NDEBUG`.

The contents of the standard header files provided are *exactly* as described by the ANSI document X3.159-1989 dated December 14, 1989.

A.2 ANSI C Library Sub-routines

There are three library images provided with each compiler,

- `lib56cx.clb` — x memory model
- `lib56cy.clb` — y memory model
- `lib56cl.clb` — l memory model

Each library contains all of the ANSI defined sub-routines compiled and/or assembled for the specific memory model. The C and assembly language source for each library routine is distributed free of charge with the compiler.

A.2.1 Hosted vs. Non-Hosted Library Routines

Some of the standard ANSI defined routines perform I/O. Programs that use these functions will not encounter problems when run using `gdb56` or `run56`. Extra work may be needed to port hosted I/O routines to custom hardware configurations. Non-hosted routines will not encounter any problems on custom hardware.

For a description of `run56`, see Appendix C, “Utilities.”

A.3 Forcing Library Routines Out-of-line

For performance reasons, several run-time library routines have been created to be used *in-line* by the compiler. The compiler *in-lines* a sub-routine by replacing the occurrence of the sub-routine call with the body of the sub-routine itself. This provides execution time benefits:

1. Eliminates sub-routine call overhead. This is a substantial portion of the run-time for small library routines.
2. Exposes more optimization opportunities to the optimizer.

The library routines in Table A-2 are automatically in-lined by the compiler when their header file is included:

Table A-2. In-Line Library Routines

Header: ctype.h				Header: math.h	Header: stdlib.h	Header: string.h
isalnum	isalpha	iscntrl	isdigit	ceil	abs	strcmp
isgraph	islower	isprint	ispunct	fabs	labs	strcpy
isspace	isupper	isxdigit	tolower	floor	—	—
toupper	—	—	—	fmod	—	—

When it is necessary to disable this feature, possibly for debugging or decreasing program size, simply do one of the following.

1. Add the following line to each C module (or once to a common header file)

```
#undef ROUTINE_NAME
```

where ROUTINE_NAME is the library routine that must be forced out-of-line. For example, to force the library routine `ceil` out-of-line:

```
#undef ceil
```

2. Use the command-line option `-U`, see Chapter 3, *Control Program Options*. This will force the library routine to be called for this compilation. If the code is re-compiled, the `-U` option must be used again.

```
C:\> g56k -Uceil file.c
```

3. Do *not* include the header file.

A.4 Function Descriptions

Table A-3 describes each function in complete detail. The synopsis provides the syntax of the function and the options section discusses each option in detail. Many function descriptions also include references to related functions and an example of how to use the function. The list shown below provides an abbreviated description of each function.

Table A-3. Function Descriptions

Function	Description
abort	Force abnormal program termination.
abs	Absolute value of integer.
acos	Arc cosine.
asin	Arc sine.
atan	Arc tangent.
atan2	Arc tangent of angle defined by point y/x.
atexit	Register a function for execution upon normal program termination.
atof	String to floating point.
atoi	String to integer.
atol	String to long integer.
bsearch	Perform binary search.
calloc	Dynamically allocate zero-initialized storage for objects.
ceil	Ceiling function.
clearerr	Clear error indicators associated with a stream.
cos	Cosine.
cosh	Hyperbolic cosine.
div	integer division with remainder.
exit	Terminate program normally.
exp	Exponential, e^x .
fabs	Absolute value of a double.
fclose	Close a stream.
feof	Test the end-of-file indicator of a stream.
ferror	Test the error indicator of a stream.

Table A-3. Function Descriptions (Continued)

Function	Description
fflush	Flush all output pending on a stream.
fgetc	Read a character from a stream.
fgetpos	Retrieve the current value of the file position indicator of a stream.
fgets	Read a string from a stream.
floor	Floor function.
fmod	Floating point remainder.
fopen	Open a named file on the disk, to be accessed via a stream.
fprintf	Write formatted output to a stream.
fputc	Write a character to a stream.
fputs	Write a string to a stream.
fread	Read unformatted input from a stream.
free	Free storage allocated by calloc, malloc, and realloc.
freopen	Open a named file on the disk, to be accessed via a stream.
frexp	Break a floating point number into mantissa and exponent.
fscanf	Read formatted input from a stream.
fseek	Set a stream's file position indicator.
fsetpos	Set a stream's file position indicator.
ftell	Retrieve the current value of a stream's file position indicator.
fwrite	Write unformatted output to a stream.
fgetc	Read a character from a stream (this may be a macro).
getchar	Read a character from the stream stdin (this may be a macro).
gets	Read a string from the stream stdin.
isalnum	Test for alphanumeric character.
isalpha	Test for alphabetic character.
iscntrl	Test for control character.
isdigit	Test for numeric character.
isgraph	Test for printing character, excluding space and tab.
islower	Test for lower-case alphabetic characters.

Table A-3. Function Descriptions (Continued)

Function	Description
isprint	Test for printing character, excluding '\t'.
ispunct	Test for punctuation character.
isspace	Test for white-space character.
isupper	Test for upper-case alphabetic character.
isxdigit	Test for hexadecimal numeric character.
labs	Absolute value of a long integer.
ldexp	Multiply floating point number by a power of two.
ldiv	Long integer division with remainder.
log	Natural logarithm, base e.
log10	Base ten logarithm.
longjmp	Execute a non-local jump.
malloc	Dynamically allocate uninitialized storage.
mblen	Length of a multibyte character.
mbstowcs	Convert multibyte string to wide character string.
mbtowc	Convert a multibyte character to a wide character.
memchr	Find a character in a memory area.
memcmp	Compare portion of two memory areas.
memcpy	Copy from one area to another.
memmove	Copy from one area to another (source and destination may overlap).
memset	Initialize memory area.
modf	Break a double into it's integral and fractional parts.
perror	Print error message indicated by errno.
pow	Raise a double to a power.
printf	Write formatted output to the stream stdout.
putc	Write a character to a stream (this may be a macro).
putchar	Write a character to the stream stdout (this may be a macro).
puts	Write a string to the stream stdout.
qsort	Quick sort.

Table A-3. Function Descriptions (Continued)

Function	Description
raise	Raise a signal.
rand	Pseudo- random number generator.
realloc	Change size of dynamically allocated storage area.
remove	Remove a file from the disk.
rename	Rename a file on the disk.
rewind	Reset the file position indicator of a stream to the beginning of the file on the disk.
scanf	Read formatted input from the stream stdin.
setjmp	Save a reference of the current calling environment for later use by longjmp.
setbuf	Associate a buffer with a stream.
setvbuf	Associate a buffer with a stream, while also specifying the buffering mode and buffer size.
signal	Set up signal handler.
sin	Sine.
sinh	Hyperbolic Sine.
sprintf	Write formatted output to a string.
sqrt	Square root.
srand	Seed the pseudo-random number generator.
sscanf	Read formatted input from a string.
strcat	Concatenate two strings.
strchr	Find first occurrence of a character in a string.
strcmp	Compare two strings.
strcoll	Compare two strings based on current locale.
strcpy	Copy one string into another.
strcspn	Compute the length of the prefix of a string not contained in a second string.
strerror	Map error code into an error message string.
strlen	Determine length of a string.
strncat	Concatenate a portion of one string to another.
strncmp	Compare a portions of two strings.
strncpy	Copy a portion of one string into another.

Table A-3. Function Descriptions (Continued)

Function	Description
strpbrk	Find the first occurrence of a character from one string in another.
strrchr	Find the last occurrence of a character in a string.
strspn	Compute the length of the prefix of a string contained in a second string.
strstr	Find the first occurrence of one string in another.
strtod	String to double.
strtok	Break string into tokens.
strtol	String to long integer.
strtoul	String to unsigned long integer.
strxfrm	Transform a string into locale-independent form.
tan	Tangent.
tanh	Hyperbolic tangent.
tmpfile	Create a temporary binary file on the disk to be referenced via a stream.
tmpnam	Generate a unique, valid file name.
tolower	Convert uppercase character to lowercase.
toupper	Convert lowercase character to uppercase.
ungetc	Push a character back onto a specified input stream.
vfprintf	Write formatted output to a stream, using a va_list.
vprintf	Write formatted output to the stream stdout, using a va_list.
vsprintf	Write formatted output to a string, using a va_list.
wcstombs	Convert wchar_t array to multibyte string.
wctomb	Convert wchar_t character to multibyte character.

A.4.1 abort—Force abnormal program termination

```
#include <stdlib.h>
void abort(void);
```

The `abort` function causes the program to terminate abnormally unless the signal `SIGABRT` is being caught and the signal handler does not return. The unsuccessful termination value, `-1`, is returned to the host environment (`gdb56`, `run56`). The `abort` function will not return to its caller.

Refer to the following sections for more information:

- Section A.4.18, “exit—Terminate program normally,” on page A-25
- Section A.4.90, “signal—Set up signal handler,” on page A-94

Example A-1. abort

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

void main()
{
printf("-- make abort call --\n");
abort();
printf("this line not reached\n");
}
```

prints to standard output:
-- make abort call --

A.4.2 abs—Absolute value of integer

```
#include <stdlib.h>
int abs(int j);
```

The `abs` function returns the absolute value of `j`. When the header file `stdlib.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Refer to the following sections for more information:

- Section A.4.20, “fabs—Absolute value of a double,” on page A-26
- Section A.4.57, “labs—Absolute value of a long integer,” on page A-58

Example A-2. abs

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

void main()
{
int neg = -709;
printf("-- abs(%d) == %d --\n", neg,abs( neg ));
}
```

prints to standard output:
-- abs(-709) == 709 --

A.4.3 acos—Arc cosine

```
#include <math.h>
double acos( double x );
```

The `acos` function computes the principle value of the arc cosine of `x` in the range $[0.0, \pi]$, where `x` is in radians. If `x` is not in the range $[-1, +1]$, a domain error occurs, `errno` is set to `EDOM` and a value of `0.0` is returned.

Refer to the following sections for more information:

- Section A.4.15, “cos—Cosine,” on page A-23

Example A-3. `acos`

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

void main()
{
double point;

for ( point = -0.8 ; point < 1.0 ; point += 0.2 )
{
printf( "%f ", acos( point ) );
if ( point >= 0.0 ) printf( "\n" );
}
}
```

prints to standard output:

```
2.498090 2.214290 1.982310 1.772150 1.570790 1.369430
1.159270
0.927295
0.643501
0.000345
```

A.4.4 asin—Arc sine

```
#include <math.h>
double asin( double x );
```

The `asin` function computes the principle value of the arc sine of `x` in the range $[-\pi/2, +\pi/2]$, where `x` is in radians. If `x` is not in the range $[-1, +1]$, a domain error occurs, `errno` is set to `EDOM` and a value of `0.0` is returned.

Refer to the following sections for more information:

- Section A.4.91, “sin—Sine,” on page A-95

Example A-4. asin

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

void main()
{
double point;

for ( point = -0.8 ; point < 1.0 ; point += 0.2 )
{
printf( "%f ", asin( point ) );
if ( point >= 0.0 ) printf( "\n" );
}
}
```

prints to standard output:

```
-0.927295 -0.643501 -0.411516 -0.201357 -0.000000 0.201357
0.411516
0.643501
0.927295
1.570450
```

A.4.5 atan—Arc tangent

```
#include <math.h>
double atan( double x );
```

The `atan` function computes the principle value of the arc tangent of `x` in the range $[-\pi/2, +\pi/2]$, where `x` is in radians.

Refer to the following sections for more information:

- Section A.4.117, “tan—Tangent,” on page A-120

Example A-5. atan

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

void main()
{
double point;

for ( point = -0.8 ; point < 1.0 ; point += 0.2 )
{
printf( "%f ", atan( point ) );
if ( point >= 0.0 ) printf( "\n" );
}
}
```

prints to standard output:

```
-0.674740 -0.540419 -0.380506 -0.197395 -0.000000 0.197395
0.380506
0.540419
0.674740
0.785398
```

A.4.6 atan2—Arc tangent of angle defined by point y/x

```
#include <math.h>
double atan2( double y, double x );
```

The `atan2` function computes the principle value of the arc tangent of y/x using the signs of both arguments to determine the quadrant of the return value. If both arguments are zero, `errno` is set to `EDOM` and 0.0 is returned. See Table A-4.

Table A-4. Argument Range/Output Range

<u>Argument Range</u>	<u>Output Range</u>
$y \geq 0.0, x \geq 0.0$	$[0.0, \pi/2]$
$y \geq 0.0, x < 0.0$	$[\pi/2, \pi]$
$y < 0.0, x < 0.0$	$[-\pi, -\pi/2]$
$y < 0.0, x \geq 0.0$	$[-\pi/2, 0.0]$

Refer to the following sections for more information:

- Section A.4.5, “atan—Arc tangent,” on page A-13
- Section A.4.117, “tan—Tangent,” on page A-120

Example A-6. `attan2`

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

void main()
{
    printf("atan2(7.09,7.09) == %f\n", atan2(7.09,7.09));
    printf("atan2(-7.09,7.09) == %f\n", atan2(-7.09,7.09));
    printf("atan2(7.09,-7.09) == %f\n", atan2(7.09,-7.09));
    printf("atan2(-7.09,-7.09) == %f\n", atan2(-7.09,-7.09));
}
```

prints to standard output:

```
atan2( 7.09, 7.09 ) == .785398
atan2( -7.09, 7.09 ) == -.785398
atan2( 7.09, -7.09 ) == 2.356194
atan2( -7.09, -7.09 ) == -2.356194
```

A.4.7 atexit—Register a function for execution at normal program termination

```
#include <stdlib.h>
int atexit( void (*func)(void) );
```

The `atexit` registers a function `func` that will be called at normal program execution. The registered function is called without arguments and returns nothing.

A total of 32 functions may be registered and will be called in the order in which they were registered. The `atexit` function returns zero if registration succeeds and a non-zero value for failure.

Refer to the following sections for more information:

- Section A.4.18, “`exit`—Terminate program normally,” on page A-25

Example A-7. `atexit`

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

void main()
{
    atexit( func_1 );
    atexit( func_2 );
    printf( "-- testing atexit --\n" );
}

void func_1( void )
{
    printf ( "first function called\n" );
}

void func_2( void )
{
    printf ( "second function called\n" );
}
```

prints to standard output:
-- testing atexit --
second function called
first function called

A.4.8 atof—String to floating point

```
#include <stdlib.h>
double atof( const char* nptr );
```

The `atof` function converts the string pointed to by `nptr` to a `double`. If the result can not be represented, the behavior is undefined. This is exactly equivalent to:

```
strtod( nptr, (char **) NULL );
```

Refer to the following sections for more information:

- Section A.4.9, “atoi—String to integer,” on page A-17
- Section A.4.10, “atol—String to long integer,” on page A-18
- Section A.4.112, “strtod—String to double,” on page A-114
- Section A.4.114, “strtol—String to long integer,” on page A-116
- Section A.4.115, “strtoul—String to unsigned long integer,” on page A-118

Example A-8. atof

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

void main()
{
printf( "atof( \"7.09\" ) == %f\n", atof( "7.09" ) );
}
```

prints to standard output:
atof("7.09") == 7.090000

A.4.9 atoi—String to integer

```
#include <stdlib.h>
int atoi( const char* nptr );
```

The `atoi` function converts the string pointed to by `nptr` to an integer. If the result can not be represented, the behavior is undefined. This is exactly equivalent to:

```
(int) strtol( nptr, (char **) NULL, 10 );
```

Refer to the following sections for more information:

- Section A.4.8, “`atof`—String to floating point,” on page A-16
- Section A.4.10, “`atol`—String to long integer,” on page A-18
- Section A.4.112, “`strtod`—String to double,” on page A-114
- Section A.4.114, “`strtol`—String to long integer,” on page A-116
- Section A.4.115, “`strtoul`—String to unsigned long integer,” on page A-118

Example A-9. `atoi`

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

void main()
{
printf( "atoi( \"709\" ) == %d\n", atoi( "709" ) );
}
```

prints to standard output:
`atoi("709") == 709`

A.4.10 atol—String to long integer

```
#include <stdlib.h>
long atol( const char* nptr );
```

The `atol` function converts the string pointed to by `nptr` to a long integer. If the result can not be represented, the behavior is undefined. This is exactly equivalent to:

```
strtol( nptr, (char **) NULL, 10 );
```

Refer to the following sections for more information:

- Section A.4.8, “`atof`—String to floating point,” on page A-16
- Section A.4.9, “`atoi`—String to integer,” on page A-17
- Section A.4.112, “`strtod`—String to double,” on page A-114
- Section A.4.114, “`strtol`—String to long integer,” on page A-116
- Section A.4.115, “`strtoul`—String to unsigned long integer,” on page A-118

Example A-10. `atol`

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

void main()
{
printf( "atol( \"709\" ) == %ld\n", atol( "709" ) );
}
```

prints to standard output:
`atol("709") == 709`

A.4.11 bsearch—Perform binary search

```
#include <stdlib.h>
void bsearch(const void* key, const void* base,
size_t nmem, size_t size,
int (*compare) (const void*, const void* ) );
```

The `bsearch` function searches an array on `nmem` objects (the initial element of which is pointed to by **base**) for an element that matches the object pointed to by **key**. The size of each element is specified by **size**. The contents of the array must be in ascending order according to a user supplied comparison function, `compare`. The `compare` function is called with two arguments that must point to the **key** object and to an array member, in that order. Also, `compare` must return an integer that is less than, equal to, or greater than zero for the **key** object to be respectively considered less than, equal to or greater than the array element.

Refer to the following sections for more information:

- Section A.4.79, “qsort—Quick sort,” on page A-86

Example A-11. bsearch

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

char* stuff[6] =
{
    "bald", "driving", "feet", "flintstone", "fred", "on"
};

int compare ( const void *key, const void *aelement )
{
    return ( strcmp( *(char*) key, *(char*) aelement ) );
}

void main()
{
    char* p;
    char* key = "bald";
    p = bsearch( &key, stuff, 6, sizeof(char*), compare );
    if ( p )
    {
        printf( "YES, fred flintstone drives on bald feet\n" );
    }

    else
    {
        printf( "NO, sam the butcher brings alice the meat\n" );
    }
}
```

prints to standard output:
YES, fred flintstone drives on bald feet

A.4.12 `calloc`—Dynamically allocate zero-initialized storage for objects

```
#include <stdlib.h>
void* calloc( size_t nmemb, size_t size );
```

The `calloc` function allocates space for an array of `nmemb` objects, each of whose size is **size**. The space is initialized to all bits equal zero. If space can not be allocated, `calloc` returns a `NULL` pointer.

Refer to the following sections for more information:

- Section A.4.35, “`free`—Free storage allocated by `calloc`, `malloc`, and `realloc`,” on page A-39
- Section A.4.63, “`malloc`—Dynamically allocate uninitialized storage,” on page A-64
- Section A.4.82, “`realloc`—change size of dynamically allocated storage area,” on page A-88
- **`alloca`**—Allocate storage in the stack frame.

Example A-12. calloc

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

int* iptr;

/* allocate space for 709 integers */
void main()
{
    iptr= (int*) calloc( 709, sizeof(int) );

    if ( iptr != NULL )
    {
        /* check first entry for zero initialization */
        if ( *iptr != 0 )
        {
            printf( "error: calloc failed to initialize\n" );
        }

        else
        {
            printf( "success: calloc ok\n" );
        }

        else
        {
            printf( "error: calloc failed\n" );
        }
    }

    prints to standard output:
    success: calloc ok

```

A.4.13 `ceil`—Ceiling function

```
#include <math.h>
double ceil( double x );
```

The `ceil` function returns the smallest integer greater than or equal to `x`.

When the header file `math.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Refer to the following sections for more information:

- Section A.4.28, “`floor`—Floor function,” on page A-33

Example A-13. `ceil`

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

void main()
{
printf( "ceil( 7.09 ) == %f\n", ceil( 7.09 ) );
}
```

prints to standard output:
`ceil(7.09) == 8.000000`

A.4.14 `clearerr`—Clear any error indicators associated with a specified stream

```
#include <stdio.h>
void clearerr ( FILE *stream );
```

The `clearerr` function clears the end-of-file and error indicators for the specified stream.

Example A-14. `clearerr`

```
#include <stdio.h>

void main ()
{
FILE *stream = tmpfile (); /* initially empty. */
clearerr ( stream );
printf ( "end-of-file indicator is: %d\n", feof ( stream ));
}
```

prints to standard output:
`end-of-file indicator is: 0`

A.4.15 cos—Cosine

```
#include <math.h>
double cos( double x );
```

The `cos` function computes and returns the cosine of `x`, measured in radians.

Refer to the following sections for more information:

- Section A.4.3, “acos—Arc cosine,” on page A-11

Example A-15. cos

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

void main()
{
    printf( "cos( 45.0 ) == %f\n", cos( 45.0 ) );
}
```

prints to standard output:
cos(45.0) == 0.525322

A.4.16 cosh—Hyperbolic cosine

```
#include <math.h>
double cosh( double x );
```

The `cosh` function computes and returns the hyperbolic cosine of `x`. If the value of `x` is too large, a range error occurs, setting `errno` to `ERANGE` and causes `cosh` to return `HUGE_VAL`.

Refer to the following sections for more information:

- Section A.4.92, “sinh—Hyperbolic Sine,” on page A-96
- Section A.4.118, “tanh—Hyperbolic tangent,” on page A-121

Example A-16. cosh

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

void main()
{
    printf( "cosh( 3.1415 ) == %f\n", cosh( 3.1415 ) );
}
```

prints to standard output:
cosh(3.1415) == 11.590883

A.4.17 `div`—integer division with remainder

```
#include <stdlib.h>
div_t div( int numer, int denom );
```

The `div` function computes the quotient and remainder of the division of the numerator **numer** by the denominator **denom** and returns them in a structure of type `div_t`. If the result can not be represented, the behavior is undefined.

Refer to the following sections for more information:

- Section A.4.59, “`ldiv`—Long integer division with remainder,” on page A-60

Example A-17. `div`

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

void main()
{
    div_t result;
    int numer = 709, denom = 56;

    result = div( numer, denom );
    printf( "quotient == %d\t", result.quot );
    printf( "remainder == %d\n", result.rem);
}
```

prints to standard output:
quotient == 12 remainder == 37

A.4.18 exit—Terminate program normally

```
#include <stdlib.h>
void exit( int status );
```

The `exit` function causes normal program termination to occur. Any functions registered with the function `atexit` are called in the order in which they were registered. **status** is returned to the environment (`gdb56`, `run56`). If more than one call is made to `exit`, the result is undefined.

Refer to the following sections for more information:

- Section A.4.1, “abort—Force abnormal program termination,” on page A-9
- Section A.4.7, “atexit—Register a function for execution at normal program termination,” on page A-15

Example A-18. exit

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

void main()
{
printf( "-- exit test --\n" );
exit ( 0 ); /* return with 0 status */
printf( "Error: exit made this unreachable\n" );
}
```

prints to standard output:
-- exit test --

A.4.19 `exp`—Exponential, e^x

```
#include <math.h>
double exp( double x );
```

The `exp` function computes and returns e^x . If the value of `x` is too large, a range error occurs with `errno` being set to `ERANGE` and `exp` returns `HUGE_VAL`. If the value of `x` is too small, a range error will also occur with `errno` being set to `ERANGE` and `exp` returns 0.0.

Refer to the following sections for more information:

- Section A.4.58, “`ldexp`—Multiply floating point number by 2,” on page A-59
- Section A.4.74, “`pow`—Raise a double to a power,” on page A-77

Example A-19. `exp`

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

void main()
{
printf( "exp( 7.09 ) == %f\n", exp( 7.09 ) );
}
```

prints to standard output:
`exp(7.09) == 1199.907686`

A.4.20 `fabs`—Absolute value of a double

```
#include <math.h>
double fabs( double x );
```

The `fabs` function computes and returns the absolute value of double `x`.

When the header file `math.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Refer to the following sections for more information:

- Section A.4.2, “`abs`—Absolute value of integer,” on page A-10
- Section A.4.57, “`labs`—Absolute value of a long integer,” on page A-58

Example A-20. fabs

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

void main()
{
double pos, neg = -7.09;

pos = fabs( neg );
printf( "-- absolute value of %d == %d --\n", neg, pos );
}
```

prints to standard output:
-- absolute value of -7.090000 == 7.090000 --

A.4.21 fclose—Close a stream

```
#include <stdio.h>
int fclose ( FILE * );
```

The function `fclose` flushes all output on the specified stream, and disassociates the stream from the file on the host. The function `fclose` returns `EOF` if there are any problems, otherwise `0`.

Refer to the following sections for more information:

- Section A.4.31, “`fprintf`—Write formatted output to a stream,” on page A-36

Example A-21. fclose

```
#include <stdio.h>

void main()
{
fprintf ( stdout, "see me second" );
fprintf ( stderr, "see me first\n" );
fclose ( stdout );
}
```

prints to combined standard error and standard output:
see me first
see me second

Note that `stdout` is by default line buffered, while `stderr` is not. The call to `fclose` causes the pending output on `stdout` to be flushed.

A.4.22 feof—Test the end-of-file indicator of a specified stream

```
#include <stdio.h>
int feof ( FILE * );
```

The function `feof` function tests the end-of-file indicator associated with the specified stream. It returns non-zero if and only if the end-of-file indicator is set.

Refer to the following sections for more information:

- Section A.4.30, “fopen—Open a named file on the host’s disk,” on page A-34
- Section A.4.39, “fseek—Set the file position indicator associated with a stream,” on page A-44
- Section A.4.31, “fprintf—Write formatted output to a stream,” on page A-36

Example A-22. feof

```
#include <stdio.h>

void main()
{
FILE *somefile = fopen ( "somefile", "rb+" );
(void) fseek ( somefile, 0L, SEEK_END );
(void) fgetc ( somefile );
fprintf ( stdout, "are we at the file's end? %s\n",
feof ( somefile ) ? "yes" : "no" );
}
prints to standard output:
are we at the file's end? yes
```

A.4.23 `ferror`—Test the error indicator of a stream

```
#include <stdio.h>
int ferror ( FILE * );
```

The function `ferror` function tests the error indicator associated with the specified stream. It returns non-zero if and only if the error indicator is set. `ferror` should be used following one or more stream I/O function calls.

A.4.24 `fflush`—Flush all pending output associated with a stream

```
#include <stdio.h>
void fflush ( FILE* );
```

The function `fflush` causes any pending output associated with the specified stream to be written to the output device.

Refer to the following sections for more information:

- Section A.4.31, “`fprintf`—Write formatted output to a stream,” on page A-36

Example A-23. `fflush`

```
#include <stdio.h>

void main()
{
    fprintf ( stdout, "see me second" );
    fprintf ( stderr, "see me first\n" );
    fflush ( stdout );
}
```

prints to combined standard error and standard output:
see me first
see me second

Note that `stdout` is by default line buffered, while `stderr` is not. The call to `fflush` causes the pending output on `stdout` to be flushed.

A.4.25 fgetc—Read a character from the specified stream

```
#include <stdio.h>
int fgetc ( FILE *stream );
```

The function `fgetc` will retrieve the next input character from the specified stream. If the stream is associated with a file on the disk, then the file position indicator is advanced. On error, `fgetc` returns EOF.

Refer to the following sections for more information:

- Section A.4.32, “fputc—Write a single character to a stream,” on page A-36

Example A-24. fgetc

```
#include <stdio.h>

void main ()
{
char value = (char) fgetc ( stdin );

while ( EOF != value )
{
fputc ( value, stdout );
value = (char) fgetc ( stdin );
}
}
```

will echo all characters from standard input to standard output until the input is exhausted.

A.4.26 `fgetpos`—Get value of file position indicator associated with a stream

```
#include <stdio.h>
int fgetpos ( FILE *stream, fpos_t *pos );
```

The function `fgetpos` fetches the value of the file position indicator associated with `stream` and stores it in the object pointed to by `pos`. `fgetpos` returns zero if it was successful. The value of the file position indicator is meaningless except as an argument to the function `fsetpos`.

Refer to the following sections for more information:

- Section A.4.30, “`fopen`—Open a named file on the host’s disk,” on page A-34
- Section A.4.39, “`fseek`—Set the file position indicator associated with a stream,” on page A-44
- **`fsetpos`**—Set the value of the file position indicator associated with a stream. `fsetpos` is equivalent to `fseek (stream, pos, seek, set)`.

Example A-25. `fgetpos`

```
#include <stdio.h>

void main ()
{
FILE *preexisting = fopen ( "already.here", "r" );
fpos_t pos;

(void) fgetpos ( preexisting, & pos );
(void) fseek ( preexisting, 0L, SEEK_END );
(void) fsetpos ( preexisting, & pos );
}
```

will open a hypothetical pre-existing file on the disk, record the initial position in `pos`, seek to the end of the file, and finally restore the initial value of the file position indicator.

A.4.27 fgets—Read a string from the specified stream

```
#include <stdio.h>
char *fgets ( char *s, int n, FILE *stream );
```

The function `fgets` will read at most `n - 1` characters from the specified stream. `fgets` will not read past a newline character. The characters are stored in memory starting at the location pointed to by `s`. `fgets` returns `s` if it was successful, `NULL` otherwise. `fgets` will update the file position indicator for any characters read.

Refer to the following sections for more information:

- Section A.4.33, “`fputs`—Write a string to a stream,” on page A-37
- Section A.4.85, “`rewind`—Reset the file position indicator to the beginning of the file,” on page A-90

Example A-26. fgets

```
#include <stdio.h>

void main ()
{
FILE *disk_file = fopen ( "newfile", "a+" );
char one_line[64];

fputs ( "read this line\n" "but not this line.\n", disk_file );
rewind ( disk_file );
fgets ( one_line, 64, disk_file );
fputs ( one_line, stdout );
}
```

will open a new file on the disk named “newfile”. Two lines will be written to the new file, and the first line will be read back using `fgets`. The retrieved line is then printed to standard output as follows:

```
read this line
```

A.4.28 floor—Floor function

```
#include <math.h>
double floor( double x );
```

The `floor` function returns the largest integer not greater than `x`.

When the header file `math.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Refer to the following sections for more information:

- Section A.4.13, “ceil—Ceiling function,” on page A-22

Example A-27. floor

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

void main()
{
printf( "floor( 7.09 ) == %f\n", floor( 7.09 ) );
}
prints to standard output:
floor( 7.09 ) == 7.000000
```

A.4.29 fmod—Floating point remainder

```
#include <math.h>
double fmod( double x, double y );
```

The `fmod` function computes and returns the floating point remainder `r` of `x/y`. The remainder, `r`, has the same sign as `x` and `x == i * y + r`, where $|r| < |y|$. If `x` and `y` can not be represented, the result is undefined.

When the header file `math.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-28. fmod

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

void main()
{
printf( "fmod( -10.0, 3.0 ) == %f\n", fmod( -10.0, 3.0 ) );
}

prints to standard output:
fmod( -10.0, 3.0 ) == -1.00000000
```

A.4.30 fopen—Open a named file on the host’s disk

```
#include <stdio.h>
void fopen ( const char *filename, const char *mode );
```

The `fopen` function attempts to open a named file for access via a stream. If `fopen` is able to open the specified file, then the new stream associated with that file is returned. If `fopen` fails, then it returns `NULL`. The `mode` argument indicates the type of the stream, and how the stream may be accessed as shown in Table A-5:

Table A-5. Mode Arguments

Argument	Stream
“r”	text type, read only,
“w”	text type, write only,
“a”	text type, append only (write after end of current file),
“rb”	binary type, read only,
“wb”	binary type, write only,
“ab”	binary type, append only (write after end of current file),
“r+”	text type, read and write,
“w+”	text type, read and write (any pre-existing file is destroyed),
“a+”	text type, append only (read/write after end of current file),
“rb+”	binary type, read and write,
“wb+”	binary type, read and write (any pre-existing file is destroyed),
“ab+”	binary type, append only (read/write after end of current file).

Note: Opening a file that does not exist will fail if **r** is the first character in the mode string. When opened, the stream is initially line buffered.

Refer to the following sections for more information:

- Section A.4.33, “fputs—Write a string to a stream,” on page A-37
- Section A.4.27, “fgets—Read a string from the specified stream,” on page A-32
- Section A.4.31, “fprintf—Write formatted output to a stream,” on page A-36

Example A-29. fopen

```
#include <stdio.h>

void main ()
{
FILE *stream = fopen ( "file.new", "w" );
char data[64];

fprintf ( stream, "verify this\n" );
fclose ( stream );

stream = fopen ( "file.new", "r" );
fgets ( data, 64, stream );
fputs ( data, stdout );
}
```

This example opens an new text file, writes to the associated stream, and closes that stream. It then reopens the file and displays the first line of that file on standard output:

```
verify this
```

A.4.31 fprintf—Write formatted output to a stream

```
#include <stdio.h>
int fprintf ( FILE *stream, const char *format, ... );
```

The function `fprintf` functions exactly like the function `printf`, except that the output is directed to the specified stream rather than being automatically directed to standard output. Please use the description of argument values in the description of the `printf` function.

Refer to the following sections for more information:

- Section A.4.31, “fprintf—Write formatted output to a stream,” on page A-36

Example A-30. fprintf

```
#include <stdio.h>

void main ()
{
fprintf ( stdout, "hello world\n" );
}
```

Will cause the following output to be printed to standard output:

```
hello world
```

A.4.32 fputc—Write a single character to a stream

```
#include <stdio.h>
int fputc ( int c, FILE *stream );
```

The function `fputc` writes the character `c` to the specified stream.

Example A-31. fputc

```
#include <stdio.h>

void main ()
{
fputc ( (int) 'S', stdout );
fputc ( (int) 'h', stdout );
fputc ( (int) 'a', stdout );
fputc ( (int) 'd', stdout );
fputc ( (int) 'r', stdout );
fputc ( (int) 'a', stdout );
fputc ( (int) 'c', stdout );
fputc ( (int) 'k', stdout );
fputc ( (int) '\n', stdout );
}
```

Will cause the following output to be printed to standard output:

```
Shadrack
```

A.4.33 fputs—Write a string to a stream

```
#include <stdio.h>
int fputs ( const char *s, FILE *stream );
```

The function `fputc` writes the string `s` to the specified stream. The trailing `'\0'` in `s` is not written to the stream.

Example A-32. fputs

```
#include <stdio.h>

void main ()
{
fputs ( "hand me down pumas\n", stdout );
}
```

Will cause the following output to be printed to standard output:

```
hand me down pumas
```

A.4.34 fread—Read data directly from a stream

```
#include <stdio.h>
size_t fread ( void *ptr, size_t size, size_t nmem, FILE *stream );
```

The function `fread` reads raw data from the specified stream. The data is stored in memory starting with the location pointed to by `ptr`. The quantity of data is `size * nmem`. `fread` returns the number of elements successfully read.

Refer to the following sections for more information:

- Section A.4.31, “`fprintf`—Write formatted output to a stream,” on page A-36
- Section A.4.30, “`fopen`—Open a named file on the host’s disk,” on page A-34

Example A-33. `fread`

Assume that the disk file “professor” has as its contents the following string, including the trailing `‘\0’`:

```
“What’s another word for pirate treasure?”
```

The following C program uses `fread`:

```
#include <stdio.h>

void main ()
{
FILE *booty = fopen ( “professor”, “r” );
char buffer[64];

fread ( buffer, 63, sizeof ( char ), booty );
fprintf ( stdout, buffer );
}
```

and will cause the following output to be printed to standard output:

```
What’s another word for pirate treasure?”
```

A.4.35 free—Free storage allocated by calloc, malloc, and realloc

```
#include <stdlib.h>
void free( void* ptr);
```

The `free` function causes the space pointed to by `ptr` to be deallocated. Once deallocated it is available to be used again by future dynamic memory allocation requests. If `ptr` is `NULL`, `free` returns immediately. If the space pointed to by `ptr` has already been deallocated by a previous call to `free` or `realloc`, the behavior is undefined.

Refer to the following sections for more information:

- Section A.4.12, “calloc—Dynamically allocate zero-initialized storage for objects,” on page A-20
- Section A.4.63, “malloc—Dynamically allocate uninitialized storage,” on page A-64
- Section A.4.82, “realloc—change size of dynamically allocated storage area,” on page A-88

Example A-34. free

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

void main()
{
char* alloc;
if ( ( alloc = (char*) malloc( 709 ) ) == NULL )
{
printf( "malloc error\n" );
exit ( -1 );
}

/* free 709 words of memory */
free( alloc );
}
```

A.4.36 freopen—Open named file on disk, to be accessed via specified stream

```
#include <stdio.h>
FILE *freopen ( const char *filename, const char *mode, FILE *stream );
```

The function `freopen` opens the named disk file in the same manner as `fopen`. However, `freopen` associates that file with the specified stream. If successful, `freopen` returns its argument `stream`, and if unsuccessful, it returns `NULL`. The range of acceptable values for the mode argument are the same as those for `fopen`.

