

Series 32000[®]

GNX — Version 4.0 C Optimizing Compiler Reference Manual

REVISION RECORD

VERSION	RELEASE DATE	SUMMARY OF CHANGES
4.0	June 1990	First Release.
		Support for the Application Specific Instruction Set (ASIS). Function prototypes. Introduction of a source level profiler. Support for the <i>Series 32000/EP</i> .

PREFACE

This is a reference manual for National Semiconductor Corporation's GNX—Version 4 C optimizing compiler. The C optimizing compiler generates high-quality code for the *Series 32000*® architecture, therefore improving the performance of the *Series 32000* system.

The main difference between the C optimizing compiler and other compilers is the advanced optimizing component of the compiler. The optimizer uses advanced optimization techniques to improve speed or save space. This reference manual provides guidelines for using the optimizer as well as a discussion of the compiler's optimization techniques. In addition, this reference manual provides information regarding the compilation process, extensions to the C programming language, and implementation issues.

This manual corresponds to the GNX—Version 4 C compiler.

A complete list of National Semiconductor's international offices may be found on the inside back cover of this manual.

The information contained in this manual is for reference only and is subject to change without notice.

No part of this document may be reproduced in any form or by any means without the prior written consent of National Semiconductor Corporation.

GENIX, GNX, ISE, ISE16, ISE32 and SYS32 are trademarks of National Semiconductor Corporation.

Series 32000 is a registered trademark of National Semiconductor Corporation.

UNIX is a registered trademark of AT&T.

VAX, VMS, and DEC are trademarks of Digital Equipment Corporation.

CONTENTS

Chapter 1	OVERVIEW	
1.1	INTRODUCTION	1-1
1.2	INTENDED AUDIENCE	1-2
1.3	FEATURES AND SUPPORTED LANGUAGE EXTENSIONS	1-3
	1.3.1 Compiler Features	1-3
	1.3.2 Supported C Language Features	1-3
1.4	DOCUMENTATION CONVENTIONS	1-4
	1.4.1 General Conventions	1-4
	1.4.2 Conventions in Syntax Descriptions	1-4
	1.4.3 Example Conventions	1-5
1.5	Incompatibilities With GNX C Compiler Version 3	1-5
Chapter 2	2 COMPILATION PROCESS	
2.1	INTRODUCTION	2-1
2.2	COMPILER STRUCTURE	2-1
2.3	COMMAND LINE OPTIONS	2-2
	2.3.1 UNIX Compilation Options	2-2
	2.3.2 VMS Compilation Qualifiers	2-9
2.4	TARGET MACHINE SPECIFICATION	2-13
2.5	RUN-TIME CHECKS	2-14
	2.5.1 Parameter Check	2-15
	2.5.2 Array Checks	2-15
	2.5.3 NIL_POINTER Checks	2-16
2.6	FLOATING-POINT EMULATION	2-17
	2.6.1 Floating-point Emulation — Native Configuration	2-17
	2.6.2 Floating-point Emulation — Cross-Configuration	2-17
	2.6.3 Floating-Point Emulation — VAX/VMS System	2-18
2.7	ENVIRONMENT VARIABLES (FOR UNIX ONLY)	2-18
Chapter 3	EXTENSIONS TO THE C LANGUAGE	
3.1	INTRODUCTION	3-1
3.2	ANSI FEATURES	3-1
	3.2.1 Function Prototypes	3-1
	3.2.2 Volatile and Const Qualifiers	3-2
	3.2.3 Void Data Type	3-2
	3.2.4 Signed Keyword	3-2

	3.2.5	The #pragma Directive	3-2
	3.2.6	Single-Precision Floating Constants	3-2
	3.2.7	Unsigned Constants	3-3
	3.2.8	Enumerated Types	3-3
	3.2.9	Structure Handling	3-3
	3.2.10	Concatenation of Adjacent String Literals	3-3
	3.2.11	Obsolesce of the Old Fashioned Compound Assignment	3-4
	3.2.12	Obsolesce of the Old Fashioned Initialization	3-4
3.3		DED SUPPORT EXTENSIONS	3-4
	3.3.1	Interrupt/Trap Routines Support	3-4
	3.3.2	Asm Keyword	3-6
	3.3.3	Intrinsic Routines	3-7
3.4	OTHER	EXTENSIONS	3-7
	3.4.1	\$ Sign in Identifiers	3-7
	3.4.2	Bitfields	3-7
	3.4.3	Ident Preprocessor Command	3-7
Chapter 4	IMPL	EMENTATION ISSUES	
4.1		OUCTION	4-1
4.2		MENTATION ASPECTS	4-1
	4.2.1	Memory Representation	4-1
	4.2.2	External Linkage	4-2
	4.2.3	Types and Conversions	4-2
	4.2.4	Variable and Structure Alignment	4-2
	4.2.5	Structure Returning Functions	4-8
	4.2.6	Calling Sequence	4-8
	4.2.7	Mixed-Language Programming	4-8
	4.2.8	Order of Evaluation	4-9
	4.2.9	Order of Allocation of Memory	4-9
	4.2.10	Register Variables	4-9
	4.2.11	Floating-Point Arithmetic	4-10
4.3	UNDEF:	INED BEHAVIOR	4-10
Chapter 5	OPTI	MIZATION TECHNIQUES	
5.1	INTROD	OUCTION	5-1
5.2	THE OP	TIMIZER	5-2
5.3	THE CO	DE GENERATOR	5-9
5.4	MEMOR	RY LAYOUT OPTIMIZATIONS	5-10
5.5	RUNTIN	ME FEEDBACK	5-11
Chapter 6	GUID	DELINES ON USING THE OPTIMIZER	
6.1	INTROL	OUCTION	6-1
-			

6.2	OPTIM	IIZATION FLAGS	6-1
	6.2.1	Optimization Options on the Command Line — UNIX	
	Sy	ystems	6-3
	6.2.2	Optimization Options on the Command Line — VMS	
	Sy	ystems	6-3
	6.2.3	Changing Default Optimization Options	6-4
6.3	PORTI	NG EXISTING C PROGRAMS	6-5
	6.3.1	Undetected Program Errors	6-5
	6.3.2	Compiling System Code	6-6
	6.3.3	Timing Assumptions	6-7
	6.3.4	Low-Level Interface	6-7
	6.3.5	Using Nonstandard Library Routines	6-7
	6.3.6	Reliance on Naive Algebraic Relations	6-8
6.4	DEBUG	GGING OF OPTIMIZED CODE	6-9
6.5	IMPRO	OVED ANNOTATION	6-10
6.6		IONAL GUIDELINES FOR IMPROVING CODE	
0.0		TY	6-11
	6.6.1	Static Functions	6-11
	6.6.2	Integer Variables	6-11
	6.6.3	Local Variables	6-11
	6.6.4	Floating-Point Computations	6-12
	6.6.5	Pointer Usage	6-12
	6.6.6	Asm Statements	6-14
	6.6.7	Register Allocation	6-15
	6.6.8	setjmp()	6-15
	6.6.9	Optimizing for Space	6-16
	6.6.10	Using /NOOPT (-Oo) option	6-16
	6.6.11	Runtime Feedback Optimization	6-16
6.7	COMPI	ILATION TIME REQUIREMENTS	6-17
		•	
Chapter '	7 PRO	FILE INFORMATION	
7.1	INTRO	DUCTION	7-1
7.2	GATHE	ERING PROFILE INFORMATION	7-2
	7.2.1	The Profile Information	7-2
	7.2.2	Code Compilation	7-2
	7.2.3	Pgen	7-3
	7.2.4	The PIT File	7-3
	7.2.5	The pfb_exit.o(pfb_exit.obj)File	7-3
	7.2.6	Compilation in the UNIX Environment	7-4
	7.2.7	Compilation in the VMS Environment	7-4
	7.2.8	Code Execution	7-5
	7.2.9	Disabling Profile Information Accumulation	7-5
	7.2.10	Redefining Standard libc Symbols	7-5

	7.2.11 7.2.12	Execution Time Considerations	7-6 7-6
7 0			
7.3		THE GNX SOURCE PROFILER	7-7 7-7
	7.3.1 7.3.2	Example	7-1 7-8
	7.3.3	Running SPROF	7-8 7-9
	7.3.4	Counts and Basic Blocks	7-10
- .			
7.4		IE FEEDBACK OPTIMIZATION	7-11
	7.4.1	Profile Information Gathering	7-11 7-11
	7.4.2	Runtime Feedback Compilation	7-11
Chapter 8	INTR	INSIC FUNCTIONS	
8.1	INTROD	UCTION	8-1
	8.1.1	Using Intrinsic Functions	8-1
8.2	General S	Series 32000 Intrinsic Functions	8-2
	8.2.1	Single Bit Instructions	8-3
	8.2.2	_ffs (Find First Set)	8-4
	8.2.3	_exti (Extract bit-field)	8-5
	8.2.4	_ins (Insert Bit-field)	8-7
	8.2.5	_cvtp (Convert to Bit Pointer)	8-9
	8.2.6	abs (Absolute Value)	8-10
8.3	CG-Core	Intrinsic Functions	8-11
	8.3.1	_extblt (External Bit Aligned Block Transfer)	8-13
	8.3.2	BITBLT instructions	8-15
	8.3.3	_bitwt (Bit Aligned Word Transfer)	8-19
	8.3.4	_movmp (Move Multiple Pattern)	8-20
	8.3.5	_sbits (Set Bit String)	8-21
	8.3.6	_sbitps (Set Bit Perpendicular String)	8-22
	8.3.7	_tbits (Test Bit String)	8-23
8.4	NS32GX	320 Intrinsic Functions	8-25
	8.4.1	NS32GX320 typedefs	8-26
	8.4.2	_mulwd (Multiply Word to Double)	8-27
	8.4.3	_cmuld (Complex Multiply Double)	8-28
	8.4.4	_cmacd (Complex Multiply and Accumulate Double)	8-29
	8.4.5	_mactd (Multiply and Accumulate Twice Double)	8-30
Appendix	A SERI	ES 32000 STANDARD CALLING CONVENTIONS	
A.1	INTROD	UCTION	A-1
A.2	CALLIN	G CONVENTION ELEMENTS	A-1
Appendix	B MIXE	D-LANGUAGE PROGRAMMING	
B.1			D 1
D.I		Weiting Mined Language Programs	B-1
	B.1.1	Writing Mixed-Language Programs	B-1

	B.1.2 B.1.3 B.1.4	Compilation on UNIX Operating Systems	B-8 B-6 B-7
ъ.		Compilation on VMS Operating Systems	
B.2		LING THE MIXED-LANGUAGE EXAMPLE	B-7
	B.2.1 B.2.2	Compiling the Example on a UNIX System	B-7
		Compiling the Example on a VMS System	B-8
B.3	PROGR	RAM MODULE LISTINGS	В-8
Appendi		OR DIAGNOSTICS	
C.1	INTRO	DUCTION	C -1
C.2	ERROR	R MESSAGES	C-:
	C.2.1	Error Messages Format	C-1
	C.2.2	System Errors	C-2
	C.2.3	Limitation Errors	C-2
	C.2.4	Syntax Errors	C-3
	C.2.5	Severe Errors	C-5
	C.2.6	Caution Errors	C-5
	C.2.7	Warnings	C-6
Appendi	x D COM	IPILER OPTIONS	
D.1	INTRO	DUCTION	D-1
Appendi	x E EMB	BEDDED PROGRAMMING HINTS	
E.1	INTRO	DUCTION	E -1
E.2	VOLAT	'ILE AND CONST	E-1
	E.2.1	Const Type Qualifier	E-1
	E.2.2	Volatile Type qualifier	E-2
	E.2.3	Memory Allocation	E-4
	E.2.4	Initialized C Variables	E-4
	E.2.5	Programming Memory Mapped Devices	E-5
E.3	ASM ST	PATEMENTS	E-6
E.4	EXAMP	PLES OF PROGRAMMING WITH INTRINSIC	
		TIONS	E-7
	E.4.1	NS32CG16 bit instructions	E-7
	E.4.2	NSGX320 specific instructions	E-11
E.5	PROGR	AMMING TRAP/INTERRUPT ROUTINES	E-12
Appendi	x F GLO	SSARY	
FIGURE	\mathbf{S}		
Figure 4-1	1. Bitfie	eld Padding	4-6

Figure 4-2.	Alignment on Bitfields	4-6
Figure 5-1.	Relationship Between Various Optimizations	5-3
Figure 5-2.	Flow Graph	5-4
Figure 5-3.	Example of Loop Unrolling	5-5
Figure 5-4.	Example of Partial Redundancy Elimination	5-7
Figure 7-1.	Example of sprof Output	7-7
Figure 7-2.	sprof Data Flow Description	7-8
Figure B-1.	Cross-Language Pairs	B-2
Figure E-1.	Example of Linker Directive File	E-4
Figure E-2.	The Image	E-8
Figure E-3.	The Image with the Reversed Shape	E-9
TABLES		
Table 2-1.	Filename Conventions	2-3
Table 2-2.	Target Selection Parameters	2-13
Table 2-3.	Run-time Check Flags	2-14
Table 4-1.	Variable Alignment	4-3
Table 6-1.	Optimization Options	6-2
Table 6-2.	Changing Default Optimization Options	6-4
Table 6-3.	Recognized Library Routines	6-8
Table 8-1.	Effect of thits on PSR L and F flags	8-24
Table B-1.	Compilers and their Associated Libraries	B-6
Table D-1.	UNIX Operating System Options	D-2
Table D-2.	VMS Operating System Options	D-4
Table D-3.	Options Passed to the Preprocessor — UNIX Systems	D-5
Table D-4.	Options Passed to the Preprocessor — VMS Systems	D-6
Table D-5.	Options Recognized and Passed to the Linker	D-6

INDEX

OVERVIEW

1.1 INTRODUCTION

This manual describes National Semiconductor's GNX—Version 4 C Optimizing Compiler. The compiler is one of a family of compatible optimizing compilers for the Series 32000 family of microprocessors.* The GNX—Version 4 C Compiler replaces and makes obsoletes the previous GNX—Version 4 C Compiler. It implements the C language as described in C Programming Language by Kernighan and Ritchie, together with most of the important features in the ANSI C standard like function prototypes, const and volatile type qualifiers, and the void data type (see Section 1.3.1). The compiler is fully compatible with the System V C compiler, a compiler derived from the portable C compiler (pcc).