Refer to the following sections for more information:

- Section A.4.75, “printf—Print to standard output,” on page A-78

Example A-35. freopen

```
#include <stdio.h>

void main ()
{
    freopen ( "diskfile", "w", stdout );
    printf ( "hello world\n" );
}
```

This example redirects standard output to a file on the disk via **freopen**. The “hello world” output will not appear on the normal standard output device, but rather in the file “diskfile”.

A.4.37 frexp—Break a floating point number into mantissa and exponent

```
#include <math.h>
double frexp( double value, int* exp );
```

The `frexp` function breaks a floating point number into a normalized fraction (the mantissa) and an integral power of 2 (the exponent). The `frexp` function returns the mantissa and stores the exponent in the integer object pointed to by `exp`. The mantissa is returned with a magnitude in the range $[1/2, 1]$ or zero, such that value equals the mantissa times 2^{exp} . If value is zero, both the exponent and mantissa are zero.

Refer to the following sections for more information:

- Section A.4.72, “`modf`—Break a double into it’s integral and fractional parts,” on page A-75

Example A-36. `frexp`

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

void main()
{
    int exp;
    double mant;

    mantissa = frexp( 70.9, &exponent );
    printf( "mantissa == %f\texponent == %d\n", mant, exp );
}
```

prints to standard output:
mantissa == 0.553906exponent == 7

A.4.38 fscanf—Read formatted input from a stream

```
#include <stdio.h>
int fscanf ( FILE *stream, const char *format, ... );
```

The function `fscanf` reads input from the specified stream. It uses the string `format` as a guide for interpreting the input and storing values in memory. Subsequent arguments to `fscanf` are used as pointers to objects in memory that will receive the input values read from the stream.

The format string is composed of directives. These are parsed from left to right from the format string, and indicate how the input from the specified stream should be processed.

If `fscanf` fails to apply a directive, then it returns. Directives are composed of either white-space characters or normal characters. White space character sequences indicate that `fscanf` should read input from the specified stream up to the first non-white-space character. Directives that consist of non-white-space characters are processed in the following manner: input is read from the specified stream until the input character is not a member of the set of characters comprising the directive. Finally, directives can be conversion specifications; these directives begin with the ‘%’ character. They describe how `fscanf` should parse input, and how `fscanf` should synthesize a value to be stored in memory.

Conversion specifications are processed as follows. First, the stream is read until all white-space characters have been exhausted, unless ‘[’, ‘c’, or ‘n’ is part of the conversion specification. Second, a value is derived from the input stream according to the conversion specifier. A conversion specifier may be one of the following as shown in Table A-6:

Table A-6. Conversion Specifiers

Conversion Specifier	Description
‘d’	match a signed, decimal integer.
‘i’	match a signed integer, whose base is determined in the same manner as a C integer constant.
‘o’	match an octal integer.
‘u’	match an unsigned, decimal integer.
‘x’	match a signed, hexadecimal integer. (‘X’ is also valid).
‘e’, ‘f’, ‘g’	match a floating-point number. (‘E’, ‘F’ and ‘G’ are also valid).
‘s’	match a sequence of non-white-space characters, essentially scan a token string.
‘[’	match a non-empty sequence of characters from a set of expected characters, which are bounded by a following ‘]’.

Table A-6. Conversion Specifiers (Continued)

Conversion Specifier	Description
'c'	match a sequence of characters, as specified by the field width. As a default, scan only one character.
'n'	don't match anything, just store the number of characters read from the input stream during this call to <code>fscanf</code> .
'%'	match a '%' character.

`fscanf` returns EOF if an input failure is detected before any conversions take place. Otherwise it returns the number of assignments made. Note that an optional assignment suppression character '*' may follow the initial '%'. This character will cause `fscanf` to discard the converted value without advancing along the list of object pointers.

Refer to the following sections for more information:

- Section A.4.86, “scanf—Read formatted input from standard input,” on page A-91
- Section A.4.96, “sscanf—Read formatted input from a string,” on page A-98

Example A-37. fscanf

(a) The following program will, assuming that the input pending on standard input is “my 98”, store the three characters ‘m’, ‘y’, ‘\0’ will be stored in `word[]`, and 98 will be stored in `number`.

```
#include <stdio.h>

void main ()
{
char word[8];
int number;
fscanf ( stdin, " %s %d", word, & number );
}
```

(b) The following program will, assuming that the input pending on standard input is “yall come”, store the following five characters in the array `word`: ‘y’, ‘a’, ‘l’, ‘l’, ‘\0’.

```
#include <stdio.h>

void main ()
{
char word[8];
fscanf ( stdin, "%[lay]", word );
}
```

A.4.39 fseek—Set the file position indicator associated with a stream

```
#include <stdio.h>
int fseek ( FILE *stream, long int offset, int whence );
```

The function `fseek` will set the file position indicator associated with the specified stream, according to the values of `offset` plus an initial value indicated by `whence`. Initial values derived of `whence` values are as follows:

SEEK_SET—The initial value is the beginning of the file.

SEEK_CUR—The initial value is the current position in the file.

SEEK_END—The initial value is the end of the file.

The value of `offset` must either be zero or the value returned by a call to `ftell`.

Refer to the following sections for more information:

- Section A.4.30, “`fopen`—Open a named file on the host’s disk,” on page A-34
- Section A.4.25, “`fgetc`—Read a character from the specified stream,” on page A-30
- Section A.4.21, “`fclose`—Close a stream,” on page A-27

Example A-38. `fseek`

The following function will read the last character in a text file specified by the parameter `name`.

```
#include <stdio.h>

char last_in_file ( char *name )
{
FILE *tmp = fopen ( name, "r" );
char return_value;

(void) fseek ( tmp, -1L, SEEK_END );
return_value = (char) fgetc ( tmp );
fclose ( tmp );
return return_value;
}
```

A.4.40 **fsetpos**—Set the file position indicator associated with a stream

```
#include <stdio.h>
int fsetpos ( FILE *stream, const fpos_t *pos );
```

The function `fsetpos` will change the file position indicator associated with the specified stream. The value of `pos` must be the return value from a prior call to `fgetpos`. Note that a successful call to `fsetpos` will undo any effect of an immediately preceding `ungetc` on the same stream. If it is successful, `fsetpos` returns zero.

Refer to the following sections for more information:

- Section A.4.26, “`fgetpos`—Get value of file position indicator associated with a stream,” on page A-31

Example A-39. fsetpos

```
#include <stdio.h>

void main ()
{
FILE *preexisting = fopen ( "already.here", "r" );
fpos_t pos;

(void) fgetpos ( preexisting, & pos );
(void) fseek ( preexisting, 0L, SEEK_END );
(void) fsetpos ( preexisting, & pos );
}
```

will open a hypothetical pre-existing file on the disk, record the initial position in `pos`, seek to the end of the file, and finally restore the initial value of the file position indicator.

A.4.41 `ftell`—Get the file position indicator associated with a stream

```
#include <stdio.h>
long int ftell ( FILE *stream );
```

The function `ftell` will return the value of the file position indicator for the specified stream. The return value is usable only by the function `fseek`. `ftell` returns `-1L`.

Refer to the following sections for more information:

- Section A.4.30, “`fopen`—Open a named file on the host’s disk,” on page A-34
- Section A.4.34, “`fread`—Read data directly from a stream,” on page A-38
- Section A.4.77, “`putchar`—Write a character to standard output,” on page A-84
- Section A.4.39, “`fseek`—Set the file position indicator associated with a stream,” on page A-44

Example A-40. `ftell`

```
#include <stdio.h>

main ()
{
FILE *stream = fopen ( "file.abc", "r" );
long int beginning = ftell ( stream );
char send;

(void) fread ( & send, sizeof ( char ), 1, stream );
putchar ( send );
fseek ( stream, beginning, SEEK_SET );
(void) fread ( & send, sizeof ( char ), 1, stream );
putchar ( send );
}
```

will read and print the first character in the file “file.abc” twice. **`ftell`** is used to reset the file position indicator to the beginning of the file.

A.4.42 fwrite—Write data directly to a stream

```
#include <stdio.h>
size_t fwrite ( const void *ptr, size_t size, size_t nmemb, FILE *stream
);
```

The function `fwrite` writes raw data to the specified stream. The data is drawn from memory starting with the location pointed to by `ptr`. The quantity of data is `size * nmemb`. `fwrite` returns the number of elements successfully written.

Example A-41. fwrite

```
#include <stdio.h>

main ()
{
char message[] = "hand me down pumas";

fwrite ( message, sizeof ( char ), strlen ( message ), stdout );
}
```

will write the message "hand me down pumas" onto standard output.

A.4.43 `getc`—Read a character from the specified stream

```
#include <stdout.h>
int fgetc ( FILE *stream );
```

`getc` is equivalent to `fgetc`, except that `getc` may be implemented as a macro. If such is the case, the argument `stream` may be evaluated more than once. This only becomes a problem if evaluation of the argument has side effects.

Refer to the following sections for more information:

- Section A.4.25, “`fgetc`—Read a character from the specified stream,” on page A-30
- Section A.4.32, “`fputc`—Write a single character to a stream,” on page A-36

Example A-42. `getc`

```
#include <stdio.h>

void main ()
{
char value = (char) getc ( stdin );

while ( EOF != value )
{
fputc ( value, stdout );
value = (char) getc ( stdin );
}
}
```

will echo all characters from standard input to standard output until the input is exhausted.

A.4.44 `getchar`—Read a character from standard input

```
#include <stdio.h>
int getchar ( void );
```

The function `getchar` reads the next character from standard input. If there is no pending input, `getchar` returns `EOF`, otherwise the read character is cast to type `int` and returned.

Refer to the following sections for more information:

- Section A.4.32, “`fputc`—Write a single character to a stream,” on page A-36

Example A-43. `getchar`

```
#include <stdio.h>

void main ()
{
char value = (char) getchar ();

while ( EOF != value )
{
fputc ( value, stdout );
value = (char) getchar ( );
}
}
```

will echo all characters from standard input to standard output until the input is exhausted.

A.4.45 `gets`—Read a string from standard input

```
#include <stdio.h>
char *gets ( char *s );
```

The function `fgets` will read characters from standard input, and place them sequentially in memory, starting with the location pointed to by `s`. `gets` will not read past a newline character, or past the end of the file. The characters are stored in memory starting at the location pointed to by `s`. `gets` returns `s` if it was successful, `NULL` otherwise. `gets` will update the file position indicator for any characters read. If `gets` encounters the end of file on its first read of standard input, `NULL` is returned. `gets` does not append a ‘\0’ to the array of characters read.

Refer to the following sections for more information:

- Section A.4.27, “`fgets`—Read a string from the specified stream,” on page A-32
- Section A.4.78, “`puts`—Write a string to standard output,” on page A-85

Example A-44. `gets`

```
#include <stdio.h>

void main ()
{
char line_buffer[BUFSIZ];

puts ( gets ( line_buffer ) );
}
```

will echo a newline-terminated line of input from standard input onto standard output. Note that **`gets`** doesn’t read past the newline; **`puts`** supplies one by definition.

A.4.46 isalnum—Test for alphanumeric character

```
#include <ctype.h>
int isalnum( int c );
```

The `isalnum` function returns a nonzero value for any alphabetic or numeric character; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-45. islnum

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

void main()
{
if ( ( isalnum( 'c' ) ) && ( isalnum( '1' ) ) )
{
printf( "c, 1 -- alpha and numeric\n" );
}

if ( ! ( isalnum( '@' ) ) )
{
printf( "@ -- neither alpha nor numeric\n" );
}
}
```

prints to standard output:
c, 1 -- alpha and numeric
@ -- neither alpha nor numeric

A.4.47 isalpha—Test for alphabetic character

```
#include <ctype.h>
int isalpha( int c );
```

The `isalpha` function returns a nonzero value for any alphabetic character; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-46. isalpha

```
#include <stdio.h>
#include <ctype.h>
void main()
{
if ( ( isalpha( 'c' ) ) )
{
printf( "c -- alpha\n" );
}
if ( ! ( isalpha( '@' ) ) && ! ( isalpha( '1' ) ) )
{
printf( "@, 1 -- non alpha\n" );
}
}
prints to standard output:
c -- alpha
@,1 -- non alpha
```

A.4.48 iscntrl—Test for control character

```
#include <ctype.h>
int iscntrl( int c );
```

The `iscntrl` function returns a nonzero value for any control character; zero is returned in the false case. A control character is any character that is *NOT* a letter, digit, punctuation, or ' ' (the space character). This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-47. iscntrl

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

void main()
{
    /* check the "beep" character */
    if ( ( iscntrl( 0x07 ) ) )
    {
        printf( "\"beep\" (0x07) -- control character\n" );
    }

    if ( ! ( iscntrl( '@' ) ) )
    {
        printf( "@ -- not control character\n" );
    }
}
```

prints to standard output:
"beep" (0x07) -- control character
@ -- not control character

A.4.49 isdigit—Test for numeric character

```
#include <ctype.h>
int isdigit( int c );
```

The `isdigit` function returns a nonzero value for any *decimal* character; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-48. isdigit

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

void main()
{
if ( ( isdigit( '1' ) ) )
{
printf( "1 -- is a decimal character\n" );
}

if ( ! ( isdigit( 'f' ) ) )
{
printf( "f -- not a decimal character\n" );
}
}
```

prints to standard output:
1 -- is a decimal character
f -- not a decimal character

A.4.50 isgraph—Test for printing character, excluding space and tab

```
#include <ctype.h>
int isgraph( int c );
```

The `isgraph` function returns a nonzero value for any printable character, excluding space (‘ ’) and tab (‘\t’); zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”].

Example A-49. isgraph

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

void main()
{
    /* check the "beep" character */
    if ( ! ( isgraph( ' ' ) ) )
    {
        printf( "space -- not \"graph\" character\n" );
    }

    if ( ( isgraph( 'f' ) ) )
    {
        printf( "f -- \"graph\" character\n" );
    }
}
```

prints to standard output:
space -- not "graph" character
f -- "graph" character

A.4.51 islower—Test for lower-case alphabetic characters

```
#include <ctype.h>
int islower( int c );
```

The `islower` function returns a nonzero value for any lower-case alphabetic character; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-50. islower

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

void main()
{
if ( ( islower( 'a' ) ) )
{
printf( "a -- lower case character\n" );
}

if ( ! ( islower( 'F' ) ) )
{
printf( "F -- not a lower case character\n" );
}
}
```

prints to standard output:
a -- lower case character
F -- not a lower case character

A.4.52 isprint—Test for printing character, excluding '\t'

```
#include <ctype.h>
int isprint( int c );
```

The `isprint` function returns a nonzero value for any printable character, excluding the tab character ('\t'); zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, "Forcing Library Routines Out-of-line."]

Example A-51. isprint

```
#include <stdio.h>
#include <ctype.h>
void main()
{
if ( ( isprint( ' ' ) ) )
{
printf( "space -- \"print\" character\n" );
}
if ( ! ( isprint( '\t' ) ) )
{
printf( "tab -- not \"print\" character\n" );
}
}
```

prints to standard output:
space -- "print" character
tab -- not "print" character

A.4.53 `ispunct`—Test for punctuation character

```
#include <ctype.h>
int ispunct( int c );
```

The `ispunct` function returns a nonzero value for any punctuation character; zero is returned in the false case. A punctuation character is one that is printable, not a digit, not a letter, and not a space (‘ ’ or ‘\t’). This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-52. `ispunct`

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

void main()
{
    if ( ( ispunct( ' ' ) ) )
    {
        printf( "space -- \"print\" character\n" );
    }

    if ( ! ( ispunct( '\t' ) ) )
    {
        printf( "tab -- not \"print\" character\n" );
    }
}
```

```
prints to standard output:
space -- "print" character
tab -- not "print" character
```

A.4.54 `isspace`—Test for white-space character

```
#include <ctype.h>
int isspace( int c );
```

The `isspace` function returns a nonzero value for any standard white-space character; zero is returned in the false case. The standard white-space characters are space (‘ ’), form feed (‘\f’), new-line (‘\n’), carriage return (‘\r’), horizontal tab (‘\t’), and vertical tab (‘\v’). This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”]

Example A-53. isspace

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

void main()
{
if ( ( isspace( ' ' ) ) )
{
printf( "space -- white-space character\n" );
}

if ( ! ( isspace( '@' ) ) )
{
printf( "@ -- not white-space character\n" );
}
}
```

prints to standard output:
space -- white-space character
@ -- not white-space character

A.4.55 isupper—Test for upper-case alphabetic character

```
#include <ctype.h>
int isupper( int c );
```

The `isupper` function returns a nonzero value for any upper-case alphabetic character; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”].

Example A-54. isupper

```
#include <stdio.h>
#include <ctype.h>
void main()
{
if ( ( isupper( 'F' ) ) )
{
printf( "F -- upper-case character\n" );
}
if ( ! ( isupper( 'f' ) ) )
{
printf( "f -- not an upper-case character\n" );
}
}

prints to standard output:
F -- upper-case character
f -- not an upper-case character
```

A.4.56 isxdigit—Test for hexadecimal numeric character

```
#include <ctype.h>
int isxdigit( int c );
```

The `isxdigit` function returns a nonzero value for any hexadecimal digit; zero is returned in the false case. This function is provided both as an in-line and out-of-line function.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”].

Example A-55. isxdigit

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

void main()
{
if ( ( isxdigit ( 'F' ) ) )
{
printf( "F -- hexadecimal character\n" );
}

if ( ! ( isxdigit ( 'G' ) ) )
{
printf( "G -- not a hexadecimal character\n" );
}
}
```

prints to standard output:
F -- hexadecimal character
G -- not a hexadecimal character

A.4.57 labs—Absolute value of a long integer

```
#include <stdlib.h>
long int labs( long int j );
```

The `labs` function returns the absolute value of the long integer `j`.

Refer to the following sections for more information:

- Section A.4.2, “abs—Absolute value of integer,” on page A-10

Example A-56. labs

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

void main()
{
long int j = -190897091;

printf( "labs( -190897091 ) == %ld\n", labs( j ) );
}
```

prints to standard output:
labs(-190897091) == 19089709

A.4.58 ldexp—Multiply floating point number by 2

```
#include <math.h>
double ldexp( double x, int exp );
```

The `ldexp` function returns $x * 2^{\text{exp}}$. If the result exceeds `HUGE_VAL`, `errno` is set to `ERANGE` and the value `HUGE_VAL` is returned with the same sign as `x`.

Refer to the following sections for more information:

- Section A.4.19, “exp—Exponential, ex,” on page A-26
- Section A.4.74, “pow—Raise a double to a power,” on page A-77
- Section A.4.19, “exp—Exponential, ex,” on page A-26
- Section A.4.74, “pow—Raise a double to a power,” on page A-77

Example A-57. ldexp

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

void main()
{
printf( "ldexp(7.09,4) == %f\n", ldexp(7.09,4) );
}
```

prints to standard output:
ldexp(7.09,4) == 113.440000

A.4.59 ldiv—Long integer division with remainder

```
#include <stdlib.h>
ldiv_t ldiv( long int numer, long int denom );
```

The `ldiv` function computes the quotient and remainder of `numer / denom` and returns the result in a structure of type `ldiv_t`. If the result cannot be represented, the result is undefined.

Refer to the following sections for more information:

- Section A.4.17, “`div`—integer division with remainder,” on page A-24

Example A-58. ldiv

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

void main()
{
    ldiv_t result;
    long numer = 709, denom = 56;

    result = ldiv( numer, denom );
    printf( "quotient == %ld\t", result.quot );
    printf( "remainder == %ld\n", result.rem);
}
```

prints to standard output:
quotient == 12 remainder == 37

A.4.60 log—Natural logarithm, base e

```
#include <math.h>
double log( double x );
```

The `log` function computes the natural logarithm of `x`. If the value of `x` is less than zero, `errno` is set to `EDOM` and the value `HUGE_VAL` is returned. If `x` is equal to zero, `errno` is set to `ERANGE` and the value `HUGE_VAL` is returned.

Refer to the following sections for more information:

- Section A.4.19, “exp—Exponential, ex,” on page A-26
- Section A.4.61, “log10—Base ten logarithm,” on page A-62

Example A-59. log

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

void main()
{
printf( "log( 7.09 ) == %f\n", log( 7.09 ) );
}
```

prints to standard output:
log(7.09) == 1.958685

A.4.61 log10—Base ten logarithm

```
#include <math.h>
double log10( double x );
```

The `log10` function computes the natural logarithm of **x**. If the value of **x** is less than zero, `errno` is set to `EDOM` and the value `HUGE_VAL` is returned. If **x** is equal to zero, `errno` is set to `ERANGE` and the value `HUGE_VAL` is returned.

Refer to the following sections for more information:

- Section A.4.19, “`exp`—Exponential, `ex`,” on page A-26
- Section A.4.60, “`log`—Natural logarithm, base `e`,” on page A-61

Example A-60. log10

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

void main()
{
printf( "log10( 7.09 ) == %f\n", log10( 7.09 ) );
}
```

prints to standard output:
`log10(7.09) == 0.850646`

A.4.62 longjmp—Execute a non-local jump

```
#include <setjmp.h>
void longjmp( jmp_buf env, int val );
```

The `longjmp` function restores the calling environment referenced by `env` which must have been initialized by a previous call to `setjmp`. If there has been no invocation of `setjmp`, or if the function containing the call to `setjmp` has terminated execution before the call to `longjmp`, the behavior is undefined.

Upon completion of `longjmp`, program execution continues as if the corresponding call to `setjmp` had returned with a value **val**; if **val** is zero, 1 is returned.

Global and *volatile-type* variables have values as of the time `longjmp` was called; all register and non-volatile automatic variables will have indeterminate values.

For more information, see Chapter 6, “Software-Hardware Integration.”

Refer to the following sections for more information:

- Section A.4.87, “setjmp—Save reference of current calling environment for later use by longjmp,” on page A-92

Example A-61. longjmp

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

jmp_buf env;

void func( void )
{
    longjmp( env, -709 );
}

main()
{
    if ( setjmp( env ) != 0 )
    {
        printf( "-- longjmp has been called --\n" );
        exit( 1 );
    }
    printf( "-- setjmp called --\n" );
    func();
}
```

prints to standard output:

```
-- setjmp called --
-- longjmp called --
```

A.4.63 malloc—Dynamically allocate uninitialized storage

```
#include <stdlib.h>
void* malloc( size_t size );
```

The `malloc` function returns a pointer to the lowest word of a size word block of storage; when **size** exceeds the amount of memory available, `malloc` returns `NULL`.

Refer to the following sections for more information:

- Section A.4.12, “`calloc`—Dynamically allocate zero-initialized storage for objects,” on page A-20
- Section A.4.35, “`free`—Free storage allocated by `calloc`, `malloc`, and `realloc`,” on page A-39
- Section A.4.82, “`realloc`—change size of dynamically allocated storage area,” on page A-88
- **alloca**—Allocate storage in the stack frame.

Example A-62. malloc

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

void main()
{
char *char_array;

if((char_array=(char*) malloc(709*sizeof(char))) == NULL)
{
printf( "error: not enough memory\n" );
exit( 1 );
}
else
{
printf("-- space for 709 chars allocated OK --\n" );
}
}
```

prints to standard output:
-- space for 709 chars allocated OK --

A.4.64 mblen—Length of a multibyte character

```
#include <stdlib.h>
int mblen( const char* s, size_t n );
```

The `mblen` function determines the number of characters in the *multibyte* character pointed to by `s`. The `mblen` function is equivalent to

```
mbtowc( (wchar_t*) 0, s, n );
```

-
1. If **s** is a NULL pointer, **mblen** returns a zero. If **s** is not NULL, **mblen** returns
 2. zero if **s** points to a NULL character,
 3. the number of characters that comprise the multibyte character, or
 4. -1 if an invalid multi-byte character is formed.

In no case will the return value exceed **n** or the **MB_CUR_MAX** macro.

Refer to the following sections for more information:

- Section A.4.66, “mbtowc—Convert a multibyte character to a wide character,” on page A-68
- Section A.4.128, “wctomb—Convert wchar_t character to multibyte character,” on page A-128

Note: The DSP56000/DSP56001 does not provide byte addressing, thus characters always require an entire word of memory each. One way to better utilize data memory (with a run-time cost) is to combine the ANSI data type `wchar_t` and the special ANSI *multibyte* and *wide character* library routines.

Example A-63. `mblen`

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

char* gstr = NULL;

void main()
{
int max = MB_CUR_MAX;
char* strnull = gstr;
char* str1 = "709";

printf("mblen(strnull,5)==%d\n", mblen( strnull,5));
printf("mblen(str1, max ) == %d\n", mblen(str1,max));
printf("mblen(\"abcdef\",5) == %d\n", mblen("abcdef",5));
printf("mblen(\"abcdef\",2) == %d\n", mblen("abcdef",2));
}
```

prints to standard output:

```
mblen( strnull, 5 ) == 3
mblen( str1, 5 ) == 3
mblen( "abcdef", 5 ) == 3
mblen( "abcdef", 2 ) == 2
```

A.4.65 mbstowcs—Convert multibyte string to wide character string

```
#include <stdlib.h>
int mbstowcs( wchar_t* pwcs, const char* s, size_t n );
```

The `mbstowcs` function converts the character string pointed to by `s` into a wide character string pointed to by `pwcs`. Each character of the *multibyte* string is converted as if by the `mbtowc` function. At most, `n` characters will be converted and stored in the wide character string. Multibyte characters that follow a `NULL` character will not be examined or converted. If `s` and `pwcs` overlap, the behavior is undefined.

If an invalid character is encountered, `mbstowcs` returns `(size_t) -1`. Otherwise, `mbstowcs` returns the number of characters converted, not including the terminating `NULL` character.

Refer to the following sections for more information:

- Section A.4.127, “`wcstombs`—Convert `wchar_t` array to multibyte string,” on page A-126

Note: The DSP56000/DSP56001 does not provide byte addressing, thus characters always require an entire word of memory each. One way to better utilize data memory (with a run-time cost) is to combine the ANSI data type `wchar_t` and the special ANSI multibyte and wide character library routines.

Example A-64. mbstowcs

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

wchar_t warray[10];

void main()
{
char array = "abcdedgh";
char* ptr = array;
int convert;

convert = mbstowcs( warray, array, 10 );

printf( "unpacked array looks like:\n" );
while ( *ptr != 0 )
{
printf( "%0.6x ", *ptr++ );
}
printf( "\n\n" );

printf( "%d chars packed, packed array looks like:\n" );
ptr = warray;
while ( *ptr != 0 )
{
printf( "%0.6x ", *ptr++ );
}
printf( "\n" );
}
```

prints to standard output:
unpacked array looks like:
000061 000062 000063 000064 000065 000066 000067 000068

8 chars packed, packed array looks like:
616263 646566 676800

A.4.66 mbtowc—Convert a multibyte character to a wide character

```
#include <stdlib.h>
int mbtowc( wchar_t* pwc, const char* s, size_t n );
```

The `mbtowc` function examines the *multibyte* (i.e., multi-character) string pointed to by `s` and converts it into a wide character (`wchar_t`). At most, `n` and never more than `MB_CUR_MAX` characters from `s` will be examined and converted.

If `s` is a `NULL` pointer, `mbtowc` returns zero. If `s` is not `NULL`, `mbtowc` returns

1. zero if `s` points to a `NULL` character,
2. the number of characters that comprise the multibyte character, or
3. -1 if an invalid multibyte character is formed.

In no case will the return value exceed `n` or the `MB_CUR_MAX` macro.

Refer to the following sections for more information:

- Section A.4.64, “`mblen`—Length of a multibyte character,” on page A-64
- Section A.4.65, “`mbstowcs`—Convert multibyte string to wide character string,” on page A-66
- Section A.4.128, “`wctomb`—Convert `wchar_t` character to multibyte character,” on page A-128

Note: The DSP56000/DSP56001 does not provide byte addressing, thus characters always require an entire word of memory each. One way to better utilize data memory (with a run-time cost) is to combine the ANSI data type `wchar_t` and the special ANSI multibyte and wide character library routines.

Example A-65. mbtowc

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

void main()
{
    wchar_t wide = 0;
    char* mbstr = "abcde";
    int convert;

    convert = mbtowc( (wchar_t*) NULL, mbstr, 2 );
    printf("%d chars packed. wide == %0.6x\n", convert, wide);

    convert = mbtowc( &wide, mbstr, strlen( mbstr ) );
    printf("%d chars packed. wide == %0.6x\n", convert, wide);

    convert = mbtowc( &wide, mbstr, 2 );
    printf("%d chars packed. wide == %0.6x\n", convert, wide);
}
```

prints to standard output:
2 chars packed. wide == 000000
3 chars packed. wide == 616263
2 chars packed. wide == 006162

A.4.67 memchr—Find a character in a memory area

```
#include <string.h>
int memchr( const void* s, int c, size_t n );
```

The `memchr` function finds the first occurrence of `c` (converted to an unsigned char) in the memory area pointed to by `s`. The terminating null character is considered to be part of the string. The `memchr` function returns a pointer to the located char or a `NULL` pointer if the character is not found.

Refer to the following sections for more information:

- Section A.4.109, “`strrchr`—Find last occurrence of a character from one string in another,” on page A-111
- Section A.4.102, “`strcspn`—Compute length of prefix of one string consisting entirely of characters not from another,” on page A-104
- Section A.4.108, “`strpbrk`—Find first occurrence of a character from one string in another,” on page A-110
- Section A.4.109, “`strrchr`—Find last occurrence of a character from one string in another,” on page A-111
- Appendix A.4.110, “`strspn`—Compute length of prefix of one string consisting entirely of characters from another,” on page A-112

Note: Due to the word orientation of the DSP56000/DSP56001, there is no way to accurately implement an ANSI version of `memchr` for the **I** memory model.

Example A-66. memchr

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

void main()
{
char*string = "fred flintstone driving on bald feet";
char*result;

/* locate the occurrence of 'b' */
result = memchr( string, 'b', strlen( string ) );
printf( "-- %s --\n", result );
}
```

prints to standard output:
-- bald feet --

A.4.68 memcmp—Compare portion of two memory areas

```
#include <string.h>
int memcmp( const void* s1, const void* s2, size_t n );
```

The `memcmp` function compares the first `n` words of the object pointed to by `s1` with the first `n` words of the object pointed to by `s2`, comparison is lexicographical. The `memcmp` function returns zero if the two areas compared are equal, a value greater than zero if `s1` is greater, or a value less than zero if `s1` is smaller.

Refer to the following sections for more information:

- Section A.4.106, “`strncmp`—Compare a portions of two strings,” on page A-108

Note: Due to the word orientation of the DSP56000/DSP56001, there is no way to accurately implement an ANSI version of `memcmp` for the `I` memory model.

Example A-67. `memcmp`

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

struct test
{
char cartoon[20];
int value;
}g1 = { "flintstones", 709 },
g2 = { "flintstones", 709 },
g3 = { "jetsons", 709 };

void main()
{
if ( memcmp( &g1, &g2, sizeof( struct test ) ) != 0 )
{
printf( "error: flintstones differ\n" );
}
else
{
printf( "-- flintstones are flintstones --\n" );
}
if ( memcmp( &g1, &g3, sizeof( struct test ) ) != 0 )
{
printf( "-- flintstones are not jetsons --\n" );
}
else
{
printf( "error: flintstones are NOT jetsons\n" );
}
}
prints to standard output:
-- flintstones are flintstones --
-- flintstones are not jetsons --
```

A.4.69 memcpy—Copy from one area to another

```
#include <string.h>
int memcpy( void* s1, const void* s2, size_t n );
```

The `memcpy` function copies `n` words from the area referenced by `s2` into the area specified by `s1`. If the source and destination areas overlap, the results are undefined. The `memcpy` function returns the value of `s1`.

Refer to the following sections for more information:

- Section A.4.107, “`strncpy`—Copy a portion of one string into another,” on page A-109
- Section A.4.101, “`strcpy`—Copy one string into another,” on page A-102

Example A-68. `memcpy`

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

struct test
{
char cartoon[20];
int value;
}g1, g2 = { "flintstones", 709 };

void main()
{
memcpy( &g1, &g2, sizeof( struct test ) );
printf( "-- I watch the %s --\n", g1.cartoon );
}
```

prints to standard output:
-- I watch the flintstones --

A.4.70 memmove—Copy storage

```
#include <string.h>
int memmove( void* s1, const void* s2, size_t n );
```

The memmove function copies **n** words from the area referenced by **s2** into the area specified by **s1**. The copy is done by first placing the **n** words into a temporary buffer and then moving the temporary buffer into the final location, this allows the source and destination areas to overlap.

Refer to the following sections for more information:

- Section A.4.69, “memcpy—Copy from one area to another,” on page A-72

Example A-69. memmove

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

struct test
{
char cartoon[20];
int value;
} g1, g2 = {"flintstones", 709 };

void main()
{
memmove( &g1, &g2, sizeof( struct test ) );
printf( "-- I watch the %s --\n", g1.cartoon );
}
```

prints to standard output:
-- I watch the flintstones --

A.4.71 memset—Initialize memory area

```
#include <string.h>
int memset( void* s, int c, size_t n );
```

The `memset` function copies the value `c` (converted to an **unsigned char**) into the first `n` words of the object pointed to by `s`.

Refer to the following sections for more information:

- Section A.4.69, “`memcpy`—Copy from one area to another,” on page A-72

Note: Due to the word orientation of the DSP56000/DSP56001, the **l** memory model version of `memset` may not act as expected when attempting to initialize double size elements (the `x` memory portion is unaffected).

Example A-70. `memset`

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

struct test
{
    char cartoon[20];
    int value;
}

void main()
{
    struct test local;

    /* auto struct local is initialized to all nines */
    memset( &local, 9, sizeof( struct test ) );

    /* random check */
    if ( local.cartoon[7] != 9 )
    {
        printf( "error: memset busted\n" );
    }
    else
    {
        printf( "-- memset OK --\n" );
    }
}
```

prints to standard output:
-- memset OK --

A.4.72 `modf`—Break a double into its integral and fractional parts

```
#include <math.h>
double modf( double value, double* iptr );
```

The `modf` function breaks `value` into its fractional and integral parts. The `modf` function returns the fractional portion of **value** and stores the integral portion in the double object pointed to by `iptr`.

Refer to the following sections for more information:

- Section A.4.37, “`frexp`—Break a floating point number into mantissa and exponent,” on page A-41

Example A-71. `modf`

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

void main()
{
double result;

printf( "-- fractional == %f\t", modf( 7.09, &result );
printf( "integral == %f --\n", result );
}
```

prints to standard output:

```
-- fractional == 0.090000 integral == 7.000000 --
```

A.4.73 perror—Print error message

```
#include <stdio.h>
void perror( const char* s );
```

The `perror` function prints out the string `s` followed by “:” and the error message associated with `errno`.

Refer to the following sections for more information:

- Section A.4.103, “`strerror`—Map error code into an error message string,” on page A-105

Example A-72. `perror`

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

void main()
{
double result;

result = asin( 7.09 );

if ( result == 0.0 && errno == EDOM )
{
perror( "asin perror test" );
}
}
```

prints to standard output:
asin perror test:domain error

A.4.74 pow—Raise a double to a power

```
#include <math.h>
double pow( double x, double y );
```

The `pow` function computes and returns x^y . If x is zero and y is less than zero, a domain error occurs setting `errno` to `EDOM` and returning 0.0. If $|x^y|$ is greater than `HUGE_VAL`, `errno` is set to `ERANGE` and `HUGE_VAL` is returned.