In addition, the GNX—Version 4 C Optimizing Compiler includes important extensions for programming embedded applications like ASIS (Application Specific Instructions) support, interrupt/trap handling in C, and asm statement. The compiler is available as a cross-support compiler running on VMSTM and UNIX® operating systems as well as a native compiler running on Series 32000 operating systems derived from UNIX System V, Release 3. Additional information on other tools in the Series 32000 family can be found in the GNX—Version 4 Commands and Operations Manual.

This manual is organized as follows:

- Introduction (Chapter 1)
- Compilation Process (Chapter 2)
- Extensions to the C language(Chapter 3)
- Implementation Issues (Chapter 4)
- Optimization Techniques (Chapter 5)
- Guidelines on Using the Optimizer (Chapter 6)
- Profile Information (Chapter 7)
- Intrinsic Routines (Chapter 8)
- Series 32000 Calling Standard Conventions (Appendix A)

^{*} At this writing, the family consists of a C Optimizing Compiler, Pascal Optimizing Compiler, and FORTRAN 77 Optimizing Compiler.

- Mixed-Language Programming (Appendix B)
- Error Messages (Appendix C)
- Compiler Options (Appendix D)
- Embedded Programming Hints (Appendix E)
- Glossary (Appendix F)

1.2 INTENDED AUDIENCE

This manual is for experienced C programmers. The information provided covers compiler options, extensions to the standard C programming language, and implementation issues. Knowledge of optimization techniques is useful but not essential; Chapter 5 provides an overview of optimization techniques used by the optimizer. And Chapter 6 provides further guidelines to help the programmer avoid problems that can occur when using the optimizer.

Recommended C reference books include:

ANSI C standard (ANSI X3.159-1989).

Harbison, Samuel and Steele, Guy. C, A Reference Manual, 2nd. ed., Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1984.

Kernighan, Brian and Ritchie, Dennis. *The C Programming Language*, 2nd, ed., Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1989.

1.3 FEATURES AND SUPPORTED LANGUAGE EXTENSIONS

1.3.1 Compiler Features

The following are the main features of the C Optimizing Compiler:

- · Accepts most ANSI C features.
- pcc compatible.
- Allows for use of Application Specific Instructions in C via intrinsic routines.
- Allows for programming of interrupt/trap handlers in C (see Section 3.3.1).
- Optimizations can be tuned to either improve speed or save space.
- Optimization level is controlled by the user.
- Code can be generated that is tuned to the specific target system.
- Full support of mixed-language programming.
- User controlled alignment of variables and structure members.
- · Improved structure handling.
- Assembly output can be annotated with source lines.
- Fast compilation mode.
- Advanced error handling, recovers from simple syntax errors.

1.3.2 Supported C Language Features

The compiler implements the full C language as defined in Appendix A of C Programming Language by Kernighan and Ritchie. In addition most of the ANSI C standard features and important extensions for embedded programming are supported. The following extensions are supported:

- ANSI C features:
 - const for defining read-only entities.
 - volatile for sensitive variables.
 - Function prototypes.
 - Signed keyword.
 - pragma preprocessor command (specifically a pragma that enables marking a trap/interrupt routine).
 - void data type.
 - Structures may be assigned, passed as arguments and returned from functions.

- Initialization of automatic aggregated types.
- Structure/union member names need not be globally unique.
- Structure and union size is not limited.
- Unsigned constants (to save run-time conversions).
- Single-precision floating constants (to save run-time conversions).
- Enumeration datatypes can be used as "int".
- Unsigned or signed bitfields.

All of the above extensions are discussed in Chapter 3.

1.4 DOCUMENTATION CONVENTIONS

The following documentation conventions are used in text, syntax descriptions, and examples in describing commands and parameters.

1.4.1 General Conventions

Nonprinting characters are indicated by enclosing a name for the character in angle brackets <>. For example, <CR> indicates the RETURN key, <ctrl/B> indicates the character input by simultaneously pressing the control key and the B key.

Constant-width type is used within text for filenames, directories, command names and program listings; it is also used to highlight individual numbers and letters. For example,

the C preprocessor, cpp, resides in the GNXDIR/lib directory.

1.4.2 Conventions in Syntax Descriptions

The following conventions are used in syntax descriptions:

Constant-width boldface type indicates actual user input.

Italics indicate user-supplied items. The italicized word is a generic term for the actual operand that the user enters. For example,

```
cc [[option] ... [filename] ...] ...
```

Spaces or blanks, when present, are significant; they must be entered as shown. Multiple blanks or horizontal tabs may be used in place of a single blank.

{ } Large braces enclose two or more items of which one, and only one, must be used. The items are separated from each other by a logical OR sign " | ."

- [] Large brackets enclose optional item(s).
- Logical OR sign separates items of which one, and only one, may be used.
- ... Three consecutive periods indicate optional repetition of the preceding item(s). If a group of items can be repeated, the group is enclosed in large parentheses "()."
- ,,, Three consecutive commas indicate optional repetition of the preceding item. Items must be separated by commas. If a group of items can be repeated, the group is enclosed in large parentheses "()."
- () Large parentheses enclose items which need to be grouped together for optional repetition. If three consecutive commas or periods follow an item, only that item may be repeated. The parentheses indicate that the group may be repeated.
- indicates a space. is only used to indicate a specific number of required spaces.

All other characters or symbols appearing in the syntax must be entered as shown. Brackets, parentheses, or braces which must be entered, are smaller than the symbols used to describe the syntax. (Compare user-entered [], with [] which show optional items.)

1.4.3 Example Conventions

In interactive examples where both user input and system responses are shown, the machine output is in constant-width regular type; user-entered input is in constant-width boldface type. Output from the machine which varies (e.g., the date) is in italic type. For example,

```
--> g <CR>
Breakpoint 2 reached at filename _main: .3
    printf("hello\r\n");
```

1.5 Incompatibilities With GNX C Compiler Version 3

The incompatibilities with the GNX-Version 3 C compiler are:

1. signed keyword added.

GNX version 4 is extended to support the ANSI C signed type specifier. Programs using this keyword in another context (i.e. as identifier or typedef names) are not legal and will be viewed as an error.

```
Example: float signed;
```

The above example defines an identifier whose name is signed and will not pass compilation.

2. Const and volatile type qualifiers fully suported.

The GNX version 3 C compiler includes a partial implementation of the ANSI C const and volatile type qualifiers. In GNX version 4 C compiler const and volatile are fully implemented as defined in ANSI C standard. Programs which rely on version 3 partial implementation may compile differently in GNX—Version 4 C.

3. -N(/TABLES on VMS) was deleted.

The compiler internal tables are allocated dynamically. Hence the -N (/TABLE_SIZE on VMS) no longer has any meaning. The compiler will silently ignore this option when used.

COMPILATION PROCESS

2.1 INTRODUCTION

The GNX—Version 4 C Compiler is a modular language processor consisting of five separate programs. This chapter describes the five programs, the GNX—Version 4 C Compiler invocation, options available to the user, and file generation during compilation.

2.2 COMPILER STRUCTURE

The GNX—Version 4 C Compiler's five programs are:

- Driver
- Macro preprocessor
- C language parser (front end)
- Optimizer
- Code generator

The driver is a program that parses and interprets the command line and then sequentially calls each of the other programs, depending on its input programs and the command line options.

The Macro preprocessor is the C preprocessor, cpp. Its input is a program file optionally containing preprocessing commands.

The C language parser is the compiler's front end, cc fe. Its input is a C program. Its output is the same program in a proprietary intermediate form.*

The optimizer, opt, is a true global, language-independent optimizer that uses advanced optimization techniques to improve the code. Both its input and output are in the same intermediate form. See Chapter 5 for more detailed information.

The intermediate form is language-independent. This allows the same optimizer and code generator to be used by all National Semiconductor GNX — Version 4 Compilers, i.e., the FORTRAN 77 Compiler, the Pascal Compiler, and the C Compiler.

The Series 32000 code generator, cgen_coff, generates an assembly program from a program in the intermediate form.

The assembly program produced by the code generator must be assembled by the *Series 32000* assembler to produce an object code program. The assembler is automatically called by the driver program.

The user produces an executable program by running the *Series 32000* linker on one or more object code programs with run-time library archives. On UNIX systems the linker is automatically called by the driver program. On VMS systems it must be called separately (for further details on invocation see the *GNX Commands and Operations Manual*).

2.3 COMMAND LINE OPTIONS

The GNX—Version 4 C Compiler operation is controlled by a large number of compilation parameters. Many of these parameters, such as the target system specification can be permanently set by means of the GNX Target Setup (GTS) facility. For details on how to use GTS, see the *Series 32000 GNX — Version 4 Commands and Operations* manual. All compilation parameters can be specified as command line options which override any previously existing default values.

Command line options and default values are the same for all supported host systems, but their syntax varies from host to host. Two host systems are currently supported: the UNIX operating system (in both cross-support and native variants) and the VMS operating system (cross-support only). The next two sections provide details on the various compilation parameters and their syntax on these host systems. The tables in Appendix D summarize the various compilation options of both operating systems and can be used as a quick reference.

2.3.1 UNIX Compilation Options

The invocation syntax of the GNX—Version 4 C Compiler under UNIX is:

```
cc [ [option]... [filename]... ]... (native configuration)

nmcc [ [option]... [filename]... ]... (cross-support configuration)
```

The compiler accepts a variable number of file arguments and compilation options. It produces an executable file, object file(s), or assembly file(s), according to the options specified. The files compiled are normally C program sources, but other types of files are also recognized. A file type is recognized by its suffix. A compilation option is recognized by the UNIX convention of a minus-sign prefix.

Filename Conventions

Files are identified by the compiler according to their suffix. Files with names ending with .c or .i are C source programs.

Files ending with .c, pass through the macro preprocessor (cpp) before compilation. Files ending with .i compile directly and assemble to produce object programs left in files whose names are those of the source files with .o substituted for the given suffix.

The intermediate .o file is deleted if a C program consisting of a single file is compiled and linked at the same time.

In support of mixed-language programming, the compiler also recognizes and compiles appropriate files written in other programming languages. Files with a .s suffix are assembly source programs and may be assembled (to produce .o files) and linked. Pascal and FORTRAN 77 source files are also recognized, and compile appropriately if your system includes the National Semiconductor GNX — Version 4 Compiler for those languages. The suffixes for these files are listed in Table 2-1. See Appendix B for details on mixed-language programming.

All other files (normally .o or .a files) are compatible object programs or archives of object programs, typically produced by previous runs of the GNX-Version 4 C Compiler, and pass directly to the linker. The object files link into one executable file with the default name a32.out (or a.out in a native-support environment).

Table 2-1. Filename Conventions

FILE NAME SUFFIX	FILE TYPE
.c	C source file
.i	Preprocessed C source file
.f, .for	FORTRAN 77 source file
.F, .FOR	FORTRAN 77 source with cpp directives
.p, .pas	Pascal source file
.P, .PAS	Pascal source with cpp directives
.s	Assembly source file
other (.o, .a, etc.)	Object code or library-archive file

Compiler Options

The following is a list of the compilation options which may be specified on the invocation line.* The tables in Appendix D summarize the various compilation options and can be used as a quick reference.

- -o (PERFORMS OPTIMIZATIONS)
- -**F**flags (SPECIFIES OPTIMIZATION FLAGS)
- -oflags (PERFORMS OPTIMIZATIONS ACCORDING TO FLAGS)

The -0 option directs the GNX—Version 4 C Compiler to perform global optimizations. The optimizer uses a variety of optimization techniques which ensure the fastest possible code. In certain cases, such as when code density is of greater importance than code speed, it is necessary to specify optimizations. Using the -F option with the optimization flags listed in Chapter 6 sets the selected optimization flags. Using the -F option by itself will do nothing. -Oflags is a shorthand notation for -0 -Fflags. A detailed discussion of optimization techniques is found in Chapter 5 and Chapter 6.

- COMPILES QUICK, NO CODE)
 This option allows for a quick error-checking compilation. No code is generated.
- -a (GENERATES RUN-TIME CHECKS)

 This option controls the generation of code that checks for run-time errors. See Section 2.5 for more details.
- -aflags (GENERATES RUN-TIME CHECKS)

 This option controls the generation of code for selective run-time error checks. See Section 2.5 for more details.
- -g (PREPARES SYMBOLIC DEBUGGING INFORMATION)

 The -g option instructs the GNX—Version 4 C Compiler to prepare symbolic debugging information for symbolic debuggers, such as dbug. See the discussion on debugging of optimized code in Section 6.4. ERROR line431contains a . No matching
- -p (PREPARES PROFILE INFORMATION FOR A PROGRAM PRO-FILER)
 This option prepares profile information for a program profiler, such as
- -B (GATHER PROFILE INFORMATION)

 This option instructs the compiler to add special code for profile

^{*} The GNX—Version 4 C compiler supports the System V Interface Definition (SVID) for C compilers. Where possible, space is allowed between an option and its following flags, i.e., -oout is the same as -o out, and -J2 is the same as -J 2. Similarly, -DHOST is equivalent to -D HOST. The notation in this section follows traditional UNIX conventions.

information gathering. See Chapter 7 for more details.

(COMPILES BUT DOES NOT LINK) -c

The -c option directs the GNX-Version 4 C Compiler to perform the compilation process up to, but not including, linking. Output is left in object files whose names end with .o. This option is useful when compiling only a portion of a program's modules. For example,

cc -c sample.c

creates the file sample.o. No executable file is created.

(COMPILES AND LEAVES ASSEMBLY FILES) -s

The -S option directs the GNX-Version 4 C Compiler to terminate the compilation process before assembly. The assembly output is left in files whose names are those of the source, with .s substituted for the original suffix. For example,

cc -S sample.c utils.c

creates the files sample.s and utils.s. No executable or object file is created.

(EMBEDS C SOURCE LINES AS COMMENTS IN ASSEMBLY) -n

This option puts the C source lines into the assembly output file as comments. If the optimizer is enabled, explanatory optimizer comments are also put into the assembly output file. Note that the -n option is useful only in conjunction with the -S option.

(LEAVES COMMENTS IN) -c

The preprocessor normally removes the comments from its output. The -C option prevents this. This option can be useful when cpp's output must be examined or when the -n option is used and C comments are required in the assembly file.

(PUTS LITERAL STRINGS IN READ-ONLY MEMORY) -R

C literal strings are, by default, writable and are thus allocated in the writable data space. The -R option allocates literal strings in a readonly area.

(RENAMES THE OUTPUT FILE) -o out

The -o option redirects the output file from the compilation process to a file named out. For example,

cc sample.c utils.c -o sample

generates the executable file sample from the two source files, and

cc -S sample.c -o new sample.s

generates the assembly file new sample.s.