Refer to the following sections for more information:

- Section A.4.19, “exp—Exponential, ex,” on page A-26
- Section A.4.58, “ldexp—Multiply floating point number by 2,” on page A-59

Example A-73. pow

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

void main()
{
printf( "-- pow( 2.0, 2.0 ) == %f --\n", pow( 2.0, 2.0 ) );
}
```

prints to standard output:
-- pow(2.0, 2.0) == 4.000000 --

A.4.75 printf—Print to standard output

```
#include <stdio.h>
int printf( const char* format, ... );
```

The `printf` function formats and writes a string to the standard output. Interpreting the format specifier `format` left to right.

The format specifier, `format`, consists of ordinary characters, escape sequences, and conversion specifications. The conversion specifications describe how arguments passed to `printf` are converted for output. All non-conversion specifying portions of `format` are sent directly to the standard output. If the number of arguments passed is less than specified by the `format` string, `printf` will write non-deterministic garbage to the standard output. If too many arguments are provided to `printf`, the extras will be evaluated but otherwise ignored.

A conversion specification is introduced by the character `%`, and has the following form:

```
 %[flags][field width][.precision][size]conversion character
```

where `flags`, `field width`, `precision`, `h`, `l`, `L` are optional.

Flags are for justification of output and printing of signs, blanks, decimal points, octal and hexadecimal prefixes. Multiple flags may be utilized at once. The ANSI flags are as shown in Table A-7:

Table A-7. ANSsbl Flags

Flag	Description
-	Left justify the result within the field. <u>The default is right justified.</u>
+	The result of a signed conversion will always have a sign (+ or -). <u>The default case provides only for -.</u>
space	If the first character of a signed conversion is not a sign, or if a signed conversion results in no characters, a space character will be prefixed to the result. If the space and the + flags both appear, the space flag is ignored. <u>The default mode is no space.</u>
#	The result is converted to an alternate form specified by the conversion character. For o conversion, it forces the first digit of the result to be a zero. For x (or X) conversion, the non-zero result will have 0x (0X) prefixed to it. For e , E , f , g , and G conversions, the result will always contain a decimal point character, even if no digits follow it. Additionally for g and G , trailing zeros will not be removed.
0	For d , i , o , u , x , X , e , E , f , g , and G conversions, leading zeros (following any indication of sign or base) are used to pad to the <i>field width</i> ; no space padding is performed. If the 0 and - flags both appear, the 0 flag will be ignored.

Each conversion takes place in *character fields*. The minimum size of the field can be specified with the field width. If the converted result contains fewer characters than specified by field width, the result will be left padded with spaces by default (see *flags* above). The field width takes the form of a decimal integer or an asterisk '*'. When the field width is an asterisk, the value is to be taken from an integer argument that precedes the argument to be converted.

Precision specifies the minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x**, **X** conversions, the number of digits appear after the decimal point character for **e**, **E**, and **f** conversions, the maximum number of significant digits for the **g** and **G** conversions, or the maximum number of characters to be written from a string in the **s** conversion. *The precision takes the form of a '.' followed by '*', or by an optional decimal integer; if only the period is specified, the precision is taken to be zero.* If precision appears with any other conversion character, the behavior is undefined.

Size specifies the argument size expected. There are three size specifiers defined by ANSI. The **h** specifies that the argument for the conversion characters **d**, **i**, **o**, **u**, **x**, or **X** will be unsigned short. The **l** specifies that the argument for the conversion characters **d**, **i**, **o**, **u**, **x**, or **X** will be long integer. The **L** specifies that the argument for the conversion characters **e**, **E**, **f**, **g**, or **G** will be long double.

There are 16 conversion characters; each is described in Table A-7.

Table A-8. Conversion Characters

Conversion Character	Description
d, i	The <code>int</code> argument is printed as a signed decimal number. The precision specifies the minimum number of digits to appear; if the value being printed can be represented in fewer digits, it is expanded with leading zeros. <u>The default precision is 1.</u> The result of printing a zero with precision zero is no characters (this is independent of padding specified by field width).
o	The <code>unsigned int</code> argument is printed as an unsigned octal number. When used in association with the <code>#</code> flag, 0 will be prefixed to non-zero results. The precision specifies the minimum number of digits to appear; if the value can be represented in fewer digits, it will be expanded with leading zeros. <u>The default precision is 1.</u> The result of printing a zero with precision zero is no characters (this is independent of padding specified by field width).
u	The <code>unsigned int</code> argument is printed as an unsigned decimal number. The precision specifies the minimum number of digits to appear; if the value can be represented in fewer digits, it will be expanded with leading zeros. <u>The default precision is 1.</u> The result of printing a zero with precision zero is no characters (this is independent of padding specified by field width).
x, X	The <code>unsigned int</code> argument is printed as an unsigned hexadecimal number. Hexadecimal alpha characters (a,b,c,d,e,f) will be printed in lower case when <code>x</code> is used and in upper case when <code>X</code> is used. When used in association with the <code>#</code> flag, 0x will be prefixed to the result (0X in the <code>X</code> case). Precision specifies the minimum number of digits to appear; if the value can be represented in fewer digits, it will be expanded with leading zeros. <u>The default precision is 1.</u> The result of printing a zero with precision zero is no characters (this is independent of padding specified by field width).
f	The <code>double</code> argument is printed out in decimal notation of the form <code>[-]ddd.ddd</code> , where precision specifies the number of digits to follow the decimal point. <u>The default precision 6.</u> When precision is 0 and the <code>#</code> flag is not specified, no decimal point character will be printed. A decimal digit will always follow at least one digit. The value printed is rounded to the appropriate number of digits.
e, E	The <code>double</code> argument is printed out in the form <code>[-] d.ddde±dd</code> , where precision specifies the number of digits to follow the decimal point. <u>The default precision 6.</u> When precision is 0 and the <code>#</code> flag is not specified, no decimal point character will be printed. A decimal digit will always follow at least one digit. The exponent always contains at least two digits.
g, G	The <code>double</code> argument is printed in the <code>f</code> , <code>e</code> , or <code>E</code> form. The <code>f</code> form is used unless the exponent to be printed is less than <code>-4</code> or greater than the precision . If precision is zero, the printed value consists of no characters (this is independent of padding specified by field width). Trailing zeros are removed from the fractional portion of the result; a decimal point character is printed only if it is followed by a digit.
c	The <code>int</code> argument is printed as an unsigned character.

Table A-8. Conversion Characters (Continued)

Conversion Character	Description
s	The argument is a pointer to a character string ((char*)). Characters from the string are printed up to (but not including) a terminating null character or until precision characters have been printed. If precision is not explicitly specified or is greater than the length of the string, the string will be printed until the null character is encountered.
p	The argument is a pointer to the void data type ((void*)). The value of the of the pointer is printed out as a hexadecimal digit.
n	The argument is a pointer to an integer ((int*)) which is the number of characters printed so far by the current call to <code>printf</code> .
%	Print the percent character, <code>%</code> . Note that the complete specifier is <code>%%</code> .

On successful completion, `printf` returns the an integer equal to the number of characters printed. On failure, `printf` returns an integer less than 0.

Refer to the following sections for more information:

- Section A.4.86, “scanf—Read formatted input from standard input,” on page A-91
- Section A.4.96, “sscanf—Read formatted input from a string,” on page A-98
- Section A.4.93, “sprintf—Print to a string,” on page A-96

A.4.76 `putc`—Write a single character to a stream

```
#include <stdio.h>
int putc ( int c, FILE *stream );
```

The function `putc` writes the character `c` to the specified stream. It is identical to the function `fputc`, except that `putc` may be implemented as a macro. This means that arguments to `putc` may be evaluated more than once. This is only a problem for function arguments that have side effects when evaluated.

Refer to the following sections for more information:

- Section A.4.32, “`fputc`—Write a single character to a stream,” on page A-36

Example A-75. `putc`

```
#include <stdio.h>

void main ()
{
    putc ( (int) 'S', stdout );
    putc ( (int) 'h', stdout );
    putc ( (int) 'a', stdout );
    putc ( (int) 'd', stdout );
    putc ( (int) 'r', stdout );
    putc ( (int) 'a', stdout );
    putc ( (int) 'c', stdout );
    putc ( (int) 'k', stdout );
    putc ( (int) '\n', stdout );
}
```

Will cause the following output to be printed to standard output:

```
Shadrack
```

A.4.77 putchar—Write a character to standard output

```
#include <stdio.h>
int putchar( int c );
```

The `putchar` function prints a character to standard output.

Refer to the following sections for more information:

- Section A.4.44, “`getchar`—Read a character from standard input,” on page A-48
- Section A.4.45, “`gets`—Read a string from standard input,” on page A-49
- Section A.4.78, “`puts`—Write a string to standard output,” on page A-85

Example A-76. `putchar`

```
#include <stdio.h>

char* str = "bald feet\n";

void main()
{
while ( *str != '\0' )
{
putchar ( *str++ );
}
}
```

prints to standard output:
bald feet

A.4.78 puts—Write a string to standard output

```
#include <stdio.h>
int puts( const char* s );
```

The `puts` function prints a string to standard output, appending a newline character. The `puts` function returns a zero if operation is successful and a non-zero value on failure.

Refer to the following sections for more information:

- Section A.4.44, “`getchar`—Read a character from standard input,” on page A-48
- Section A.4.45, “`gets`—Read a string from standard input,” on page A-49
- Section A.4.77, “`putchar`—Write a character to standard output,” on page A-84

Example A-77. `puts`

```
#include <stdio.h>

char* str = "bald feet";

void main()
{
puts ( str );
}
```

prints to standard output:
bald feet

A.4.79 qsort—Quick sort

```
#include <stdlib.h>
void qsort(void* base, size_t nmem, size_t size,
int (*compar) (const void*, const void* ) );
```

The `qsort` function sorts an array of `nmem` objects of size `size`, pointed to by **base**.

The array is sorted in ascending order according to a comparison function pointed to by `compar` which is called with two pointers to the array members. The `compar` function must return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second argument.

Example A-78. qsort

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char* stuff[] = {
"fred", "flintstone", "driving", "bald", "on", "feet"
};

static int compare( const char** a1, const char** a2 )
{
return( strcmp( *a1, *a2 ) );
}

main()
{
int i;

qsort(stuff, (size_t)6, (size_t)sizeof(char*), compare);

for ( i = 0 ; i < 6 ; i++ )
{
printf( "%s\t", stuff[i] );
}

printf( "\n" );
}
```

prints to standard output:
balddrivingfeetflintstonefredon

A.4.80 raise—Raise a signal

```
#include <signal.h>
int raise( int sig );
```

The `raise` function sends the signal `sig` to the executing program and returns 0 if successful or non-zero if unsuccessful. See `signal.h` for list of available signals and their default actions.

For more information, see Chapter 6, “Software-Hardware Integration.”

Refer to the following sections for more information:

- Section A.4.90, “signal—Set up signal handler,” on page A-94

Example A-79. raise

```
#include <stdio.h>

void main()
{
int onintr();

signal (SIGINT, onintr);
raise( SIGINT );
}

onintr()
{
printf( "caught SIGINT, see ya ...\\n" );
exit( -9 );
}
```

prints to standard output:
caught SIGINT, see ya ...

A.4.81 rand—Pseudo- random number generator

```
#include <stdlib.h>
int rand( void );
```

The `rand` function computes and returns a sequence of pseudo-random integers in the range of 0 to 32767.

Refer to the following sections for more information:

- Section A.4.95, “`srand`—Seed the pseudo-random number generator,” on page A-98

Example A-80. `rand`

```
#include <stdio.h>

void main()
{
/* seed the random number sequence */
srand( 1638 );

/* spew out random numbers in the range 0 to 709 */
for ( ; ; )
{
printf( "%d", ( rand() ) % 709 );
}
}
```

prints to standard output:

```
545 174 307 657 220 267 66 47 651 268 284 338 35 615 403 etc.
```

A.4.82 realloc—change size of dynamically allocated storage area

```
#include <stdlib.h>
int realloc( void* ptr, size_t size );
```

The `realloc` function changes the size of the storage area pointed to by `ptr` to `size`. The contents of the storage area are unchanged. If the new storage area is larger, the value of the new area is indeterminate. If `ptr` is null, `realloc` acts like `malloc`. If `ptr` was not dynamically allocated or the area was previously deallocated by a call to `free`, the behavior is undefined.

If `realloc` is unable to allocate the new size storage area, a null pointer is returned and the original storage area is unchanged.

Refer to the following sections for more information:

- Section A.4.12, “`calloc`—Dynamically allocate zero-initialized storage for objects,” on page A-20

-
- Section A.4.35, “free—Free storage allocated by calloc, malloc, and realloc,” on page A-39
 - Section A.4.63, “malloc—Dynamically allocate uninitialized storage,” on page A-64
 - **alloca**—Allocate storage in the stack frame.

Example A-81. realloc

```
#include <stdio.h>

void main()
{
char* str;

if ( ( str = (char*) malloc( (size_t) 15 ) ) == NULL )
{
error( "malloc failed" );
exit (-8);
}

strcpy( str, "short string" );
printf( "%s\n", string );

/* allocate space for 40 character string */
if ((str = (char*)realloc(str, 40*sizeof(char)) )== NULL )
{
perror( "realloc test" );
exit ( -9 );
}

strcat( str, " becomes a long string" );
printf( "%s\n", str );
}

prints to standard output:
short string
short string becomes a long string
```

A.4.83 `remove`—Remove a file from the disk

```
#include <stdio.h>
int remove ( char *filename );
```

The function `remove` will eliminate the file associated with the specified filename. The effect of this call on open files may vary from host to host, and is considered undefined.

Example A-82. `remove`

```
#include <stdio.h>

void main()
{
remove ( "foo.exe" );
}
```

will remove the file "foo.exe" on the disk, if such a file exists.

A.4.84 `rename` —Rename a file on the disk

```
#include <stdio.h>
int rename ( const char *old, const char *new );
```

The function `rename` disassociates the a disk file from the name `old`, and associates it with the name `new`. The behavior of this call is undefined if there already exists a file associated with the name `new`. `rename` returns zero if it is successful. If it fails, the file remains associated with the old name, and is not altered in any way.

Example A-83. `rename`

```
#include <stdio.h>

void main()
{
rename ( "old.exe", "new.exe" );
}
```

will rename the file "old.exe" to "new.exe", provided that "old.exe" actually exists on the disk. Note that "old.exe" will cease to exist.

A.4.85 `rewind`—Reset the file position indicator to the beginning of the file

```
#include <stdio.h>
void rewind ( FILE *stream );
```

The function `rewind` will reset the file position indicator associated with the specified stream. Any pending error is also cleared.

Refer to the following sections for more information:

-
- Section A.4.26, “fgetpos—Get value of file position indicator associated with a stream,” on page A-31
 - **fsetpos**—Set the file position indicator value associated with a stream.

Example A-84. `rewind`

```
#include <stdio.h>

void main ()
{
FILE *preexisting = fopen ( "already.here", "r" );

putchar ( fgetc ( preexisting ));
rewind ( preexisting );
putchar ( fgetc ( preexisting ));
}
will print the first character in the file "already.here" onto standard
output twice.
```

A.4.86 `scanf`—Read formatted input from standard input

```
#include <stdio.h>
int scanf (char *format, ... );
```

The function `scanf` is equivalent to the `fscanf` function, except that input is always read from standard input. Please use the description of argument values in the description of the `fscanf` function.

Refer to the following sections for more information:

- Section A.4.38, “fscanf—Read formatted input from a stream,” on page A-42

Example A-85. `scanf`

See the manual entry for **fscanf** for examples. The only difference between **scanf** and **fscanf** is that **scanf** does not require a **FILE*** argument; **stdin** is implied.

A.4.87 setjmp—Save reference of current calling environment for later use by longjmp

```
#include <setjmp.h>
int setjmp( jmp_buf env );
```

The `setjmp` function saves its calling environment in `env` for later use by `longjmp`. If the return is direct from `setjmp`, the value zero is returned. If the return is from the `longjmp` function, the value returned is non-zero.

For more information, see Chapter 6, “Software-Hardware Integration.” .

Refer to the following sections for more information:

- Section A.4.62, “`longjmp`—Execute a non-local jump,” on page A-63

Example A-86. setjmp

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

jmp_buf env;

void func( void )
{
    longjmp( env, -709 );
}

void main()
{
    if ( setjmp( env ) != 0 )
    {
        printf( "-- longjmp has been called --\n" );
        exit( 1 );
    }
    printf( "-- setjmp called --\n" );
    func();
}
```

prints to standard output:

```
-- setjmp called --
-- longjmp called --
```

A.4.88 `setbuf`—Read formatted input from standard input

```
#include <stdio.h>
void setbuf ( FILE *stream, char *buf );
```

If `buf` is `NULL`, the specified stream will be unbuffered. If `buf` is non-`NULL`, then the stream will be fully buffered with a buffer of size `BUFSIZ`. Note that `setbuf` must be used only before any other operations are performed on the specified stream, and that the stream argument must be associated with an opened file. Calling `setbuf` is equivalent to calling `setvbuf` using `_IOBUF` for the mode argument, and `BUFSIZ` for the `size` argument.

Refer to the following sections for more information:

- Section A.4.89, “`setvbuf`—Alter stream buffering,” on page A-94

A.4.89 setvbuf—Alter stream buffering

```
#include <stdio.h>
int setvbuf ( FILE *stream, char *buf, int mode, size_t size );
```

The function `setvbuf` is used to alter the way a specified stream is buffered. It must only be used before any other operation is performed on the specified stream. The argument `mode` determines the buffering policy:

_IOFBF—Use the full size of the buffer in the most efficient way.

_IOLBF—Use a line buffering policy: flush on newlines.

_IONBF—Do not buffer the stream at all.

The argument `size` specifies the buffer size for this stream. The pointer `buf`, if non-NULL, may be used for stream buffering. If `buf` is NULL, then `setvbuf` will allocate any needed buffer.

Refer to the following sections for more information:

- Section A.4.88, “`setbuf`—Read formatted input from standard input,” on page A-93

A.4.90 signal—Set up signal handler

```
#include <setjmp.h>
void (*signal( int sig, void (*func)(int) ) ) (int);
```

The `signal` function chooses one of three ways in which to handle the receipt of signal `sig`.

1. If the value of `func` is the macro `SIG_DFL`, default handling for the signal will occur.
2. If the value of `func` is the macro `SIG_IGN`, the signal is ignored.
3. Otherwise, `func` is a pointer to a function that will be called on the receipt of signal `sig`.

When a signal occurs, the signal handler for `sig` is reset to `SIG_DFL`; this is equivalent to making the call `signal(sig, SIG_DFL)`. The function `func` terminates by executing the `return` statement or by calling the `abort`, `exit`, or `longjmp` function. If the function `func` terminates with a `return` statement, execution continues at the point the signal was caught. Note that if the value of `sig` was `SIGFPE`, the behavior is undefined.

Also note that in order to continue catching signal `sig`, the signal handler must reissue the signal call.

For more information, see Chapter 6, “Software-Hardware Integration.”

Refer to the following sections for more information:

- Section A.4.80, “raise—Raise a signal,” on page A-87

Example A-87. signal

```
#include <stdio.h>

void main()
{
int onintr();

signal (SIGINT, onintr);
raise( SIGINT );
}

onintr()
{
printf( "caught SIGINT, see ya ...\\n" );
exit( -9 );
}
```

prints to standard output:
caught SIGINT, see ya ...

A.4.91 sin—Sine

```
#include <math.h>
double sin( double x );
```

The `sin` function computes and returns the sine of `x`, measured in radians.

Refer to the following sections for more information:

- Section A.4.4, “asin—Arc sine,” on page A-12

Example A-88. sin

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

void main()
{
printf( "sin( 45.0 ) == %f\\n", sin( 45.0 ) );
}
```

prints to standard output:
sin(45.0) == 0.850904

A.4.92 sinh—Hyperbolic Sine

```
#include <math.h>
double sinh( double x );
```

The `sinh` function computes and returns the hyperbolic sine of `x`, measured in radians. When the value of `x` is too large, `errno` will be set to `ERANGE` and the return value will be `HUGE_VAL` with the sign of `x`.

Refer to the following sections for more information:

- Section A.4.16, “cosh—Hyperbolic cosine,” on page A-23
- Section A.4.118, “tanh—Hyperbolic tangent,” on page A-121

Example A-89. sinh

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

void main()
{
    printf( "sinh( 3.1415 ) == %f\n", sinh( 3.1415 ) );
}
```

prints to standard output:
sinh(3.1415) == 11.547661

A.4.93 sprintf—Print to a string

```
#include <stdio.h>
int sprintf( char* s, const char* format, ... );
```

The `sprintf` function is equivalent to `printf` except that `s` specifies a string that the generated output is printed to rather than standard output. A null character is written at the end of the string. The `sprintf` function returns the number of characters written to the string.

Refer to the following sections for more information:

- Section A.4.75, “printf—Print to standard output,” on page A-78

Example A-90. sprintf

```
#include <stdio.h>

void main()
{
char buffer[256];
char* bptr = buffer;
char* str = "strings";
int i = 709, count;
double d = 7.09;

bptr += sprintf( bptr,"testing sprintf with:\n" );
sprintf( bptr, "\tstrings\t(%s)\n%n", str, &count );
bptr += count;
bptr += sprintf( bptr, "\thex digits\t%x\n", i );
bptr += sprintf( bptr, "\tfloating point\t%f\n", d );

puts( buffer );
}
```

```
prints to standard output:
testing sprintf with:
strings(strings)
hex digits2c5
floating point7.090000
```

A.4.94 sqrt—Square root

```
#include <math.h>
double sqrt( double x );
```

The `sqrt` function computes and returns the nonnegative square root of `x`. If `x` is less than zero, `errno` is set to `EDOM` and `0.0` is returned.

Example A-91. sqrt

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

void main()
{
double = 50.2681;

printf( "sqrt( 50.2681 ) == %.2f\n", sqrt( d ) );
}
```

```
prints to standard output:
sqrt( 50.2681 ) == 7.09
```

A.4.95 srand—Seed the pseudo-random number generator

```
#include <stdlib.h>
void sqrt( unsigned int seed );
```

The `srand` function uses the argument as a seed for a new sequence of pseudo-random numbers to be returned by `rand`. When `srand` is called with the same argument, the sequence of pseudo-random numbers will be repeated. If `srand` is not called, the default seed is 1.

Refer to the following sections for more information:

- Section A.4.81, “`rand`—Pseudo- random number generator,” on page A-88

Example A-92. `srand`

```
#include <stdio.h>

void main()
{
/* seed the random number sequence */
srand( 1638 );

/* spew out random numbers in the range 0 to 709 */
for ( ; ; )
{
printf( "%d", ( rand() ) % 709 );
}
}

prints to standard output:
545 174 307 657 220 267 66 47 651 268 284 338 35 615 403 etc.
```

A.4.96 sscanf—Read formatted input from a string

```
#include <stdio.h>
int sscanf ( const char *s, const char *format, ... );
```

The function `sscanf` reads formatted input from the string argument `s`, according to the format string `format`. The operation of `sscanf` is identical to `fscanf` except that input is read from a string.

Refer to the following sections for more information:

- Section A.4.38, “`fscanf`—Read formatted input from a stream,” on page A-42

A.4.97 strcat—Concatenate two strings

```
#include <string.h>
char* strcat( char* s1, const char* s2 );
```

The `strcat` function appends a copy of the string pointed to by `s2` (including the terminating null character) to the end of the string pointed to by `s1`. The first character of the second string is written over the first string's terminating character. The `strcat` function returns a pointer to `s1`.

Refer to the following sections for more information:

- Section A.4.105, “strncat—Concatenate a portion of one string to another,” on page A-106

Example A-93. strcat

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

void main()
{
char bigstr[80] = "string 1";
char smallstr[20] = " string 2";

printf("concatenate (%s) and (%s)\n", bigstr, smallstr);
(void) strcat( bigstr, smallstr );
puts( bigstr );
}
```

prints to standard output:
concatenate (string 1) and (string 2)
string 1 string 2

A.4.98 strchr—Find first occurrence of a character in a string

```
#include <string.h>
char* strchr( const char* s, int c );
```

The `strchr` function locates the first occurrence of `c` (converted to a `char`) in the string pointed to by `s`. The terminating null character is considered part of the string. The `strchr` function returns a pointer to the located character or a null pointer if the character is not found in the string.

Refer to the following sections for more information:

- Section A.4.67, “`memchr`—Find a character in a memory area,” on page A-70
- Section A.4.102, “`strcspn`—Compute length of prefix of one string consisting entirely of characters not from another,” on page A-104
- Section A.4.108, “`strpbrk`—Find first occurrence of a character from one string in another,” on page A-110
- Appendix A.4.110, “`strspn`—Compute length of prefix of one string consisting entirely of characters from another,” on page A-112.

Example A-94. `strchr`

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

void main()
{
char* string = "fred flintstone driving on bald feet";
char* found;

found = strchr( string, 'b' );
puts( found );
}
```

prints to standard output:
bald feet

A.4.99 strcmp—Compare two strings

```
#include <string.h>
int strcmp( const char* s1, const char* s2 );
```

The `strcmp` function compares the string pointed to by `s1` to the string pointed to by `s2`. If string `s1` is lexicographically greater than, equal to, or less than `s2`; an integer respectively greater than, equal to, or less than zero will be returned. The comparison of two strings of unequal length in which the longer string contains the smaller string yields the results that the longer string compares *greater than*.

i.e. `strcmp("xxx", "xxxxyz") < 0` or `strcmp("xxxxyz", "xxx") > 0`

When the header file `string.h` is included, the default case will be in-line [see section A.3, “Forcing Library Routines Out-of-line.”]

Refer to the following sections for more information:

- Section A.4.68, “`memcmp`—Compare portion of two memory areas,” on page A-71
- Section A.4.100, “`strcoll`—Compare two strings based on current locale,” on page A-102
- Section A.4.106, “`strncmp`—Compare a portions of two strings,” on page A-108

Example A-95. `strcmp`

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

void main()
{
if ( strcmp( "xxx", "xxxxyz" ) < 0 )
{
puts( "xxx is less than xxxxyz" );
}
else
{
puts( "xxx is greater than xxxxyz" );
}

if ( strcmp( "xxxxyz", "xxx" ) < 0 )
{
puts( "xxxxyz is less than xxx" );
}
else
{
puts( "xxxxyz is greater than xxx" );
}

if ( strcmp( "xxxxyz", "xxxxyz" ) == 0 )
{
puts( "xxxxyz is equal to xxxxyz" );
}
}
```

prints to standard output:
xxx is less than xxxxyz
xxxxyz is greater than xxx
xxxxyz is equal to xxxxyz

A.4.100 `strcoll`—Compare two strings based on current locale

```
#include <string.h>
int strcoll( const char* s1, const char* s2 );
```

The `strcoll` function compares the string pointed to by `s1` to the string pointed to by `s2`, both strings are interpreted using the `LC_COLLATE` category of the current locale. If string `s1` is lexicographically greater than, equal to, or less than `s2`, an integer greater than, equal to, or less than zero will be returned. The comparison of two strings of unequal length in which the longer string contains the smaller string yields the result that the longer string compares *greater than*.

For DSP56KCC, `strcoll` functions exactly like `strcmp`.

Refer to the following sections for more information:

- Section A.4.116, “`strxfrm`—Transform a string into locale-independent form,” on page A-120
- Section A.4.99, “`strcmp`—Compare two strings,” on page A-100

A.4.101 `strcpy`—Copy one string into another

```
#include <string.h>
int strcpy( char* s1, const char* s2 );
```

The `strcpy` function copies the characters of string `s2`, including the terminating character, into the string pointed to by `s1`. If the strings overlap, the behavior is undefined. The value of `s1` is returned.

When the header file `string.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”].

Refer to the following sections for more information:

- Section A.4.69, “`memcpy`—Copy from one area to another,” on page A-72
- Section A.4.71, “`memset`—Initialize memory area,” on page A-74
- Section A.4.107, “`strncpy`—Copy a portion of one string into another,” on page A-109

Example A-96. strcpy

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

void main()
{
char string[80];
char* sptr;

strcpy( string, "-- no bald feet for george jetson --" );
puts( sptr );
}
```

prints to standard output:
-- no bald feet for george jetson --

A.4.102 `strcspn`—Compute length of prefix of one string consisting entirely of characters not from another

```
#include <string.h>
int strcspn( const char* s1, const char* s2 );
```

The `strcspn` function computes and returns the length of the prefix of the string pointed to by `s1` that consists entirely of characters not found in the string pointed to by `s2`.

Refer to the following sections for more information:

- Section A.4.67, “`memchr`—Find a character in a memory area,” on page A-70
- Section A.4.98, “`strchr`—Find first occurrence of a character in a string,” on page A-100
- Section A.4.108, “`strpbrk`—Find first occurrence of a character from one string in another,” on page A-110
- Section A.4.109, “`strrchr`—Find last occurrence of a character from one string in another,” on page A-111

Example A-97. `strcspn`

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

void main()
{
  int i;

  i = strcspn( "azbyfghjki", "fkjeughtrg" );
  printf( "-- prefix length == %d --\n", i );
}
```

prints to standard output:
-- prefix length == 4 --

A.4.103 `strerror`—Map error code into an error message string

```
#include <string.h>
char* strerror( int errnum );
```

The `strerror` function maps `errnum` to an error message string. A pointer to the string is returned. The string returned *should not be modified* by the programmer.

Refer to the following sections for more information:

- Section A.4.73, “`perror`—Print error message,” on page A-76

Example A-98. `strerror`

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

void main()
{
  int i = 0;
  char* s;

  while ( ( s = strerror( i ++ ) ) != NULL )
  {
    printf( "\tmessage %d:\t%s\n", i, s );
  }
}
```

prints to standard output:
message 1:domain error
message 2:range error
message 3:out of heap memory
message 4:bad format for conversion string

A.4.104 `strlen`—Determine length of a string

```
#include <string.h>
size_t strlen( const char* s );
```

The `strlen` function computes and returns the number of characters *proceeding* the terminating character.

Example A-99. `strlen`

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

void main()
{
char* s = "is your name michael diamond?";

printf( "strlen( \"%s\" ) == %d\n", s, strlen( s ) );
}
```

prints to standard output:

```
strlen( "is your name michael diamond?" ) == 29
```

A.4.105 `strncat`—Concatenate a portion of one string to another

```
#include <string.h>
char* strncat( char* s1, const char* s2, size_t n );
```

The `strncat` function appends, at most, **n** characters from the string pointed by **s2** to the end of the string pointed to by **s1**. The first character of the second string is written over the first string's terminating character and a new terminating character is appended. The `strncat` function returns a pointer to **s1**. If **s1** does not have **n** words allocated past the terminating character, the behavior is undefined.

Refer to the following sections for more information:

- Section A.4.97, “`strcat`—Concatenate two strings,” on page A-99

Example A-100. strncat

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

void main()
{
char bstr[80] = "string 1";
char sstr[20] = " string 2";

printf("paste 5 chars of (%s) on to (%s)\n", sstr, bstr);
(void) strncat( bstr, sstr, 5 );
puts( bigstr );
}
```

prints to standard output:
paste 5 chars of (string 2) on to (string 1)
string 1 stri

A.4.106 strncmp—Compare a portions of two strings

```
#include <string.h>
int strncmp( const char* s1, const char* s2, size_t n );
```

The `strncmp` function compares `n` characters of the string pointed to by `s2` with the string pointed to by `s1`. If string `s1` is lexicographically greater than, equal to, or less than `s2`; an integer respectively greater than, equal to, or less than zero will be returned. This is similar to `strcmp`.

Refer to the following sections for more information:

- Section A.4.99, “`strcmp`—Compare two strings,” on page A-100

Example A-101. strncmp

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

void main()
{
char bigstr[80] = "string 1";
char smallstr[20] = "string 2";

if ( strncmp( bigstr, smallstr, 5 ) == 0 )
{
printf( "-- strncmp ok --\n" );
}
else
{
printf( "?? strncmp error ??\n" );
}
}
```

prints to standard output:

```
-- strncmp ok --
```

A.4.107 strncpy—Copy a portion of one string into another

```
#include <string.h>
char* strncpy( char* s1, const char* s2, size_t n );
```

The `strncpy` function copies exactly **n** characters from a string pointed to by **s2** into a string pointed to by **s1**. If `strlen(s2)` is less than **n**, the string **s1** is null padded. If `strlen(s2)` is greater than or equal to **n**, no null termination character is copied to **s1**. The **s1** pointer is returned.

Note: The behavior of non null terminated strings is undefined.

Refer to the following sections for more information:

- Section A.4.69, “memcpy—Copy from one area to another,” on page A-72
- Section A.4.101, “strcpy—Copy one string into another,” on page A-102

Example A-102. strncpy

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

void main()
{
char bigstr[80] = "string 1";
char smallstr[20] = "spanky 2";

( void ) strncpy( bigstr, smallstr, 6 );
puts( bigstr );
}
```

prints to standard output:
spanky 1

A.4.108 `strpbrk`—Find first occurrence of a character from one string in another

```
#include <string.h>
char* strpbrk( char* s1, const char* s2 );
```

The `strpbrk` function finds the first occurrence of *any* character in the string pointed to by `s2` in the string pointed to by `s1`. If a character is found, a pointer to the character is returned. If a character is not found, a null pointer is returned.

Refer to the following sections for more information:

- Section A.4.67, “`memchr`—Find a character in a memory area,” on page A-70
- Section A.4.98, “`strchr`—Find first occurrence of a character in a string,” on page A-100
- Section A.4.102, “`strcspn`—Compute length of prefix of one string consisting entirely of characters not from another,” on page A-104
- Section A.4.109, “`strrchr`—Find last occurrence of a character from one string in another,” on page A-111
- **`strrchr`**—Find the last occurrence of a character in a string.
- Appendix A.4.110, “`strspn`—Compute length of prefix of one string consisting entirely of characters from another,” on page A-112

Example A-103. `strbrk`

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

void main()
{
char* string = "abcde random characters fghijkl";
char* fndstr = "klmnopqr";
char* found;

if ( ( found = strpbrk( string, fndstr ) ) != NULL )
{
puts( found );
}

else
{
puts( "can't find a character" );
}

prints to standard output:
random characters fghijkl
```

A.4.109 `strrchr`—Find last occurrence of a character from one string in another

```
#include <string.h>
char* strrchr( char* s1, const char* s2 );
```

The `strrchr` function locates the last occurrence of `c` (converted to a `char`) in the string pointed to by `s`. The terminating null character is considered part of the string. `strrchr` returns a pointer to the located character, or a null pointer if the character is not found in the string.

Refer to the following sections for more information:

- Section A.4.67, “`memchr`—Find a character in a memory area,” on page A-70
- Section A.4.98, “`strchr`—Find first occurrence of a character in a string,” on page A-100
- Section A.4.102, “`strcspn`—Compute length of prefix of one string consisting entirely of characters not from another,” on page A-104
- Section A.4.108, “`strpbrk`—Find first occurrence of a character from one string in another,” on page A-110
- Appendix A.4.110, “`strspn`—Compute length of prefix of one string consisting entirely of characters from another,” on page A-112

Example A-104. `strrchr`

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

void main()
{
char* string = "fred flintstone driving on bald feet";
char* found;

found = strrchr( string, 'f' );
puts( found );
}
```

prints to standard output:
feet

A.4.110 `strspn`—Compute length of prefix of one string consisting entirely of characters from another

```
#include <string.h>
int strspn( const char* s1, const char* s2 );
```

The `strspn` function computes and returns the length of the prefix of the string pointed to by **s1** that consists entirely of characters found in the string pointed to by **s2**.

Refer to the following sections for more information:

- Section A.4.67, “`memchr`—Find a character in a memory area,” on page A-70
- Section A.4.98, “`strchr`—Find first occurrence of a character in a string,” on page A-100
- Section A.4.108, “`strpbrk`—Find first occurrence of a character from one string in another,” on page A-110
- Section A.4.109, “`strrchr`—Find last occurrence of a character from one string in another,” on page A-111

Example A-105. `strspn`

```
#include "stdio.h"
#include "string.h"
void main()
{
    int i;
    i= strspn("fghjkiazby", "fkjeughtrg");
    printf("%d\n", i);
}
```

```
/*result will be 5 */
```

A.4.111 `strstr`—Find the first occurrence of one string in another

```
#include <string.h>
char* strstr( const char* s1, const char* s2 );
```

The `strstr` function locates the first occurrence of the string pointed to by **s2** (excluding the termination character) in the string pointed to by **s1**. If the string **s2** is found, a pointer to it is returned. If the string **s2** is not found, a null pointer is returned. If **s2** has a length of zero, a pointer to **s1** is returned.

Example A-106. strstr

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

void main()
{
char* string = "abcdef random characters ghijkl";
char* fndstr = "random";
char* found;

if ( ( found = strstr( string, fndstr ) ) != NULL )
{
puts( found );
}
else
{
puts( "can't find the string" );
}
}
```

prints to standard output:
random characters ghijkl

A.4.112 strtod—String to double

```
#include <stdlib.h>
double strtod( const char* nptr, char** endptr );
```

The `strtod` function converts and returns the string pointed to by `nptr` to floating point number. First `strtod` decomposes `nptr` into three sections;

1. an initial, possibly empty, sequence of white space characters,
2. a subject in the form of a floating point constant; and
3. a final string of one or more unrecognized characters, including the terminating null character of the input string.

If the first unrecognized character is *not* `NULL`, a pointer to that character is stored into the object that `endptr` points to. If the string is empty or the subject contains no floating point constant description, zero is returned.

Refer to the following sections for more information:

- Section A.4.8, “`atof`—String to floating point,” on page A-16
- Section A.4.9, “`atoi`—String to integer,” on page A-17
- Section A.4.10, “`atol`—String to long integer,” on page A-18
- Section A.4.114, “`strtol`—String to long integer,” on page A-116
- Section A.4.115, “`strtoul`—String to unsigned long integer,” on page A-118

Example A-107. `strtod`

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

void main()
{
char* string = "7.09strtod stopped";
char* stopped;
double result;

result = strtod( string, &stopped );
printf( "string == (%s)\n", string );
printf( "result == %f\n", result );
printf( "stop string == (%s)\n", stopped );
}
```

prints to standard output:
string == (7.09strtod stopped)
result == 7.089998
stop string == (strtod stopped)

A.4.113 strtok—Break string into tokens

```
#include <stdlib.h>
char* strtok( char* s1, const char* s2 );
```

The `strtok` function breaks the string pointed to by `s1` into tokens and each token is delimited by characters from the string pointed to by `s2`. The first call in the sequence has `s1` as its first argument and is followed by calls with a null pointer as the first argument. The separator string, `s2`, may be different from call to call. If a token is not found, a null pointer is returned. If a token is found, a null terminated token is returned.

Example A-108. strtok

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

void main()
{
char* str1 = "$%^this#is string\tnumber!one.";
char str2[] = "?a???b,,,#c";
char* token;

while ( ( token = strtok( str1, "$%^#\t! " ) ) != NULL )
{
printf( "%s ", token );
str1 = NULL;
}

printf( "\n" );

token = strtok( str2, "?" ); printf( "%s ", token );
token = strtok( NULL, "," ); printf( "%s ", token );
token = strtok( NULL, "#," ); printf( "%s\n", token );

if ( ( token = strtok( NULL, "?" ) ) != NULL )
{
printf( "error: strtok busted\n" );
}
}
```

prints to standard output:
this is string number one.
a ??b c

A.4.114 `strtol`—String to long integer

```
#include <stdlib.h>
long int strtol( const char* nptr, char** endptr, int base );
```

The `strtol` function converts and returns the string pointed to by `nptr` to a long integer.

First `strtol` decomposes `nptr` into three sections;

1. an initial, possibly empty, sequence of white space characters,
2. a subject in the form of an integer constant; and
3. a final string of one or more unrecognized characters, including the terminating null character of the input string.

If the first unrecognized character is *not* `NULL`, a pointer to that character is stored in to the object that `endptr` points to. If the string is empty or the subject contains no floating-point constant description, zero is returned.

If **base** is between 2 and 36, the expected form of the long integer subject is a sequence of letters and digits with the radix specified by **base**. The letters **a** (or **A**) through **z** (or **Z**) are ascribed values 10 to 35; only letters whose value is less than **base** are valid. If **base** is 16, **0x** or **0X** may optionally proceed the long integer subject. If **base** is zero, the long integer subject determines its own base.

Leading `0x`, or `0Xbase == 16`

Leading `0base == 8`

otherwise `base == 10`

If the value of the return value is too large to be expressed by a long int, `errno` is set to `ERANGE` and `LONG_MAX` is returned.

Refer to the following sections for more information:

- Section A.4.8, “`atof`—String to floating point,” on page A-16
- Section A.4.9, “`atoi`—String to integer,” on page A-17
- Section A.4.10, “`atol`—String to long integer,” on page A-18
- Section A.4.112, “`strtod`—String to double,” on page A-114
- Section A.4.115, “`strtoul`—String to unsigned long integer,” on page A-118

Example A-109. strtol

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

void main()
{
char* hexstr = "0xabcdef709hexstr stopped";
char* decstr = "709709709decstr stopped";
char* octstr = "012341234octstr stopped";
char* stopped;
long result;

printf( "result\t\tstop string\n" );
result = strtol( hexstr, &stopped, 16 );
printf( "%lx\t\t%s", result, stopped );

result = strtol( decstr, &stopped, 10 );
printf( "%ld\t\t%s", result, stopped );

result = strtol( octstr, &stopped, 8 );
printf( "%lo\t\t%s", result, stopped );
}
```

prints to standard output:

```
resultstop string
abcdef709hexstr stopped
709709709decstr stopped
12341234octstr stopped
```

A.4.115 strtoul—String to unsigned long integer

```
#include <stdlib.h>
unsigned long int strtoul( const char* nptr, char** endptr, int base );
```

The `strtoul` function converts and returns the string pointed to by `nptr` to a long integer. First `strtoul` decomposes `nptr` into three sections; an initial, possibly empty, sequence of white space characters, a subject in the form of an integer constant; and a final string of one or more unrecognized characters, including the terminating null character of the input string.

If the first unrecognized character is *not* `NULL`, a pointer to that character is stored in to the object that `endptr` points to. If the string is empty or the subject contains no floating point constant description, zero is returned.

If **base** is between 2 and 36, the expected form of the long integer subject is a sequence of letters and digits with the radix specified by `base`. The letters **a** (or **A**) through **z** (or **Z**) are ascribed values 10 to 35; only letters whose value is less than **base** are valid. If **base** is 16, **0x** or **0X** may optionally proceed the long integer subject. If **base** is zero, the long integer subject determines its own base.

Leading `0x`, or `0Xbase == 16`

Leading `0base == 8`

otherwise `base == 10`

If the value of the return value is too large to be expressed by a long int, `errno` is set to `ERANGE`, and `ULONG_MAX` is returned.

Refer to the following sections for more information:

- Section A.4.8, “`atof`—String to floating point,” on page A-16
- Section A.4.9, “`atoi`—String to integer,” on page A-17
- Section A.4.10, “`atol`—String to long integer,” on page A-18
- Section A.4.114, “`strtol`—String to long integer,” on page A-116
- Section A.4.112, “`strtod`—String to double,” on page A-114

Example A-110. strtoul

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

void main()
{
char* hexstr = "0xabcdef709hexstr stopped";
char* decstr = "709709709decstr stopped";
char* octstr = "012341234octstr stopped";
char* stopped;
unsigned long result;

printf( "result\t\tstop string\n" );
result = strtoul( hexstr, &stopped, 16 );
printf( "%lu\t\t%s", result, stopped );

result = strtoul( decstr, &stopped, 10 );
printf( "%lu\t\t%s", result, stopped );

result = strtoul( octstr, &stopped, 8 );
printf( "%lu\t\t%s", result, stopped );
}
```

prints to standard output:

```
resultstop string
2883516169hexstr stopped
709709709decstr stopped
12341234octstr stopped
```

A.4.116 `strxfrm`—Transform a string into locale-independent form

```
#include <string.h>
size_t strxfrm( char* s1, const char* s2, size_t n );
```

The `strxfrm` function transforms the string pointed to by `s2` and places the resulting string in the array pointed to by `s1`. The transformation is such that if the `strcmp` function is applied to the two transformed strings, it returns a value greater than, equal to, or less than zero, corresponding to the result of the `strcoll` function applied to the same two original strings. No more than `n` characters are placed into `s1`, including the terminating null character. If `s1` and `s2` overlap, the behavior is undefined.

The `strxfrm` function returns the length of the transformed string excluding the terminating null character. If the value returned is `n` or more, the contents of `s1` are indeterminate.

Refer to the following sections for more information:

- Section A.4.100, “`strcoll`—Compare two strings based on current locale,” on page A-102
- Section A.4.106, “`strncmp`—Compare a portions of two strings,” on page A-108

Note: DSP56KCC only supports the standard locale, so no transformation is done.

A.4.117 `tan`—Tangent

```
#include <math.h>
double tan( double x );
```

The `tan` function computes and returns the tangent of `x`, where `x` is in radians.

Refer to the following sections for more information:

- Section A.4.5, “`atan`—Arc tangent,” on page A-13
- Section A.4.6, “`atan2`—Arc tangent of angle defined by point `y/x`,” on page A-14

Example A-111. `tan`

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

void main()
{
printf( "tan( 45 ) == %f\n", tan( 45 ) );
}
```

prints to standard output:
`tan(45) == 1.619775`

A.4.118 tanh—Hyperbolic tangent

```
#include <math.h>
double tanh( double x );
```

The tanh function computes and returns the hyperbolic tangent of **x**,

$$\tanh(x) == \sinh(x) / \cosh(x)$$

If the value of **x** is too large, `errno` is set to `ERANGE` and the value `HUGE_VAL` is returned with the sign of **x**.

Refer to the following sections for more information:

- Section A.4.16, “cosh—Hyperbolic cosine,” on page A-23
- Section A.4.92, “sinh—Hyperbolic Sine,” on page A-96

Example A-112. tanh

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

void main()
{
printf( "tanh( 45 ) == %f\n", tanh( 45 ) );
}
```

prints to standard output:
tanh(45) == 1.000000

A.4.119 tmpfile—Create a temporary binary file

```
#include <stdio.h>
FILE *tmpfile ( void );
```

The function `tmpfile` will create a temporary file on the disk. The file will be automatically removed when the program terminates. The file will be opened with the mode “**wb+**”. If `tmpfile` fails, it returns a `NULL` pointer.

Refer to the following sections for more information:

- Section A.4.120, “tmpnam—Create a temporary file name,” on page A-122

A.4.120 tmpnam—Create a temporary file name

```
#include <stdio.h>
char *tmpnam ( char *s );
```

The function `tmpnam` will create a string that could be used as a unique temporary file name. This function may be called as many as `TMP_MAX` times. Each time it will return a different string. If the argument `s` is `NULL`, then `tmpnam` will return an internal static buffer that may be clobbered by subsequent calls. If `s` is non-`NULL`, then it must point to a writable buffer of at least `L_tmpnam` characters.

Refer to the following sections for more information:

- Section A.4.119, “`tmpfile`—Create a temporary binary file,” on page A-121

Example A-113. tmpnam

```
#include <stdio.h>

void main ()
{
char buffer[L_tmpnam];

(void) fopen ( tmpnam ( buffer ), "w+" );
}
```

will create a temporary *text* file on the disk. Note that unlike when `tmpfile` is called, one must remove any files created using `fopen` and `tmpnam`.

A.4.121 tolower—Convert uppercase character to lowercase

```
#include <ctype.h>
int tolower( int c );
```

The `tolower` function converts uppercase to lowercase. If `c` is an uppercase letter, return the corresponding lowercase letter; otherwise return `c`.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, “Forcing Library Routines Out-of-line.”].

Example A-114. tolower

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

void main()
{
printf( "tolower( 'A' ) == %c\n", tolower( 'A' ) );
printf( "tolower( 'z' ) == %c\n", tolower( 'z' ) );
printf( "tolower( '#' ) == %c\n", tolower( '#' ) );
}
```

prints to standard output:

```
tolower( 'A' ) == a
tolower( 'z' ) == z
tolower( '#' ) == #
```

A.4.122 toupper—Convert lowercase character to uppercase

```
#include <ctype.h>
int toupper( int c );
```

The `toupper` function converts lowercase to uppercase. If `c` is a lowercase letter, return the corresponding uppercase letter; otherwise return `c`.

When the header file `ctype.h` is included, the default case will be **in-line** [see section A.3, *Forcing Library Routines Out-of-line*].

Example A-115. toupper

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

void main()
{
printf( "toupper( 'A' ) == %c\n", toupper( 'A' ) );
printf( "toupper( 'z' ) == %c\n", toupper( 'z' ) );
printf( "toupper( '#' ) == %c\n", toupper( '#' ) );
}
```

prints to standard output:

```
toupper( 'A' ) == A
toupper( 'z' ) == Z
toupper( '#' ) == #
```

A.4.123 ungetc—Push a character back onto an input stream

```
#include <stdio.h>
int ungetc ( int c, FILE *stream );
```

The function `ungetc` converts the argument `c` to an unsigned `char`, and pushes it back onto the specified input stream. Pushed characters will be read back in reverse order by any functions reading from said stream. If a call is made to a file positioning function, such as `fseek`, all pushed characters will be lost. Only one call to `ungetc` before a read from the stream is allowed. EOF cannot be pushed. `ungetc` returns EOF upon failure, while the converted value is returned upon success.

Refer to the following sections for more information:

- Section A.4.119, “`tmpfile`—Create a temporary binary file,” on page A-121

Example A-116. ungetc

```
#include <stdio.h>

void main ()
{
char peek = getchar ();
putchar ( peek );
ungetc ( peek, stdin );
putchar ( getchar ());
}
```

will print the first character from standard input twice on standard output.

A.4.124 vfprintf—Write formatted output to a stream using a va_list

```
#include <stdio.h>
int vfprintf ( FILE *stream, const char *format, va_list arg );
```

The function `vfprintf` is exactly the same as the function `fprintf` except that an existing `va_list` is used in place of a series of arguments. The macro `va_start` must have been invoked on the argument `arg` before the call to `vfprintf` is made. `vfprintf` returns the number of characters printed. On error, `vfprintf` returns a negative value.

Refer to the following sections for more information:

- Section A.4.31, “`fprintf`—Write formatted output to a stream,” on page A-36

Example A-117. `vfprintf`

```
#include <stdio.h>

int printf ( const char *format, ... )
{
va_list ap;
int result;

va_start ( ap, format );
result = vfprintf ( stdout, format, ap );
va_end ( ap );
return result;
}
```

is essentially the library function `printf`.

A.4.125 `vprintf`—Write formatted output to standard output using a `va_list`

```
#include <stdio.h>
int vprintf ( const char *format, va_list arg );
```

The function `vprintf` is exactly the same as the function `printf` except that an existing `va_list` is used in place of a series of arguments. The macro `va_start` must have been invoked on the argument `arg` before the call to `vprintf` is made. `vprintf` returns the number of characters printed. On error, `vprintf` returns a negative value.

Refer to the following sections for more information:

- Section A.4.75, “`printf`—Print to standard output,” on page A-78

Example A-118. `vprintf`

```
#include <stdio.h>

int printf ( const char *format, ... )
{
va_list ap;
int result;

va_start ( ap, format );
result = vprintf ( format, ap );
va_end ( ap );
return result;
}
```

is essentially the library function `printf`.

A.4.126 vsprintf—Write formatted output to a string using a va_list

```
#include <stdio.h>
int vsprintf ( char *s, const char *format, va_list arg );
```

The function `vsprintf` is exactly the same as the function `printf` except that an existing `va_list` is used in place of a series of arguments. The macro `va_start` must have been invoked on the argument `arg` before the call to `vsprintf` is made. `vsprintf` returns the number of characters printed. On error, `vsprintf` returns a negative value.

Refer to the following sections for more information:

- Section A.4.93, “`sprintf`—Print to a string,” on page A-96

Example A-119. vsprintf

```
#include <stdio.h>

int sprintf ( char *s, const char *format, ... )
{
    va_list ap;
    int result;

    va_start ( ap, format );
    result = vsprintf ( s, format, ap );
    va_end ( ap );
    return result;
}
```

is essentially the library function `sprintf`.

A.4.127 wcstombs—Convert wchar_t array to multibyte string

```
#include <stdlib.h>
size_t wcstombs( char* s, const wchar_t* pwcs, size_t n );
```

The `wcstombs` function converts a wide character string pointed to by `pwcs` into the character string pointed to by `s`. Each character of the wide character string is converted into the corresponding *multibyte* character as if by the `wctomb` function. Conversion will stop when `n` total characters have been converted or a null character is encountered. If `s` and `pwcs` overlap, the behavior is undefined.

If an invalid character is encountered, `wcstombs` returns `(size_t) -1`. Otherwise, `wcstombs` returns the number of characters converted not including the terminating `NULL` character, if any.

Note: The DSP56000/DSP56001 does not provide byte addressing, thus characters always require an entire word of memory each. One way to better utilize data memory (with a run-time cost) is to combine the ANSI data type `wchar_t` and the special ANSI multibyte and wide character library routines.

Example A-120. wcstombs

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

char array[16];
wchar_t wstr[10] = L"abcdefgh";

void main()
{
char* ptr = (char*) wstr;
int convert;

convert = wcstombs( array, wstr, 10 );

printf( "packed array:\n" );
while ( *ptr != 0 )
{
printf( "%0.6x ", *ptr++ );
}
printf( "\n\n" );

printf( "%d chars extracted, unpacked array:\n" );
ptr = array;
while ( *ptr != 0 )
{
printf( "%0.6x ", *ptr++ );
}
printf( "\n" );
}

prints to standard output:
packed array:
616263 646566 676800

8 chars extracted, unpacked array:
000061 000062 000063 000064 000065 000066 000067 000068
```

A.4.128 wctomb—Convert wchar_t character to multibyte character

```
#include <stdlib.h>
int wctomb( char* s, wchar_t wchar );
```

The `wctomb` function examines and converts the wide character `wchar` into a string of characters pointed to by `s`. At most, `MB_CUR_MAX` characters will be stored in `s`.

If `s` is a `NULL` pointer, `wctomb` returns zero. If `s` is not `NULL`, `wctomb` returns the number of characters that comprise the converted multibyte character unless an invalid multibyte character is detected in which case a -1 will be returned.

Refer to the following sections for more information:

- Section A.4.64, “`mblen`—Length of a multibyte character,” on page A-64
- Section A.4.65, “`mbstowcs`—Convert multibyte string to wide character string,” on page A-66
- Section A.4.66, “`mbtowlc`—Convert a multibyte character to a wide character,” on page A-68

Note: The DSP56000/DSP56001 does not provide byte addressing, thus characters always require an entire word of memory each. One way to better utilize data memory (with a run-time cost) is to combine the ANSI data type `wchar_t` and the special ANSI multibyte and wide character library routines.

Example A-121. wctomb

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

char mbararray[8];

void main()
{
    wchar_t wide = L'abc';
    char* ptr = array;
    int convert;

    convert = wctomb( array, wide );

    printf( "packed char looks like:\n" );
    printf( "%0.6x\n\n", wide );

    printf( "%d extracted chars looks like:\n", convert );
    while ( *ptr != 0 )
    {
        printf( "%0.6x ", *ptr++ );
    }
    printf( "\n" );
}
```

```
prints to standard output:
packed char looks like:
616263
```

```
3 extracted chars looks like:
000061 000062 000063
```



Appendix B

DSP56000/DSP56001 Instruction Set and Assembler Directive Summary

B.1 Overview

This appendix is intended to assist the C programmer in understanding the C compiler's output. For a more in-depth discussion of assembler and assembly language issues, see the *DSP56000/DSP56001 Digital Signal Processor User's Manual* or the *DSP56000/DSP56001 Assembler User's Manual*.

B.2 Instruction Set

Instructions can be grouped by function into six types:

1. Arithmetic instructions
2. Logical instructions
3. Bit manipulation instructions
4. Loop instructions
5. Move instructions
6. Program control instructions

B.2.1 Arithmetic Instructions

The instructions used for arithmetic operations are as shown in Table B-1:

Table B-1. Arithmetic Operations

Arithmetic Operations	Description
ABS	Absolute value
ADC	Add long with carry
ADD	Add

Table B-1. Arithmetic Operations

Arithmetic Operations	Description
ADDL	Shift left then add
ADDR	Shift right then add
ASL	Arithmetic shift accumulator left
ASR	Arithmetic shift accumulator right
CLR	Clear accumulator
CMP	Compare
CMPM	Compare magnitude
DIV	Divide iteration
MAC	Signed multiply—accumulate
MACR	Signed multiply—accumulate and round
MPY	Signed multiply
MPYR	Signed multiply and round
NEG	Negate accumulator
NORM	Normalize accumulator iteration
RND	Round accumulator
SUB	Subtract
SUBL	Shift left then subtract
SUBR	Shift right then subtract
Tcc	Transfer conditionally
TFR	Transfer data ALU register
TST	Test

B.2.2 Logical Instructions

The instructions used for logical operations are as shown in Table B-2:

Table B-2. Logical Operations

Logical Operations	Description
AND	Logical AND
ANDI	AND Immediate with control register

Table B-2. Logical Operations

Logical Operations	Description
EOR	Logical exclusive OR
LSL	Logical shift accumulator left
LSR	Logical shift accumulator right
NOT	Logical complement on accumulator
OR	Logical inclusive OR
ORI	OR immediate with control register
ROL	Rotate accumulator left
ROR	Rotate accumulator right

B.2.3 Bit Manipulation Instructions

The instructions used for bit manipulation are as shown in Table B-3:

Table B-3. Bit Manipulation Instructions

Instruction	Description
BCLR	Bit test and clear
BSET	Bit test and set
BCHG	Bit test and change
BTST	Bit test on memory
JCLR	Jump if bit clear
JSET	Jump if bit set
JSCLR	Jump to subroutine if bit clear
JSSET	Jump to subroutine if bit set

B.2.4 Loop Instructions

The instructions used for loop operations are as shown in Table B-4:

Table B-4. Loop Operation Instructions

Instruction	Description
DO	Start hardware loop
ENDDO	Exit from hardware loop

B.2.5 Move Instructions

The instructions used for move operations are as shown in Table B-5:

Table B-5. Move Operation Instructions

Instruction	Description
LUA	Load updated address
MOVE	Move data
MOVEC	Move control register
MOVEM	Move program memory
MOVEP	Move peripheral data

B.2.6 Program Control Instructions

The instructions used for program control are as shown in Table B-6:

Table B-6. Program Control Instructions

Instruction	Description
Jcc	Jump conditionally
JMP	Jump
JScC	Jump to subroutine conditionally
JSR	Jump to subroutine
NOP	No operation
REP	Repeat next instruction
RESET	Reset on-chip peripheral devices
RTI	Return from interrupt
RTS	Return from subroutine
STOP	Stop processing (low power stand-by)
SWI	Software interrupt
WAIT	Wait for interrupt (low power stand-by)

B.3 Directive Summary

Directives can be grouped by function into seven types:

1. Assembly control
2. Symbol definition
3. Data definition / storage allocation
4. Listing control and options
5. Object file control
6. Macros and conditional assembly
7. Structured programming

B.3.1 Assembly Control

The directives used for assembly control are as shown in Table B-7:

Table B-7. Assembly Control Directives

Directive	Description
COMMENT	Comment line(s)
DEFINE	Define substitution strings
END	End of source program
FAIL	Programmer generated error message
INCLUDE	Include secondary file
MSG	Program generated message
ORG	Initialize memory space and location counters
RADIX	Change input radix for constants
RDIRECT	Remove directive / mnemonic from assembler table
UNDEF	Undefine DEFINE symbol
WARN	Programmer generated warning message

B.3.2 Symbol Definition

The directives used to control symbol definition are as shown in Table B-8:

Table B-8. Control Symbol Directives

Directive	Description
ENDSEC	End section
EQU	Equate symbol to a value
SET	Set symbol to a value
SECTION	Start section
XDEF	External section symbol definition
XREF	External section symbol reference

B.3.3 Data Definition/Storage Allocation

The directives used to control constant data definition and storage allocation are as shown in Table B-9:

Table B-9. Storage Allocation Directives

Directive	Description
BSC	Block storage of a constant
DC	Define constant
DS	Define storage
DSM	Define modulo storage
DSR	Define reverse carry storage

B.3.4 Listing Control and Options

The directives used to control the output listing are as shown in Table B-10:

Table B-10. Output Listing Directives

Directive	Description
LIST	List the assembly
LSTCOL	Set listing field widths
NOLIST	Stop assembly listing
OPT	Assembler options

Table B-10. Output Listing Directives

Directive	Description
PAGE	Top of page / define page size
PRCTL	Send control string to printer
STITLE	Initialize program sub—title
TITLE	Initialize program title

B.3.5 Object File Control

The directives used to control the object file are as shown in Table B-11:

Table B-11. Object file Directives

Directive	Description
COBJ	Comment object code
IDENT	Object code identification record
SYMOBJ	Write symbol information to object file

B.3.6 Macros and Conditional Assembly

The directives used for macros and conditional assembly are:

Table B-12. Conditional Assembly Directives

DUP	Duplicate a sequence of source lines
DUPA	Duplicate sequence with arguments
DUPC	Duplicate sequence with characters
ENDIF	End of conditional assembly
ENDM	End of macro definition
EXITM	Exit macro
IF	Conditional assembly directive
MACLIB	Macro library
MACRO	Macro definition
PMACRO	Purge a macro definition

B.3.7 Structured Programming

The directives used for structured programming are as shown in Table B-13:

Table B-13. Programming Directives

Directive	Description
.BREAK	Exit from structured loop construct
.CONTINUE	Continue next iteration of structured loop
.ELSE	Perform following statements when .IF false
.ENDF	End of .FOR loop
.ENDI	End of .IF condition
.ENDL	End of hardware loop
.ENDW	End of .WHILE loop
.FOR	Begin .FOR loop
.IF	Begin .IF condition
.LOOP	Begin hardware loop
.REPEAT	Begin .REPEAT loop
.UNTIL	End of .REPEAT loop
.WHILE	Begin .WHILE loop

Appendix C

Utilities

C.1 Utility Programs

There are nine utility programs available with the DSP56KCC Rev. g1.11 compiler. They are:

1. asm56000
2. cldinfo
3. cldlod
4. cofdmp
5. dsplib
6. dsplnk
7. gdb56
8. run56
9. srec

These programs are described in detail in the following pages.

C.1.1 asm56000—Motorola DSP56000/DSP56001 COFF Assembler

```
asm56000 [-A] [-B[<objfil>]] [-D <symbol> <string>] [-F <argfil>]
[-G][ -I <ipath>] [-L[<lstfil>]] [-M <mpath>] [-O <opt>[, <opt> ... ]]
[-R <rev> [, <rev>...]] [-V] [-Z] <files...>
```

asm56000 is a program that processes source program statements written in DSP56000 assembly language, translating these source statements into object programs compatible with other DSP56000 software and hardware products.

files is a list of operating system compatible file names including optional pathnames. If no extension is supplied for a given file, the assembler will first attempt to open the file using the file name as supplied. If that is not successful, the assembler appends.asm to the file name and tries to open the file again. If no path is given for a particular file, the assembler will look for that file in the current directory. The list of files will be processed

sequentially in the order given and all files will be used to generate the output listing and object file.

The assembler will redirect the output listing to the standard output if it is not redirected via the `-L` command line option described below. Error messages will always appear on the standard output regardless of any option settings. Note that some options (`-B` and `-L`) allow a hyphen as an optional argument which indicates that the corresponding output should be sent to the standard output stream. Unpredictable results may occur if, for example, the object file is explicitly routed to standard output while the listing file is allowed to default to the same output stream.

C.1.1.1 `asm5600` Options

Any of the following command line options may be specified. These can be in any order but must precede the list of source file names. Option letters may be entered in either upper or lower case.

Option arguments may immediately follow the option letter or may be separated from the option letter by blanks or tabs. However, an ambiguity arises if an option takes an optional argument. Consider the following command line:

```
asm56000 -b main io
```

In this example it is not clear whether the file `main` is a source file or is meant to be an argument to the `-B` option. If the ambiguity is not resolved, the assembler will assume that `main` is a source file and attempt to open it for reading. This may not be what the programmer intended.

There are several ways to avoid this ambiguity. If `main` is supposed to be an argument to the `-B` option, it can be placed immediately after the option letter, without intervening white space:

```
asm56000 -bmain io
```

If there are other options on the command line besides those that take optional arguments, the other options can be placed between the ambiguous option and the list of source file names:

```
asm56000 -b main -v io
```

Alternatively, two successive hyphens may be used to indicate the end of the option list:

asm56000 -b -- main io

In this case the assembler interprets *main* as a source file name and uses the default naming conventions for the -B option. Table C-1 lists the different Command Line Options that can be used with asm56000:

Table C-1. asm56000 Command Line Options

Option	Description
-A	indicates that the assembler should operate in absolute mode, creating a load file (. <i>cl</i> d) if the -B option is given. By default, the assembler produces a link file (. <i>cl</i> n) which is subsequently processed by the Motorola DSP linker.
-B[< <i>objfil</i> >]	<p>This option specifies that an object file is to be created for assembler output. objfil can be any legal operating system file name, including an optional pathname. A hyphen may also be used as an argument to indicate that the object file should be sent to the standard output.</p> <p>If a path is not specified, the file will be created in the current directory. If no file name is supplied, the assembler will use the basename (file name without extension) of the first file name encountered in the source input file list. The resulting output file will have an extension of .<i>cl</i>n unless the -A option is given in which case the file will have a .<i>cl</i>d extension. If the -B option is not given, then the assembler will not generate an object file. The -B option should be specified only once.</p>
-D< <i>symbol</i> > < <i>string</i> >	<p>This is equivalent to a source statement of the form:</p> <p>DEFINE <<i>symbol</i>> <<i>string</i>></p> <p>string must be enclosed in quotes if it contains any embedded blanks. Note that if quotes are used, they must be passed to the assembler intact, e.g. some host command interpreters will strip quotes from around arguments. The -D<<i>symbol</i>> <<i>string</i>> sequence may be repeated as often as desired.</p>
-F< <i>argfil</i> >	<p>This option indicates that an external file should be read for further command arguments. It is useful in host environments where the command line length is restricted. argfil must be present on the command line, but can be any legal operating system file name including an optional pathname.</p> <p>The file may contain any legal command line options, including the -F option itself. The arguments need be separated only by white space (spaces, tabs, or newlines). A semicolon (;) on a line following white space causes the rest of the line in the file to be treated as a comment.</p>
-G	Send source file line number information to the object file. This option is valid only in conjunction with -B command line option. The generated line number information can be used by debuggers to provide source-level debugging.
-I< <i>ipath</i> >	When the assembler encounters <i>include</i> files, the current directory (or the directory specified in the INCLUDE directive) is first searched for the file. If it is not found and the -I option is supplied, the assembler prefixes the file name (and optional pathname) specified in the INCLUDE directive with ipath and searches the newly formed directory pathname for the file. The -I< <i>ipath</i> > sequence may be repeated as many times as desired. The directories will be searched in the order given on the command line.

Table C-1. asm56000 Command Line Options

Option	Description
-L[<lstfil>]	<p>This option specifies that a listing file is to be created for the assembler output. <i>lstfil</i> can be any legal operating system file name including an optional pathname. A hyphen also may be used as an argument to indicate that the listing file should be sent to the standard output.</p> <p>If a path is not specified, the file will be created in the current directory. If no file name is supplied, the assembler will use the basename (file name without extension) of the first file name encountered in the source input file list. The resulting output file will have an extension of <i>.lst</i>. If the -L option is not given, then the assembler will route the listing output to the standard output. The -L option should be specified only once.</p>
-M<mpath>	<p>This is equivalent to a source statement of the form:</p> <p>MACLIB <mpath></p> <p>The -M<mpath> sequence may be repeated as many times as desired. The directories will be searched in the order specified on the command line.</p>
-O<opt>[,<opt>...]	<p>opt can be any of the options that are available with the assembler OPT directive. If multiple options are supplied, they must be separated by commas. The -O<opt> sequence may be repeated for as many options as desired.</p>
-R<rev>[,<rev>...]	<p>Run the assembler without the specified processor revision level enhancements. This is for backward compatibility so that the assembler will flag new constructions as illegal. <i>rev</i> can be any of the revision specifiers given below, but must be appropriate for the target processor.</p> <p>ProcessorRevision DSP56000C</p> <p>If multiple revision levels are supplied, they must be separated by commas. The -R<rev> sequence may be repeated as many times as desired.</p>
-Z	<p>This option causes the assembler to strip symbol information from the absolute load file. Normally symbol information is retained in the object file for symbolic reference purposes. Note that this option is valid only when the assembler is in absolute mode via the -A command line option and when an object file is created (-B option).</p>
-V	<p>Indicates that the assembler should be verbose during processing, displaying a progress report as it assembles the input files. The assembler will show the beginning of each pass and when files are opened and closed. The information is sent to the standard error output stream.</p>

C.1.2 cldinfo—Memory size information from Motorola DSP COFF object file.

cldinfo file

cldinfo is a utility that reads an absolute or relocatable Common Object File Format (COFF) file and produces a formatted display of the program memory size, data memory size and the programs starting address.

`file` is an operating system compatible file name. Only a single file name may be supplied, and it must be explicit; there is no default extension for the input file.

C.1.3 `cldlod`—Motorola COFF to LOD Format converter

```
cldlod cldfile > lodfile
```

`cldlod` is a utility that converts a binary COFF object file into an ascii LOD file.

`cldfile` is an operating system compatible file name which contains COFF information. Only a single file name may be supplied, and it must be explicit; there is no default extension for the input file.

`lodfile` is an LOD file.

C.1.4 `cofdmp`—Motorola DSP COFF File Dump Utility

```
cofdmp [ -cfhlorstv ] [ -d file ] files
```

`cofdmp` is a utility that reads an absolute or relocatable COFF file and produces a formatted display of the object file contents. The entire file or only selected portions may be processed depending on command line options. The program also can generate either codes or symbolic references to entities such as symbol type or storage class.

`file` is an operating system compatible file name. Only a single file name may be supplied, and it must be explicit; there is no default extension for the input file.

C.1.4.1 `cofdmp` Options

Any of the following command line options listed in Table C-2, may be given. Option letters may be entered in either upper or lower case. If no option is specified, the entire object file is dumped.

Table C-2. `cofdmp` Command Line Options

Option	Description
-c	Dump the object file string table. This information may not be available if the object file has been stripped.
-d	Dump to output file.
-f	Dump the file header of the object file.
-h	Dump the object file section headers.
-l	Dump the object file line number information. This information may not be available if the object file has been stripped.
-o	Dump the object file optional header.

Table C-2. cofdmp Command Line Options

Option	Description
-r	Dump the object file relocation information. This information is available only in relocatable object files.
-s	Dump the object file raw data contents.
-t	Dump the object file symbol table.
-v	Dump the object file symbolically, expanding bit flag, symbol type, and storage class names.

C.1.5 dsplib—Motorola DSP COFF Librarian

dsplib [-a|-c|-d|-l|-r|-u|-v|-x] [-f<argfil>] library [files...]

`dsplib` is a utility that allows separate files to be grouped together into a single file. The resulting library file can then be used for linking by the Motorola DSP Cross Linker program or for general-purpose archival storage.

library is an operating system compatible file name (including optional pathname) indicating the library file to create or access. If no extension is supplied, the librarian will automatically append *.clb* to the file name. If no pathname is specified, the librarian will look for the library in the current directory.

files is a list of operating system compatible file names. For input operations the file names may also contain an optional pathname; the path is stripped when the file is written to the library. For output operations only the file name should be used to refer to library modules.

If no arguments are given on the command line, the librarian enters an interactive mode where multiple commands may be entered without exiting the program. The syntax for the interactive mode is *command library [files...]* where *command* is an action corresponding to one of the options listed below, *library* is the library name, and *files* is the optional (based on the action requested) list of files/modules upon which to operate. For example the command *add foo bar.cln* adds the module *bar.cln* to the library *foo*. Because interactive input is taken from the standard input channel of the host environment, it is possible to create a batch of librarian commands and feed them to the program for execution via redirection. For more information on interactive commands, invoke the librarian without any arguments and enter *help*.

C.1.5.1 dsplib Options

Only one of the following command line options listed in Table C-3 may be given for each invocation of the librarian. Option letters may be entered in either upper or lower case. If no option is given, the librarian operates as if the **-U** option were specified.

Table C-3. cldinfo Command Line Options

Option	Description
-a	This option adds the modules in the file list to the named library. The library file must exist and the modules must not already be in the library.
-c	Create a new library file and add any specified modules to it. If the library file already exists, an error is issued.
-d	Delete named modules from the library. If the module is not in the library, an error is issued.
-f<argfil>	This option indicates that an external file should be read for further command arguments. It is useful in host environments where the command line length is restricted. argfil must be present on the command line but can be any legal operating system file name, including an optional pathname. argfil is a text file containing module names to be passed to the librarian. The names need be separated only by white space (spaces, tabs, or newlines). A semicolon (;) on a line following white space causes the rest of the line to be treated as a comment.
-l	List library contents. This option lists the module name as contained in the library header, the module size (minus library overhead), and the date and time the file was stored into the library. The listing output is routed to standard output so that it may be redirected to a file if desired.
-r	This option replaces the named modules in the given library. The modules must already be present in the library file.
-u	This option updates the specified modules if they exist in the library; otherwise it adds them to the end of the library file.
-v	This option displays the librarian version number and copyright notice.
-x	Extract named modules from the library. The resulting files are given the name of the modules as stored in the library module header. All files are created in the current working directory.

C.1.6 dsplnk—Motorola DSP COFF Linker

```
dsplnk [-B[<lodfil>]] [-F<argfil>] [-I [-L<library>]] [-M<mapfil>] [-N ]
[-O<mem>[<ctr>][<map>]:<origin>] [-P<lpath>] [-R<ctlfil>] [-Z]
[-U<symbol>] [-V] [-X<opt>[, <opt>...]] [-Z] <files...>
```

dsplnk is a program that processes relocatable link files produced by the DSP assemblers, generating an absolute load file which can be

1. loaded directly into the Motorola DSP simulator or
2. converted to Motorola S-record format for PROM burning.

files is a list of operating system compatible file names including optional pathnames. If no extension is supplied for a given file, the linker will first attempt to open the file using the file name as supplied. If that is not successful the linker appends *.cln* to the file name and tries to open the file again. If no pathname is supplied for a given file, the linker will

look for that file in the current directory. The list of files will be processed sequentially in the order given and all files will be used to generate the load file and map listing.

Note that some options (**-B** and **-M**) allow a hyphen as an optional argument which indicates that the corresponding output should be sent to the standard output stream. Unpredictable results may occur if, for example, the object file is explicitly routed to standard output while the listing file is allowed to default to the same output stream.

C.1.6.1 dsplnk Options

Any of the following command line options may be specified. These can be in any order but must precede the list of link file names (except for the **-L** option). Option letters may be specified in either upper or lower case.

Option arguments may immediately follow the option letter or may be separated from the option letter by blanks or tabs. However, an ambiguity arises if an option takes an optional argument. Consider the following command line:

```
dsplnk -b main io
```

In this example it is not clear whether the file *main* is a link file or is meant to be an argument to the **-B** option. If the ambiguity is not resolved, the linker will assume that *main* is a link file and attempt to open it for reading. This may not be what the programmer intended. There are several ways to avoid this ambiguity. If *main* is supposed to be an argument to the **-B** option, it can be placed immediately after the option letter without intervening white space:

```
dsplnk -bmain io
```

If there are other options on the command line besides those that take optional arguments the other options can be placed between the ambiguous option and the list of link file names.