-Jwidth (ALIGNMENT WITHIN STRUCTURES)

This option allows the user to set structure-member alignment on bytes (width = 1), words (width = 2), or double-words (width = 4). Default value for width is 4 (double-word-aligned).

-w (NO WARNING DIAGNOSTICS)

The GNX—Version 4 C Optimizing Compiler normally prints warnings regarding inconsistencies in the input program. The -w option suppresses these warning diagnostics. See Appendix C for a complete list of the warning diagnostics.

-w66 (SUPPRESSES FORTRAN 66 WARNINGS)

This is only useful when compiling FORTRAN 77 programs.

-T (UNDEFINED VARIABLE TYPE)

This is only useful when compiling FORTRAN 77 programs.

-A (ALLOCATES VARIABLES AS STANDARD)

This option directs the compiler to adhere to the ANSI C standard, with respect to the declaration and allocation of global variables. When this option is used, there must be exactly one declaration of each global variable without the keyword *extern* within the entire program. This declaration is considered the *definition* of the variable.

-m (USES THE m4 PREPROCESSOR)

With this option, the m4 preprocessor is used on assembly and FOR-TRAN 77 files before assembling and compiling them.

-d (CASE SENSITIVITY)

This is only useful when compiling Pascal and FORTRAN 77 programs.

-N[parameter][size](SET INTERNAL TABLE SIZE)

This option is only useful for FORTRAN programs.

-v (VERBOSE)

This option lists the subprograms of the GNX—Version 4 C Compiler as they are executed by the driver program.

-vn (SHOWS BUT DOES NOT ACTUALLY EXECUTE)

This option lists the compiler subprograms that are called by the compiler's driver program, without actually executing them. This option can be used to verify how other compiler options work.

-Kparameter (SETS TARGET CPU, FPU, OR BUSWIDTH)

The -K option allows the user to "tune" the GNX—Version 4 C Compiler by specifying the CPU, the FPU (or absence of), and/or buswidth of the target system. See Sections 2.4 and 2.6 for more details.

-**z**c (USES ALTERNATIVE LIBRARY)

This option directs the compiler to link an alternative library and

initialization file, determined by the character which follows the option. For example,

cc -Z2 unix.c

links unix.o with crt2.o and lib2.a.

-x (GENERATES MODULAR CODE)

This option directs the compiler to generate code that conforms to the Series 32000 architectural feature of modularity (which allows the modular use of external references). For further information see the Series 32000 GNX — Version 4 Language Tools Technical Notes and the Series 32000 Programmer's Reference Manual.

-f (FLOATING-POINT EMULATION)

This option tells the compiler driver that there is no FPU on the target and floating-point emulation is desired. See Section 2.6 for a discussion of this option and floating-point emulation.

The compiler accepts the following options and passes them to the C preprocessor.

-D name[=def] (DEFINES)

The -D switch defines *name* equal to *def* to the preprocessor. If no explicit value is given, *name* is defined as having the value 1. The use of this option is equivalent to putting a "#define *name def*" at the beginning of each C source file.

For example:

cc -DHOST=VAX sample.c

works as if the following define was at the head of sample.c:

#define HOST VAX

-E (RUNS cpp ONLY)

This option terminates the compilation after preprocessing; only the cpp preprocessor is invoked, and its output is sent to the standard output, stdout.

-Idir (SPECIFIES DIRECTORY FOR INCLUDED FILES)

This option tells to use the specified directory as the default directory for included files. Include files that are called using double quotes, e.g., #include "filename", are sought first in the directory of the compiled file, then in the directories specified by -I, and finally in directories on a standard list (/usr/include). If the user explicitly names the file to be included by using the complete path, e.g., "/a/mydir/filename", the named file is sought directly. If angle brackets instead of double are used quotes, #include <filename>, the file is sought in the directories on a standard list (/usr/include).

-M (RUNS cpp ONLY, GENERATES MAKEFILE DEPENDENCIES)

This option runs only the cpp macro preprocessor on the named C programs, requests it to generate makefile dependencies and then sends the result to the standard output, stdout. For example:

runs cpp on all of the C programs in the current directory and generates all makefile dependencies. These dependencies are then sent to the file new.makefile.

-P (RUNS cpp ONLY, REDIRECTS OUTPUT TO .i FILE)

This option is similar to -E, except that the output of cpp is sent to a file with a .i extension. For example:

cc -P sample.c utils.c

creates the files sample.i and utils.i.

(UNDEFINES) -Uname

> Using this option is equivalent to putting "#undef name" at the beginning of each C source file.

In addition, the compiler accepts the following compiler options and passes them to the linker. See the GNX — Version 4 Linker User's Guide manual for details.

-v	(LINKER VERSION)
-1 lib	(SPECIFIES A PROGRAM LIBRARY)
-s	(STRIPS THE EXECUTABLE FILE OF SYMBOL TABLE AND RELOCATION BITS)
-r	(RETAINS RELOCATION)
-u symname	(UNDEFINES SYMBOL IN SYMBOL TABLE)
-e epname	(DEFINES ENTRY POINT)
-x	(NO LOCAL SYMBOLS IN OUTPUT SYMBOL TABLE)
-i	(RUN-TIME INITIALIZATIONS)

The following option can be used as an "escape" to pass additional options (not recognized by the GNX—Version 4 C Compiler) to the C preprocessor, assembler, or linker.

-wx, options (PASSES OPTIONS TO COMPILATION PHASE x)

This option passes options to the C preprocessor (x = p), the assembler (x = a), or the linker (x = 1). The options must be a single argument (no embedded space, unless quoted). For example, the command,

cc -Wl, -mmu382 sample.c

passes the option -mmu382 to the linker.

2.3.2 VMS Compilation Qualifiers

The command line invocation syntax of the GNX—Version 4 C Compiler is as follows:

```
nmcc [qualifier]... filename
```

The normal operation of the GNX—Version 4 C Compiler compiles and assembles a file specified on the command line to create an object file. Command qualifiers (preceded by a /) are applied as necessary. Most qualifiers can be preceded by NO to reverse their function. The usual VMS conventions regarding default filename extensions, case insensitivity, qualifier syntax and abbreviation rules apply. The GNX— Version 4 C Compiler accepts only one C source file as input and produces an object file with optional intermediate results (such as an assembly file). If the source file has no extension, a . C extension is assumed.

The following is a list of the compilation qualifiers which may be specified on the invo-

The tables in Appendix D summarize the various compilation qualifiers and can be used as a quick reference.

/[NO]OBJECT[=filename]

This qualifier directs the compiler to leave the object code in a file named filename. If filename has no suffix, .OBJ is added as a suffix. If filename is not specified, the object code is placed in a file with the source's filename, with the .OBJ suffix substituted for the original suffix. Default of this qualifier is /OBJECT. For example,

```
NMCC/OBJ=NEW UTILS.OBJ UTILS.C
```

compiles the file utils.c, and leaves the result in a file called new utils.obj.

The command:

```
NMCC/NOOBJ/ASM/OPT/ANNO SAMPLE.C
```

results in an annotated, optimized assembly translation of sample.c and does not generate an object file.

The command NMCC/NOOBJ x.c results in a quick compilation of x.c without producing any output. This is useful for error checking.

/[NO]OPTIMIZE[=(flags)]

This qualifier directs the GNX-Version 4 C Compiler to perform global optimizations. A detailed discussion of the GNX-Version 4 C Compiler optimization techniques is located in Chapter 5 and Chapter 6. Default is /NOOPTIMIZE.

/ [NO]CHECK

This qualifier controls the generation of code that checks for run-time errors. Default is /NOCHECK. See Section 2.5 for more details.

/ NO DEBUG

The /DEBUG qualifier instructs the GNX—Version 4 C Compiler to prepare symbolic debugging information for symbolic debuggers, such as DBUG. See the discussion on debugging of optimized code in Section 6.4. Default is /NODEBUG.

/[NO]GATHER (GATHER PROFILE INFORMATION)

This qualifier instructs the compiler to add special code for profile information gathering. The default is /NOGATHER. See Chapter 7 for more details.

/[NO]ASM[=filename]

This qualifier directs the compiler to leave the intermediate assembly file in a file named *filename*. If *filename* has no suffix, .ASM is added as a suffix. If *filename* is not given, the source filename is used substituting the .ASM suffix with the source filename's suffix. Default of this qualifier is /NOASM. For example,

NMCC/ASM=NEW UTILS.ASM UTILS.C

compiles the file UTILS.C, and produces NEW_UTILS.ASM and UTILS.OBJ.

/[NO]ANNOTATE

This qualifier directs the compiler to put GNX—Version 4 C source lines as comments into the assembly output file. If the optimizer is enabled, explanatory optimizer comments are also added into the assembly output. Note that this qualifier is useful only in conjunction with the /ASM qualifier. Default is /NOANNOTATE.

/[NO]ROM_STRINGS

C literal strings are, by default, writable and are thus allocated in the writable data space. This qualifier directs the compiler to put all literal strings in read-only memory.

/ALIGN [=width]

This qualifier allows the user to set structure member alignment on bytes (width = 1), words (width = 2), or double-words (width = 4). Default value for width is 4 (double-word-aligned). See Section 4.2.4 for details of the GNX—Version 4 C Compiler's alignment scheme.

/[NO]WARNING

The GNX—Version 4 C Compiler prints warnings regarding inconsistencies found in the input program. The /NOWARNING qualifier suppresses these warning diagnostics. Default is /WARNING. See Appendix C for details on warning diagnostics.

/ [NO]STANDARD

This qualifier directs the compiler to adhere to the draft-proposed ANSI C standard, with respect to the declaration and allocation of global variables. When /STANDARD is used, there must be exactly only one declaration of each global variable without the keyword extern within the entire program. This declaration is considered the "definition" of the variable. Default is /NOSTANDARD.

/TABLE SIZE=(table_name=size[,...])

This option is only useful for compiling FORTRAN programs.

/ NO VERBOSE

This qualifier lists the parts of the GNX—Version 4 C Compiler as they are called by the driver program. Default is /NOVERBOSE.

/[no]vn

With this qualifier, the compiler lists the subprograms that are called by the driver program, without actually executing them. This qualifier can be used to verify how the other qualifiers work. Default is /NOVN.

/TARGET= (CPU=cpu, FPU=fpu, BUSWIDTH=bus)

The /TARGET qualifier allows the user to "tune" the GNX—Version 4 C Compiler by specifying the CPU, the FPU (or absence of), and/or buswidth of the target system. See Sections 2.4 and 2.6 for more details.

/ [NO]MODULAR

This qualifier directs the compiler to generate code that conforms to the Series 32000 architectural feature of modularity (which allows the use of external references). For further information see the Series 32000 GNX — Version 4 Language Tools Technical Notes and the Series 32000 Programmer's Reference Manual. Default is /NOMO-DULAR.

/[NO]ERROR[=filename]

The /ERROR qualifier instructs the GNX—Version 4 C Compiler to direct compilation error messages to an error log file in addition to the standard output. If filename has no suffix, the suffix .ERR is added. If no destination file is given, the source filename is used, substituting .ERR for the source filename's suffix. Default sends the errors to the standard output only. For instance,

NMCC /ERROR=FILE1 FILE1.C

creates an error log file named FILE1.ERR.

/[NO]PRE PROCESSOR

This qualifier causes the source file to be passed to the GNX C preprocessor before the normal processing by the GNX—Version 4 C language parser.

```
Default is /PRE PROCESSOR.
```

In addition, the compiler recognizes the following compiler qualifiers and passes them to the C preprocessor. These qualifiers must be used in conjunction with the $\protect{PRCESSOR}$ qualifier.

```
/DEFINE = (name [=def] [,...])
```

The use of this option is equivalent to putting a #define name def at the beginning of the C source file. The /DEFINE switch defines name equal to the value def to the preprocessor. If no explicit value is given, name is defined as having the value 1. For example:

```
NMCC/PRE PROCESSOR/DEFINE=("VAX", "TARGET IS NS32000") SAMPLE.C
```

works as if the following two defines were at the head of SAMPLE.C:

```
#define VAX 1
#define TARGET IS NS32000
```

/ [NO]COMMENT

The preprocessor normally removes the comments from its output. The /COMMENT qualifier prevents this. This qualifier is useful when cpp's output must be examined or when the /ANNOTATE qualifier is used and C comments are required in the assembly file. Default is /NOCOMMENT.

/ [NO]EXPAND [=filename]

This qualifier controls whether the output of the preprocessor is saved to a file. If *filename* has no suffix, the suffix .MAC is added. If *filename* is not given, the source file name is used substituting the suffix .MAC for the source file name's suffix. (Default is /NOEXPAND.)

/INCLUDE=(include_dir [,...])

This qualifier tells the cpp preprocessor to use the specified directory as the default directory for included files. Include files that are specified using double quotes, e.g., #include "filename", are sought first in the directory of the compiled file, then in the directories specified by the /INCLUDE option, and finally in directories on a standard list (GNXDIR:INCLUDE). If the user explicitly names the file to be included by using the complete path, i.e., #include "[MYDIR] filename", the named file is sought directly. If angle brackets are used instead of double quotes, e.g., #include <filename>, the file is sought in the directories on a standard list (GNXDIR:INCLUDE).

/UNDEFINE=(name [,...])

Using this qualifier is equivalent to putting #undef name at the beginning of each C source file.

2.4 TARGET MACHINE SPECIFICATION

The compiler provides a way for the user to tune the code for a specific target system by specifying its CPU, FPU and buswidth. This tuning is performed by setting permanent defaults using the GNX Target Setup (GTS) facility, or by specifying -K (/TARGET on VMS) on the command line. Table 2-2 lists the flags and the possible settings. The values for the CPU and FPU can either be the complete device name e.g., NS32332 or NS32081, or the last characters of the device name, e.g. 332 or cq16. The absence of an FPU on the target system can be indicated by specifying emulation or nofpu (for more details see Section 2.6). The buswidth is specified in bytes.

Table 2-2. Target Selection Parameters

CPU (C)	FPU (F)	BUSWIDTH (B)
[NS32]008	[NS32]081	1
[NS32]016	[NS32]181	2
[NS32]cg16	[NS32]381	4
[NS32]fx16	[NS32]580	
[NS32]cg160	emulation	
[NS32]032	nofpu	
[NS32]332		
[NS32]532		
[NS32]gx32		
[NS32]gx320		

Example:

The following example specifies an NS32CG16 CPU, an NS32081 FPU, and a buswidth of 4 bytes.