```
dsplnk -b main -v io
```

Alternatively, two successive hyphens may be used to indicate the end of the option list:

```
dsplnk -b -- main io
```

In this case the linker interprets *main* as a link file name and uses the default naming conventions for the **-B** option. Table C-4 lists the different Command Line Options that can be used with *dsplnk*:

Table C-4. dsplnk Command Line Options

Option	Description
-B[<objfil>]	<p>This option specifies a name for the object file generated by the linker. <i>objfil</i> can be any legal operating system file name including an optional pathname. A hyphen may also be used as an argument to indicate that the object file should be sent to the standard output.</p> <p>If a pathname is not given, the file will be created in the current directory. If no file name is supplied or if the -B option is not given, the linker will use the basename (file name without extension) of the first file name encountered in the link input file list. The resulting output file will have an extension of <i>.cld</i>. The -B option should be specified only once.</p>
-F<argfil>	<p>This option indicates that an external file should be read for further command arguments. It is useful in host environments where the command line length is restricted. <i>argfil</i> must be present on the command line but can be any legal operating system file name, including an optional pathname.</p> <p>The file may contain any legal command line options including the -F option itself. The arguments need be separated only by white space (spaces, tabs, or newlines). A semicolon (;) on a line following white space causes the rest of the line to be treated as a comment.</p>
-l	<p>Under normal operation, the linker produces an absolute load file as output. If the -l option appears on the command line, the linker combines the input files into a single relocatable link file suitable for a subsequent linker pass. No absolute addresses are assigned and no errors are issued for unresolved external references.</p>
-L<library>	<p>The linker ordinarily processes a list of link files which each contain a single relocatable code module. If the -L option is encountered, the linker treats the accompanying pathname as a library file, and searches the file for any outstanding unresolved references.</p> <p>If a module is found in the library that resolves an outstanding external reference, the module is read from the library and included in the load file output. The linker continues to search a library until all external references are resolved or no more references can be satisfied within the current library. The linker searches a library only once: when it is encountered on the command line. Therefore, the position of the -L option on the command line is significant.</p>
-M[<mapfil>]	<p>This option specifies that a map file is to be created. <i>mapfil</i> can be any legal operating system file name including an optional pathname.</p> <p>If a pathname is not given, the file will be created in the current directory. If no file name is supplied, the linker will use the basename (file name without extension) of the first file name encountered in the link input file list. The resulting output file will have an extension of <i>.map</i>. The linker will not generate a map file if the -M option is not specified. The -M option should be specified only once.</p>
-N	<p>Indicates that the linker should ignore case in symbol names. Ordinarily the linker is sensitive to upper and lower case letters in symbol names. If the -N option is supplied the linker maps all symbol characters to lower case.</p>

Table C-4. dsplnk Command Line Options

Option	Description														
-O<mem>[<ctr>][<map>]:<origin>	By default, the linker generates instructions and data for the load file beginning at absolute location zero for all DSP memory spaces. This option allows the programmer to redefine the start address for any memory space and associated location counter. <i>mem</i> is one of the single-character memory space identifiers (X, Y, L, and P). The letter may be upper or lower case. The optional <i>ctr</i> is a letter indicating the high (H) or low (L) location counters. If no counter is specified, the default counter is used. <i>map</i> is also optional and signifies the desired physical mapping for all relocatable code in the given memory space. It may be I for internal memory, E for external memory, or B for bootstrap memory (valid only in P program memory space). If <i>map</i> is not supplied, then no explicit mapping is presumed. The <i>origin</i> is a hexadecimal number signifying the new relocation address for the given memory space. The -O option may be specified as many times as needed on the command line.														
-P< path>	When the linker encounters a library specification on the command line, the current directory (or the directory given in the library specification) is first searched for the library file. If it is not found and the -P option is supplied, the linker prefixes the file name (and optional pathname) provided in the library specification with <i>lpath</i> and searches the newly formed directory pathname for the file. The directories will be searched in the order given on the command line.														
-R[<ctlfil>]	This option indicates that a memory control file is to be read to determine the absolute placement of sections in DSP memory. <i>ctlfil</i> can be any legal operating system file name including an optional pathname. If a pathname is not given, an attempt will be made to open the file in the current directory. If no file name is supplied, the linker will use the basename (file name without extension) of the first file name encountered in the link input file list, appending an extension of <i>.ctl</i> . If the -R option is not specified, then the linker will not use a memory map file. The -R option should be specified only once.														
-U<symbol>	Causes <i>symbol</i> to be entered into the unresolved external reference table. This is useful when the initial or only link file is a library. Since there are no external references when the linker is invoked, the -U option may be used to force inclusion of a library module that resolves the undefined reference. The -U option may be specified as often as desired.														
-V	Indicates that the linker should be verbose during processing, displaying a progress report as it links the input files. The linker will show the beginning of each pass and when files are opened and closed. The information is sent to the standard error output stream.														
-X<opt>[,<opt>,...,<opt>]	<p>The -X option directs the linker to perform a little bit of different operation than standard operation of the linker. The options are described below with their different operations performed. All options may be preceded by NO to reverse their meaning. The -X<opt> sequence can be repeated for as many options as desired.</p> <table border="0" data-bbox="493 1556 1386 1755"> <thead> <tr> <th style="text-align: left;">Option</th> <th style="text-align: left;">Meaning</th> </tr> </thead> <tbody> <tr> <td>XC</td> <td>Relative terms from different sections used in an expression cause an error</td> </tr> <tr> <td>RSV</td> <td>Reserve special target processor memory areas (e.g. DSP96000 DMA)</td> </tr> <tr> <td>AEC</td> <td>Check form of address expressions</td> </tr> <tr> <td>RO</td> <td>Allow region overlap</td> </tr> <tr> <td>ESO</td> <td>Do not allocate memory below ordered sections</td> </tr> <tr> <td>ASC</td> <td>Enable absolute section bounds checking</td> </tr> </tbody> </table>	Option	Meaning	XC	Relative terms from different sections used in an expression cause an error	RSV	Reserve special target processor memory areas (e.g. DSP96000 DMA)	AEC	Check form of address expressions	RO	Allow region overlap	ESO	Do not allocate memory below ordered sections	ASC	Enable absolute section bounds checking
Option	Meaning														
XC	Relative terms from different sections used in an expression cause an error														
RSV	Reserve special target processor memory areas (e.g. DSP96000 DMA)														
AEC	Check form of address expressions														
RO	Allow region overlap														
ESO	Do not allocate memory below ordered sections														
ASC	Enable absolute section bounds checking														
-Z	The linker strips source file line number and symbol information from the input file. Symbol information normally is retained for debugging purposes. This options has no effect if incremental linking is being done (see the -I option).														

C.1.7 gdb56—GNU Source-level Debugger for DSP56000/DSP56001

```
gdb56 [-s SYMFILE] [-e EXECFILE] [-se FILE] [-c COREFILE] [-x FILE]
[-d directory] [-nx] [-q] [-batch] [-fullname] [-help] [-tty TTY]
[-cd DIR] <name> <core>
```

gdb56 is a source-level symbolic debugger for C programs created by Richard Stallman for the GNU Project and distributed by the Free Software Foundation. **gdb56** has something of the flavor of the UNIX source-level debugger **dbx** but has more features and power.

gdb56 is invoked with the shell command *gdb56*. Once started, it reads commands from the terminal until given the quit command. *name* is the name of the executable program, and *core*, if specified, is the name of the DSP56000/DSP56001 *state file* to be examined or have execution resumed.

gdb56 has been modified to run with both software and hardware DSP56000/DSP56001 execution devices.

C.1.7.1 gdb56 Options

The following command-line options listed in Table C-5 can be used to more precisely specify the files to be debugged. All the options and command line arguments given are processed in sequential order. The order makes a difference when the **-x** command is used.

Table C-5. gdb56 Command Line Options

Option	Description
-s SYMFILE	Read symbol table from file SYMFILE.
-e EXECFILE	Use EXECFILE as the executable file to execute when appropriate and for examining pure data in conjunction with a core dump.
-se FILE	Read symbol table from FILE and use it as the executable file.
-c COREFILE	Analyze the core dump COREFILE.
-x file	Execute GDB commands from file.
-d directory	Add directory to the path to search for source files.
-batch	Run in batch mode. Exit with code 1 after processing all the command files specified with -x (and .gdbinit, if not inhibited). Exit also if, due to an error, <i>gdb56</i> would otherwise attempt to read a command from the terminal.
-fullname	This option is used when Emacs runs GDB as a subprocess. It tells GDB to produce the full file name and line number each time a stack frame is displayed (which includes each time the program stops).
-help	Print command-line argument summary.
-tty TTY	Use TTY for input/output by the program being debugged.

Table C-5. gdb56 Command Line Options

Option	Description
-cd DIR	Change current directory to DIR .
The following additional command-line options can be used to affect certain aspects of the behavior of gdb56 :	
-nx	Don't execute commands from the init files .gdbinit. Normally, the commands in these files are executed after all the command options and arguments have been processed.
-q	Do not print the version number on start-up.

SEE ALSO

Appendix E — GNU Debugger (GDB)

C.1.8 run56—Motorola DSP56000/DSP56001 Simulator Based Execution Device.

run56 [-b BCR_VALUE] [-s STATE_FILE] [-tx] [file]

run56 is a COFF object file execution utility that provides operating system style hooks to support *hosted* ANSI run-time library routines such as `printf()`.

file is an operating system compatible file name. Only a single file name may be supplied and it must be explicit; there is no default extension for the input file.

C.1.8.1 run56 Options

Table C-6 lists the different Command Line Options that can be used with run56:

Table C-6. run56 Command Line Options

Option	Description
-b BCR_VALUE	Use <i>BCR_VALUE</i> to set the DSP56000/DSP56001 bus control register. Default bcr value is 0.
-s STATE_FILE	Resume the execution of the DSP56000/DSP56001 state file <i>STATE_FILE</i> .
-t	Update global C variable <code>__time</code> at each clock tick. Used to benchmark code.
-x	Indicates the input file was generated from the x memory model of C compiler. The C compiler may use x, y or l memory model for the compilation (refer to the option -m of the C compiler), and this option directs the run56 program to use the X memory space for data buffer. Without this option, y memory model or L memory model is assumed.

C.1.9 srec—Motorola DSP S-Record Conversion Utility

srec [-b|-w] [-m|-s] [-l] [-r] <files ...>

srec converts Motorola DSP .cld and .lod format files into Motorola S-record files. The S-record format was devised for the purpose of encoding programs or data files in a printable form for transportation between computer systems. Motorola S-record format is recognized by many PROM programming systems.

files is a list of operating system compatible file names. If no pathname is specified for a given file, *srec* will look for that file in the current directory. If the special character '-' is used as a file name *srec* will read from the standard input stream. The list of files will be processed sequentially in the order given.

C.1.9.1 srec Options

Only one of the following command line options listed in Table C-7 may be given for each invocation of the librarian. Option letters may be entered in either upper or lower case. If no option is given, the librarian operates as if the -U option were specified.

Table C-7. srec Command Line Options

Option	Description
-b	Use byte addressing when transferring load addresses to S-record addresses. This means that load file DATA record start addresses are multiplied by the number of bytes per target DSP word and subsequent S1/S3 record addresses are computed based on the data byte count. The -b and -w options are mutually exclusive.
-l	Use double-word addressing when transferring load addresses from L space to S-record addresses. This means that load file DATA records for L space data are moved unchanged and subsequent S1/S3 record addresses are computed based on the data word count divided by 2. This option should always be used when the source load file contains DATA records in L memory space.
-m	Split each DSP word into bytes and store the bytes in parallel S-records. The -m and -s options are mutually exclusive.
-r	Write bytes high to low, rather than low to high. This option has no effect when used with the -m option.
-s	Write data to a single file, putting memory space information into the address field of the S0 header record. The -m and -s options are mutually exclusive.
-w	Use word addressing when transferring load addresses to S-record addresses. This means that load file DATA record start addresses are moved unchanged and subsequent S1/S3 record addresses are computed based on the data word count.

Appendix D

GNU Debugger (GDB)

D.1 Introduction

Appendix E contains the GNU Debugger Manual. This preface to the GDB manual describes the differences between the standard `gdb` source level debugger and the `gdb56` debugger that uses the Motorola DSP56000 family simulator.

A short description of the way the simulator is embedded in `gdb56` will help explain some of the differences. The DSP56000 device is simulated by the same non-display version of the simulator program that is supplied with the DSP56000 Development Software. The simulator is linked with the `gdb` program. When a program is executed on the simulated DSP56000, it is done via direct function calls to the simulator library functions rather than via `ptrace` calls to an inferior process as in other versions of `gdb`.

D.1.1 Commands Not Implemented

1. The `attach` and `detach` commands are not implemented because the simulator program is linked directly to the `gdb56` program. As a result, there is no ability to attach or release a separately running simulation.
2. The `core-file` command is not implemented at this time.

All other commands work as documented in the *GDB Manual for the gnu Source-Level Debugger*, Third Edition, October 1989. Some enhancements, listed below, were necessary to accommodate operation using the DSP56000.

D.1.2 DSP56000 Family Differences / Specific Requirements

1. Since the DSP56000 has three memory spaces, it will sometimes be necessary to use one of the prefixes “`p:`”, “`x:`”, or “`y:`” in front of an address in order to examine or modify the specific memory space. The debugger symbol information already has the proper memory space designator accompanying each symbol address so symbolic names *do not* require a memory space prefix; however,

addresses entered as numeric constants *require* the prefix for x or y memory. *P memory is the default.*

2. The gdb56 debugger operates with COFF format “.cld” files generated by the DSP56000 COFF Assembler and Linker programs. The gdb56 subdirectory “m56kinc” contains header files which define the COFF structures used by gdb56 as well as other device-specific header files.
3. The gdb56 evaluator and display routines have been enhanced to handle the 48 bit long values and special format floating-point values used by the DSP56000.
4. The simulation is able to halt at breakpoints without actually breaking the pipeline activity of the device and without actually inserting breakpoint code into the device memory. As a result, the simulator can maintain an accurate record of the device execution time regardless of the number of inserted breakpoints. The cycle count is maintained in the variable \$cyc.
5. The normal gdb program uses.gdbinit as the default initialization file at start-up whereas gdb56 uses.gdbinit56.

GDB Manual

The GNU Source-Level Debugger

Third Edition, GDB version 3.4

October 1989

Richard M. Stallman

Copyright © 1988, 1989 Free Software Foundation, Inc.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the section entitled “GNU General Public License” is included exactly as in the original, and provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that the section entitled “GNU General Public License” may be included in a translation approved by the author instead of in the original English.

D.1.3 Summary of GDB

The purpose of a debugger such as GDB is to allow you to execute another program while examining what is going on inside it. We call the other program “your program” or “the program being debugged”.

GDB can do four kinds of things (plus other things in support of these):

1. Start the program, specifying anything that might affect its behavior.
2. Make the program stop on specified conditions.
3. Examine what has happened, when the program has stopped, so that you can see bugs happen.
4. Change things in the program, so you can correct the effects of one bug and go on to learn about another without having to re-compile first.

GDB can be used to debug programs written in C and C++. Pascal support is being implemented, and Fortran support will be added when a GNU Fortran compiler is written.

D.1.4 GNU General Public License

Version 1, February 1989

Copyright ©1989 Free Software Foundation, Inc.
675 Mass Ave., Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

D.1.5 Preamble

The license agreements of most software companies try to keep users at the mercy of those companies. By contrast, our General Public License is intended to guarantee your freedom to share and change free software---to make sure the software is free for all its users. The General Public License applies to the Free Software Foundation’s software and to any other program whose authors commit to using it. You can use it for your programs, too.

When we speak of free software, we are referring to freedom, not price. Specifically, the General Public License is designed to make sure that you have the freedom to give away or sell copies of free software, that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of a such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must tell them their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

The precise terms and conditions for copying, distribution and modification follow.

D.1.6 Terms and Conditions

1. This License Agreement applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any work containing the Program or a portion of it, either verbatim or with modifications. Each licensee is addressed as "you".
2. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this General Public License and to the absence of any warranty; and give any other recipients of the Program a copy of this General Public License along with the Program. You may charge a fee for the physical act of transferring a copy.
3. You may modify your copy or copies of the Program or any portion of it, and copy and distribute such modifications under the terms of Paragraph 1 above, provided that you also do the following:
 - cause the modified files to carry prominent notices stating that you changed the files and the date of any change; and

- cause the whole of any work that you distribute or publish, that in whole or in part contains the Program or any part thereof, either with or without modifications, to be licensed at no charge to all third parties under the terms of this General Public License (except that you may choose to grant warranty protection to some or all third parties, at your option).
- If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the simplest and most usual way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this General Public License.
- You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

Mere aggregation of another independent work with the Program (or its derivative) on a volume of a storage or distribution medium does not bring the other work under the scope of these terms.

4. You may copy and distribute the Program (or a portion or derivative of it, under Paragraph 2) in object code or executable form under the terms of Paragraphs 1 and 2 above provided that you also do one of the following:
 - accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Paragraphs 1 and 2 above; or,
 - accompany it with a written offer, valid for at least three years, to give any third party free (except for a nominal charge for the cost of distribution) a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Paragraphs 1 and 2 above; or,
 - accompany it with the information you received as to where the corresponding source code may be obtained. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form alone.)

Source code for a work means the preferred form of the work for making modifications to it. For an executable file, complete source code means all the source code for all modules it contains; but, as a special exception, it need not include source code for modules which are standard libraries that accompany the operating system on which the executable file runs, or for standard header files or definitions files that accompany that operating system.

5. You may not copy, modify, sublicense, distribute or transfer the Program except as expressly provided under this General Public License. Any attempt otherwise to copy, modify, sublicense, distribute or transfer the Program is void, and will automatically terminate your rights to use the Program under this License. However, parties who have received copies, or rights to use copies, from you under this General Public License will not have their licenses terminated so long as such parties remain in full compliance.
6. By copying, distributing or modifying the Program (or any work based on the Program) you indicate your acceptance of this license to do so, and all its terms and conditions.
7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein.

The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

8. Each version is given a distinguishing version number. If the Program specifies a version number of the license which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the license, you may choose any version ever published by the Free Software Foundation.
9. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

D.1.7 No Warranty

1. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE

IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

2. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

D.1.8 End of Terms and Conditions

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to humanity, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

one line to give the program’s name and a brief idea of what it does.

Copyright (C) 19yy *name of author*

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 1, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave., Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) 19yy *name of author*

Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type ‘show w’.

This is free software, and you are welcome to redistribute it under certain conditions; type ‘show c’ for details.

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ‘show w’ and ‘show c’; they could even be mouse-clicks or menu items---whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program ‘Gnomovision’ (a program to direct compilers to make passes at assemblers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

That’s all there is to it!

D.2 Input and Output Conventions

GDB is invoked with the shell command ‘gdb’. Once started, it reads commands from the terminal until you tell it to exit.

A GDB command is a single line of input. There is no limit on how long it can be. It starts with a command name, which is followed by arguments whose meaning depends on the command name. For example, the command *step* accepts an argument which is the number of times to step, as in **step 5**. You can also use the *step* command with no arguments. Some command names do not allow any arguments.

GDB command names may always be abbreviated if the abbreviation is unambiguous. Sometimes even ambiguous abbreviations are allowed; for example, ‘s’ is specially defined as equivalent to *step* even though there are other commands whose names start with ‘s’. Possible command abbreviations are often stated in the documentation of the individual commands.

A blank line as input to GDB means to repeat the previous command verbatim. Certain commands do not allow themselves to be repeated this way; these are commands for which unintentional repetition might cause trouble and which you are unlikely to want to repeat. Certain others (*list* and ‘x’) act differently when repeated because that is more useful.

A line of input starting with ‘#’ is a comment; it does nothing. This is useful mainly in command files (See section Command Files).

GDB indicates its readiness to read a command by printing a string called the prompt. This string is normally ‘(gdb)’. You can change the prompt string with the ‘set prompt’ command. For instance, when debugging GDB with GDB, it is useful to change the prompt in one of the GDBs so that you tell which one you are talking to.

```
set prompt newprompt
```

Directs GDB to use *newprompt* as its prompt string henceforth.

To exit GDB, use the *quit* command (abbreviated ‘q’). Ctrl-c will not exit from GDB, but rather will terminate the action of any GDB command that is in progress and return to GDB command level. It is safe to type Ctrl-c at any time because GDB does not allow it to take effect until a time when it is safe.

Certain commands to GDB may produce large amounts of information output to the screen. To help you read all of it, GDB pauses and asks you for input at the end of each page of output. Type RET when you want to continue the output. Normally GDB knows the size of the screen from on the termcap data base together with the value of the TERM environment variable; if this is not correct, you can override it with the ‘set screensize’ command:

```
set screensize lpp  
set screensize lpp cpl
```

Specify a screen height of lpp lines and (optionally) a width of cpl characters. If you omit cpl, the width does not change.

If you specify a height of zero lines, GDB will not pause during output no matter how long the output is. This is useful if output is to a file or to an editor buffer.

Also, GDB may at times produce more information about its own workings than is of interest to the user. Some of these informational messages can be turned on and off with the 'set verbose' command:

```
set verbose off
```

Disables GDB's output of certain informational messages.

```
set verbose on
```

Re-enables GDB's output of certain informational messages.

Currently, the messages controlled by 'set verbose' are those which announce that the symbol table for a source file is being read (see section File Commands, in the description of the command 'symbol-file').

D.3 Specifying GDB's Files

GDB needs to know the file name of the program to be debugged, both in order to read its symbol table and in order to start the program. To debug a core dump of a previous run, GDB must be told the file name of the core dump.

D.3.1 Specifying Files with Arguments

The usual way to specify the executable and core dump file names is with two command arguments given when you start GDB. The first argument is used as the file for execution and symbols, and the second argument (if any) is used as the core dump file name. Thus,

```
gdb prog core
```

specifies 'prog' as the executable program and 'core' as a core dump file to examine. (You do not need to have a core dump file if what you plan to do is debug the program interactively.)

See section Options, for full information on options and arguments for invoking GDB.

D.3.2 Specifying Files with Commands

Usually you specify the files for GDB to work with by giving arguments when you invoke GDB. But occasionally it is necessary to change to a different file during a GDB session. Or you may run GDB and forget to specify the files you want to use. In these situations the GDB commands to specify new files are useful.

Table D-1. GDB File Commands

Command	Description
<code>exec-file filename</code>	Specify that the program to be run is found in filename. If you do not specify a directory and the file is not found in GDB's working directory, GDB will use the environment variable PATH as a list of directories to search, just as the shell does when looking for a program to run.
<code>symbol-file filename</code>	<p>Read symbol table information from file filename. PATH is searched when necessary. Most of the time you will use both the 'exec-file' and 'symbol-file' commands on the same file.</p> <p>'symbol-file' with no argument clears out GDB's symbol table.</p> <p>The 'symbol-file' command does not actually read the symbol table in full right away. Instead, it scans the symbol table quickly to find which source files and which symbols are present. The details are read later, one source file at a time, when they are needed.</p> <p>The purpose of this two-stage reading strategy is to make GDB start up faster. For the most part, it is invisible except for occasional messages telling you that the symbol table details for a particular source file are being read. (The 'set verbose' command controls whether these messages are printed; see section User Interface).</p> <p>However, you will sometimes see in backtraces lines for functions in source files whose data has not been read in; these lines omit some of the information, such as argument values, which cannot be printed without full details of the symbol table.</p> <p>When the symbol table is stored in COFF format, 'symbol-file' does read the symbol table data in full right away. We haven't bothered to implement the two-stage strategy for COFF.</p>
<code>core-file filename</code>	<p>Specify the whereabouts of a core dump file to be used as the "contents of memory". Note that the core dump contains only the writable parts of memory; the read-only parts must come from the executable file.</p> <p>'core-file' with no argument specifies that no core file is to be used.</p> <p>Note that the core file is ignored when your program is actually running under GDB. So, if you have been running the program and you wish to debug a core file instead, you must kill the sub-process in which the program is running. To do this, use the 'kill' command (see section Kill Process).</p>
<code>add-file filename address</code>	<p>The 'add-file' command reads additional symbol table information from the file filename. You would use this when that file has been dynamically loaded into the program that is running. address should be the memory address at which the file has been loaded; GDB cannot figure this out for itself.</p> <p>The symbol table of the file filename is added to the symbol table originally read with the 'symbol-file' command. You can use the 'add-file' command any number of times; the new symbol data thus read keeps adding to the old. The 'symbol-file' command forgets all the symbol data GDB has read; that is the only time symbol data is forgotten in GDB.</p>

Table D-1. GDB File Commands

Command	Description
<code>info files</code>	Print the names of the executable and core dump files currently in use by GDB, and the file from which symbols were loaded. While all three file-specifying commands allow both absolute and relative file names as arguments, GDB always converts the file name to an absolute one and remembers it that way. The 'symbol-file' command causes GDB to forget the contents of its convenience variables, the value history, and all breakpoints and auto-display expressions. This is because they may contain pointers to the internal data recording symbols and data types, which are part of the old symbol table data being discarded inside GDB.

D.4 Compiling Your Program for Debugging

In order to debug a program effectively, you need to ask for debugging information when you compile it. This information in the object file describes the data type of each variable or function and the correspondence between source line numbers and addresses in the executable code.

To request debugging information, specify the '-g' option when you run the compiler.

The Unix C compiler is unable to handle the '-g' and '-O' options together. This means that you cannot ask for optimization if you ask for debugger information.

The GNU C compiler supports '-g' with or without '-O', making it possible to debug optimized code. We recommend that you always use '-g' whenever you compile a program. You may think the program is correct, but there's no sense in pushing your luck.

GDB no longer supports the debugging information produced by giving the GNU C compiler the '-gg' option, so do not use this option.

D.5 Running Your Program Under GDB

To start your program under GDB, use the 'run' command. The program must already have been specified using the 'exec-file' command or with an argument to GDB (see section Files); what 'run' does is create an inferior process, load the program into it, and set it in motion.

The execution of a program is affected by certain information it receives from its superior. GDB provides ways to specify this information, which you must do before starting the program. (You can change it after starting the program, but such changes do not affect the program unless you start it over again.) This information may be divided into three categories:

- The *arguments*.

You specify the arguments to give the program as the arguments of the ‘run’ command.

- The *environment*.

The program normally inherits its environment from GDB, but you can use the GDB commands ‘set environment’ and ‘unset environment’ to change parts of the environment that will be given to the program.

- The *working directory*.

The program inherits its working directory from GDB. You can set GDB’s working directory with the ‘cd’ command in GDB.

After the ‘run’ command, the debugger does nothing but wait for your program to stop. See section Stopping.

Note that once your program has been started by the ‘run’ command, you may evaluate expressions that involve calls to functions in the inferior. See section Expressions. If you wish to evaluate a function simply for its side effects, you may use the ‘set’ command. See section Assignment.

D.5.1 Your Program’s Arguments

The arguments to your program are specified by the arguments of the ‘run’ command. They are passed to a shell, which expands wildcard characters and performs redirection of I/O, and thence to the program.

‘run’ with no arguments uses the same arguments used by the previous ‘run’.

The command ‘set args’ can be used to specify the arguments to be used the next time the program is run. If ‘set args’ has no arguments, it means to use no arguments the next time the program is run. If you have run your program with arguments and want to run it again with no arguments, this is the only way to do so.

D.5.2 Your Program’s Environment

The environment consists of a set of environment variables and their values. Environment variables conventionally record such things as your user name, your home directory, your terminal type, and your search path for programs to run. Usually you set up environment variables with the shell and they are inherited by all the other programs you run. When debugging, it can be useful to try running the program with different environments without having to start the debugger over again.

Table D-2. Environment Variables

Variable	Description
<code>info environment varname</code>	Print the value of environment variable <i>varname</i> to be given to your program when it is started. This command can be abbreviated 'i env varname'.
<code>info environment</code>	Print the names and values of all environment variables to be given to your program when it is started. This command can be abbreviated 'i env'.
<code>set environment varname value</code> <code>set environment varname = value</code>	Sets environment variable <i>varname</i> to <i>value</i> , for your program only, not for GDB itself. <i>value</i> may be any string; the values of environment variables are just strings, and any interpretation is supplied by your program itself. The <i>value</i> parameter is optional; if it is eliminated, the variable is set to a null value. This command can be abbreviated as short as 'set e'. For example, this command: <code>set env USER = foo</code> tells the program, when subsequently run, to assume it is being run on behalf of the user named 'foo'.
<code>delete environment varname</code> <code>unset environment varname</code>	Remove variable <i>varname</i> from the environment to be passed to your program. This is different from 'set env varname =' because 'delete environment' leaves the variable with no value, which is distinguishable from an empty value. This command can be abbreviated 'd e'.

D.5.3 Your Program's Working Directory

Each time you start your program with 'run', it inherits its working directory from the current working directory of GDB. GDB's working directory is initially whatever it inherited from its parent process (typically the shell), but you can specify a new working directory in GDB with the 'cd' command.

The GDB working directory also serves as a default for the commands that specify files for GDB to operate on as listed in Table D-3. See section Files.

Table D-3. Working Directory Commands

Command	Description
<code>cd directory</code>	Set GDB's working directory to <i>directory</i> .
<code>pwd</code>	Print GDB's working directory.

D.5.4 Your Program's Input and Output

By default, the program you run under GDB does input and output to the same terminal that GDB uses.

You can redirect the program's input and/or output using 'sh'-style redirection commands in the 'run' command. For example,

```
run > outfile
```

starts the program, diverting its output to the file 'outfile'.

Another way to specify where the program should do input and output is with the 'tty' command. This command accepts a file name as argument, and causes this file to be the default for future 'run' commands. It also resets the controlling terminal for the child process, for future 'run' commands. For example,

```
tty /dev/ttyb
```

directs that processes started with subsequent 'run' commands default to do input and output on the terminal '/dev/ttyb' and have that as their controlling terminal.

An explicit redirection in 'run' overrides the 'tty' command's effect on input/output redirection, but not its effect on the controlling terminal.

When you use the 'tty' command or redirect input in the 'run' command, only the *input for your program* is affected. The input for GDB still comes from your terminal.

D.5.5 Debugging an Already-Running Process

Some operating systems allow GDB to debug an already running process that was started outside of GDB. To do this, you use the 'attach' command instead of the 'run' command.

The 'attach' command requires one argument, which is the process-id of the process you want to debug. (The usual way to find out the process-id of the process is with the ps utility.)

The first thing GDB does after arranging to debug the process is to stop it. You can examine and modify an attached process with all the GDB commands that ordinarily available when you start processes with 'run'. You can insert breakpoints; you can step and continue; you can modify storage. If you would rather the process continue running, you may use the 'continue' command after attaching GDB to the process.

When you have finished debugging the attached process, you can use the 'detach' command to release it from GDB's control. Detaching the process continues its execution. After the 'detach' command, that process and GDB become completely independent once more, and you are ready to 'attach' another process or start one with 'run'.

If you exit GDB or use the 'run' command while you have an attached process, you kill that process. You will be asked for confirmation if you try to do either of these things.

The ‘attach’ command is also used to debug a remote machine via a serial connection. See section Attach, for more info.

D.5.6 Killing the Child Process

`kill`

Kill the child process in which the program being debugged is running under GDB.

This command is useful if you wish to debug a core dump instead. GDB ignores any core dump file if it is actually running the program, so the ‘kill’ command is the only sure way to make sure the core dump file is used once again.

It is also useful if you wish to run the program outside the debugger for once and then go back to debugging it.

The ‘kill’ command is also useful if you wish to re-compile and re-link the program, since on many systems it is impossible to modify an executable file which is running in a process. But, in this case, it is just as good to exit GDB, since you will need to read a new symbol table after the program is re-compiled if you wish to debug the new version, and restarting GDB is the easiest way to do that.

D.6 Stopping and Continuing

When you run a program normally, it runs until it terminates. The principal purpose of using a debugger is so that you can stop it before that point; or so that if the program runs into trouble you can investigate and find out why.

D.6.1 Signals

A signal is an asynchronous event that can happen in a program. The operating system defines the possible kinds of signals, and gives each kind a name and a number. For example, SIGINT is the signal a program gets when you type Ctrl-c; SIGSEGV is the signal a program gets from referencing a place in memory far away from all the areas in use; SIGALRM occurs when the alarm clock timer goes off (which happens only if the program has requested an alarm).

Some signals, including SIGALRM, are a normal part of the functioning of the program. Others, such as SIGSEGV, indicate errors; these signals are fatal (kill the program immediately) if the program has not specified in advance some other way to handle the signal. SIGINT does not indicate an error in the program, but it is normally fatal so it can carry out the purpose of Ctrl-c: to kill the program.

GDB has the ability to detect any occurrence of a signal in the program running under GDB's control. You can tell GDB in advance what to do for each kind of signal.

Normally, GDB is set up to ignore non-erroneous signals like SIGALRM (so as not to interfere with their role in the functioning of the program) but to stop the program immediately whenever an error signal happens. You can change these settings with the 'handle' command. You must specify which signal you are talking about with its number.