UNIX

```
nmcc -KCcg16 -KF081 -KB4 temp.c(cross-support)
or cc -KCcg16 -KF081 -KB4 temp.c
```

VMS

NMCC /TARGET=(CPU=cg16, FPU=081, BUS=4) TEMP.C

2.5 RUN-TIME CHECKS

Run-time checks detect and report run-time errors. The compiler by default does not generate code to perform run-time checks. If run-time checks are required, they can be turned on selectively or all at once on the command line by using the -a option on UNIX systems (/CHECK qualifier on VMS).

The -a option (/CHECK qualifier on VMS) causes all run-time checks to be performed. The full syntax for UNIX is:

-aflags

And for VMS:

/[NO]CHECK[=(
$$flags[,...]$$
)]

By adding flags, only specified checks are performed. Table 2-3 lists the flags for each run-time error check.

Table 2-3. Run-time Check Flags

UNIX	VMS	CHECK PERFORMED
р	PARAMETER	Intrinsic routines parameters
i	INDEX	Index exceeding array bounds
n	NIL_POINTER	Dereferencing through a pointer to the 0 address

An example for generating all checks in the UNIX environment is:

An example for generating only index and NIL pointer checks is:

When a run-time error occurs, a detailed message is displayed describing the file, error, and line at which the error occurred is displayed. The program terminates after the error information is displayed.

2.5.1 Parameter Check

The parameter check option generates code to check for incorrect parameter values on calls to intrinsic routines (see Chapter 9). The following calls are checked:

- mask1 and mask2 parameters in bitblt routines. The value of the actual mask1 and mask2 parameters in a call to a bitblt routine must be in the range of 0 to the maximum unsigned value of a word (65535).
- shift val parameter in bitwt routines. The value of the actual shift val parameter in a call to a bitwt routine must be in the range of 0 to 15.
- length in ext and ins routines. The value of the actual length parameter in calls to ext and ins routines must be in the range of 1 to 32.
- src addr and dest addr in extblt routines. The value of the actual src addr and dest addr parameters in calls to a extblt routine must be an even number.
- width in extblt routines. The value of the actual width in calls to a extblt routine must be an even number and a multiple of the value of the actual horiz inc parameter.
- horiz_inc in extblt routines. The value of the actual horiz_inc parameter in calls to a extblt routine must be either (+2) or (-2).

2.5.2 Array Checks

Each array index is checked to be within the array bounds (i.e. greater or equal to 0 and less then the array's dimension).

For example, the following code:

```
main() {
    int index,array[5];
    index = 6;
    array[index] = 1;
}
```

will result in run-time in the error message

```
"bad.c", line 5: value of 6 is out of bounds
```

NOTE: Index run-time checks are generated only for arrays whose dimensions are known during the compilation of the file.

2.5.3 NIL_POINTER Checks

Whenever a pointer is dereferenced, a check is performed for NIL pointers. If a NIL pointer is dereferenced, an error message results.

For example, the following code:

```
main() {
      char *ptr; = ((char *) 0);
      *ptr = 1;
}
```

results in the error message in run-time

```
"badptr.c", line 4: trying to dereference through a NIL
pointer
```

2.6 FLOATING-POINT EMULATION

Two different floating point emulation options are available with the GNX-Version 4 C Compiler: Hfp and fpee. Additional information, such as the difference between these options and the way they are implemented, can be found in Chapter 6 of the Series 32000 GNX-Version 4 Support Libraries Reference Manual. The use of the Hfp package is indicated by the -KFemulation compiler (/TARGET=(FPU=emulation) on VMS). The Hfp package may be used for cross configuration only. The use of the fpee package is indicated by the -f or -KFnofpu compiler option (/TARGET=(FPU=nofpu) on VMS). The fpee package may be used for cross configuration and for IEEE compatibility in native configuration. These options may also be set permanently by using the GTS facility.

2.6.1 Floating-point Emulation — Native Configuration

There is no way to unconfigure the FPU on the SYS32/50 and no floating-point emulation is therefore required. To use the fpee library you must do the following:

- 1. Include a call to the library routine fpinit at the beginning of the main module.
- 2. Include a -1fpe field after the source and object module in the "compile" command. For example,

```
cc file1.c -lfpe -lm
```

where file1.c is the input source file.

2.6.2 Floating-point Emulation — Cross-Configuration

In Cross-Configuration (UNIX system), floating-point emulation is achieved by using either the -f option on the nmcc invocation line or including a call to the INIT routine prior to any floating-point operations and explicitly linking files and libraries.

When -f is used on the nmcc invocation line the cross-compiler driver:

- · assumes there is no FPU on the target system
- assumes that the user wants to use floating-point emulation
- generates the correct command line and passes this to the linker

For example:

```
nmcc -f file1.c
```

The following is an example of explicitly linking files and libraries:

In cross host:

```
nmcc -c file1.c
nmeld GNXDIR/lib/fcrt0.o file1.o -lfpe -lm -lc
```

In native host (Series 32000/UNIX system):

```
nmcc -c file1.c
ld GNXDIR/lib/db fcrt0.o file1.o -ldb_c -ldb_fpe
```

2.6.3 Floating-Point Emulation — VAX/VMS System

Files and libraries must be explicitly linked to achieve floating-point emulation on a VAX/VMS system. This is a two-step process:

```
nmcc file1.c
nmeld gnxdir:fcrt0.obj, file1.obj, gnxdir:libfpe.a, gnxdir:libc.a
```

2.7 ENVIRONMENT VARIABLES (FOR UNIX ONLY)

On UNIX systems, in addition to the command line options, the compiler accepts several implicit options. These can be set through the environment variables CMDDIR, TMPDIR, LIBPATH, PITFILE, and INCLUDEPATH which are described below:

CMDDIR

The environment variable CMDDIR can be given the value of a directory name, in which the driver looks for the indirectly called programs (cpp, cc_fe, opt, etc.). For example, if CMDDIR = ''/usr/nsc/lib'', the driver will look for /usr/nsc/lib/cpp, /usr/nsc/lib/cc fe, etc.

TMPDIR

This environment variable redefines the location at which temporary files are created in the compilation process: Default is /tmp. This environment variable should be used on small systems with tiny /tmp partitions, which overflow when compiling huge files.

LIBPATH

The environment variable LIBPATH can be defined to contain one or more directories (separated by ":"). If LIBPATH is defined, then libraries will be taken from one of these directories. For example, if LIBPATH = /usr/mylib:/usr/vourlib. then libraries will be in either /usr/mylib or /usr/yourlib.

PITFILE

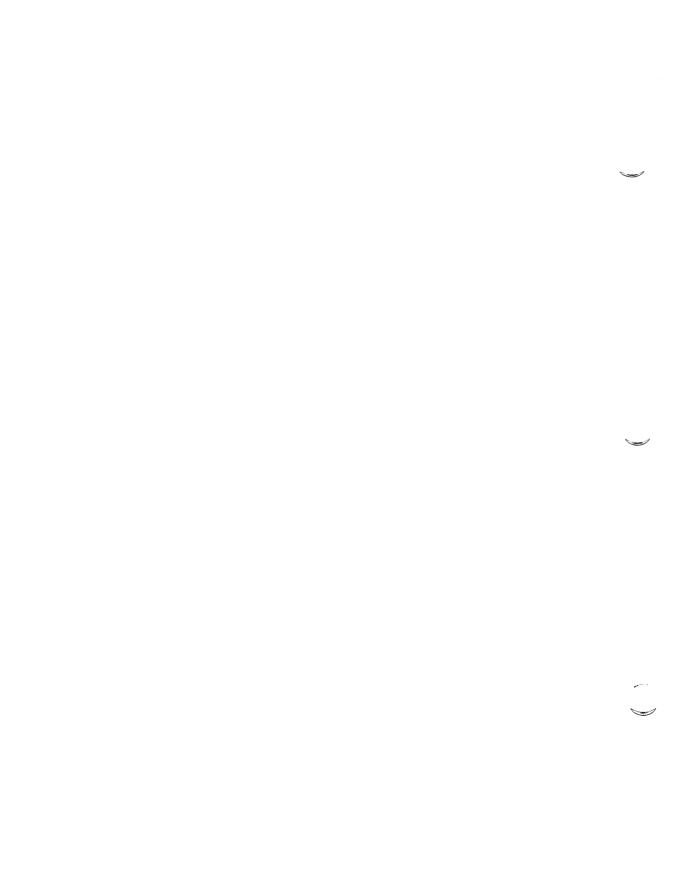
The environment variable PITFILE is used to redefine the default filename for profile information table file (PIT) used by sprof and the compiler. See Chapter 7 for more details.

INCLUDEPATH

If the INCLUDEPATH variable is defined (in a similar format as LIB-PATH), the standard include files (such as <stdio.h>) will be searched for in its directories

AVAIL_SWAP

The environment variable AVAIL_SWAP sets the maximum swap sapce of the optimizer in megabyte units. AVAIL_SWAP should be set to the number of megabytes to be used. See Section 6.7 for use of the AVAIL SWAP environment variable.



EXTENSIONS TO THE C LANGUAGE

3.1 INTRODUCTION

The GNX—Version 4 C compiler is based on the UNIX portable C compiler, pcc. All pcc extensions to the C language (as defined by Kernighan and Ritchie) are implemented by the GNX-Version 4 C compiler. In addition, the compiler includes two main types of extensions:

- 1. ANSI C features Most non pre-processor features of the ANSI C Standard are implemented.
- 2. Embedded support extensions Special features to assist programming embedded applications.

This chapter describes the extensions implemented by the GNX—Version 4 C compiler. Section 3.2 reviews the ANSI C extensions. Section 3.3 describes the embedded support extensions. All other extensions are presented in Section 3.4.

3.2 ANSI FEATURES

This section describes ANSI C features implemented in the GNX - Version 4 C compiler. For more details see C - A Reference Manual (second edition) by Harbison and Steele, and the ANSI C standard.

3.2.1 Function Prototypes

Function prototypes are fully implemented.

3.2.2 Volatile and Const Qualifiers

volatile and const type qualifiers are fully supported. See Appendix E for more details.

3.2.3 Void Data Type

The void data type is used as the type mark for a function that returns no result. It may also be used in any context where the value of an expression is discarded to explicitly indicate that a value is ignored. This is done by writing a cast to void.

The type void * is used for the generic pointer and is compatible with other pointer types.

3.2.4 Signed Keyword

The signed keyword is recognized by the compiler.

3.2.5 The #pragma Directive

The #pragma directive is recognized by the preprocessor and by the compiler. However, only the use of #pragma for interrupt/trap routines will be recognized by the compiler. Any other use of the #pragma directive will be ignored by the compiler.

3.2.6 Single-Precision Floating Constants

These floating constants allow the explicit specification of constants as single-precision in order to eliminate wasteful run-time conversions. This is accomplished by appending an f suffix to a float constant.

Example:

```
fmax += 17.0f
```

The same effect can be achieved by casting the constant to float, as in fmax += (float) 17.0;. Not using either the suffix or the cast results in both fmax and the value 17.0 being converted to double-precision for a double-precision addition; with the result being converted back to single-precision.

3.2.7 Unsigned Constants

Unsigned constants allow the explicit specification of unsigned constants. This is accomplished by appending a u suffix to a positive integer constant.

Example: "65u"

As with single-precision floating constants, unsigned constants eliminate wasteful run-time conversions.

3.2.8 Enumerated Types

Enumerated types as defined in ANSI C standard are fully supported. In addition, a warning is issued on assignment of different enumeration.

3.2.9 Structure Handling

The GNX—Version 4 C compiler implements the following improvements to structure handling:

- structure assignment
- structures as function arguments and return values
- reuse of structure and union member names
- initialization of first member of a union.
- initialization of auto storage class structures

Unlike initialization of automatic scalar variables, initialization of automatic variables is limited to initializers known at compile time.

3.2.10 Concatenation of Adjacent String Literals

According to the ANSI C standard, string literals that are adjacent tokens are concatenated into one character string literal.

For example the following code:

```
char s = "hello "
              "world";
printf(s);
```

prints the message:

```
hello world
```

3.2.11 Obsolesce of the Old Fashioned Compound Assignment

Since old fashioned compound assignment syntax is obsolete in ANSI C, it is no longer recognized by the GNX compiler.

For example, the following line:

```
int_var =+ 5; /* used to be equivalent to 'int_var += 5' */
```

is flagged as an error by the compiler.

3.2.12 Obsolesce of the Old Fashioned Initialization

Since the old fashioned initialization syntax is obsolete in ANSI C, it is no longer recognized by the GNX compiler.

For example, the following code:

```
int int_var 14;  /* used to be equivalent to 'int int_var = 14;'
*/
```

is flagged as an error by the compiler.

3.3 EMBEDDED SUPPORT EXTENSIONS

3.3.1 Interrupt/Trap Routines Support

As part of the embedded support, the GNX C compiler enables programming of trap and interrupt handlers in C. Handlers are defined as functions in the regular C syntax, preceded by a #pragma directive used to mark these functions as trap/interrupt handler routines.

Special code is produced by the compiler for the enter and exit sequence of routines marked as interrupt/trap handlers. This code is responsible for saving the proper registers (i.e. all registers used by the routine and scratch registers if the routine calls another routine) when entering an interrupt/trap routine. When the routine is exited, the saved registers are restored and RETI (for interrupts) or RETT (for traps) is performed (see the Series 32000 Programmer's Reference Manual for further details).

This section describes the syntax and semantics of writing interrupt/traps handlers in the GNX C Compiler. See Appendix E for more details.

INTERRUPT/TRAP HANDLER DEFINITION

The interrupt/trap handler is written as a regular C routine in the usual C function definition syntax. For example:

```
void hndlr foo(void)
        printf("division by zero");
        exit (1);
  };
```

The function is designated as an interrupt/trap handler in the following manner (the #pragma is used to mark an interrupt/trap handler).

Syntax for interrupts:

```
#pragma interrupt (function name [,save_regs={int_regs|all_regs}])
```

Syntax for traps:

```
#pragma trap (function name [,save_regs={int_regs|all_regs}])
```

function name is the name of the function to be marked as an interrupt/trap save regs can be either all regs (save all registers for general purpose and floating point), or int regs (save only general purpose registers).

In many applications the interrupt/trap handlers do not perform floating-point operations. In such applications there is no need to save the scratch floating point registers. The option save regs enables you to specify the register type to be saved (when the handler calls another routine). The default (if save_regs is omitted) is int_regs. Options different from all regs or int regs are considered errors.

Only the registers used in the interrupt/trap routine (and the scratch registers if the interrupt/trap calls another function) are saved.

A warning is issued by the compiler if a function is marked as an interrupt/trap handler using the #pragma directive, but no definition of the function was found in the compiled module.

Multiple #pragma directives with the same function name are considered errors, unless they are identical.

Restriction

The #pragma directive must appear before any declaration or definition of the function. The placement of the #pragma interrupt/trap in any other location results in an error message.

USING INTERRUPT/TRAP HANDLERS

It is your responsibility to install the address (or descriptor) of the interrupt/trap handler in the proper entry of the interrupt dispatch table (see the Series 32000 Development Board Monitor Reference Manual and the example presented in Appendix E for further information)

Calling an interrupt/trap handler directly from within the C code is not permitted. Any attempt to do so causes an error. This is because different instructions are used for returning from the interrupt/trap routine (RETI/RETT) and for returning from a regular routine (RET).

Attempts to call an interrupt/trap routine from within the C code is detected by the compiler only for calls in the same module in which the interrupt/trap routine was defined. All other calls are not detected by the compiler.

The default exception mode is direct-exception. In order to make non direct-exception mode possible, you should insert an asm statement ".module" before the function definition. It is your responsibility to do so. For more details see Appendix E.

3.3.2 Asm Keyword

The keyword asm is recognized for the insertion of assembly instructions directly into the generated instruction stream. The syntax is

asm(constant-string);

where *constant-string* is a double-quoted character string.

The keyword asm can be used within functions as a statement and outside of functions as a global declaration. A newline character will be appended to the given string without causing any change in the assembly code. See Appendix E for a detailed example.

3.3.3 Intrinsic Routines

The compiler enables the use of Series 32000/EP application specific instructions without the need of the asm keyword, by recognizing a set of intrinsic functions known internally to the C compiler. These intrinsic functions are used in the code as regular C functions, but are translated to an instruction sequence containing the special instructions and not to a function call. See Appendix E for more details.

3.4 OTHER EXTENSIONS

3.4.1 \$ Sign in Identifiers

The GNX—Version 4 C compiler allows the use of \$ signs in identifier names.

3.4.2 Bitfields

The GNX—Version 4 C compiler implements signed, unsigned, int, short, and char bitfields. Due to the Series 32000 architecture, the code for unsigned bitfields is more efficient than the code for signed bitfields.

3.4.3 Ident Preprocessor Command

A new cpp-style directive is recognized for placing strings into the .comment section of the object file. The syntax is

#ident constant-string

where constant-string is a double-quoted character string. The string is passed to the assembly file with a .ident directive and placed by the assembler in the .comment section of the object file.*

See the Series 32000 GNX - Version 4 COFF Programmer's Guide and the Series 32000 GNX -Version 4 Assembler Reference Manual for a description of the .comment section and the .ident directive.

IMPLEMENTATION ISSUES

4.1 INTRODUCTION

This chapter describes compiler implementation aspects which may differ from other compilers and which may affect code portability.

Portability issues are recognized by the C standard as issues that may differ from one implementation to another. The following two sections discuss portability issues. Section 4.2 defines how the GNX—Version 4 C compiler behaves under the listed issues. Section 4.3 lists issues that cause an undefined behavior of the GNX-Version 4 C compiler.

4.2 IMPLEMENTATION ASPECTS

The following cases are aspects of this implementation.

4.2.1 Memory Representation

• The representation of the various C types in this compiler are

С ТҮРЕ	SERIES 32000 DATE TYPE
int long short char float double	32-bit double-word 32-bit double-word 16-bit word 8-bit byte 32-bit single-precision floating-point 64-bit double-precision floating-point

- The set of values stored in a char object is signed.
- The padding and alignment of members of structures as described in Section 4.2.4.
- · A field of a structure can generally straddle storage unit boundaries.
- While signed bitfields are implemented, it is not recommended to use them since their implementation is slow. Bitfields are not allowed to straddle a double-word boundary.

4.2.2 External Linkage

- There is no limit to the number of characters in external names.
- Case distinctions are significant in an identifier with external linkage.

4.2.3 Types and Conversions

- A right shift of a signed integral type is arithmetic, i.e., the sign is maintained.
- When a negative floating-point number is converted to an integer, it is truncated to the nearest integer that is less than or equal to it in absolute value. The result is returned as a signed integer.
- When a double-precision entity is converted to a single-precision entity, it is converted to the nearest representation that will fit in a float with default rounding performed to the nearest value.
- The presence of a float operand in an operation not containing double operands causes a conversion of the other operand to float and the use of single-precision arithmetic. If double operands are present, conversion to double occurs.

4.2.4 Variable and Structure Alignment

The alignment of entities in a program is a trade-off issue. Most *Series 32000* CPUs are more efficient when dealing with entities aligned to a double-word boundary. This normally makes it necessary to have some amount of padding added to a program. This padding represents an overhead in storage space.

The GNX—Version 4 C compiler allows the user to tailor the alignment of structures/unions and their members and, independently, the alignment of other variables. Function parameters are always double-word aligned. This allows the calling of functions across modules without dealing with alignment issues.

Alignment of Variables

Extern, static, and auto variables are aligned in memory according to their size and the buswidth setting. Table 4-1 lists variable size, buswidth, and the alignment determined by these two parameters.

A buswidth setting of 1 means "align to 1 byte." Variables start on a byte boundary, in other words, there is no alignment and no padding. When allocating storage for variables, bytes are allocated sequentially with no padding between bytes.

Variables of size 1 are of the C type char, variables of size 2 are of the C type short, and variables of size 4 or greater are of the C types int, long, float, and double (size 8).

Table 4-1. Variable Alignment

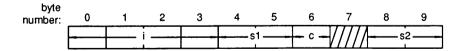
BUS WIDTH	VARIABLE SIZE (BYTES)			
BUS WIDTH	1	2	>= 4	
1	byte	byte	byte	
2	byte	word	word	
4	byte	word	double-word	

A buswidth setting of 2 means "align to an even byte." Variables that are larger than 1 byte start on a word boundary. This means that there may be padding of single bytes.

A buswidth setting of 4 means "align to a double-word boundary" (a byte whose address is divisible by four). Variables that are 2 bytes long start on a word boundary; variables that are 4 bytes or larger in size start on a double-word boundary. This means that there may be padding of up to three bytes.

Arrays are aligned as the alignment of their element type. Structures are aligned according to the alignment of the largest structure members. This is affected by the -J (/ALIGN) option. See "Structure/Union Alignment" and "Allocation of Bit-Fields" for more details.

Example: The arrangement of int i; short s1; char c; short s2; with a buswidth of 2 or 4 is



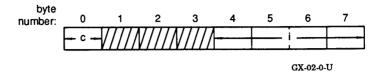
GX-01-0-U

Note that to align s2 to a word boundary, padding space of one byte is needed after c. This padding does not exist with a buswidth of 1.

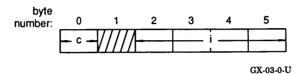
Example: The arrangement of

char c; int i;

with a buswidth of 4 is



With a buswidth of 2, the arrangement is



With a buswidth of 1, there is no padding.

It is important to note that the order in memory is the same as the declaration order only for extern and static variables. The optimizer may reorder auto variables in order to minimize padding space.

Fastest code is achieved by setting the default alignment to that of the data buswidth of the CPU (for all but the NS32008, the NS32CG16, the NS32FX16, the NS32CG160 and the NS32016). This can be accomplished by setting the BUS parameter in the target specification file, or by overwriting that file on the command line with the -KB (/TARGET) option.

Structure/Union Alignment

Structure members are aligned within the structure, relative to the beginning of the structure, in the same way that variables are aligned in memory. In order to maintain the alignment of the members relative to memory, the structure itself is aligned in memory according to the alignment of its largest members. This alignment may be controlled by putting -J (/ALIGN) on the command line.

In addition, the total size of a structure is such that it also ends on an alignment boundary of its largest member. This maintains the alignment of individual members in arrays of structures. This is illustrated in the FILE struct example at the end of this section.

For unions, there is no padding. The alignment of the union's largest members determine the alignment of the union itself.

Allocation of Bit-Fields

To understand the way bit-fields are handled, think of the situation where a field is fetched from memory. The number of bits fetched is determined by buswidth. For instance, if a bus is 2-bytes wide, then 2 bytes are fetched, even if only the first few bits are needed. For convenience, the number of bits fetched is called the "fetching unit."

Note that for the purpose of structure member alignment, the align switch value (1 byte, 2 bytes, or 4 bytes) is taken as a "virtual buswidth," even if it is different from the actual buswidth.

A complication exists when allocating bit-fields. The complication arises from the fact that different base types for bit-fields (char short, and int) are supported. The maximum length of a bit-field is the size of its base type; therefore, there may be times when a bit-field is larger than the buswidth. When the size of the base type is larger than the buswidth, the size of the fetching unit is considered to be the base-type size.

The precise rules for determining the start of the fetching unit are quite complicated. In general, it is determined by the current position in the allocation of structure members and by the base-type of the first bit-field in a group of consecutive bit-fields.

An attempt is made to pack consecutive bit-fields as much as possible, as long as the bit-fields remain in the same fetching unit. As soon as a field "spills over" into the next fetching unit, the alignment is set to the next memory unit (byte, word, or double-word, according to the align switch value and the base type of the field). A hole of padding bits remains, and the beginning of the spill-over field determines the start of a new fetching unit for following bit-fields. Using this method, bit-fields are packed as much as possible while still maintaining the alignment.

If, because of the bit-fields, the structure as a whole does not terminate on a byte boundary, padding bits are added to it to fill up to the end of the last byte it occupies. Additional padding bytes may be needed to fill to the alignment boundary of the largest structure member. This is seen in Figure 4-1. The bit-field does not quite reach the byte boundary; therefore, padding bits are added until the byte boundary is reached. Additional padding bytes are added to fill to the alignment boundary of the doubleword structure member.

```
Example:
            struct A {
              int i;
              unsigned bitfield: 4;
            } a:
```

The arrangement of a's fields in memory will be:

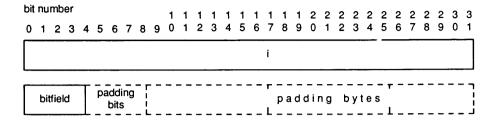


Figure 4-1. Bitfield Padding

Figure 4-2 is an example of the alignment on bit-fields given the different align switch settings. To summarize, the -J (/ALIGN) switch affects:

- the alignment and padding used for structure members and the alignment of variables of the structure type.
- the total storage allocated to a structure by determining if, and how many, padding bytes will be added after its last field.

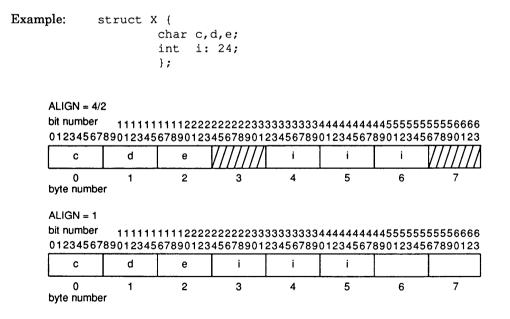


Figure 4-2. Alignment on Bitfields

CAUTION

The user must make sure that all parts of the program use the same alignment for the same structures; otherwise, problems result. The following example illustrates this point.

Suppose the example program includes "foo.h". The file "foo.h" contains the following definitions:

```
typedef struct {
         int
                          counter;
        unsigned char
                          *pointer;
        char
                          flag1;
        char
                          flag2;
} XXX;
extern XXX array[10];
```

Note that XXX has two char members at its end. If align=4, any variable declared to be of type XXX will have two padding bytes added at its end in order to make it occupy an integral number of double-words. When align=1 or align=2, no padding is performed.

If a module using "foo.h" is compiled with align=4 and later linked with a module compiled with align=1 or align=2 that tries to use array[n] where n > 0, the result will be wrong. This is because the two modules disagree on the size of the elements in the array.

The solution to this problem is to make sure all modules are compiled using either the same alignment setting or a revised header file that has been made insensitive to the setting of the alignment switch. This is performed by including the necessary padding to enforce equal sizes and offsets. If the latter solution is chosen, XXX is revised to look like:

```
typedef struct {
         int
                          counter;
         unsigned char
                          *pointer;
        char
                          flag1;
        char
                          flag2;
         short
                          padding;
} XXX;
```

No padding is added by the compiler, and the size of the structure is the same for all switch settings.

4.2.5 Structure Returning Functions

In the GNX—Version 4 C compiler, structure returning functions have a hidden argument which is the address of an area the size of the returned structure. This area is allocated by the caller and its address is passed as a first argument to the structure returning function. Structure returning functions are, therefore, re-entrant and interruptible.

NOTE: At the optimizer's discretion, small structures (less than 5 bytes) may be passed and/or returned in a register.

4.2.6 Calling Sequence

The standard Series 32000 calling conventions are used by the GNX-Version 4 C compiler for calls to external routines of all languages. It is, therefore, unnecessary to use the fortran keyword in C programs (if present, the keyword is ignored).

However, local or internal routines (functions which in C are preceded by the static keyword) are called by more efficient calling sequences.

The standard Series 32000 calling conventions are described in Appendix A.

NOTE: Code using the Series 32000 modularity features cannot be mixed with code not using those features. By default, the GNX—Version 4 tools assume no modularity.

4.2.7 Mixed-Language Programming

Mixed-language programs are frequently used for a couple of reasons. First, one language may be more convenient than another for certain tasks. Second, code sections already written in another language (e.g., an already existing library function) can be reused simply by calling them.

A programmer who wishes to mix several programming languages needs to be aware of subtle differences between the compilation of the various languages. Appendix B describes the issues one needs to be aware of when writing mixed-language programs and compiling and linking such programs successfully.

4.2.8 Order of Evaluation

The evaluation order of expressions and actual parameters in the GNX—Version 4 C compiler differs from those of other compilers. Therefore, programs that rely on a specific order of evaluation may not run correctly when compiled. In particular, the following orders of evaluation are unspecified:

- The order in which expressions are evaluated.
- The order in which function arguments are evaluated.
- The order in which side effects take place. For instance, a[i++] = i may be evaluated as

```
a[i] = i;
   i++;
or as
  t = i;
   i++;
   a[t] = i;
```

4.2.9 Order of Allocation of Memory

The order of allocation of local variables in memory is compiler-dependent. After the optimizer of the GNX—Version 4 C compiler performs register allocation, it reorders the local variables left in memory. This reordering reduces memory space requirements and minimizes displacement length. User programs that rely on any order of allocation of local variables may not run correctly. See Chapter 6.

4.2.10 Register Variables

By default, register variables, as well as other local variables, are equal candidates for register allocation. When given complete freedom, the optimizer generally performs a better job of register allocation than when forced to follow the programmer's allocation suggestions. For programs which make assumptions about variables which reside in specific registers, an optimization flag (-Ou or -O -Fu on UNIX and USER REGISTERS on VMS) is available to enforce the pcc allocation scheme for register variables of scalar types and of type double. See also Section 6.6.7.