```
info signal
```

Print a table of all the kinds of signals and how GDB has been told to handle each one. You can use this to see the signal numbers of all the defined types of signals.

```
handle signalnum keywords...
```

Change the way GDB handles signal *signalnum*. The keywords say what change to make.

To use the 'handle' command you must know the code number of the signal you are concerned with. To find the code number, type 'info signal' which prints a table of signal names and numbers.

The keywords allowed by the handle command can be abbreviated. Table D-4 lists their full names:

Table D-4. Keywords

Keyword	Description
stop	GDB should stop the program when this signal happens. This implies the 'print' keyword as well.
print	GDB should print a message when this signal happens.
nostop	GDB should not stop the program when this signal happens. It may still print a message telling you that the signal has come in.
noprint	GDB should not mention the occurrence of the signal at all. This implies the 'nostop' keyword as well.
pass	GDB should allow the program to see this signal; the program will be able to handle the signal, or may be terminated if the signal is fatal and not handled.
nopass	GDB should not allow the program to see this signal. When a signal has been set to stop the program, the program cannot see the signal until you continue. It will see the signal then, if 'pass' is in effect for the signal in question at that time. In other words, after GDB reports a signal, you can use the 'handle' command with 'pass' or 'nopass' to control whether that signal will be seen by the program when you later continue it. You can also use the 'signal' command to prevent the program from seeing a signal, or cause it to see a signal it normally would not see, or to give it any signal at any time. See section Signaling.

D.6.2 Breakpoints

A *breakpoint* makes your program stop whenever a certain point in the program is reached. You set breakpoints explicitly with GDB commands, specifying the place where the program should stop by line number, function name or exact address in the program. You can add various other conditions to control whether the program will stop.

Each breakpoint is assigned a number when it is created; these numbers are successive integers starting with 1. In many of the commands for controlling various features of breakpoints you use the breakpoint number to say which breakpoint you want to change. Each breakpoint may be *enabled* or *disabled*; if disabled, it has no effect on the program until you enable it again.

The command ‘info break’ prints a list of all breakpoints set and not deleted, showing their numbers, where in the program they are, and any special features in use for them. Disabled breakpoints are included in the list, but marked as disabled. ‘info break’ with a breakpoint number as argument lists only that breakpoint. The convenience variable `$_` and the default examining-address for the ‘x’ command are set to the address of the last breakpoint listed (see section Memory).

D.6.2.1 Setting Breakpoints

Breakpoints are set with the ‘break’ command (abbreviated ‘b’). You have several ways to say where the breakpoint should go as shown in Table D-5.

Table D-5. Breakpoints

Command	Description
<code>break function</code>	Set a breakpoint at entry to function <i>function</i> .
<code>break +offset</code> <code>break -offset</code>	Set a breakpoint some number of lines forward or back from the position at which execution stopped in the currently selected frame.
<code>break linenum</code>	Set a breakpoint at line <i>linenum</i> in the current source file. That file is the last file whose source text was printed. This breakpoint will stop the program just before it executes any of the code on that line.
<code>break filename:linenum</code>	Set a breakpoint at line <i>linenum</i> in source file <i>filename</i> .
<code>break filename:function</code>	Set a breakpoint at entry to function <i>function</i> found in file <i>filename</i> . Specifying a file name as well as a function name is superfluous except when multiple files contain similarly named functions.
<code>break *address</code>	Set a breakpoint at address <i>address</i> . You can use this to set breakpoints in parts of the program which do not have debugging information or source files.

Table D-5. Breakpoints

Command	Description
<code>break</code>	Set a breakpoint at the next instruction to be executed in the selected stack frame (see section Stack). In any selected frame but the innermost, this will cause the program to stop as soon as control returns to that frame. This is equivalent to a 'finish' command in the frame inside the selected frame. If this is done in the innermost frame, GDB will stop the next time it reaches the current location; this may be useful inside of loops. GDB normally ignores breakpoints when it resumes execution, until at least one instruction has been executed. If it did not do this, you would be unable to proceed past a breakpoint without first disabling the breakpoint. This rule applies whether or not the breakpoint already existed when the program stopped.
<code>break ... if cond</code>	Set a breakpoint with condition <code>cond</code> ; evaluate the expression <code>cond</code> each time the breakpoint is reached, and stop only if the value is nonzero. '...' stands for one of the possible arguments described above (or no argument) specifying where to break. See section Conditions, for more information on breakpoint conditions.
<code>tbreak args</code>	Set a breakpoint enabled only for one stop. <code>args</code> are the same as in the 'break' command, and the breakpoint is set in the same way, but the breakpoint is automatically disabled the first time it is hit. See section Disabling. GDB allows you to set any number of breakpoints at the same place in the program. There is nothing silly or meaningless about this. When the breakpoints are conditional, this is even useful (see section Conditions).

D.6.2.2 Deleting Breakpoints

It is often necessary to eliminate a breakpoint once it has done its job and you no longer want the program to stop there. This is called *deleting* the breakpoint. A breakpoint that has been deleted no longer exists in any sense; it is forgotten.

With the 'clear' command you can delete breakpoints according to where they are in the program. With the 'delete' command you can delete individual breakpoints by specifying their breakpoint numbers. See Table D-6.

It is not necessary to delete a breakpoint to proceed past it. GDB automatically ignores breakpoints in the first instruction to be executed when you continue execution without changing the execution address.

Table D-6. Deleting Breakpoints Commands

Command	Description
<code>clear</code>	Delete any breakpoints at the next instruction to be executed in the selected stack frame (see section Selection). When the innermost frame is selected, this is a good way to delete a breakpoint that the program just stopped at.
<code>clear function</code> <code>clear filename:function</code>	Delete any breakpoints set at entry to the function <i>function</i> .
<code>clear linenum</code> <code>clear filename:linenum</code>	Delete any breakpoints set at or within the code of the specified line.
<code>delete bnums...</code>	Delete the breakpoints of the numbers specified as arguments.

D.6.2.3 Disabling Breakpoints

Rather than deleting a breakpoint, you might prefer to *disable* it. This makes the breakpoint inoperative as if it had been deleted, but remembers the information on the breakpoint so that you can *enable* it again later.

You disable and enable breakpoints with the ‘enable’ and ‘disable’ commands, specifying one or more breakpoint numbers as arguments. Use ‘info break’ to print a list of breakpoints if you don’t know which breakpoint numbers to use.

A breakpoint can have any of four different states of enablement:

- Enabled. The breakpoint will stop the program. A breakpoint made with the ‘break’ command starts out in this state.
- Disabled. The breakpoint has no effect on the program.
- Enabled once. The breakpoint will stop the program, but when it does so it will become disabled. A breakpoint made with the ‘tbreak’ command starts out in this state.
- Enabled for deletion. The breakpoint will stop the program, but immediately after it does so it will be deleted permanently.

You change the state of enablement of a breakpoint with the following commands as shown in Table D-7:

Table D-7. Disabling Breakpoint Commands

Command	Description
<code>disable breakpoints bnums...</code> <code>disable bnums...</code>	Disable the specified breakpoints. A disabled breakpoint has no effect but is not forgotten. All options such as ignore-counts, conditions and commands are remembered in case the breakpoint is enabled again later.

Table D-7. Disabling Breakpoint Commands

Command	Description
<code>enable breakpoints bnums...</code> <code>enable bnums...</code>	Enable the specified breakpoints. They become effective once again in stopping the program, until you specify otherwise.
<code>enable breakpoints once bnums...</code> <code>enable once bnums...</code>	Enable the specified breakpoints temporarily. Each will be disabled again the next time it stops the program (unless you have used one of these commands to specify a different state before that time comes).
<code>enable breakpoints delete bnums...</code> <code>enable delete bnums...</code>	Enable the specified breakpoints to work once and then die. Each of the breakpoints will be deleted the next time it stops the program (unless you have used one of these commands to specify a different state before that time comes).

Aside from the automatic disablement or deletion of a breakpoint when it stops the program, which happens only in certain states, the state of enablement of a breakpoint changes only when one of the commands above is used.

D.6.2.4 Break Conditions

The simplest sort of breakpoint breaks every time the program reaches a specified place. You can also specify a *condition* for a breakpoint. A condition is just a boolean expression in your programming language. (See section Expressions). A breakpoint with a condition evaluates the expression each time the program reaches it, and the program stops only if the condition is true.

Break conditions may have side effects, and may even call functions in your program. These may sound like strange things to do, but their effects are completely predictable unless there is another enabled breakpoint at the same address. (In that case, GDB might see the other breakpoint first and stop the program without checking the condition of this one.) Note that breakpoint commands are usually more convenient and flexible for the purpose of performing side effects when a breakpoint is reached (see section Break Commands).

Break conditions can be specified when a breakpoint is set, by using ‘if’ in the arguments to the ‘break’ command. See section Set Breaks. They can also be changed at any time with the ‘condition’ command:

```
condition bnum expression
```

Specify *expression* as the break condition for breakpoint number *bnum*. From now on, this breakpoint will stop the program only if the value of *expression* is true (nonzero, in C). *expression* is not evaluated at the time the ‘condition’ command is given. See section Expressions.

```
condition bnum
```

Remove the condition from breakpoint number *bnum*. It becomes an ordinary unconditional breakpoint.

A special case of a breakpoint condition is to stop only when the breakpoint has been reached a certain number of times. This is so useful that there is a special way to do it, using the ignore count of the breakpoint. Every breakpoint has an ignore count, which is an integer. Most of the time, the ignore count is zero, and therefore has no effect. But if the program reaches a breakpoint whose ignore count is positive, then instead of stopping, it just decrements the ignore count by one and continues. As a result, if the ignore count value is *n*, the breakpoint will not stop the next *n* times it is reached.

```
ignore bnum count
```

Set the ignore count of breakpoint number *bnum* to *count*. The next *count* times the breakpoint is reached, it will not stop.

To make the breakpoint stop the next time it is reached, specify a count of zero.

```
cont count
```

Continue execution of the program, setting the ignore count of the breakpoint that the program stopped at to *count* minus one. Thus, the program will not stop at this breakpoint until the *count*'th time it is reached.

This command is allowed only when the program stopped due to a breakpoint. At other times, the argument to 'cont' is ignored.

If a breakpoint has a positive ignore count and a condition, the condition is not checked. Once the ignore count reaches zero, the condition will start to be checked.

Note that you could achieve the effect of the ignore count with a condition such as '\$foo-<= 0' using a debugger convenience variable that is decremented each time. See section Convenience Vars.

D.6.2.5 Commands Executed on Breaking

You can give any breakpoint a series of commands to execute when the program stops due to that breakpoint. For example, you might want to print the values of certain expressions, or enable other breakpoints.

```
commands bnum
```

Specify commands for breakpoint number *bnum*. The commands themselves appear on the following lines. Type a line containing just 'end' to terminate the commands.

To remove all commands from a breakpoint, use the command ‘commands’ and follow it immediately by ‘end’; that is, give no commands.

With no arguments, ‘commands’ refers to the last breakpoint set.

It is possible for breakpoint commands to start the program up again. Simply use the ‘cont’ command, or ‘step’, or any other command to resume execution. However, any remaining breakpoint commands are ignored. When the program stops again, GDB will act according to the cause of that stop.

If the first command specified is ‘silent’, the usual message about stopping at a breakpoint is not printed. This may be desirable for breakpoints that are to print a specific message and then continue. If the remaining commands too print nothing, you will see no sign that the breakpoint was reached at all. ‘silent’ is not really a command; it is meaningful only at the beginning of the commands for a breakpoint.

The commands ‘echo’ and ‘output’ that allow you to print precisely controlled output are often useful in silent breakpoints. See section Output.

For example, here is how you could use breakpoint commands to print the value of x at entry to foo whenever it is positive.

```
break foo if x>0
commands
silent
echo x is\040
output x
echo \n
cont
end
```

One application for breakpoint commands is to correct one bug so you can test another. Put a breakpoint just after the erroneous line of code, give it a condition to detect the case in which something erroneous has been done, and give it commands to assign correct values to any variables that need them. End with the ‘cont’ command so that the program does not stop, and start with the ‘silent’ command so that no output is produced.

Here is an example:

```
break 403
commands
silent
set x = y + 4
cont
end
```

One deficiency in the operation of automatically continuing breakpoints under Unix appears when your program uses raw mode for the terminal. GDB switches back to its own terminal modes (not raw) before executing commands, and then must switch back to raw mode when your program is continued. This causes any pending terminal input to be lost.

In the GNU system, this will be fixed by changing the behavior of terminal modes.

Under Unix, when you have this problem, you might be able to get around it by putting your actions into the breakpoint condition instead of commands. For example

```
condition 5 (x = y + 4), 0
```

specifies a condition expression (See section Expressions) that will change `x` as needed, then always have the value 0 so the program will not stop. Loss of input is avoided here because break conditions are evaluated without changing the terminal modes. When you want to have nontrivial conditions for performing the side effects, the operators ‘&&’, ‘||’ and ‘?...:’ may be useful.

D.6.2.6 “Cannot Insert Breakpoints” Error

Under some operating systems, breakpoints cannot be used in a program if any other process is running that program. Attempting to run or continue the program with a breakpoint in this case will cause GDB to stop it.

When this happens, you have three ways to proceed:

1. Remove or disable the breakpoints, then continue.
2. Suspend GDB, and copy the file containing the program to a new name. Resume GDB and use the ‘exec-file’ command to specify that GDB should run the program under that name. Then start the program again.
3. Re-link the program so that the text segment is non-sharable, using the linker option ‘-N’. The operating system limitation may not apply to non-sharable executables.

D.6.3 Continuing

After your program stops, most likely you will want it to run some more if the bug you are looking for has not happened yet.

`cont`

Continue running the program at the place where it stopped.

If the program stopped at a breakpoint, the place to continue running is the address of the breakpoint. You might expect that continuing would just stop at the same breakpoint immediately. In fact, ‘cont’ takes special care to prevent that from happening. You do not need to delete the breakpoint to proceed through it after stopping at it.

You can, however, specify an ignore-count for the breakpoint that the program stopped at, by means of an argument to the ‘cont’ command. See section Conditions.

If the program stopped because of a signal other than SIGINT or SIGTRAP, continuing will cause the program to see that signal. You may not want this to happen. For example, if the program stopped due to some sort of memory reference error, you might store correct values into the erroneous variables and continue, hoping to see more execution; but the program would probably terminate immediately as a result of the fatal signal once it sees the signal. To prevent this, you can continue with ‘signal 0’. See section Signaling. You can also act in advance to prevent the program from seeing certain kinds of signals, using the ‘handle’ command (see section Signals).

D.6.4 Stepping

Stepping means setting your program in motion for a limited time, so that control will return automatically to the debugger after one line of code or one machine instruction. Breakpoints are active during stepping and the program will stop for them even if it has not gone as far as the stepping command specifies.

Table D-8. Stepping Commands

Command	Description
<code>step</code>	Continue running the program until control reaches a different line, then stop it and return control to the debugger. This command is abbreviated ‘s’. This command may be given when control is within a function for which there is no debugging information. In that case, execution will proceed until control reaches a different function, or is about to return from this function. An argument repeats this action.
<code>step count</code>	Continue running as in ‘step’, but do so count times. If a breakpoint is reached or a signal not related to stepping occurs before count steps, stepping stops right away.

Table D-8. Stepping Commands (Continued)

Command	Description
next	<p>Similar to 'step', but any function calls appearing within the line of code are executed without stopping. Execution stops when control reaches a different line of code at the stack level which was executing when the 'next' command was given. This command is abbreviated 'n'. An argument is a repeat count, as in 'step'.</p> <p>'next' within a function without debugging information acts as does 'step', but any function calls appearing within the code of the function are executed without stopping.</p>
finish	<p>Continue running until just after the selected stack frame returns (or until there is some other reason to stop, such as a fatal signal or a breakpoint). Print value returned by the selected stack frame (if any).</p> <p>Contrast this with the 'return' command (see section Returning).</p>
until	<p>This command is used to avoid single stepping through a loop more than once. It is like the 'next' command, except that when 'until' encounters a jump, it automatically continues execution until the program counter is greater than the address of the jump.</p> <p>This means that when you reach the end of a loop after single stepping though it, 'until' will cause the program to continue execution until the loop is exited. In contrast, a 'next' command at the end of a loop will simply step back to the beginning of the loop, which would force you to step through the next iteration.</p> <p>'until' always stops the program if it attempts to exit the current stack frame.</p> <p>'until' may produce somewhat counterintuitive results if the order of the source lines does not match the actual order of execution. For example, in a typical C for-loop, the third expression in the for-statement (the loop-step expression) is executed after the statements in the body of the loop, but is written before them. Therefore, the 'until' command would appear to step back to the beginning of the loop when it advances to this expression. However, it has not really done so, not in terms of the actual machine code.</p> <p>Note that 'until' with no argument works by means of single instruction stepping, and hence is slower than 'until' with an argument.</p>
until location	<p>Continue running the program until either the specified location is reached, or the current (innermost) stack frame returns. This form of the command uses breakpoints, and hence is quicker than 'until' without an argument.</p>
stepi si	<p>Execute one machine instruction, then stop and return to the debugger.</p> <p>It is often useful to do 'display/i \$pc' when stepping by machine instructions. This will cause the next instruction to be executed to be displayed automatically at each stop. See section Auto Display.</p> <p>An argument is a repeat count, as in 'step'.</p>
nexti ni	<p>Execute one machine instruction, but if it is a subroutine call, proceed until the subroutine returns.</p> <p>An argument is a repeat count, as in 'next'.</p> <p>A typical technique for using stepping is to put a breakpoint (see section Breakpoints) at the beginning of the function or the section of the program in which a problem is believed to lie, and then step through the suspect area, examining the variables that are interesting, until the problem happens.</p> <p>The 'cont' command can be used after stepping to resume execution until the next breakpoint or signal.</p>

D.7 Examining the Stack

When your program has stopped, the first thing you need to know is where it stopped and how it got there.

Each time your program performs a function call, the information about where in the program the call was made from is saved in a block of data called a *stack frame*. The frame also contains the arguments of the call and the local variables of the function that was called. All the stack frames are allocated in a region of memory called the *call stack*.

When your program stops, the GDB commands for examining the stack allow you to see all of this information.

One of the stack frames is *selected* by GDB and many GDB commands refer implicitly to the selected frame. In particular, whenever you ask GDB for the value of a variable in the program, the value is found in the selected frame. There are special GDB commands to select whichever frame you are interested in.

When the program stops, GDB automatically selects the currently executing frame and describes it briefly as the ‘frame’ command does (see section Frame Info, Info).

D.7.1 Stack Frames

The call stack is divided up into contiguous pieces called *stack frames*, or *frames* for short; each frame is the data associated with one call to one function. The frame contains the arguments given to the function, the function’s local variables, and the address at which the function is executing.

When your program is started, the stack has only one frame, that of the function *main*. This is called the *initial* frame or the *outermost* frame. Each time a function is called, a new frame is made. Each time a function returns, the frame for that function invocation is eliminated. If a function is recursive, there can be many frames for the same function. The frame for the function in which execution is actually occurring is called the *innermost* frame. This is the most recently created of all the stack frames that still exist.

Inside your program, stack frames are identified by their addresses. A stack frame consists of many bytes, each of which has its own address; each kind of computer has a convention for choosing one of those bytes whose address serves as the address of the frame. Usually this address is kept in a register called the *frame pointer register* while execution is going on in that frame.

GDB assigns numbers to all existing stack frames, starting with zero for the innermost frame, one for the frame that called it, and so on upward. These numbers do not really

exist in your program; they are to give you a way of talking about stack frames in GDB commands.

Many GDB commands refer implicitly to one stack frame. GDB records a stack frame that is called the *selected* stack frame; you can select any frame using one set of GDB commands, and then other commands will operate on that frame. When your program stops, GDB automatically selects the innermost frame.

Some functions can be compiled to run without a frame reserved for them on the stack. This is occasionally done with heavily used library functions to save the frame setup time. GDB has limited facilities for dealing with these function invocations; if the innermost function invocation has no stack frame, GDB will give it a virtual stack frame of 0 and correctly allow tracing of the function call chain. Results are undefined if a function invocation besides the innermost one is frameless.

D.7.2 Backtraces

A backtrace is a summary of how the program got where it is. It shows one line per frame, for many frames, starting with the currently executing frame (frame zero), followed by its caller (frame one), and on up the stack. See Table D-6.

Table D-9. Backtraces

Backtrace	Description
<code>backtrace</code> <code>bt</code>	Print a backtrace of the entire stack: one line per frame for all frames in the stack. You can stop the backtrace at any time by typing the system interrupt character, normally Control-C.
<code>backtrace n</code> <code>bt n</code>	Similar, but print only the innermost <i>n</i> frames.
<code>backtrace -n</code> <code>bt -n</code>	Similar, but print only the outermost <i>n</i> frames. The names 'where' and 'info stack' are additional aliases for 'backtrace'. Every line in the backtrace shows the frame number, the function name and the program counter value.

If the function is in a source file whose symbol table data has been fully read, the backtrace shows the source file name and line number, as well as the arguments to the function. (The program counter value is omitted if it is at the beginning of the code for that line number.)

If the source file's symbol data has not been fully read, just scanned, this extra information is replaced with an ellipsis. You can force the symbol data for that frame's source file to be read by selecting the frame. (See section Selection).

Here is an example of a backtrace. It was made with the command ‘bt 3’, so it shows the innermost three frames.

```
#0 rtx_equal_p (x=(rtx) 0x8e58c, y=(rtx) 0x1086c4)
(/gp/rms/cc/rtlanal.c line 337)
#1 0x246b0 in expand_call (...) (...)
#2 0x21cfc in expand_expr (...) (...)
(More stack frames follow...)
```

The functions `expand_call` and `expand_expr` are in a file whose symbol details have not been fully read. Full detail is available for the function `rtx_equal_p`, which is in the file ‘rtlanal.c’. Its arguments, named `x` and `y`, are shown with their typed values.

D.7.3 Selecting a Frame

Most commands for examining the stack and other data in the program work on whichever stack frame is selected at the moment. Here are the commands for selecting a stack frame; all of them finish by printing a brief description of the stack frame just selected. See Table D-10.

Table D-10. Frame Commands

Frame	Description
<code>frame n</code>	Select frame number <code>n</code> . Recall that frame zero is the innermost (currently executing) frame, frame one is the frame that called the innermost one, and so on. The highest-numbered frame is main’s frame.
<code>frame addr</code>	Select the frame at address <code>addr</code> . This is useful mainly if the chaining of stack frames has been damaged by a bug, making it impossible for GDB to assign numbers properly to all frames. In addition, this can be useful when the program has multiple stacks and switches between them.
<code>up n</code>	Select the frame <code>n</code> frames up from the frame previously selected. For positive numbers <code>n</code> , this advances toward the outermost frame, to higher frame numbers, to frames that have existed longer. <code>n</code> defaults to one.
<code>down n</code>	Select the frame <code>n</code> frames down from the frame previously selected. For positive numbers <code>n</code> , this advances toward the innermost frame, to lower frame numbers, to frames that were created more recently. <code>n</code> defaults to one.

All of these commands end by printing some information on the frame that has been selected: the frame number, the function name, the arguments, the source file and line number of execution in that frame, and the text of that source line. For example:

```
#3 main (argc=3, argv=??, env=??) at main.c, line 67
67 read_input_file (argv[i]);
```

After such a printout, the ‘list’ command with no arguments will print ten lines centered on the point of execution in the frame. See section List.

D.7.4 Information on a Frame

There are several other commands to print information about the selected stack frame as listed in Table D-11.

Table D-11. Frame Command

Command	Description
<code>frame</code>	This command prints a brief description of the selected stack frame. It can be abbreviated 'f'. With an argument, this command is used to select a stack frame; with no argument, it does not change which frame is selected, but still prints the same information.
<code>info frame</code>	This command prints a verbose description of the selected stack frame, including the address of the frame, the addresses of the next frame in (called by this frame) and the next frame out (caller of this frame), the address of the frame's arguments, the program counter saved in it (the address of execution in the caller frame), and which registers were saved in the frame. The verbose description is useful when something has gone wrong that has made the stack format fail to fit the usual conventions.
<code>info frame addr</code>	Print a verbose description of the frame at address <i>addr</i> , without selecting that frame. The selected frame remains unchanged by this command.
<code>info args</code>	Print the arguments of the selected frame, each on a separate line.
<code>info locals</code>	Print the local variables of the selected frame, each on a separate line. These are all variables declared static or automatic within all program blocks that execution in this frame is currently inside of.

D.8 Examining Source Files

GDB knows which source files your program was compiled from, and can print parts of their text. When your program stops, GDB spontaneously prints the line it stopped in. Likewise, when you select a stack frame (see section Selection), GDB prints the line which execution in that frame has stopped in. You can also print parts of source files by explicit command.

D.8.1 Printing Source Lines

To print lines from a source file, use the 'list' command (abbreviated 'l'). There are several ways to specify what part of the file you want to print.

Table D-12 shows the forms of the ‘list’ command most commonly used:

Table D-12. List Commands

Command	Description
<code>list linenum</code>	Print ten lines centered around line number <i>linenum</i> in the current source file.
<code>list function</code>	Print ten lines centered around the beginning of function <i>function</i> .
<code>list</code>	Print ten more lines. If the last lines printed were printed with a ‘list’ command, this prints ten lines following the last lines printed; however, if the last line printed was a solitary line printed as part of displaying a stack frame (see section Stack), this prints ten lines centered around that line.
<code>list -</code>	Print ten lines just before the lines last printed. Repeating a ‘list’ command with RET discards the argument, so it is equivalent to typing just ‘list’. This is more useful than listing the same lines again. An exception is made for an argument of ‘-’; that argument is preserved in repetition so that each repetition moves up in the file. In general, the ‘list’ command expects you to supply zero, one or two linespecs. Linespecs specify source lines; there are several ways of writing them but the effect is always to specify some source line. Here is a complete description of the possible arguments for ‘list’:
<code>list linespec</code>	Print ten lines centered around the line specified by <i>linespec</i> .
<code>list first, last</code>	Print lines from first to last. Both arguments are linespecs.
<code>list , last</code>	Print ten lines ending with last.
<code>list first,</code>	Print ten lines starting with first.
<code>list +</code>	Print ten lines just after the lines last printed.
<code>list -</code>	Print ten lines just before the lines last printed.
<code>list</code>	As described in the preceding table.

Table D-13 shows the ways of specifying a single source line--all the kinds of linespec.

Table D-13. Single Source Line Commands

Command	Description
<code>linenum</code>	Specifies line <i>linenum</i> of the current source file. When a ‘list’ command has two linespecs, this refers to the same source file as the first linespec.
<code>+offset</code>	Specifies the line <i>offset</i> lines after the last line printed. When used as the second linespec in a ‘list’ command that has two, this specifies the line offset lines down from the first linespec.
<code>-offset</code>	Specifies the line <i>offset</i> lines before the last line printed.
<code>filename:linenum</code>	Specifies line <i>linenum</i> in the source file <i>filename</i> .
<code>function</code>	Specifies the line of the open-brace that begins the body of the function <i>function</i> .

Table D-13. Single Source Line Commands

Command	Description
<i>filename: function</i>	Specifies the line of the open-brace that begins the body of the function <i>function</i> in the file <i>filename</i> . The file name is needed with a function name only for disambiguation of identically named functions in different source files.
<i>* address</i>	Specifies the line containing the program address <i>address</i> . <i>address</i> may be any expression. One other command is used to map source lines to program addresses.
<i>info line linenum</i>	Print the starting and ending addresses of the compiled code for source line <i>linenum</i> .

The default examine address for the ‘x’ command is changed to the starting address of the line, so that ‘x/i’ is sufficient to begin examining the machine code (see section Memory). Also, this address is saved as the value of the convenience variable \$_ (see section Convenience Vars).

D.8.2 Searching Source Files

There are two commands for searching through the current source file for a regular expression.

The command ‘forward-search *regex*’ checks each line, starting with the one following the last line listed, for a match for *regex*. It lists the line that is found. You can abbreviate the command name as ‘fo’.

The command ‘reverse-search *regex*’ checks each line, starting with the one before the last line listed and going backward, for a match for *regex*. It lists the line that is found. You can abbreviate this command with as little as ‘rev’.

D.8.3 Specifying Source Directories

Executable programs do not record the directories of the source files from which they were compiled, just the names. GDB remembers a list of directories to search for source files; this is called the source path. Each time GDB wants a source file, it tries all the directories in the list, in the order they are present in the list, until it finds a file with the desired name.

Note: The executable search path is *not* used for this purpose. Neither is the **current** working directory, unless it happens to be in the source path.

When you start GDB, its source path contains just the current working directory. To add other directories, use the ‘directory’ command as shown in Table D-14.

Table D-14. Directory Commands

Command	Description
<code>directory dirnames...</code>	Add directory <i>dirname</i> to the end of the source path. Several directory names may be given to this command, separated by whitespace or ':'.
<code>directory</code>	Reset the source path to just the current working directory of GDB. This requires confirmation. Since this command deletes directories from the search path, it may change the directory in which a previously read source file will be discovered. To make this work correctly, this command also clears out the tables GDB maintains about the source files it has already found.
<code>info directories</code>	Print the source path: show which directories it contains.

Because the ‘`directory`’ command adds to the end of the source path, it does not affect any file that GDB has already found. If the source path contains directories that you do not want, and these directories contain misleading files with names matching your source files, the way to correct the situation is as follows:

1. Choose the directory you want at the beginning of the source path. Use the ‘`cd`’ command to make that the current working directory.
2. Use ‘`directory`’ with no argument to reset the source path to just that directory.
3. Use ‘`directory`’ with suitable arguments to add any other directories you want in the source path.

D.9 Examining Data

The usual way to examine data in your program is with the ‘`print`’ command (abbreviated ‘`p`’). It evaluates and prints the value of any valid expression of the language the program is written in (for now, C). You type

```
print exp
```

where *exp* is any valid expression, and the value of *exp* is printed in a format appropriate to its data type.

A more low-level way of examining data is with the ‘`x`’ command. It examines data in memory at a specified address and prints it in a specified format.

D.9.1 Expressions

Many different GDB commands accept an expression and compute its value. Any kind of constant, variable or operator defined by the programming language you are using is legal in an expression in GDB. This includes conditional expressions, function calls, casts and

string constants. It unfortunately does not include symbols defined by preprocessor `#define` commands.

Casts are supported in all languages, not just in C, because it is so useful to cast a number into a pointer so as to examine a structure at that address in memory.

GDB supports three kinds of operator in addition to those of programming languages as shown in Table D-15:

Table D-15. GDB Operators

Operator	Description
@	'@' is a binary operator for treating parts of memory as arrays. See section Arrays, for more information.
::	::' allows you to specify a variable in terms of the file or function it is defined in. See section Variables.
{type}	addr Refers to an object of type <i>type</i> stored at address <i>addr</i> in memory. <i>addr</i> may be any expression whose value is an integer or pointer (but parentheses are required around non-unary operators, just as in a cast). This construct is allowed regardless of what kind of data is officially supposed to reside at <i>addr</i> .

D.9.2 Program Variables

The most common kind of expression to use is the name of a variable in your program.

Variables in expressions are understood in the selected stack frame (see section Selection); they must either be global (or static) or be visible according to the scope rules of the programming language from the point of execution in that frame. This means that in the function

```
foo (a)
int a;
{
    bar (a);
    {
        int b = test ();
        bar (b);
    }
}
```

the variable `a` is usable whenever the program is executing within the function `foo`, but the variable `b` is visible only while the program is executing inside the block in which `b` is declared.

As a special exception, you can refer to a variable or function whose scope is a single source file even if the current execution point is not in this file. But it is possible to have more than one such variable or function with the same name (if they are in different source files). In such a case, it is not defined which one you will get. If you wish, you can specify any one of them using the colon colon construct:

```
block::variable
```

Here `block` is the name of the source file whose variable you want.

D.9.3 Artificial Arrays

It is often useful to print out several successive objects of the same type in memory; a section of an array, or an array of dynamically determined size for which only a pointer exists in the program.

This can be done by constructing an artificial array with the binary operator '@'. The left operand of '@' should be the first element of the desired array, as an individual object. The right operand should be the length of the array. The result is an array value whose elements are all of the type of the left argument. The first element is actually the left argument; the second element comes from bytes of memory immediately following those that hold the first element, and so on. Here is an example. If a program says

```
int *array = (int *) malloc (len * sizeof (int));
```

you can print the contents of array with

```
p *array@len
```

The left operand of '@' must reside in memory. Array values made with '@' in this way behave just like other arrays in terms of subscripting, and are coerced to pointers when used in expressions. (It would probably appear in an expression via the value history, after you had printed it out.)

D.9.4 Format options

GDB provides a few ways to control how arrays and structures are printed.

info format

Display the current settings for the format options.

set array-max number-of-elements

If GDB is printing a large array, it will stop printing after it has printed the number of elements set by the 'set array-max' command. This limit also applies to the display of strings.

set prettyprint on

Cause GDB to print structures in an indented format with one member per line, like this:

```
$1 = {
    next = 0x0,
    flags = {
        sweet = 1,
        sour = 1
    },
    meat = 0x54 "Pork"
}
```

```
set prettyprint off
```

Cause GDB to print structures in a compact format, like this:

```
$1 = {next = 0x0, flags = {sweet = 1, sour = 1}, meat = 0x54 "Pork"}
```

This is the default format.

```
set unionprint on
```

Tell GDB to print unions which are contained in structures. This is the default setting.