4.2.11 Floating-Point Arithmetic

The floating-point arithmetic conversion rules of the GNX—Version 4 C compiler complies with the ANSI C standard and may differ from other C compilers.

In an operation not containing double operands, if one of two operands is of type float, the other operand is converted to type float and single-precision arithmetic is used. The result of the operation is of type float. Some other compilers perform such operations in double precision.

In old C compilers, the result of float-returning functions was actually returned in double format and placed in the F0-F1 register pair. When compiled by the GNX—Version 4 C compiler, such functions return the result in float format and place the result in the F0 register. Note that assembly programs that interface with float-returning functions may now incorrectly expect a double precision result.

Float parameters, however, are passed as double because the C language semantics do not require type identity between actual and formal parameters. Code is generated in the called function to convert these actual double values back to float if necessary.

Floating-point constants are of type double, unless they are typecast to float or are suffixed by the letter f or F. By preference, constants of type float should be used in float expressions to avoid the unnecessary casting of other operands to double precision. For example,

```
fmax += 17.5f;
```

is more efficient than

```
fmax += 17.5;
```

The following examples are of double constants and float constants.

Example: double constants float constants

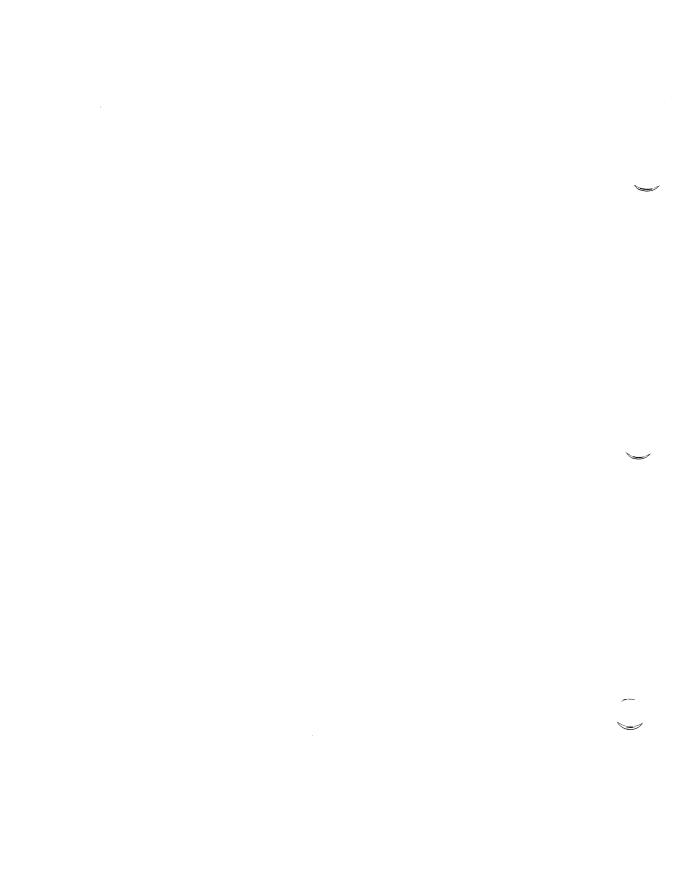
14.5e6 14.5e6f 14.5 (float) 14.5

4.3 UNDEFINED BEHAVIOR

In the following cases, the behavior of the GNX—Version 4 C compiler is undefined:

- The value of a floating-point or integer constant is not representable.
- An arithmetic conversion produces a result that cannot be represented in the space provided.
- A volatile object is referred to by means of a pointer to a type without the volatile attribute.

- An arithmetic operation is invalid, such as division by 0, or produces a result that cannot be represented in the space provided, such as overflow or underflow.
- A member of a union object is accessed using a member of a different type.
- An object is assigned to an overlapping object.
- The value of a register variable has been changed between a setimp call and a longjmp call.



OPTIMIZATION TECHNIQUES

5.1 INTRODUCTION

The main difference between the GNX—Version 4 C Compiler and other compilers is the optimizer. Recompiling and optimizing with the GNX—Version 4 C Compiler will result in a 10 percent to 200 percent speedup for most programs, with the mean above 30 percent.

This chapter describes some of the advanced optimization techniques used by the GNX-Version 4 C Compiler to improve speed or save space. The most important techniques are:

- Value propagation
- Constant folding
- Redundant assignment elimination
- · Partial redundancy elimination
- Common subexpression elimination
- Flow optimizations
- Loop unrolling
- Dead-code removal
- Loop-invariant code motion
- Strength reduction
- Induction variable elimination
- · Register-allocation by coloring
- Peephole optimizations
- · Memory-layout optimizations
- · Fixed frame
- Runtime feedback optimization

The following sections describe these techniques in more detail. For coding suggestions and other practical guidelines on how to make best use of the optimizing aspects of the compiler, see Chapter 6.

5.2 THE OPTIMIZER

The optimizer, shared by all the GNX — Version 4 Compilers, is based on advanced optimization theory, developed over the past 15 years. Central to the optimizer is an innovative global-data-flow-analysis technique which simplifies the optimizer's implementation. It allows the optimizer to perform some unique optimizations in addition to standard optimizations found in other compilers. Optimizations are performed globally on the code of a whole procedure at a time and not just in a local context.

The optimizer can be regarded as a multi-step process. Each step performs its particular optimizations and provides new opportunities for the optimizations of the next step.

STEP ONE

The first step in the optimization process is to read in the source program one procedure at a time and to partition this procedure into basic blocks. A basic block is a straight line sequence of code with a branch only at the entry or exit. Some of the optimizations performed during this step are:

Value Propagation

Value propagation (or copy propagation) is the attempt to replace a variable with the most recent value that has been assigned to it. This optimization is primarily useful in the special case of constant propagation. It is important because it creates opportunities for other optimizations. Value propagation can be turned off by the CODE MOTION optimization flag (-Om on UNIX systems).

Constant Folding

If an expression or condition consists of constants only, it is evaluated by the optimizer into one constant, thereby avoiding this computation at run-time. The optimizer, using algebraic properties such as the commutative, associative and distributive law, sometimes rearranges expressions to allow constant folding of part of an expression.

The GNX—Version 4 C Compiler also folds floating-point constant expressions. This feature can be turned off using the NOFLOAT FOLD option (-Oc on UNIX systems) of the optimizer.

Redundant Assignment Elimination

The optimizer detects and eliminates assignments to variables which are not used later in the program or which are assigned again before being used. This optimization can often be applied as a result of value propagation.

Value propagation, constant folding, and redundant assignment elimination are illustrated in Figure 5-1.

```
The program sequence
                         a = 4;
                         if (a*8 < 0) b = 15;
                         else b = 20;
            ... code which uses b but not a ...
is translated by the GNX—Version 4 C Compiler front end into the
following intermediate code
                         if (a*8 >= 0) goto L1
                         b ← 15
                         goto L2
                    L1: b \leftarrow 20
                    L2: ...
which is transformed by "value propagation" into
                         a ← 4
                         if (4*8 >= 0) goto L1
                         b ← 15
                         goto L2
                    L1: b ← 20
                    L2: ...
which after "constant folding" becomes
                         a ← 4
                         if (true) goto L1
                         b ← 15
                         goto L2
                    L1: b \leftarrow 20
                    L2: ...
"dead code removal" results in
                         a ← 4
                         goto L1
                    L1: b \leftarrow 20
                    L2: ...
which is transformed by another "flow optimization" into
                         a \leftarrow 4
                         b ← 20
Since there is no further use of a, a \leftarrow 4 is a "redundant assignment:"
                         b ← 20
                         . . .
```

Figure 5-1. Relationship Between Various Optimizations

STEP TWO

The second step in the optimization process is the construction of the program's "flow graph." This is a graph in which each node represents a basic block. As mentioned in STEP ONE, a basic block is a linear segment of code with only one entry point and one exit point. If there is a path in the program that leads from one basic block to another, then an "arrow" is drawn in the graph to represent this path.

Figure 5-2 illustrates a flow graph, representing an "if-then-else" sequence.

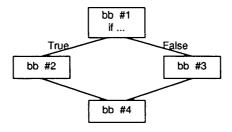


Figure 5-2. Flow Graph

During the construction of the flow graph, additional optimizations can be performed:

Flow Optimizations

Flow optimizations reduce the number of branches performed in the program. One example is to replace a branch whose target is another branch with a direct branch to the ultimate target. This often makes the second branch redundant. At other times code is reordered to eliminate unnecessary branches. Branches to "return" are replaced by the return-sequence itself.

Loop Unrolling

Loop unrolling duplicates the body of a loop. This reduces the number of times that the loop control code is executed. Loop unrolling improves performance by reducing the number of increment, comparison, and branch instructions.

This technique is particularly useful for small loops whose iteration control constitutes a significant part of loop execution time. However, because loop unrolling does involve the duplication of the loop's body, more space is needed.

When a loop is unrolled, it is replaced by code with the following structure:

- 1. Pre-Loop Code Checks whether to enter the unrolled loop body or to branch to the tail-code.
- 2. Unrolled Loop Body The loop body duplicated a number of times.
- 3. Tail-Code Performs the remaining iterations.

Based on a loop's code size, the optimizer determines whether to perform loop unrolling, and if so how many times. An example of loop unrolling is shown in Figure 5-3.

```
Code sequence for initialization of an array:
    foo(int j)
      int i;
      int a[100];
      for (i=j; i < 100; i++) {
            a[i] = 0;
         }
    }
When unrolled 5 times, is equivalent to the following code sequence:
    foo(int j)
      int i;
      int a[100];
      /* pre-loop code */
      i = j;
      if (i > 95) goto TAIL_CODE;
      /* unrolled loop body */
      for (; i < 95; i += 5) {
            a[i] = 0;
            a[i+1] = 0;
            a[i+2] = 0;
            a[i+3] = 0;
            a[i+4] = 0;
         }
      TAIL CODE:
      /* tail code */
      for (; i < 100; i++)
            a[i] = 0;
    }
```

Figure 5-3. Example of Loop Unrolling

Dead Code Removal

Flow optimizations are also designed to help the optimizer discover code which will never actually be executed. Removal of this code, called "dead code removal," results in smaller object programs.

STEP THREE

Step three of the optimization process is called "global-data-flow-analysis." It identifies

OPTIMIZATION TECHNIQUES 5-5

desirable global code transformations which speed program execution. Many of these concentrate on speeding up loop execution, since most programs spend 90 percent or more of their time in loops. Global-data-flow-analysis is the computation of a large number of properties for each expression in the procedure.

Unlike most optimizers, which employ unrelated and separate techniques, the optimizer centers around one innovative technique which involves the recognition of a situation called "partial redundancy." This technique is so powerful that many other optimizations turn out to be special cases. The central idea is that it is wasteful to compute an expression, say a*b, twice on the same path; it is often faster to save the result of the first computation and then replace the fully redundant second computation with the saved value. More common, however, is the case in which an expression is partially redundant; there is one path to an expression, which already contains a computation of that expression, but another path to that same expression does not.

The following optimizations are performed by a common technique:

• Elimination of Fully Redundant Expressions

This optimization is often called "Common Subexpression Elimination." It is relatively simple to avoid the recomputation of fully redundant expressions. The optimizer saves the result of the first computation (usually in a register variable) and uses the saved value in place of the second computation. Performance-conscious programmers sometimes do this themselves, but many cases, such as array index and structure member calculations, are recognized only by the optimizer.

• Partial Redundancy Elimination

A partially redundant expression can be eliminated in two steps. First, insert the expression on the paths in which it previously did not occur; this makes the expression fully redundant. Second, save the first computations and use the save value to replace the redundant computation. An example of this optimization is shown in Figure 5-4.

Partial redundancy elimination sometimes results in slightly larger code, but execution is not harmed, since all inserted expressions are in parallel and only one is actually executed.

Loop Invariant Code Motion

If an expression occurs within a loop and its value does not change throughout that loop, it is called "loop invariant." Loop invariant expressions are also partially redundant. This can be understood by realizing that there are two paths into the loop body: one is through the loop entry (the first time the loop is executed), and the other is from the end of the loop, while the exit condition is false. Loop invariant computations are, therefore, removed from the loop in the same way: the expression is first inserted on the entry path to the loop, and then is saved on the entry path in a register, while the redundant computation in the loop is replaced by that register.

Strength Reduction

This optimization globally replaces complex operations with simpler ones. This is

In the following code, a*b is "partially redundant" (computed twice only if C is true):

It is first transformed into a "fully redundant" expression

if
$$C = 1$$

 $x \leftarrow a*b$
else
 $b \leftarrow b+10$
 $temp \leftarrow a*b$
 $y \leftarrow a*b$

Then, as in the simple case of "redundant expression elimination," this is reduced to

```
if C = 1
      temp \leftarrow a*b
      x \leftarrow temp
      b \leftarrow b+10
      temp ← a*b
y \leftarrow temp
```

Now, the expression a*b is computed only once on any path.

Figure 5-4. Example of Partial Redundancy Elimination

primarily useful for reducing complex array-subscript computations (involving multiplication into simpler additions).

```
Example:
            static int a[15]; for (i=0; i<15; i+=1)
                 a[i] = 1;
            is transformed into:
            for (i=0, p=a; i<15; i+=1, p+=4)
                 *p = 1;
```

Induction Variable Elimination

Induction variables are variables which maintain a fixed relation to other variables. The use of such variables can often be replaced by a simple transformation. For instance, the example given for strength reduction can be reduced to the following:

```
for (p=a; p<a+60; p+=4)
*p = 1;
```

STEP FOUR

The fourth optimization step performed by the optimizer, and possibly the most profitable, is the "register allocation" phase. Register allocation places variables in machine registers instead of main memory. References to a register are always much faster and use less code space than respective memory references.

The algorithm used by the optimizer is called the "coloring algorithm." First, global-flow-analysis is performed to determine the different live ranges of variables within the procedure. A live range is the program path along which a variable has a particular value. Generally, an assignment to a variable starts a new live range; this live range terminates with the last use of that assigned value.

The optimizer subsequently constructs a graph as follows: each node represents a live range; two nodes are connected if there exists a point in the program in which the two live ranges intersect. The allocation of registers to live ranges is now the same as coloring the nodes of the graph so that two connected nodes have different colors. This is a classic problem from graph theory, for which good solutions exist. If there are not enough registers, more frequently used variables have higher priority than less frequently used ones. Loop nesting is taken into account when calculating the frequency of use, meaning that variables used inside of loops have higher priority than those that are not.

Most optimizing compilers attempt register allocation only for true local variables, for which there is no danger of "aliasing." An alias occurs when there are two different ways to access a variable. This can happen when a global variable is passed as a reference parameter; the variable can be accessed through its global name, or through the parameter alias. A common case in C is when the address of a variable is assigned to a pointer.

The optimizer takes a more general approach by considering all variables with appropriate data types as candidates for register allocation, including global variables, variables whose addresses have been taken, array elements, and items pointed to by pointers. These special candidates cannot reside in registers across procedure calls and pointer references and, therefore, normally have lower priority than local variables. However, instead of completely disqualifying the special candidates in advance, the decision is made by the coloring algorithm.

Additional important optimizations performed by the register allocator are:

• Use of Safe and Scratch Registers

The Series 32000 machine registers are, by convention, divided into two groups: registers R0 through R2 and F0 through F3, the so-called "scratch" registers which can be used as temporaries but whose values may be changed by a procedure call, and the "safe" registers (R3 through R7 and F4 through F7) which are

guaranteed to retain their value across procedure calls. The register allocator spends a special effort to maximize the use of scratch registers, since it is not necessary to save these upon entry or restore them upon exit from the current procedure. The use of scratch registers, therefore, reduces the overhead of procedure calls.

Register Parameter Allocation

The register allocator attempts to detect routines, whose parameters can be passed in registers. This is possible for static routines only, since by definition all the calls to such routines are visible to the optimizer. Calls to other (externally callable) routines are subject to the standard Series 32000 calling sequence. Passing parameters in registers is another way to reduce the overhead of procedure calls.

STEP FIVE

The last optimization step consolidates the results of all previous steps by writing out the optimized procedure in intermediate form for the separate code generator. Some reorganization takes place during this step. Local variables which have been allocated in registers are removed from the procedure's activation record (frame), which is reordered to minimize overall frame size.

5.3 THE CODE GENERATOR

The back end (code generator) attempts to match expression trees with optimal code sequences. It applies standard techniques to minimize the use of temporary registers, which are necessary for the computation of the subexpressions of a tree. The main strength of the code generator lies in the number of "peephole optimizations" it performs.

Peephole optimizations are machine-dependent code transformations that are performed by the code generator on small sequences of machine code just before emitting the code. Some of the most important peephole transformations are listed below:

- The code for maintaining the frame of routines which have no local variables, or whose variables are all allocated in registers, is removed.
- Case statements are optimized into binary search, linear search or table-indexed code (using the Series 32000 CASE instruction), in order to obtain optimal code in each situation.
- The stack and frame areas are always aligned for minimal data fetches.
- Reduction of arithmetic identities, i.e., x*1 = x, x+0 = x, etc.
- Use of the ADDR instruction instead of ADD of three operands.
- Some optimizations performed in the optimizer, such as the application of the distributive law of algebra, i.e., (10+i)*4 = 40+4*i, provide additional opportunities to the code generator to fully exploit the Series 32000's addressing modes.

- Use of ADDR instead of MOVZBD of small constant.
- Strength Reduction Optimizations. Use of MOVD instead of MOVF from memory to memory; use of index addressing mode instead of multiplication by 2, 4 or 8; use of combinations of ADDR instructions or shift and ADD sequences instead of multiplication by other constants up to 200.
- Fixed Frame Optimization. An important contribution of the code generator is its ability to precompute the stack requirements of a procedure in advance. This allows the generation of code which does not use (nor update) the FP (frame pointer), resulting in cheaper calling sequences.

This optimization is most useful when the procedure contains many procedure calls because it is not necessary to execute code to adjust the stack after every call. Parameters are moved to the pre-allocated space instead of pushing them on to the stack using the top-of-stack addressing mode. Note that when using this optimization, the run-time stack pointer stays the same throughout the procedure, and all references to local variables are relative to it and not to the FP. Also note that since parameters are not pushed on to the stack, the evaluation order of parameters is not defined solely by their original order.

While most optimizations are beneficial for both speed and space, some optimizations favor one over the other. The default setting of the optimizer switch favors speed over space in trade-off situations. The following are the effects of favoring space over speed (by an optimization flag):

- · Code is not aligned after branches.
- All returns within the code are replaced by a jump to a common return sequence.
- Certain space-expensive peephole transformations are not performed.

5.4 MEMORY LAYOUT OPTIMIZATIONS

The following memory layout optimizations are performed by the GNX—Version 4 C Compiler:

- Frame variables that are allocated in registers are removed from the frame.
- Internal, static routines whose parameters are passed in registers have smaller frames.
- The stack alignment is always maintained. Stack parameters are passed in aligned positions.
- Frame variables are allocated in aligned positions. The optimizer reorders these variables to save overall frame space.
- Code is aligned after every unconditional jump.

5.5 RUNTIME FEEDBACK

The optimizer has normally no way to determine the actual runtime behavior of a program. What looks like a loop may in reality never be executed. The GNX-Version 4 C Compiler has an option to create a statistic record of a program's execution path. This execution profile can then be used in a subsequent optimization pass of the same program, to improve the optimizer's heuristic algorithms. This technique, call runtime feedback optimization, effects mainly the following optimizations:

- · Loop Unrolling
- · Register Allocation

For more details on runtime feedback optimization see Section 7.4.

GUIDELINES ON USING THE OPTIMIZER

6.1 INTRODUCTION

The following sections are provided as guidelines on using the GNX-Version 4 C Compiler. Experienced programmers should understand this compiler's optimization techniques in order to:

- Learn how to port programs to the GNX—Version 4 C Compiler.
- Understand how to recognize and avoid nonportable code.
- Avoid using programming tricks that rely on the way ordinary compilers generate code.
- · Avoid performing "hand optimizations" that the optimizer does anyway.
- Avoid writing code that may prevent certain optimizations.
- · Understand how to select the different command line optimization flags to achieve optimal performance.

Please read Chapter 5 for a complete description of the optimization techniques.

6.2 OPTIMIZATION FLAGS

Optimization options available to the user are listed in Table 6-1. Default options are marked by (*).

Table 6-1. Optimization Options

UNIX	VMS	DESCRIPTION
0	NOOPT	does not invoke the optimizer phase.
В	RUNTIME_FEEDBACK	performs runtime feedback optimization
* b	NORUNTIME_FEEDBACK	does not perform runtime feedback optimization
С	NOFLOAT_FOLD	does not compute floating-point constant expressions at compile time.
* C	FLOAT_FOLD	performs floating-point constant folding.
* F	FIXED_FRAME	uses fixed frame references, avoids use of the FP register or the Series 32000 ENTER/EXIT instruction.
f	NOFIXED_FRAME	compiles for debugging: uses slower FP and TOS addressing modes.
I	NOVOLATILE	applies all optimizations to all variables (including global variables).
i	VOLATILE	compiles system code: assumes that all global and static memory variables and pointer dereferences are volatile.
L	STANDARD_LIBRARIES	assumes use of standard run-time library
1	NO STANDARD_LIBRARIES	assumes that all routines have corrupting side effects.
* M	CODE_MOTION	performs global code motion optimizations.
m	NOCODE_MOTION	does not perform global code motion optimizations.
N	LOOP_UNROLLING	performs loop-unrolling optimizations.
* n	NOLOOP_UNROLLING	does not perform loop-unrolling optimizations.
U	NOUSER_REGISTERS	ignores user register declarations.
u	USER_REGISTERS	allocates user-declared register variables in registers as done by pcc.
* R	REGISTER_ALLOCATION	performs the register allocation pass of the optimizer.
r	NOREGISTER_ALLOCATION	does not perform the register allocation pass of the optimizer.
* S	SPEED_OVER_SPACE	optimizes for speed only.
s	NOSPEED_OVER_SPACE	does not waste space in favor of speed.
1-9		maximal memory/swap-space available is 1 through 9 Mbytes (default: 4 Mbytes).

6.2.1 Optimization Options on the Command Line — UNIX Systems

The -O option enables the optimizer. Specifying -O on the command line results in the fastest possible code without undue increase in code size. (-ObCFIUMnLRS). In special cases, such as when compiling operating system code, there may be a need to further refine the optimization phase by specifying optimization flags. Individual optimization flags can be specified either by using the -F option or by simply appending them to -o. Table 6-2 lists reasons why a particular default option might be changed.

Even when the optimizer pass is omitted, some local optimizations are performed by the code generator. Note that specifying the compiler debug option (-g) on the command line automatically turns off the optimizer fixed frame flag (-OF), unless otherwise specified by the user.

Also note that using the compiler target option (-KB1) favors space over speed by saving alignment holes normally produced when the bus width is the default (-KB4).

6.2.2 Optimization Options on the Command Line — VMS Systems

The fastest possible code, without undue increase in code size, is generated by specifying /OPTIMIZE on the command line. This is equivalent to entering:

```
/OPTIMIZE=(FIXED FRAME, CODE MOTION, REGISTER ALLOCATION, FLOAT FOLD,
SPEED OVER SPACE, NOVOLATILE, STANDARD LIBRARIES, NOUSER REGISTERS,
NOLOOP UNROLLING, NORUNTIME FEEDBACK)
```

In special cases, such as when compiling operating system code, there may be a need to further refine the optimization phase by specifying optimization flags. Table 6-2 lists reasons why a particular default option might be changed.

Even when the optimizer pass is omitted, some local optimizations are performed by the code generator. Therefore, specifying /NOOPTIMIZE (which is the default for this qualifier) is equivalent to entering:

```
NOOPT, NOFIXED FRAME, NOCODE MOTION, NOREGISTER ALLOCATION,
/OPTIMIZE=(
             NOFLOAT FOLD, SPEED OVER SPACE, NOVOLATILE,
             NOSTANDARD_LIBRARIES, USER REGISTERS, NOLOOP UNROLLING,
             NORUNTIME FEEDBACK)
```

Note that specifying the compiler debug option (/DEBUG) on the command line automatically turns off the optimizer fixed frame option (FIXED FRAME), unless otherwise specified by the user.

Also note that using the compiler option /TARGET=(BUSWIDTH=1) favors space over speed by saving alignment holes normally produced when the bus width is the default (BUSWIDTH=4).

6.2.3 Changing Default Optimization Options

There is normally no reason to turn off any of the optimization options; the default produces the best results, see Table 6-2. Refer to Chapters 2 and 5 for more on optimization options.

with a crospt her whole the control of the control

ory to market

OPTION	REASON FOR CHANGING OPTION	SEE ALSO
NOFIXED_FRAME (-Of) -O g pats argumen \$ 2 was 5 Fr Son net fully	to debug the program or to compile nonport- able programs that assume knowledge of the run-time stack.	6.3.4, 6.4
VOLATILE (-Oi)	to compile system programs, such as device drivers, which contain variables that change or are referenced spontaneously.	6.3.2
NO_STANDARD_LIBRARIES(-01)	to compile programs which re-implement standard functions, in a way which does not agree with the optimizers assumptions (i.e., have side effects).	6.3.5
NOFLOAT_FOLD (-Oc)	to compile programs whose correct execution depends on the order in which floating-point expressions are evaluated.	6.3.6
NOCODE_MOTION (-Om)	to compile programs which contain huge func- tions, which are a drain on the system's resources and are time consuming to optimize.	
LOOP_UNROLLING (-ON)	to compile program segments containing tight loops which are most often executed.	6.6.9
USER_REGISTERS (-Ou)	to compile programs which rely on the register allocation scheme of pcc.	6.6.7
NOREGISTER_ALLOCATION (-Or)	to run programs that cease to work when performing register allocation.	6.6.7
NOSPEED_OVER_SPACE (-Os)	to compile programs which must fit as tightly as possible in memory.	6.6.9
NOOPT (-00 or use -Fflags without giving -0)	when the optimizer phase is not required and another flag needs to be turned off as well.	6.6.10
NORUNTIME_FEEDBACK (-Ob or use -Fflags without giving -O)	when run-time feedback is required to achieve better optimization results based on the typi- cal behavior of the program.	6.6.11

6.3 PORTING EXISTING C PROGRAMS

Almost every program which runs when compiled by other C Compilers, will compile and run on the GNX—Version 4 C Compiler without any changes in the sources. However, there might be a few programs which will cease to work in the same manner as before, when compiled by the GNX—Version 4 C Compiler. There might be other programs, which seem to work when compiled without the optimizer, but which cease to work when optimized. The following sections describe some of the reasons for this phenomenon.

6.3.1 Undetected Program Errors

The single most common reason for a nonfunctioning program is an undetected program error, which becomes apparent only when compiling under a different compiler or only when optimizing. Many of these errors result from the fact that the program author relied on the way the compiler compiled, and thereby created a program which is clearly nonportable.

The following partial list points out some of the most common problems:

· Uninitialized local variables.

Since the memory and register allocation algorithms of the GNX—Version 4 C Compiler are very different from those of other compilers, a local variable may wind up in a completely different place. For example, a programmer may fail to initialize a local variable, with the assumption that, upon program start, it would certainly contain zero. This may become false as a result of the register allocation phase of the GNX—Version 4 C Compiler.

· Relying on memory allocation

One cannot assume that if two variables are declared in a certain order, they will actually be allocated in that order. A program that uses address calculations to proceed from one declared variable to another declared variable might not work.

· Failing to declare a function

A char returning function will return a value in the lower-order byte of R0, without affecting the other bytes. A failure to declare that function where it is used, might result in an error. For instance, assuming that get_code() is defined to return a char, then

```
main() {
   int i;
   if ((i = get_code()) == 17)      do_something();
}
```

might never execute do_something even if get_code returns 17 since the whole register is compared to 17, not just the low-order byte.

A similar problem exists for functions which return short or float, or those which return a structure.

6.3.2 Compiling System Code

System code is distinguished from general "high-level" code, by the fact that it is machine-dependent, often contains real-time aspects and interspersed asm statements, and is often driven by asynchronous events, such as interrupts. Examples of system code are interrupt routines, device handlers and kernel code. From the optimizer's point of view, ordinary looking global variables can actually be semaphores or memory-mapped I/O, that can be affected by external events, which are not under the optimizer's control. Even so, it is still possible to optimize such code, by taking some precaution, and by activating some special optimization flags. Some of these aspects are discussed in the following sections.

Volatile variables

Volatile variables are variables, which might be used or changed by asynchronous events, such as I/O or interrupts. The /OPTIMIZE=VOLATILE (-Oi under the UNIX operating system) qualifier treats all global variables, static variables, and pointer dereferences as volatile, which means that they are not subject to any optimizations. As a result, the number and nature of memory references to them will not change. Remember that individual identifiers can be declared as volatile by using volatile type qualifiers. The following examples demonstrate the consequences of volatile variables and pointer dereferences.