```
set unionprint off
```

Tell GDB not to print unions which are contained in structures.

For example, given the declarations

```
typedef enum {Tree, Bug} Species;
typedef enum {Big_tree, Acorn, Seedling} Tree_forms;
typedef enum {Caterpillar, Cocoon, Butterfly}
Bug_forms;
struct thing {
    Species it;
    union {
        Tree_forms tree;
        Bug_forms bug;
    } form;
};
struct thing foo = {Tree, {Acorn}};
```

with ‘set unionprint on’ in effect ‘p foo’ would print

```
$1 = {it = Tree, form = {tree = Acorn, bug = Cocoon}}
```

and with ‘set unionprint off’ in effect it would print

```
$1 = {it = Tree, form = {...}}
```

D.9.5 Output formats

GDB normally prints all values according to their data types. Sometimes this is not what you want. For example, you might want to print a number in hex, or a pointer in decimal. Or you might want to view data in memory at a certain address as a character string or an instruction. These things can be done with output formats.

The simplest use of output formats is to say how to print a value already computed. This is done by starting the arguments of the ‘print’ command with a slash and a format letter. The format letters supported are listed in Table D-16:

Table D-16. Print Output Format Letters

Format Letters	Description
‘x’	Regard the bits of the value as an integer, and print the integer in hexadecimal.
‘d’	Print as integer in signed decimal.
‘u’	Print as integer in unsigned decimal.
‘o’	Print as integer in octal.
‘a’	Print as an address, both absolute in hex and then relative to a symbol defined as an address below it.
‘c’	Regard as an integer and print it as a character constant.
‘f’	Regard the bits of the value as a floating point number and print using typical floating point syntax.

For example, to print the program counter in hex (see section Registers), type

```
p/x $pc
```

Note that no space is required before the slash; this is because command names in GDB cannot contain a slash.

To reprint the last value in the value history with a different format, you can use the ‘print’ command with just a format and no expression. For example, ‘p/x’ reprints the last value in hex.

D.9.5.1 Examining Memory

The command ‘x’ (for ‘examine’) can be used to examine memory without reference to the program’s data types. The format in which you wish to examine memory is instead explicitly specified. The allowable formats are a superset of the formats described in the previous section.

‘x’ is followed by a slash and an output format specification, followed by an expression for an address. The expression need not have a pointer value (though it may); it is used as an integer, as the address of a byte of memory. See section Expressions for more information on expressions. For example, ‘x/4xw \$sp’ prints the four words of memory above the stack pointer in hexadecimal.

The output format in this case specifies both how big a unit of memory to examine and how to print the contents of that unit. It is done with one or two of the following letters:

These letters listed in Table D-17 specify just the size of unit to examine:

Table D-17. Memory Output Format Letters

Format Letters	Description
‘b’	Examine individual bytes.
‘h’	Examine halfwords (two bytes each).
‘w’	Examine words (four bytes each).

Many assemblers and cpu designers still use ‘word’ for a 16-bit quantity, as a holdover from specific predecessor machines of the 1970’s that really did use two byte words. But more generally the term ‘word’ has always referred to the size of quantity that a machine normally operates on and stores in its registers. This is 32 bits for all the machines that GDB runs on.

‘g’ Examine giant words (8 bytes).

These letters, as listed in Table D-18, specify just the way to print the contents:

Table D-18. Print Output Format Letters

Format Letter	Description
‘x’	Print as integers in unsigned hexadecimal.
‘d’	Print as integers in signed decimal.
‘u’	Print as integers in unsigned decimal.
‘o’	Print as integers in unsigned octal.

Table D-18. Print Output Format Letters (Continued)

Format Letter	Description
<code>'a'</code>	Print as an address, both absolute in hex and then relative to a symbol defined as an address below it.
<code>'c'</code>	Print as character constants.
<code>'f'</code>	Print as floating point. This works only with sizes <code>'w'</code> and <code>'g'</code> .
<code>'s'</code>	Print a null-terminated string of characters. The specified unit size is ignored; instead, the unit is however many bytes it takes to reach a null character (including the null character).
<code>'i'</code>	Print a machine instruction in assembler syntax (or nearly). The specified unit size is ignored; the number of bytes in an instruction varies depending on the type of machine, the opcode and the addressing modes used.

If either the manner of printing or the size of unit fails to be specified, the default is to use the same one that was used last. If you don't want to use any letters after the slash, you can omit the slash as well.

You can also omit the address to examine. Then the address used is just after the last unit examined. This is why string and instruction formats actually compute a unit size based on the data: so that the next string or instruction examined will start in the right place. The `'print'` command sometimes sets the default address for the `'x'` command; when the value printed resides in memory, the default is set to examine the same location. `'info line'` also sets the default for `'x'`, to the address of the start of the machine code for the specified line and `'info breakpoints'` sets it to the address of the last breakpoint listed.

When you use RET to repeat an `'x'` command, it does not repeat exactly the same: the address specified previously (if any) is ignored, so that the repeated command examines the successive locations in memory rather than the same ones.

You can examine several consecutive units of memory with one command by writing a repeat-count after the slash (before the format letters, if any). The repeat count must be a decimal integer. It has the same effect as repeating the `'x'` command that many times except that the output may be more compact with several units per line. For example,

```
x/10i $pc
```

prints ten instructions starting with the one to be executed next in the selected frame. After doing this, you could print another ten following instructions with

```
x/10
```

in which the format and address are allowed to default.

The addresses and contents printed by the ‘x’ command are not put in the value history because there is often too much of them and they would get in the way. Instead, GDB makes these values available for subsequent use in expressions as values of the convenience variables `$_` and `$_`.

After an ‘x’ command, the last address examined is available for use in expressions in the convenience variable `$_`. The contents of that address, as examined, are available in the convenience variable `$_`.

If the ‘x’ command has a repeat count, the address and contents saved are from the last memory unit printed; this is not the same as the last address printed if several units were printed on the last line of output.

The specialized command ‘disassemble’ is also provided to dump a range of memory as machine instructions. The default memory range is the function surrounding the program counter of the selected frame. A single argument to this command is a program counter value; the function surrounding this value will be dumped. Two arguments specify a range of address (first inclusive, second exclusive) to be dumped.

D.9.6 Automatic Display

If you find that you want to print the value of an expression frequently (to see how it changes), you might want to add it to the automatic display list so that GDB will print its value each time the program stops. Each expression added to the list is given a number to identify it; to remove an expression from the list, you specify that number. The automatic display looks like this:

```
2: foo = 38
```

```
3: bar[5] = (struct hack *) 0x3804
```

showing item numbers, expressions and their current values.

If the expression refers to local variables, then it does not make sense outside the lexical context for which it was set up. Such an expression is printed only when execution is inside that lexical context. For example, if you give the command ‘display name’ while inside a function with an argument name, then this argument will be displayed whenever the program stops inside that function, but not when it stops elsewhere (since this argument doesn’t exist elsewhere). See Table D-19.

Table D-19. Display Commands

Command	Description
<code>display exp</code>	Add the expression <i>exp</i> to the list of expressions to display each time the program stops. See section Expressions.
<code>display/fmt exp</code>	For <i>fmt</i> specifying only a display format and not a size or count, add the expression <i>exp</i> to the auto-display list but arranges to display it each time in the specified format <i>fmt</i> .
<code>display/fmt addr</code>	For <i>fmt</i> 'i' or 's', or including a unit-size or a number of units, add the expression <i>addr</i> as a memory address to be examined each time the program stops. Examining means in effect doing 'x/fmt addr'. See section Memory.
<code>undisplay dnums...</code> <code>delete display dnums...</code>	Remove item numbers <i>dnums</i> from the list of expressions to display.
<code>disable display dnums...</code>	Disable the display of item numbers <i>dnums</i> . A disabled display item is not printed automatically, but is not forgotten. It may be re-enabled later.
<code>enable display dnums...</code>	Enable display of item numbers <i>dnums</i> . It becomes effective once again in auto display of its expression, until you specify otherwise.
<code>display</code>	Display the current values of the expressions on the list, just as is done when the program stops.
<code>info display</code>	Print the list of expressions previously set up to display automatically, each one with its item number, but without showing the values. This includes disabled expressions, which are marked as such. It also includes expressions which would not be displayed right now because they refer to automatic variables not currently available.

D.9.7 Value History

Every value printed by the 'print' command is saved for the entire session in GDB's value history so that you can refer to it in other expressions.

The values printed are given history numbers for you to refer to them by. These are successive integers starting with 1. 'print' shows you the history number assigned to a value by printing '\$num = ' before the value; here num is the history number.

To refer to any previous value, use '\$' followed by the value's history number. The output printed by 'print' is designed to remind you of this. Just \$ refers to the most recent value in the history, and \$\$ refers to the value before that.

For example, suppose you have just printed a pointer to a structure and want to see the contents of the structure. It suffices to type

```
p *$
```

If you have a chain of structures where the component 'next' points to the next one, you can print the contents of the next one with this:

```
p *$.next
```

It might be useful to repeat this command many times by typing RET.

Note: The history records values, not expressions.

If the value of `x` is 4 and you type this command:

```
print x
set x=5
```

then the value recorded in the value history by the ‘print’ command remains 4 even though the value of `x` has changed. See Table D-20.

Table D-20. Info Commands

Command	Description
<code>info values</code>	Print the last ten values in the value history, with their item numbers. This is like ‘p \$\$9’ repeated ten times, except that ‘info values’ does not change the history.
<code>info values n</code>	Print ten history values centered on history item number <code>n</code> .
<code>info values +</code>	Print ten history values just after the values last printed.

D.9.8 Convenience Variables

GDB provides *convenience variables* that you can use within GDB to hold on to a value and refer to it later. These variables exist entirely within GDB; they are not part of your program, and setting a convenience variable has no effect on further execution of your program. That’s why you can use them freely.

Convenience variables have names starting with ‘\$’. Any name starting with ‘\$’ can be used for a convenience variable, unless it is one of the predefined set of register names (see section Registers).

You can save a value in a convenience variable with an assignment expression, just as you would set a variable in your program. Example:

```
set $foo = *object_ptr
```

would save in `$foo` the value contained in the object pointed to by `object_ptr`.

Using a convenience variable for the first time creates it; but its value is void until you assign a new value. You can alter the value with another assignment at any time.

Convenience variables have no fixed types. You can assign a convenience variable any type of value, even if it already has a value of a different type. The convenience variable as an expression has whatever type its current value has.

info convenience

Print a list of convenience variables used so far, and their values. Abbreviated ‘i con’.

One of the ways to use a convenience variable is as a counter to be incremented or a pointer to be advanced. For example:

```
set $i = 0
print bar[$i++]>contents
...repeat that command by typing RET.
```

Some convenience variables are created automatically by GDB and given values likely to be useful. See Table D-21.

Table D-21. Convenience Variables

Convenience Variable	Description
\$_	The variable \$_ is automatically set by the ‘x’ command to the last address examined (see section Memory). Other commands which provide a default address for ‘x’ to examine also set \$_ to that address; these commands include ‘info line’ and ‘info breakpoint’.
\$__	The variable \$__ is automatically set by the ‘x’ command to the value found in the last address examined.

D.9.9 Registers

Machine register contents can be referred to in expressions as variables with names starting with ‘\$’. The names of registers are different for each machine; use ‘info registers’ to see the names used on your machine. The names \$pc and \$sp are used on all machines for the program counter register and the stack pointer. Often \$fp is used for a register that contains a pointer to the current stack frame, and \$ps is used for a register that contains the processor status. These standard register names may be available on your machine even though the info registers command displays them with a different name. For example, on the SPARC, info registers displays the processor status register as \$psr but you can also refer to it as \$ps.

GDB always considers the contents of an ordinary register as an integer when the register is examined in this way. Some machines have special registers which can hold nothing but floating point; these registers are considered floating point. There is no way to refer to the contents of an ordinary register as floating point value (although you can print it as a floating point value with ‘print/f \$regname’).

Some registers have distinct “raw” and “virtual” data formats. This means that the data format in which the register contents are saved by the operating system is not the same one that your program normally sees. For example, the registers of the 68881 floating point coprocessor are always saved in “extended” format, but all C programs expect to work

with “double” format. In such cases, GDB normally works with the virtual format only (the format that makes sense for your program), but the ‘info registers’ command prints the data in both formats.

Register values are relative to the selected stack frame (see section Selection). This means that you get the value that the register would contain if all stack frames farther in were exited and their saved registers restored. In order to see the real contents of all registers, you must select the innermost frame (with ‘frame 0’).

Some registers are never saved (typically those numbered zero or one) because they are used for returning function values; for these registers, relativization makes no difference.

```
info registers
```

Print the names and relativized values of all registers.

```
info registers regname
```

Print the relativized value of register *regname*. *regname* may be any register name valid on the machine you are using, with or without the initial ‘\$’.

D.9.9.1 Examples

You could print the program counter in hex with

```
p/x $pc
```

or print the instruction to be executed next with

```
x/i $pc
```

or add four to the stack pointer with

```
set $sp += 4
```

The last is a way of removing one word from the stack, on machines where stacks grow downward in memory (most machines, nowadays). This assumes that the innermost stack frame is selected. Setting \$sp is not allowed when other stack frames are selected.

D.10 Examining the Symbol Table

The commands described in this section allow you to make inquiries for information about the symbols (names of variables, functions and types) defined in your program. This information is found by GDB in the symbol table loaded by the ‘symbol-file’ command; it is inherent in the text of your program and does not change as the program executes. See Table D-22.

Table D-22. Symbol-File Commands

Command	Description
<code>whatis exp</code>	Print the data type of expression <code>exp</code> . <code>exp</code> is not actually evaluated, and any side-effecting operations (such as assignments or function calls) inside it do not take place. See section Expressions.
<code>whatis</code>	Print the data type of <code>\$</code> , the last value in the value history.
<code>info address symbol</code>	Describe where the data for <code>symbol</code> is stored. For a register variable, this says which register it is kept in. For a non-register local variable, this prints the stack-frame offset at which the variable is always stored. Note the contrast with <code>'print &symbol'</code> , which does not work at all for a register variables, and for a stack local variable prints the exact address of the current instantiation of the variable.
<code>ptype typename</code>	Print a description of data type <code>typename</code> . <code>typename</code> may be the name of a type, or for C code it may have the form <code>'struct struct-tag'</code> , <code>'union union-tag'</code> or <code>'enum enum-tag'</code> .
<code>info sources</code>	Print the names of all source files in the program for which there is debugging information.
<code>info functions</code>	Print the names and data types of all defined functions.
<code>info functions regexp</code>	Print the names and data types of all defined functions whose names contain a match for regular expression <code>regexp</code> . Thus, <code>'info fun step'</code> finds all functions whose names include <code>'step'</code> ; <code>'info fun ^step'</code> finds those whose names start with <code>'step'</code> .
<code>info variables</code>	Print the names and data types of all variables that are declared outside of functions (i.e., except for local variables).
<code>info variables regexp</code>	Print the names and data types of all variables (except for local variables) whose names contain a match for regular expression <code>regexp</code> .
<code>info types</code>	Print all data types that are defined in the program.
<code>info types regexp</code>	Print all data types that are defined in the program whose names contain a match for regular expression <code>regexp</code> .
<code>printsyms filename</code>	Write a complete dump of the debugger's symbol data into the file <code>filename</code> .

D.11 Altering Execution

Once you think you have find an error in the program, you might want to find out for certain whether correcting the apparent error would lead to correct results in the rest of the run. You can find the answer by experiment, using the GDB features for altering execution of the program.

For example, you can store new values into variables or memory locations, give the program a signal, restart it at a different address, or even return prematurely from a function to its caller.

D.11.1 Assignment to Variables

To alter the value of a variable, evaluate an assignment expression. See section Expressions. For example,

```
print x=4
```

would store the value 4 into the variable x, and then print the value of the assignment expression (which is 4).

All the assignment operators of C are supported, including the incrementation operators ‘++’ and ‘--’, and combining assignments such as ‘+=’ and ‘<<=’.

If you are not interested in seeing the value of the assignment, use the ‘set’ command instead of the ‘print’ command. ‘set’ is really the same as ‘print’ except that the expression’s value is not printed and is not put in the value history (see section Value History). The expression is evaluated only for side effects.

Note that if the beginning of the argument string of the ‘set’ command appears identical to a ‘set’ sub-command, it may be necessary to use the ‘set variable’ command. This command is identical to ‘set’ except for its lack of sub-commands.

GDB allows more implicit conversions in assignments than C does; you can freely store an integer value into a pointer variable or vice versa, and any structure can be converted to any other structure that is the same length or shorter.

To store values into arbitrary places in memory, use the ‘{...}’ construct to generate a value of specified type at a specified address (see section Expressions). For example, {int}0x83040 would refer to memory location 0x83040 as an integer (which implies a certain size and representation in memory), and

```
set {int}0x83040 = 4
```

would store the value 4 into that memory location.

D.11.2 Continuing at a Different Address

Ordinarily, when you continue the program, you do so at the place where it stopped, with the ‘cont’ command. You can instead continue at an address of your own choosing, with the following commands:

jump *linenum*

Resume execution at line number *linenum*. Execution may stop immediately if there is a breakpoint there.

The ‘jump’ command does not change the current stack frame, or the stack pointer, or the contents of any memory location or any register other than the program counter. If line *linenum* is in a different function from the one currently executing, the results may be bizarre if the two functions expect different patterns of arguments or of local variables. For this reason, the ‘jump’ command requests confirmation if the specified line is not in the function currently executing. However, even bizarre results are predictable based on careful study of the machine-language code of the program.

```
jump *address
```

Resume execution at the instruction at address *address*.

You can get much the same effect as the jump command by storing a new value into the register \$pc. The difference is that this does not start the program running; it only changes the address where it will run when it is continued. For example,

```
set $pc = 0x485
```

causes the next ‘cont’ command or stepping command to execute at address 0x485, rather than at the address where the program stopped. See section Stepping.

The most common occasion to use the ‘jump’ command is when you have stepped across a function call with next, and found that the return value is incorrect. If all the relevant data appeared correct before the function call, the error is probably in the function that just returned.

In general, your next step would now be to rerun the program and execute up to this function call, and then step into it to see where it goes astray. But this may be time consuming. If the function did not have significant side effects, you could get the same information by resuming execution just before the function call and stepping through it. To do this, first put a breakpoint on that function; then, use the ‘jump’ command to continue on the line with the function call.

D.11.3 Giving the Program a Signal

```
signal signalnum
```

Resume execution where the program stopped, but give it immediately the signal number *signalnum*.

Alternatively, if *signalnum* is zero, continue execution without giving a signal. This is useful when the program stopped on account of a signal and would ordinarily see the signal when resumed with the ‘cont’ command; ‘signal 0’ causes it to resume without a signal.

D.11.4 Returning from a Function

You can cancel execution of a function call with the ‘return’ command. This command has the effect of discarding the selected stack frame (and all frames within it), so that control moves to the caller of that function. You can think of this as making the discarded frame return prematurely.

First select the stack frame that you wish to return from (see section Selection). Then type the ‘return’ command. If you wish to specify the value to be returned, give that as an argument.

This pops the selected stack frame (and any other frames inside of it), leaving its caller as the innermost remaining frame. That frame becomes selected. The specified value is stored in the registers used for returning values of functions.

The ‘return’ command does not resume execution; it leaves the program stopped in the state that would exist if the function had just returned. Contrast this with the ‘finish’ command (see section Stepping), which resumes execution until the selected stack frame returns naturally.

D.12 Canned Sequences of Commands

GDB provides two ways to store sequences of commands for execution as a unit: user-defined commands and command files.

D.13 User-Defined Commands

A *user-defined* command is a sequence of GDB commands to which you assign a new name as a command. This is done with the ‘define’ command. See Table D-23.

Table D-23. Define Commands

Command	Description
<code>define <i>commandname</i></code>	Define a command named <i>commandname</i> . If there is already a command by that name, you are asked to confirm that you want to redefine it. The definition of the command is made up of other GDB command lines, which are given following the ‘define’ command. The end of these commands is marked by a line containing ‘end’.
<code>document <i>commandname</i></code>	Give documentation to the user-defined command <i>commandname</i> . The command <i>commandname</i> must already be defined. This command reads lines of documentation just as ‘define’ reads the lines of the command definition, ending with ‘end’. After the ‘document’ command is finished, ‘help’ on command <i>commandname</i> will print the documentation you have specified.

You may use the ‘document’ command again to change the documentation of a command. Redefining the command with ‘define’ does not change the documentation.

User-defined commands do not take arguments. When they are executed, the commands of the definition are not printed. An error in any command stops execution of the user-defined command.

Commands that would ask for confirmation if used interactively proceed without asking when used inside a user-defined command. Many GDB commands that normally print messages to say what they are doing omit the messages when used in user-defined command.

D.13.1 Command Files

A command file for GDB is a file of lines that are GDB commands. Comments (lines starting with '#') may also be included. An empty line in a command file does nothing; it does not mean to repeat the last command, as it would from the terminal.

When GDB starts, it automatically executes its init files, command files named '.gdbinit'. GDB reads the init file (if any) in your home directory and then the init file (if any) in the current working directory. (The init files are not executed if the '-nx' option is given.) You can also request the execution of a command file with the 'source' command:

```
source filename
```

Execute the command file *filename*.

The lines in a command file are executed sequentially. They are not printed as they are executed. An error in any command terminates execution of the command file.

Commands that would ask for confirmation if used interactively proceed without asking when used in a command file. Many GDB commands that normally print messages to say what they are doing omit the messages when used in a command file.

D.13.2 Commands for Controlled Output

During the execution of a command file or a user defined command, the only output that appears is what is explicitly printed by the commands of the definition. This section describes three commands useful for generating exactly the output you want as listed in Table D-24.

Table D-24. Controlled Output Commands

Command	Description
<code>echo text</code>	<p>Print <i>text</i>. Non-printing characters can be included in text using C escape sequences, such as ‘\n’ to print a newline. No newline will be printed unless you specify one. In addition to the standard C escape sequences a backslash followed by a space stands for a space. This is useful for outputting a string with spaces at the beginning or the end, since leading and trailing spaces are trimmed from all arguments. Thus, to print “and foo = ”, use the command “echo \ and foo = \ ”. A backslash at the end of text can be used, as in C, to continue the command onto subsequent lines.</p> <pre>echo This is some text\n\ which is continued\n\ onto several lines.\n produces the same output as echo This is some text\n echo which is continued\n echo onto several lines.\n</pre>
<code>output expression</code>	<p>Print the value of expression and nothing but that value: no newlines, no ‘\$nn = ’. The value is not entered in the value history either. See section Expressions for more information on expressions.</p>
<code>output/fmt expression</code>	<p>Print the value of expression in format <i>fmt</i>. See section Output formats, for more information.</p>
<code>printf string, expressions...</code>	<p>Print the values of the expressions under the control of string. The expressions are separated by commas and may be either numbers or pointers. Their values are printed as specified by string, exactly as if the program were to execute</p> <pre>printf (string, expressions...);</pre> <p>For example, you can print two values in hex like this:</p> <pre>printf "foo, bar-foo = 0x%x, 0x%x\n", foo, bar-foo</pre> <p>only backslash-escape sequences that you can use in the string are the simple ones that consist of backslash followed by a letter.</p>

D.14 Options and Arguments for GDB

When you invoke GDB, you can specify arguments telling it what files to operate on and what other things to do.

D.14.1 Mode Options

Table D-25 lists possible Mode Options:

Table D-25. Mode Options

Mode Option	Description
<code>'-nx'</code>	Do not execute commands from the init files <code>'gdbinit'</code> . Normally, the commands in these files are executed after all the command options and arguments have been processed. See section Command Files.
<code>'-q'</code>	“Quiet”. Do not print the usual introductory messages.
<code>'-batch'</code>	Run in batch mode. Exit with code 0 after processing all the command files specified with <code>'-x'</code> (and <code>'gdbinit'</code> , if not inhibited). Exit with nonzero status if an error occurs in executing the GDB commands in the command files.
<code>'-fullname'</code>	This option is used when Emacs runs GDB as a subprocess. It tells GDB to output the full file name and line number in a standard, recognizable fashion each time a stack frame is displayed (which includes each time the program stops). This recognizable format looks like two <code>'\032'</code> characters, followed by the file name, line number and character position separated by colons, and a newline. The Emacs-to-GDB interface program uses the two <code>'\032'</code> characters as a signal to display the source code for the frame.

D.14.2 File-specifying Options

All the options and command line arguments shown in Table D-26 are processed in sequential order. The order makes a difference when the `'-x'` option is used.

Table D-26. File Specifying Options

Options	Description
<code>'-s file'</code>	Read symbol table from file <i>file</i> .
<code>'-e file'</code>	Use file <i>file</i> as the executable file to execute when appropriate, and for examining pure data in conjunction with a core dump.
<code>'-se file'</code>	Read symbol table from file <i>file</i> and use it as the executable file.
<code>'-c file'</code>	Use file <i>file</i> as a core dump to examine.
<code>'-x file'</code>	Execute GDB commands from file <i>file</i> .
<code>'-d directory'</code>	Add directory to the path to search for source files.

D.14.3 Other Arguments

If there are arguments to GDB that are not options or associated with options, the first one specifies the symbol table and executable file name (as if it were preceded by `'-se'`) and the second one specifies a core dump file name (as if it were preceded by `'-c'`).

D.15 Using GDB under GNU Emacs

A special interface allows you to use GNU Emacs to view (and edit) the source files for the program you are debugging with GDB.

To use this interface, use the command `M-x gdb` in Emacs. Give the executable file you want to debug as an argument. This command starts GDB as a subprocess of Emacs, with input and output through a newly created Emacs buffer.

Using GDB under Emacs is just like using GDB normally except for two things:

- All “terminal” input and output goes through the Emacs buffer. This applies both to GDB commands and their output, and to the input and output done by the program you are debugging.

This is useful because it means that you can copy the text of previous commands and input them again; you can even use parts of the output in this way.

All the facilities of Emacs’s Shell mode are available for this purpose.

- GDB displays source code through Emacs. Each time GDB displays a stack frame, Emacs automatically finds the source file for that frame and puts an arrow (`'=>'`) at the left margin of the current line.

Explicit GDB ‘list’ or search commands still produce output as usual, but you probably will have no reason to use them.

In the GDB I/O buffer, you can use these special Emacs commands as listed in Table D-27:

Table D-27. Emac Commands

Commands	Description
<code>M-s</code>	Execute to another source line, like the GDB ‘step’ command.
<code>M-n</code>	Execute to next source line in this function, skipping all function calls, like the GDB ‘next’ command.
<code>M-i</code>	Execute one instruction, like the GDB ‘stepi’ command.
<code>C-c C-f</code>	Execute until exit from the selected stack frame, like the GDB ‘finish’ command.
<code>M-c</code>	Continue execution of the program, like the GDB ‘cont’ command.
<code>M-u</code>	Go up the number of frames indicated by the numeric argument (see section Arguments, Numeric Arguments, emacs, The GNU Emacs Manual), like the GDB ‘up’ command.
<code>M-d</code>	Go down the number of frames indicated by the numeric argument, like the GDB ‘down’ command.

In any source file, the Emacs command `C-x SPC` (`gdb break`) tells GDB to set a breakpoint on the source line point is on.

The source files displayed in Emacs are in ordinary Emacs buffers which are visiting the source files in the usual way. You can edit the files with these buffers if you wish; but keep in mind that GDB communicates with Emacs in terms of line numbers. If you add or delete lines from the text, the line numbers that GDB knows will cease to correspond properly to the code.

Appendix E

Additional Support

User support from the conception of a design through completion is available from Motorola and third-party companies as shown in Table E-1:

Table E-1. User Support

Support	Motorola	Third Party
Design	<ul style="list-style-type: none"> - Data Sheets - Application Notes - Application Bulletins - Software Examples 	<ul style="list-style-type: none"> - Data Acquisition Packages - Filter Design Packages - Operating System Software - Simulator
Prototyping	<ul style="list-style-type: none"> - Assembler - Linker - C Compiler - Simulator - Application Development System (ADS) - In-Circuit EmulatorDebug - Cable for ADS 	<ul style="list-style-type: none"> - Logic Analyzer with DSP56000/DSP56001 ROM Packages - Data Acquisition Cards - DSP Development System Cards - Operating System Software - In-Circuit Emulator - Debug Software
Design Verification	<ul style="list-style-type: none"> - Application Development System (ADS) - In-Circuit Emulator - Simulator 	<ul style="list-style-type: none"> - Data Acquisition Packages - Logic Analyzer with DSP56000/DSP56001 ROM Packages - Data Acquisition Cards - DSP Development System Cards - Application-Specific Development Tools - Debug Software

The following is a partial list of the support available for the DSP56000/DSP56001. Additional information can be obtained through Dr. BuB or the appropriate support telephone service.

E.1 Motorola DSP Product Support

- DSP56000CLASx Design-In Software Package which includes:
 - Relocatable Macro Assembler
 - Linker
 - Simulator (simulates single or multiple DSP56000/DSP56001s)
 - Librarian
- 56KCCx Full ANSI Compliant C Compiler
- DSP320to56001 Translator Software
- DSP56000/DSP56001 Applications Development System (ADS)
- Support Integrated Circuits
- DSP Bulletin Board (Dr. BuB)
- Motorola DSP Newsletter
- Motorola Technical Service Engineers (TSEs)
See your local telephone directory for the Motorola Semiconductor Sector sales office telephone number.
- Design Hotline
- Applications Assistance
- Marketing Information
- Third-Party Support Information
- University Support Information

E.2 DSP56000CLASx Assembler/Simulator

The macro cross assembler and simulator run on:

1. IBM™ PCs and clones using an 80386 or upward compatible processor
2. Macintosh™ computers with a NU-BUS™ expansion port
3. SUN computer
4. NeXT™ computer

E.2.1 Macro Cross Assembler Features:

- Production of relocatable object modules compatible with linker program when in relocatable mode
- Production of absolute files compatible with simulator program when in absolute mode
- Supports full instruction set, memory spaces, and parallel data transfer fields of the DSP56000/DSP56001
- Modular programming features: local labels, sections, and external definition/reference directives
- Nested macro processing capability with support for macro libraries
- Complex expression evaluation including boolean operators
- Built-in functions for data conversion, string comparison, and common transcendental math functions
- Directives to define circular and bit-reversed buffers
- Extensive error checking and reporting

E.2.2 Simulator Features:

- Simulation of all DSP56000 family DSPs
- Simulation of multiple DSP56000 family DSPs
- Linkable object code modules:
 - -Nondisplay simulator library
 - -Display simulator library
- C language source code for:
 - -Screen management functions
 - -Terminal I/O functions
 - -Simulation examples
- Single stepping through object programs
- Conditional or unconditional breakpoints
- Program patching using a single-line assembler/disassembler
- Instruction, clock cycle, and histogram counters
- Session and/or command logging for later reference
- ASCII input/output files for peripherals
- Help-line display and expanded on-line help for simulator commands

-
- Loading and saving of files to/from simulator memory
 - Macro command definition and execution
 - Display enable/disable of registers and memory
 - Hexadecimal/decimal/binary calculator

E.3 DSP320to56001 Translator

E.3.1 DSP320to56001 Translator Features:

- Translates any TMS32010 linked object code to DSP56001 source assembler code
- Two modes of operation:
 - Translates to DSP56001 source assembler code for optimization and assembly using DSP56000CLASx
 - Translates and runs “as is” directly and immediately on the DSP56000ADSx
- C language DSP320to56001 source code is provided in addition to IBM PC/XT/AT object code to allow:
 - User modification for TMS32020 or TMS320C25 translation
 - User compilation to accommodate different host platforms

E.4 DSP56000ADSx Application Development Systems (ADS)

E.4.1 DSP56000ADS ADS Hardware Features:

- Full-speed 27 MHz operation
- Multiple application development module (ADM) support with programmable ADM addresses
- 8Kx24 (expandable to 32Kx24)user-configurable RAM for DSP56000/DSP56001 code development
- 1Kx24 monitor ROM expandable to 4Kx24
- 96-pin Euro-card connector making all DSP56001 pins accessible
- In-circuit emulation capabilities when used with the DSP56KEMULTRCABL cable
- Separate berg pin connectors for alternate accessing of serial or host/DMA ports
- ADM can be used in stand-alone configuration
- No external power supply needed when connected to a host platform

E.4.2 DSP56000ADSx ADS Software Features:

- Single/multiple stepping through DSP56000/DSP56001 object programs
- Up to 99 conditional or unconditional breakpoints
- Program patching using a single-line assembler/disassembler
- Session and/or command logging for later reference
- Loading and saving files to/from ADM memory
- Macro command definition and execution
- Display enable/disable of registers and memory
- Debug commands supporting multiple ADMs
- Hexadecimal/decimal/binary calculator
- Host operating system commands from within ADS user interface program
- Multiple OS I/O file access from DSP56000/DSP56001 object programs
- Fully compatible with the DSP56000CLASx design-in software package
- On-line help screens for each command and DSP56000/DSP56001 register

E.4.2.1 Support Integrated Circuits:

- 8Kx24 Static RAM
- DSP56ADC16 16-bit, 100-kHz analog-to-digital converter
- DSP56401 AES/EBU processor
- DSP56200 FIR filter

E.5 Dr. BuB Electronic Bulletin Board

Dr. BuB is an electronic bulletin board providing free source code for a large variety of topics that can be used to develop applications with Motorola DSP products. The software library includes files including FFTs, FIR filters, IIR filters, lattice filters, matrix algebra routines, companding routines, floating-point routines, a software debug monitor, and others. In addition, the latest product information and documentation (including information on new products and improvements on existing products) is posted. Questions concerning Motorola DSP products posted on Dr. BuB are answered promptly. Access to Dr. BuB is through the following phone numbers:

(212A – 300/1200 Baud) (512) 891-DSP1
 (V.22 – 1200 Baud)..... (512) 891-DSP2
 (V.22bis – 2400 Baud)..... (512) 891-DSP3
 Format: 7 data bits, even parity, 1 stop bit
 User ID=guest

Table E-2 is a partial list of the software available on Dr. BuB.

Table E-2. Available Software

Document ID	Version	Synopsis	Size
Audio:			
rvb1.asm	1.0	Easy-to-read reverberation routine	17056
rvb2.asm	1.0	Same as RVB1.ASM but optimized	15442
stereo.asm	1.0	Code for C-QUAM AM stereo decoder	4830
stereo.hlp	1.0	Help file for STEREO.ASM	620
dge.asm	1.0	Digital Graphic Equalizer code from	14880
Codec Routines:			
loglin.asm	1.0	Companded CODEC to linear PCM data conversion	4572
loglin.hlp		Help for loglin.asm	1479
loglint.asm	1.0	Test program for loglin.asm	2184
loglint.hlp		Help for loglint.asm	1993
linlog.asm	1.1	Linear PCM to companded CODEC data conversion	4847
linlog.hlp		Help for linlog.asm	1714

Table E-2. Available Software

Document ID	Version	Synopsis	Size
DTMF Routines:			
clear.cmd	1.0	Explained in read.me file	119
data.lod	1.0		421
det.asm	1.0	Subroutine used in IIR DTMF	5923
dtmf.asm	1.0	Main routine used in IIR DTMF	10685
dtmf.mem	1.0	Memory for DTMF routine	48
dtmfmstr.asm	1.0	Main routine for multichannel DTMF	7409
dtmfmstr.mem	1.0	Memory for multichannel DTMF routine	41
dtmftwo.asm	1.0		10256
ex56.bat	1.0		94
genxd.lod	1.0	Data file	183
genyd.lod	1.0	Data file	180
goertzel.asm	1.0	Goertzel routine	4393
goertzel.lnk	1.0	Link file for Goertzel routine	6954
goertzel.lst	1.0	List file for Goertzel routine	11600
load.cmd	1.0		46
tstgoert.mem	1.0	Memory for Goertzel routine	384
sub.asm	1.0	Subroutine linked for use in IIR DTMF	2491
read.me	1.0	Instructions	738
Fast Fourier Transforms:			
sincos.asm	1.2	Sine-Cosine Table Generator for FFTs	1185
sincos.hlp		Help for sincos.asm	887
sinewave.asm	1.1	Full-Cycle Sine wave Table Generator Macro	1029
sinewave.hlp		for sinewave.asm	1395
fftr2a.asm	1.1	Radix 2, In-Place, DIT FFT (smallest)	3386
fftr2a.hlp		Help for fftr2a.asm	2693
fftr2at.asm	1.1	Test Program for FFTs (fftr2a.asm)	999
fftr2at.hlp		Help for fftr2at.asm	563

Table E-2. Available Software

Document ID	Version	Synopsis	Size
fftr2b.asm	1.1	Radix 2, In-Place, DIT FFT (faster)	4290
fftr2b.hlp		Help for fftr2b.asm	3680
fftr2c.asm	1.2	Radix 2, In-Place, DIT FFT (even faster)	5991
fftr2c.hlp		Help for fftr2c.asm	3231
fftr2d.asm	1.0	Radix 2, In-Place, DIT FFT (using DSP56001 sine-cosine ROM tables)	3727
fftr2d.hlp		Help for fftr2d.asm	3457
fftr2dt.asm	1.0	Test program for fftr2d.asm	1287
fftr2dt.hlp		Help for fftr2dt.asm	614
fftr2e.asm	1.0	1024 Point, Non-In-Place, FFT (3.39ms)	8976
fftr2e.hlp		Help for fftr2e.asm	5011
fftr2et.asm	1.0	Test program for fftr2e.asm	984
fftr2et.hlp		Help for fftr2et.asm	408
dct1.asm	1.1	Discrete Cosine Transform using FFT	5493
dct1.hlp	1.1	Help file for dct1.asm	970
fftr2cc.asm	1.0	Radix 2, In-place Decimation-in-time complex FFT macro	6524
fftr2cc.hlp	1.0	Help file for fftr2cc.asm	3533
fftr2cn.asm	1.0	Radix 2, Decimation-in-time Complex FFT macro with normally ordered input/output	6584
fftr2cn.hlp	1.0	Help file for fftr2cn.asm	2468
fftr2en.asm	1.0	1024 point, not-in-place, complex FFT macro with normally ordered input/output	9723
fftr2en.hlp	1.0	Help file for fftr2en.asm	4886
dhit1.asm	1.0	Routine to compute Hilbert transform in the frequency domain	1851
dhit1.hlp	1.0	Help file for dhit1.asm	1007
fftr2bf.asm	1.0	Radix-2, decimation-in-time FFT with	13526
		block floating point	
fftr2bf.hlp	1.0	Help file for fftr2bf.asm	1578
fftr2aa.asm	1.0	FFT program for automatic scaling	3172

Table E-2. Available Software

Document ID	Version	Synopsis	Size
Filters:			
fir.asm	1.0	Direct Form FIR Filter	545
fir.hlp		Help for fir.asm	2161
firt.asm	1.0	Test program for fir.asm	1164
iir1.asm	1.0	Direct Form Second Order All Pole	656
		IIR Filter	
iir1.hlp		Help for iir1.asm	1786
iir1t.asm	1.0	Test program for iir1.asm	1157
iir2.asm	1.0	Direct Form Second Order All Pole IIR Filter with Scaling	801
iir2.hlp		Help for iir2.asm	2286
iir2t.asm	1.0	Test program for iir2.asm	1311
iir3.asm	1.0	Direct Form Arbitrary Order All Pole IIR Filter	776
iir3.hlp		Help for iir3.asm	2605
iir3t.asm	1.0	Test program for iir3.asm	1309
iir4.asm	1.0	Second Order Direct Canonic IIR Filter (Biquad IIR Filter)	713
iir4.hlp		Help for iir4.asm	2255
iir4t.asm	1.0	Test program for iir4.asm	1202
iir5.asm	1.0	Second Order Direct Canonic IIR Filter with Scaling (Biquad IIR Filter)	842
iir5.hlp		Help for iir5.asm	2803
iir5t.asm	1.0	Test program for iir5.asm	1289
iir6.asm	1.0	Arbitrary Order Direct Canonic IIR Filter	923
iir6.hlp		Help for iir6.asm	3020
iir6t.asm	1.0	Test program for iir6.asm	1377
iir7.asm	1.0	Cascaded Biquad IIR Filters	900
iir7.hlp		Help for iir7.asm	3947
iir7t.asm	1.0	Test program for iir7.asm	1432
lms.hlp	1.0	LMS Adaptive Filter Algorithm	5818
transiir.asm	1.0	Implements the transposed IIR filter	1981

Table E-2. Available Software

Document ID	Version	Synopsis	Size
transiir.hlp	1.0	Help file for transiir.asm	974
Floating-Point Routines:			
fpdef.hlp	2.0	Storage format and arithmetic representation definition	10600
fpcalls.hlp	2.1	Subroutine calling conventions	11876
fplist.asm	2.0	Test file that lists all subroutines	1601
fprevs.hlp	2.0	Latest revisions of floating-point lib	1799
fpinit.asm	2.0	Library initialization subroutine	2329
fpadd.asm	2.0	Floating point add	3860
fpsub.asm	2.1	Floating point subtract	3072
fpcmp.asm	2.1	Floating point compare	2605
fpmpy.asm	2.0	Floating point multiply	2250
fpmac.asm	2.1	Floating point multiply-accumulate	2712
fpdiv.asm	2.0	Floating point divide	3835
fpsqrt.asm	2.0	Floating point square root	2873
fpneg.asm	2.0	Floating point negate	2026
fpabs.asm	2.0	Floating point absolute value	1953
fp scale.asm	2.0	Floating point scaling	2127
fpfix.asm	2.0	Floating to fixed point conversion	3953
fpfloat.asm	2.0	Fixed to floating point conversion	2053
fpceil.asm	2.0	Floating point CEIL subroutine	1771
fpfloor.asm	2.0	Floating point FLOOR subroutine	2119
durbin.asm	1.0	Solution for LPC coefficients	5615
durbin.hlp	1.0	Help file for DURBIN.ASM	2904
fpfrac.asm	2.0	Floating point FRACTION subroutine	1862
Functions:			
log2.asm	1.0	Log base 2 by polynomial approximation	1118
log2.hlp		Help for log2.asm	719
log2t.asm	1.0	Test program for log2.asm	1018

Table E-2. Available Software

Document ID	Version	Synopsis	Size
log2nrm.asm	1.0	Normalizing base 2 logarithm macro	2262
log2nrm.hlp		Help for log2nrm.asm	676
log2nrmt.asm	1.0	Test program for log2nrm.asm	1084
exp2.asm	1.0	Exponential base 2 by polynomial approximation	926
exp2.hlp		Help for exp2.asm	759
exp2t.asm	1.0	Test program for exp2.asm	1019
sqrt1.asm	1.0	Square Root by polynomial approximation, 7 bit accuracy	991
sqrt1.hlp		Help for sqrt1.asm	779
sqrt1t.asm	1.0	Test program for sqrt1.asm	1065
sqrt2.asm	1.0	Square Root by polynomial approximation, 10 bit accuracy	899
sqrt2.hlp		Help for sqrt2.asm	776
sqrt2t.asm	1.0	Test program for sqrt2.asm	1031
sqrt3.asm	1.0	Full precision Square Root Macro	1388
sqrt3.hlp		Help for sqrt3.asm	794
sqrt3t.asm	1.0	Test program for sqrt3.asm	1053
tli.asm	1.1	Linear table lookup/interpolation routine for function generation	3253
tli.hlp	1.1	Help for tli.asm	1510
bingray.asm	1.0	Binary to Gray code conversion macro	601
bingrayt.asm	1.0	Test program for bingray.asm	991
rand1.asm	1.1	Pseudo Random Sequence Generator	2446
rand1.hlp		Help for rand1.asm	704
Lattice Filters:			
latfir1.asm	1.0	Lattice FIR Filter Macro	1156
latfir1.hlp		Help for latfir1.asm	6327
latfir1t.asm	1.0	Test program for latfir1.asm	1424
latfir2.asm	1.0	Lattice FIR Filter Macro (modified modulo count)	1174
latfir2.hlp		Help for latfir2.asm	1295
latfir2t.asm	1.0	Test program for latfir2.asm	1423

Table E-2. Available Software

Document ID	Version	Synopsis	Size
latiir.asm	1.0	Lattice IIR Filter Macro	1257
latiir.hlp		Help for latiir.asm	6402
latiirt.asm	1.0	Test program for latiir.asm	1407
latgen.asm	1.0	Generalized Lattice FIR/IIR Filter Macro	1334
latgen.hlp		Help for latgen.asm	5485
latgent.asm	1.0	Test program for latgen.asm	1269
latnrm.asm	1.0	Normalized Lattice IIR Filter Macro	1407
latnrm.hlp		Help for latnrm.asm	7475
latnrmt.asm	1.0	Test program for latnrm.asm	1595
Matrix Operations:			
matmul1.asm	1.0	[1x3][3x3]=[1x3] Matrix Multiplication	1817
matmul1.hlp		Help for matmul1.asm	527
matmul2.asm	1.0	General Matrix Multiplication, C=AB	2650
matmul2.hlp		Help for matmul2.asm	780
matmul3.asm	1.0	General Matrix Multiply-Accumulate, C=AB+Q	2815
matmul3.hlp	1.0	Help for matmul3.asm	865
Reed-Solomon Encoder:			
readme.rs	1.0	Instructions for Reed-Solomon coding	5200
rscd.asm	1.0	Reed-Solomon coder for DSP56000 simulator	5822
newc.c	1.0	Reed-Solomon coder coded in C	4075
table1.asm	1.0	Include file for R-S coder	7971
table2.asm	1.0	Include file for R-S coder	4011
Sorting Routines:			
sort1.asm	1.0	Array Sort by Straight Selection	1312
sort1.hlp		Help for sort1.asm	1908
sort1t.asm	1.0	Test program for sort1.asm	689
sort2.asm	1.1	Array Sort by Heapsort Method	2183
sort2.hlp		Help for sort2.asm	2004

Table E-2. Available Software

Document ID	Version	Synopsis	Size
sort2t.asm	1.0	Test program for sort2.asm	700
Speech:			
lgsol1.asm	2.0	Leroux-Gueguen solution for PARCOR (LPC) coefficients	4861
lgsol1.hlp		Help for lgsol1.asm	3971
durbin1.asm	1.2	Durbin Solution for PARCOR (LPC) coefficients	6360
durbin1.hlp		Help for durbin1.asm	3616
adpcm.asm	1.0	32 kbits/s CCITT ADPCM Speech Coder	
adpcm.hlp	1.0	Help file for adpcm.asm	14817
adpcmns.asm	1.0	Nonstandard ADPCM source code	54733
adpcmns.hlp	1.0	Help file for adpcmns.asm	9952
Standard I/O Equates:			
ioequ.asm	1.1	Motorola Standard I/O Equate File	8774
ioequlc.asm	1.1	Lower Case Version of ioequ.asm	8788
intequ.asm	1.0	Standard Interrupt Equate File	1082
intequlc.asm	1.0	Lower Case Version of intequ.asm	1082
Tools and Utilities:			
srec.c	4.10	Utility to convert DSP56000 OMF format to SREC.	38975
srec.doc	4.10	Manual page for srec.c.	7951
srec.h	4.10	Include file for srec.c	3472
srec.exe	4.10	Srec executable for IBM PC	22065
sloader.asm	1.1	Serial loader from the SCI port for the DSP56001	3986
sloader.hlp	1.1	Help for sloader.asm	2598
sloader.p	1.1	Serial loader s-record file for download to EPROM	736
parity.asm	1.0	Parity calculation of a 24-bit number in accumulator A	1641
parity.hlp	1.0	Help for parity.asm	936
parityt.asm	1.0	Test program for parity.asm	685
parityt.hlp	1.0	Help for parityt.asm	259

Table E-2. Available Software

Document ID	Version	Synopsis	Size
dspsbug		Ordering information for free debug monitor for DSP56000/DSP56001	882
The following is a list of current DSP56200 related software:			
p1	1.0	Information on 56200 Filter Software	6343
p2	1.0	Interrupt Driven Adaptive Filter Flowchart.	10916
p3	1.0	“C” code implementation of p2	25795
p4	1.0	Polled I/O Adaptive Filter Flowchart	10361
p5	1.0	“C” code implementation of p4	24806
p6	1.1	Interrupt Driven Dual FIR Filter Flowchart.	9535
p7	1.0	“C” code implementation of p6	28489
p8	1.0	Polled I/O Dual FIR Filter Flowchart	9656
p9	1.0	“C” code implementation of p8	28525

E.5.1 Motorola DSP News

The Motorola DSP News is a quarterly newsletter providing information on new products, application briefs, questions and answers, DSP product information, third-party product news, etc. This newsletter is free and is available upon request by calling the marketing information phone number listed below.

E.5.2 Motorola Field Application Engineers

Information and assistance for DSP applications is available through the local Motorola field office. See your local telephone directory for telephone numbers or call (512)891-2030.

E.5.3 Design Hotline – 1-800-521-6274

This is the Motorola number for information pertaining to any Motorola product.

E.5.4 DSP Applications Helpline– (512) 891-3230

Design assistance for specific DSP applications is available by calling this number.

E.5.5 DSP Marketing Information – (512) 891-2030

Marketing information including brochures, application notes, manuals, price quotes, etc. for Motorola DSP-related products are available by calling this number.

E.5.6 DSP Third-Party Support Information – (512) 891-3098

Information concerning third-party manufacturers using and supporting Motorola DSP products is available by calling this number. Third-party support includes:

- Filter design software
- Logic analyzer support
- Boards for VME, IBM-PC/XT/AT, MACII, SPARC, HP300
- Development systems
- Data conversion cards
- Operating system software
- Debug software

Additional information is available on Dr. BuB and in DSP News.

E.5.7 DSP University Support – (512) 891-3098

Information concerning university support programs and university discounts for all Motorola DSP products is available by calling this number.

E.5.8 DSP Training Courses – (602) 994-6900

There are two courses available for the DSP56000 Family:

1. Introduction to the DSP56000/DSP56001 (MTTA5) which is a 4.5-hour audio-tape course on the DSP56000/DSP56001 architecture and programming.
2. Introduction to the DSP56000/DSP56001 (MTT31) which is a four-day instructor-led course and laboratory covering the details of the DSP56000/DSP56001 architecture and programming.

Additional information is available by writing:

- Motorola SPS Training and Technical Operations
- Mail Drop HW68
- P. O. Box 21007
- Phoenix, Arizona 85036

or by calling the number above. A technical training catalog is available which describes these courses and gives the current training schedule and prices.

E.5.9 Reference Books and Manuals

A list of DSP-related books is included here as an aid for the engineer who is new to the field of DSP. This is a partial list of DSP references intended to help the new user find useful information in some of the many areas of DSP applications. Many books could be included in several categories but are not repeated.

E.5.10 General DSP:

ADVANCED TOPICS IN SIGNAL PROCESSING

Jae S. Lim and Alan V. Oppenheim
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988

APPLICATIONS OF DIGITAL SIGNAL PROCESSING

A. V. Oppenheim
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1978

DISCRETE-TIME SIGNAL PROCESSING

A. V. Oppenheim and R. W. Schafer
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1989

DIGITAL PROCESSING OF SIGNALS THEORY AND PRACTICE

Maurice Bellanger
New York, NY: John Wiley and Sons, 1984

DIGITAL SIGNAL PROCESSING

Alan V. Oppenheim and Ronald W. Schafer
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1975

DIGITAL SIGNAL PROCESSING: A SYSTEM DESIGN APPROACH

David J. DeFatta, Joseph G. Lucas, and William S. Hodgkiss
New York, NY: John Wiley and Sons, 1988

FOUNDATIONS OF DIGITAL SIGNAL PROCESSING AND DATA ANALYSIS

J. A. Cadzow
New York, NY: MacMillan Publishing Company, 1987

HANDBOOK OF DIGITAL SIGNAL PROCESSING

D. F. Elliott
San Diego, CA: Academic Press, Inc., 1987

INTRODUCTION TO DIGITAL SIGNAL PROCESSING

John G. Proakis and Dimitris G. Manolakis

New York, NY: Macmillan Publishing Company, 1988

MULTIRATE DIGITAL SIGNAL PROCESSING

R. E. Crochiere and L. R. Rabiner

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1983

SIGNAL PROCESSING ALGORITHMS

S. Stearns and R. Davis

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988

SIGNAL PROCESSING HANDBOOK

C.H. Chen

New York, NY: Marcel Dekker, Inc., 1988

SIGNAL PROCESSING – THE MODERN APPROACH

James V. Candy

New York, NY: McGraw-Hill Company, Inc., 1988

THEORY AND APPLICATION OF DIGITAL SIGNAL
PROCESSING

Rabiner, Lawrence R., Gold and Bernard

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1975

E.5.11 Digital Audio and Filters:

ADAPTIVE FILTER AND EQUALIZERS

B. Mulgrew and C. Cowan

Higham, MA: Kluwer Academic Publishers, 1988

ADAPTIVE SIGNAL PROCESSING

B. Widrow and S. D. Stearns

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1985

ART OF DIGITAL AUDIO, THE

John Watkinson

Stoneham, MA: Focal Press, 1988

DESIGNING DIGITAL FILTERS

Charles S. Williams

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1986

DIGITAL AUDIO SIGNAL PROCESSING AN
ANTHOLOGY

John Strawn
William Kaufmann, Inc., 1985

DIGITAL CODING OF WAVEFORMS
N. S. Jayant and Peter Noll
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984

DIGITAL FILTERS: ANALYSIS AND DESIGN
Andreas Antoniou
New York, NY: McGraw-Hill Company, Inc., 1979

DIGITAL FILTERS AND SIGNAL PROCESSING
Leland B. Jackson
Higham, MA: Kluwer Academic Publishers, 1986

DIGITAL SIGNAL PROCESSING
Richard A. Roberts and Clifford T. Mullis
New York, NY: Addison-Welsey Publishing Company, Inc.,
1987

INTRODUCTION TO DIGITAL SIGNAL PROCESSING
Roman Kuc
New York, NY: McGraw-Hill Company, Inc., 1988

INTRODUCTION TO ADAPTIVE FILTERS
Simon Haykin
New York, NY: MacMillan Publishing Company, 1984

MUSICAL APPLICATIONS OF MICROPROCESSORS
(Second Edition)
H. Chamberlin
Hasbrouck Heights, NJ: Hayden Book Co., 1985

E.5.12 C Programming Language:

C: A REFERENCE MANUAL
Samuel P. Harbison and Guy L. Steele
Prentice-Hall Software Series, 1987.

PROGRAMMING LANGUAGE - C
American National Standards Institute,
ANSI Document X3.159-1989
American National Standards Institute, inc., 1990

THE C PROGRAMMING LANGUAGE
Brian W. Kernighan, and Dennis M. Ritchie
Prentice-Hall, Inc., 1978.

E.5.13 Controls:

ADAPTIVE CONTROL

K. Astrom and B. Wittenmark
New York, NY: Addison-Welsey Publishing Company, Inc.,
1989

ADAPTIVE FILTERING PREDICTION & CONTROL

G. Goodwin and K. Sin
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984

AUTOMATIC CONTROL SYSTEMS

B. C. Kuo
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1987

COMPUTER CONTROLLED SYSTEMS: THEORY &
DESIGN

K. Astrom and B. Wittenmark
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984

DIGITAL CONTROL SYSTEMS

B. C. Kuo
New York, NY: Holt, Reinholt, and Winston, Inc., 1980

DIGITAL CONTROL SYSTEM ANALYSIS & DESIGN

C. Phillips and H. Nagle
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984

ISSUES IN THE IMPLEMENTATION OF DIGITAL
FEEDBACK COMPENSATORS

P. Moroney
Cambridge, MA: The MIT Press, 1983

E.5.14 Graphics:

CGM AND CGI

D. B. Arnold and P. R. Bono
New York, NY: Springer-Verlag, 1988

COMPUTER GRAPHICS (Second Edition)

D. Hearn and M. Pauline Baker
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1986

FUNDAMENTALS OF INTERACTIVE COMPUTER
GRAPHICS

J. D. Foley and A. Van Dam

Reading MA: Addison-Wesley Publishing Company Inc.,
1984

GEOMETRIC MODELING

Michael E. Morteson

New York, NY: John Wiley and Sons, Inc.

GKS THEORY AND PRACTICE

P. R. Bono and I. Herman (Eds.)

New York, NY: Springer-Verlag, 1987

ILLUMINATION AND COLOR IN COMPUTER
GENERATED IMAGERY

Roy Hall

New York, NY: Springer-Verlag

POSTSCRIPT LANGUAGE PROGRAM DESIGN

Glenn C. Reid - Adobe Systems, Inc.

Reading MA: Addison-Wesley Publishing Company, Inc.,
1988

MICROCOMPUTER DISPLAYS, GRAPHICS, AND
ANIMATION

Bruce A. Artwick

Englewood Cliffs, NJ: Prentice-Hall, Inc., 1985

PRINCIPLES OF INTERACTIVE COMPUTER GRAPHICS

William M. Newman and Roger F. Sproull

New York, NY: McGraw-Hill Company, Inc., 1979

PROCEDURAL ELEMENTS FOR COMPUTER
GRAPHICS

David F. Rogers

New York, NY: McGraw-Hill Company, Inc., 1985

RENDERMAN INTERFACE, THE

Pixar

San Rafael, CA. 94901

E.5.15 Image Processing:

DIGITAL IMAGE PROCESSING

William K. Pratt

New York, NY: John Wiley and Sons, 1978

DIGITAL IMAGE PROCESSING (Second Edition)

Rafael C. Gonzales and Paul Wintz

Reading, MA: Addison-Wesley Publishing Company, Inc.,
1977

DIGITAL IMAGE PROCESSING TECHNIQUES

M. P. Ekstrom

New York, NY: Academic Press, Inc., 1984

DIGITAL PICTURE PROCESSING

Azriel Rosenfeld and Avinash C. Kak

New York, NY: Academic Press, Inc., 1982

SCIENCE OF FRACTAL IMAGES, THE

M. F. Barnsley, R. L. Devaney, B. B. Mandelbrot, H. O.

Peitgen,

D. Saupe, and R. F. Voss

New York, NY: Springer-Verlag

E.5.16 Motorola DSP Manuals:

MOTOROLA DSP56000 LINKER/LIBRARIAN
REFERENCE MANUAL

Motorola, Inc., 1991.

MOTOROLA DSP56000 MACRO ASSEMBLER
REFERENCE MANUAL

Motorola, Inc., 1991.

MOTOROLA DSP56000 SIMULATOR REFERENCE
MANUAL

Motorola, Inc., 1991.

MOTOROLA DSP56000/DSP56001 USER'S MANUAL

Motorola, Inc., 1990.

E.5.17 Numerical Methods:

ALGORITHMS (THE CONSTRUCTION, PROOF, AND
ANALYSIS OF PROGRAMS)

P. Berliout and P. Bizard
New York, NY: John Wiley and Sons, 1986

MATRIX COMPUTATIONS
G. H. Golub and C. F. Van Loan
John Hopkins Press, 1983

NUMERICAL RECIPES IN C - THE ART OF SCIENTIFIC
PROGRAMMING
William H. Press, Brian P. Flannery,
Saul A. Teukolsky, and William T. Vetterling
Cambridge University Press, 1988

NUMBER THEORY IN SCIENCE AND
COMMUNICATION
Manfred R. Schroeder
New York, NY: Springer-Verlag, 1986

E.5.18 Pattern Recognition:

PATTERN CLASSIFICATION AND SCENE ANALYSIS
R. O. Duda and P. E. Hart
New York, NY: John Wiley and Sons, 1973

CLASSIFICATION ALGORITHMS
Mike James
New York, NY: Wiley-Interscience, 1985
Spectral Analysis:

STATISTICAL SPECTRAL ANALYSIS, A
NONPROBABILISTIC THEORY
William A. Gardner
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988

THE FAST FOURIER TRANSFORM AND ITS
APPLICATIONS
E. Oran Brigham
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988

THE FAST FOURIER TRANSFORM AND ITS
APPLICATIONS
R. N. Bracewell
New York, NY: McGraw-Hill Company, Inc., 1986

E.5.19 Speech:

ADAPTIVE FILTERS – STRUCTURES, ALGORITHMS,
AND APPLICATIONS

Michael L. Honig and David G. Messerschmitt
Higham, MA: Kluwer Academic Publishers, 1984

DIGITAL CODING OF WAVEFORMS

N. S. Jayant and P. Noll
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984

DIGITAL PROCESSING OF SPEECH SIGNALS

Lawrence R. Rabiner and R. W. Schafer
Englewood Cliffs, NJ: Prentice-Hall, Inc., 1978

LINEAR PREDICTION OF SPEECH

J. D. Markel and A. H. Gray, Jr.
New York, NY: Springer-Verlag, 1976

SPEECH ANALYSIS, SYNTHESIS, AND PERCEPTION

J. L. Flanagan
New York, NY: Springer-Verlag, 1972

SPEECH COMMUNICATION – HUMAN AND MACHINE

D. O'Shaughnessy
Reading, MA: Addison-Wesley Publishing Company, Inc.,
1987

E.5.20 Telecommunications:

DIGITAL COMMUNICATION

Edward A. Lee and David G. Messerschmitt
Higham, MA: Kluwer Academic Publishers, 1988

DIGITAL COMMUNICATIONS

John G. Proakis
New York, NY: McGraw-Hill Publishing Co., 1983

Index

Symbols

#pragma 5-22
.cld 3-1
.cln 3-1
_ 6-9
__asm
 multiple instructions 5-3
__asm() 5-2
__asm() reg_save 5-7
__c_sig_goto_dispatch 6-10
__c_sig_handlers 6-10
__DATE__ 4-2
__DSP56K__ 4-2
__FILE__ 4-2
__fp_shift 6-9
__INCLUDE_LEVEL__ 4-2
__LINE__ 4-2
__mem_limit 6-9
__receive 6-9
__send 6-9
__sig_dfl 6-11
__sig_drop_count 6-11
__sig_err 6-11
__sig_ign 6-11
__stack_safety 6-8
__STDC__ 4-2
__time 6-9
__TIME__ 4-2
__VERSION__ 4-2

Numerics

80386 2-1
80486 2-1

A

a.cld 2-4
abort A-9
abs A-10
Accumulator Registers 4-8
acos A-11
Address Offset Registers 4-8
Address Registers 4-8
affine arithmetic 4-7
-alo 3-4, 3-20, 4-21
alo561 3-2, 3-20, 4-21

ANSI 1-1
asin A-12
-asm option 3-4, 3-31
asm56000 C-1
atan A-13
atan2 A-14
atexit A-15
atof A-16
atoi A-17
atol A-18
autoexec.bat 2-1, 2-3

B

-B option 3-4, 3-5
-b option 3-4, 3-5
bar 2-3, 2-4
brk() 6-9
bsearch A-19

C

-C option 3-4, 3-7
-c option 3-4, 3-31
calloc 6-8, A-20
ceil A-22
char 4-2
cldinfo C-4
cldlod C-5
cofdmp C-5
COFF 1-3, 3-1
compiler's dsp directory tree 2-1, 2-2, 2-4
control line 3-10
control program 3-1
cos A-23
cosh A-23
counter_string 5-23
-crt option 3-4, 3-31
crt0 6-1

D

-D option 3-4, 3-8
Data ALU 4-8
DELETESWAP 2-3
denormalized numbers 4-7
div A-24
div_ba 4-16
DOS extended memory manager 2-2

DOS4GVM 2-2
DOS4GVM.SWP 2-3
dos4gw.exe 2-2
DSIZE 6-2
dsp 2-1
dsplib C-6
dsplnk C-7
DSPLOC 2-1, 2-2, 2-4

E

-E option 3-4, 3-9
errno 4-16, 6-8
exit A-25
exp A-26
exponent 4-6

F

F__uldiv 4-16
F__ulmod 4-16
fabs A-26, A-27, A-28, A-29
fadd_ba 4-16
-fcaller-saves option 3-4, 3-21
fcmp_a 4-16
-fcond-mismatch option 3-4, 3-22
fdiv_ba 4-16
-ffixed-REG option 3-4, 3-22
-fforce-addr option 3-4, 3-21
file inclusion 3-10
-finline-functions option 3-4, 3-21
-fkeep-inline-functions option 3-4, 3-21
float 4-3
floor A-29, A-30, A-31, A-32, A-33, A-48, A-49
fmod A-33
fmpy_ba 4-16
fneg_a 4-16
-fno-defer-pop option 3-4, 3-21
-fno-opt option 3-4, 3-20
-fno-peek option 3-4, 3-20
-fno-strength-reduce option 3-4, 3-20
Frame Pointer 4-9
frame pointer 4-11, 6-2
free A-34, A-36, A-37, A-38, A-39, A-40, A-47, A-83
frexp A-41, A-42
fsub_ba 4-16
-fvolatile option 3-4, 3-22
-fwritable-strings option 3-4, 3-22

G

-g option 3-4, 3-22
G56_EXEC_PREFIX 3-5, 3-31
g56lc 2-3

G56K 3-1
GDB56 3-22
gdb56 C-11
global assembler directive 5-33

H

hello.c 2-4
host port 6-2

I

-I option 3-4, 3-10, 3-11
-i option 3-4, 3-13
identifier length limits 4-1
IEEE P754-1985 4-3
in-line assembly
 instruction template 5-3
in-line assembly code 5-1
Input Registers 4-8
install.exe 2-1
int 4-2
interrupt vectors 6-1
isalnum A-45, A-46, A-50
isalpha A-51
isctrl A-52
isdigit A-52
isgraph A-54
islower A-54
isprint A-55
ispunct A-56
isspace A-56
isupper A-57
isxdigit A-58

J

-j option 3-4, 3-32

L

-l option 3-4, 3-32
l_load 5-23
l_run 5-23
labs A-58
ldexp A-59
ldiv A-60
ldiv_ba 4-16
lmod_ba 4-16
lmpy_ba 4-16
log A-61
log10 A-62
long 4-2
longjmp 3-26, 6-1, 6-12, A-63

M

- M option 3-4, 3-14
- m56002 3-23
- malloc 6-8, A-64
- Mantissa 4-5
- map files 3-32
- MAXMEM 2-2
- MB_CUR_MAX A-128
- mblen A-64
- mbstowcs A-66
- mbtowc A-68
- memchr A-70
- memcmp A-71
- memcpy A-72
- memmove A-73
- memset A-74
- MINMEM 2-2
- ml-memory option 3-4, 3-24, 3-25
- MM option 3-4, 3-15
- mno-biv-plus-linv-promotion option 3-4, 3-24
- mno-do-loop-generation option 3-4, 3-24
- mno-dsp-optimization option 3-4, 3-23, 3-24
- mod_ba 4-16
- modf A-75
- Modifier Registers 4-8
- mp-mem-switchtable 3-24
- mstack_check option 3-4, 3-24, 4-17
- mx-memory option 3-4, 3-24, 3-25
- my-memory option 3-4, 3-24, 3-25

N

- NaNs 4-7
- nostdinc option 3-4, 3-16

O

- O option 3-4, 3-23
- o option 3-4, 3-6
- omr 6-2
- Option, Assemble
 - asm string 3-31
 - c 3-31
- Option, Command line
 - Bdirectory 3-5
 - bPREFIX 3-5
 - o FILE 3-6
 - v 3-7
- Option, Compile
 - fcaller-saves 3-21
 - fcond-mismatch 3-22
 - ffixed-REG 3-22
 - fforce-addr 3-21

- finline-functions 3-21
- fkeep-inline-functions 3-21
- fno-defer-pop 3-21
- fno-opt 3-20
- fno-peephole 3-20
- fno-strength-reduce 3-20
- fvolatile 3-22
- fwritable-strings 3-22
- g 3-22
- ml-memory 3-25
- mno-biv-plus-linv-promotion 3-24
- mno-do-loop-generation 3-24
- mno-dsp-optimization 3-23, 3-24
- mstack_check 3-24
- mx-memory 3-24
- my-memory 3-25
- O 3-23
- pedantic 3-25
- Q 3-25
- S 3-25
- W 3-26
- w 3-26
- Wall 3-30
- Wcast-qual 3-30
- Wid-clash-LEN 3-30
- Wimplicit 3-28
- Wpointer-arith 3-30
- Wreturn-type 3-28
- Wshadow 3-30
- Wswitch 3-30
- Wunused 3-29
- Wwrite-strings 3-30

Option, Link

- crt file 3-31
- j string 3-32
- LIBRARY 3-32
- r MAPFILE 3-32

Option, Preprocessor

- C 3-7
- DMACRO 3-8
- DMACRO=DEFN 3-8
- E 3-9
- I- 3-11
- i FILE 3-13
- IDIR 3-10
- M 3-14
- MM 3-15
- nostdinc 3-16
- P 3-17
- pedantic 3-17
- UMACRO 3-18
- v 3-17
- Wcomment 3-19
- Wtrigraphs 3-20

out-of-line calling C routines 5-35

P

-P option 3-4, 3-17
p_load 5-23
p_run 5-23
-pedantic option 3-4, 3-17, 3-25
perror A-76
pow A-77
pragma 5-22
printf A-78
putchar A-84
puts A-85

Q

-Q option 3-4, 3-25
qsort A-86

R

-r option 3-4, 3-32
raise 6-1, A-87
rand A-88
realloc 6-8, A-88
RUN56 1-3
run56sim 2-4, C-12

S

-S option 3-4, 3-25
setjmp 3-27, 6-1, 6-12, A-90, A-92
short 4-2
SIG_ERR 6-11
SIG_IGN 6-11
signal 6-1, A-94
signal file 6-9
sin A-95
sinh A-96
sizeof 4-1
sprintf A-96
sqrt A-97, A-98
srec C-13
stack pointer 4-9, 4-11, 6-2
standard directory search list 3-2
standard include directory 3-10
strcat A-98, A-99
strchr A-100
strcmp A-100
strcoll A-102
strcpy A-102
strncpy A-104, A-112
strerror A-105
strlen A-106

strncat A-106
strncmp A-108
strncpy A-109
strpbrk A-110, A-111
strstr A-112
strtod A-114
strtok A-115
strtol A-116
strtoul A-118
strxfrm A-120
SWAPNAME 2-3

T

tan A-120
tanh A-121
TERM 2-1
TERMCAP 2-1
tolower A-121, A-122
toupper A-123

U

-U option 3-4, 3-18
udiv_ba 4-16
unsigned char 4-2
unsigned int 4-2
unsigned short 4-2

V

-v option 3-4, 3-7, 3-17
VIRTUALSIZE 2-3
volatile 3-26, 5-22

W

-W option 3-4, 3-19, 3-26
-w option 3-26
-Wall option 3-4, 3-30
-Wcast-qual option 3-4, 3-30
wctombs A-124, A-125, A-126
wctomb A-128
-Wid-clash-LEN option 3-4, 3-30
-Wimplicit option 3-4, 3-28
-Wpointer-arith option 3-4, 3-30
-Wreturn-type option 3-4, 3-28
-Wshadow option 3-4, 3-30
-Wswitch option 3-4, 3-30
-Wunused option 3-4, 3-29
-Wwrite-strings option 3-4, 3-30

X

x_load 5-23

x_run 5-23
XDEF assembler directive 5-33
XREF assembler directive 5-33

Y

y_load 5-23
y_run 5-23

DSP56KCC Optimizing C Compiler Trouble Report

DSP Applications Assistance – (512) 891-3230

Name _____ Company Name _____

Street _____ City _____ zip code _____

Phone _____ Date _____

Version Number: _____ Serial Number _____

Rate the impact of this problem (select those applicable).

Shut down system

Software development halted

Needs eventual fix

Suggested enhancement

Describe the system used with this product (include PC manufacturer, Operating system type and version, memory size, and configuration).

Describe the problem and give details on how it can be reproduced, and any suggestions on how it can or may be fixed. Use extra sheets if needed.

DSP56KCC Optimizing C Compiler User's Manual Report Form

We welcome your comments and suggestions. They help us provide you with better product documentation. Mail completed form to:

Motorola Inc.
6501 Wm. Cannon Drive West
Austin, Texas 78735-8598
Attn: Digital Signal Processors

1. Did you find errors in the manual? Please give page number and a description of each error.

2. Is anything missing or damaged?

3. Did you find the manual clear and easy to use? Please comment on specific sections that you feel need improvement.

4. What sections of this manual do you consider most important/least important.