Examples: 1. x = 17; x = 18;

If $\,\times\,$ is volatile, both of the two assignments to $\,\times\,$ are executed even though the first one seems redundant.

2.
$$x = 9;$$

 $y = x + 1;$

If x is volatile, this program segment is not optimized to y = 10;

3.
$$*p = b + c;$$

if *p is volatile, then this results in

The difference stems from the fact that the second sequence, though faster, makes two references to O(p) when the programmer may have wanted only one.

6.3.3 Timing Assumptions

Optimizing a program changes the timing of various constructs. In particular, delay-loops might now run faster than before.

6.3.4 Low-Level Interface

· Relying on register order

A program that relies on the fact that a given register variable resides in a specific register must be compiled with the /OPTIMIZE=USER_REGISTERS flag (-Ou on UNIX systems) turned on (see Section 6.6.7).

Relying on frame structure

A program, that relies on a specific frame structure, must be compiled with the FIXED_FRAME flag turned off (-Of on UNIX systems). This includes, in particular, programs that use the standard alloca() function (which allocates space on the user's frame).

Referring to variables on the frame of a different function (such as the caller of this function) by complex pointer arithmetic may also cease to work. See Appendix A for more details.

• Using asm statements

The code inserted by asm statements may cease to work because the surrounding code produced by the GNX—Version 4 C Compiler will normally be different from another compiler's code. See Section 6.6.6.

6.3.5 Using Nonstandard Library Routines

The GNX—Version 4 C Compiler assumes by default that all the C standard mathematical library routines listed in Table 6-3 are available as a standard run-time library. These library routines have absolutely no access to global variables. Therefore, calls to these routines are specially recognized and marked as calls which do not disturb optimizations of the global variables of the program. The global library variable errno is treated as volatile, so no references to it will be optimized. This is normally a safe assumption since it is unusual for a program to redefine (and thereby hide) these standard routines. The functions abs, fabs, and ffabs actually compile into in-line code and do not generate a procedure call at all.

In addition, a set of intrinsic routines known internally to the computer are supported. See Chapter 8.

The compiler generates a warning message whenever it compiles a program which does redefine one of these routines. In this case the user must decide whether the redefined routine's behavior is consistent with the previously mentioned assumption of the optimizer. If it is not, the user has the choice of renaming the redefined routine (so that calls to it are not specially recognized), or of using the NOSTANDARD_LIBRARY flag (-Ol on UNIX), which turns off the recognition of all library routines.

Table 6-3. Recognized Library Routines

abs	acos	facos	asin	fasin	atan	fatan
atan2	fatan2	cabs	fcabs	ceil	fceil	cos
fcos	cosh	fcosh	erf	ferf	erfc	ferfc
exp	fexp	fabs	ffabs	fmod		
ffmod	fmodf	ffmodf	frexp	gamma	hypot	fhypot
j0	ქ1	jn	ldexp	log	flog	log10
flog10	modf	pow	fpow	sin	fsin	sinh
fsinh	sqrt	fsqrt	tan	ftan	tanh	ftanh
у0	y1	yn				

6.3.6 Reliance on Naive Algebraic Relations

Since the optimizer performs floating-point constant folding, i.e., it rearranges expressions to evaluate constant subexpressions at compile time, some naive algebraic expressions are folded away.

```
Example:
              do {
                       a = a * 2;
              while ((a + 1.0) - 1.0 == a);
              is optimized to
              do {
                        a = a \times 2;
              while (1);
```

which was not the programmer's intention.

To maintain the program and keep the programmer's original intention, the programmer should use the NOFLOAT FOLD (-Oc on UNIX systems) optimization flag to suppress the folding optimization.

6.4 DEBUGGING OF OPTIMIZED CODE

Most of the time, the user should not need to debug an optimized program. The majority of all bugs can be found before optimization is turned on. However, there are some rare bugs which make their appearance only when the optimizer is introduced, bugs that are difficult to find without a debugger.

The problem is that code motion optimizations and register allocation obsolete most of the symbolic debugging information generated by the compiler. The following "rules of thumb" can be employed when using symbolic debug information together with the optimizer:

- Line number information is correct, but the code performed at the specified lines may be different from non-optimized code as a result of various code motion optimizations, such as moving loop invariant expressions out of loops.
- Symbolic information for global variables is normally correct, since global variables are rarely put in registers. In particular, if a global variable is not referenced within the current procedure, the value in memory is valid and the symbolic information is correct.
- Symbolic information for parameters is correct except in the following two cases:
 - 1. When a parameter is allocated a register and there is an assignment to that parameter, the symbolic information is incorrect.
 - 2. When a parameter of a local procedure is passed in a register as a result of an optimization, the symbolic information is incorrect. In this case, the symbolic information of all other parameters is incorrect because their offset within the procedure's frame is changed.
- Symbolic information of local variables is likely to be incorrect because most of the local variables are put in registers; the rest of the local variables are reordered into new frame locations, or "optimized out".
- Note that if symbolic information is requested, then slightly different code is generated. This happens because the optimizing flag FIXED FRAME (-OF on UNIX systems) is automatically disabled when the /DEBUG (-g qualifier on UNIX systems) is used. Specifically, the ENTER instruction is always generated at the entry of procedures, and frame variables are referenced by FP-relative rather than SP relative addressing mode. Without disabling this flag, symbolic debugging is almost impossible.

It is helpful to have an assembly listing of the program in question which has been compiled with the /ASM (-S on UNIX systems) and the /ANNOTATE (-n on UNIX systems) qualifiers. Such a listing contains comments from the optimizer regarding its actions (see Section 6.5).

6.5 IMPROVED ANNOTATION

The GNX C compiler has a unique annotation feature which helps in the debugging of an optimized code.

Upon invocation with the -n and -S flag (/ANNOTATE and /ASM on VMS), the compiler emits the source lines into the assembly code as comments (see Section 2.3). In addition the GNX optimizer emits annotated comments explaining its actions.

Example:

The following code accumulates the first n elements of array a into the global accumulator acc. n resides in register r4.

```
for(i = 0; i < n; i++)
acc += a[i];
```

The optimizer may generate the following annotated code

```
\#--- for(i = 0; i < n; i++)
   # temp initialized to &a
                 $(1),r2
      movad
                  (0) + a, r1
      movd
   # load (moved up) acc to r3
      movd
                 acc,r3
.LL2:
#--- acc += a[i];
                  0(r1), r3
      addd
      addqd
                  $(1),r2
   # temp = temp + 4 (temp incremented)
      addqd
                  $(4),r1
      cmpd
                  r2, r4
                  .LL2
      blt
   # store (moved down) r3 to acc
      movd
                  r3, acc
```

The actions taken by the optimizer can be inferred from the comments:

```
"load (moved up) acc to r3" (the value of the variable acc is first loaded into the register r3)
```

[&]quot;temp initialized to &a (temp initialized)" (while performing strength reduction optimization, a compiler

pointer is initialized by the first element's address)

- "temp = temp + 4 (temp incremented)" (the temporary pointer allocated by the optimizer is updated to point to the next element at the end of each iteration)
- "store (moved down) r3 to acc" (the value of acc is updated)

6.6 ADDITIONAL GUIDELINES FOR IMPROVING CODE QUALITY

Using some of the following programming guidelines results in programs that take advantage of the GNX—Version 4 C Compiler optimizations.

6.6.1 Static Functions

It is not only good software engineering practice, but also good optimization practice to declare all functions not called from outside the file as "static." This allows the optimizer to use a more efficient internal calling sequence upon calls to such routines. This internal calling sequence uses the BSR instruction instead of the JSR or CXP instruction and also passes parameters in registers rather than on the stack.

If a program consists of a single file and this is discovered by the GNX—Version 4 C Compiler (by indicating compilation and linking in one step), then all functions within that file are automatically considered static by the compiler, resulting in the same advantages.

6.6.2 Integer Variables

Many operators, including index calculations, are defined in C to operate on integers, and imply a conversion when given non-integer operands. Therefore, to avoid frequent run-time conversions from char or short to int, integer variables, particularly variables which serve as array-indices, should be defined as type int and not short or char.

6.6.3 Local Variables

Local variables should be used as much as possible, particularly when they are employed as loop counters or array indices, as they have a better chance of being placed in registers.

6.6.4 Floating-Point Computations

In programs which do not require double-precision floating-point computations, a significant run-time improvement can be achieved by paying attention to the following points:

- define all functions as returning float type, not double
- define all constants to be 'float' using the f suffix or cast expressions explicitly to float
- use the single precision version of the standard floating-point routines such as fabs() instead of abs(), fsin() instead of sin(), etc.

6.6.5 Pointer Usage

The following terms are used throughout this section.

potential definition

A statement potentially defines a memory location if the execution of the statement may change the contents of that memory location.

Example:

- 1. A call to a function potentially defines all global variables, because their value may change during the execution of that function.
- 2. Imagine the following code fragment.

The assignment statement potentially defines the memory location *q, because q may point to the same memory location as p. The location *p is defined, *i.e.*, given a new value, by the assignment. About location *q, it can only be said that it may be changed, hence the potential definition.

· potential use

A statement "potentially uses" a memory location if it *may* reference (read from) that memory location.

· address taken variable

A variable is considered "address taken" if the address operator (&) is applied to it within the file or if the variable is a global variable that is visible by other files.

volatile/non-volatile registers

By convention, the registers are divided into volatile registers (registers R0 through

R2 and F0 through F3) and non-volatile registers (registers R3 through R7 and F4 through F7). Volatile registers may be changed by a procedure call, whereas non-volatile registers are guaranteed to retain their value across procedure calls. Therefore, all non-volatile registers used within a procedure have to be saved at the entry and restored at the exit of that procedure.

The optimizer does not keep track of the contents of pointers; therefore, it cannot tell, for any given location in the program, where each pointer is pointing.

Since a pointer can point at any memory location, the optimizer makes the following assumptions concerning pointer usage:

- 1. Every assignment to a "pointer dereference," the location pointed to by a pointer, potentially defines all other pointer dereferences and all address-taken variables.
- 2. Every use of a pointer dereference (i.e, a value read through a pointer) potentially uses all other pointer dereferences and all address-taken variables. This is because any accessible memory location is potentially read.
- 3. Every function call potentially defines and potentially uses all pointer dereferences, all address taken-variables, and all global variables. This is so because that function's code may, using pointers, read and/or write any accessible memory location. Of course, any global variable may be used and/or changed.

It is advisable to keep these assumptions in mind when using pointers. In particular, using arrays is preferable to using pointers. The following example illustrates this point. Assume a is an array of char and p is a pointer to char. The two program segments perform the same function.

Example: program segment 1

```
for (i = 0 ; i != 10 ; i++) {
    a[i] = global_var; a[i+1] = global_var + 1;
}

program segment 2

for (p = &a[0] ; p != &a[10] ; p++) {
    *p = global_var; *(p+1) = global_var + 1;
}
```

In program segment 1, global_var can be put in a register. In program segment 2, however, p may point to global_var. The first statement (*p = global_var) potentially defines global var; therefore, it cannot be put in a register.

Another aspect of this same issue is that of *common subexpressions*. The optimizer normally recognizes multiple uses of the same expression and saves that expression in a temporary variable (usually a register). This cannot be performed when worst case

assumptions are made about potential definition of expressions (as described in the previous section). Expressions that contain pointer dereferences or global variables are vulnerable; therefore, if many uses of the same expression span across procedure calls, it is advisable to save them in local variables. In the following example:

```
foo1 (p -> x);
foo2 (p -> x);
```

The expression $p \rightarrow x$ cannot be recognized by the optimizer as a common subexpression because fool() may change its value. The following hand optimization may help:

```
t = p \rightarrow x; /* t is local, therefore */
foo1(t); /* not potentially defined by foo1() */
foo2(t); /* so its value is still valid for foo2()
*/
```

The programmer is using his knowledge that $p \to x$ is not changed by fool() to make this optimization. The optimizer cannot do the same because it assumes the worst case.

6.6.6 Asm Statements

Extreme care should be taken if using asm statements. If using asm statements, remember the following:

- The optimizer is not aware of the contents of an asm statement. Therefore, it assumes that an asm statement potentially defines and potentially uses all of the variables (including local variables). This means that no common subexpressions can be recognized across an asm statement.
- In order to allow an asm statement to use a specific register (e.g., asm ("save [r0,r1,r2]");), the optimizer de-allocates all the registers.
- The compiler usually generates code which is different from the code generated by other compilers. This applies particularly to allocation of local variables and parameters of static procedures.
- The code surrounding the asm statement may change as a result of changes in other parts of the procedure.
- An asm statement that contains a branch instruction or a branch target (label) may cause the optimizer to generate wrong code.

For the above mentioned reasons, it is strongly advised to look at the generated assembly code before and after inserting asm statements into a program.

6.6.7 Register Allocation

The C language is unique in that it allows the programmer to specify (or rather recommend) that some variables be allocated to machine registers. The optimizer normally ignores these recommendations, since it turns out that in most cases the optimizer's own register allocation algorithms are as good as or superior to the programmer's recommendations. There are several reasons for this fact:

- The user can use a register for one variable only. The optimizer however allocates a register along live ranges of variables, making it possible for several variables with non-conflicting live ranges to use the same register.
- The user can declare as a register, only local variables whose addresses are not taken; whereas, the optimizer allocates global variables as well as variables whose addresses are taken (where possible).
- The user can allocate variables in safe registers only. Therefore, every register which is used has to be saved/restored at the entry/exit of the procedure. The optimizer allocates variables that do not live across procedure calls in unsafe registers. Therefore, these registers need not be saved/restored.
- Because of code motion optimizations, the number of references of variables may be changed. Therefore, the choice of register variables may not be optimal. In the following example:

```
int
     j;
register int i;
i = j;
if (i == 3 || i == 4 || i == 5)
```

undesired effects result if optimized with the /USER REGISTERS flag (-Ou on UNIX systems) The reason is that j is copy propagated and replaces all occurrences of i. Therefore, i occupies a register for nothing, while j may end up in memory (because either the ordinary register allocation of the optimizer is not invoked or there are no registers left for i).

6.6.8 setimp()

Calls to set jmp() are specially recognized by the compiler. Procedures that contain calls to set jmp() are only partially optimized because procedure calls may end up in a call to longimp(). Code motion optimizations are performed only within linear code sequences (those sequences not containing branches or branch targets). Register allocation is limited to optimizer generated temporary variables, register declared variables, and variables whose live ranges do not contain function calls.

6.6.9 Optimizing for Space

The default behavior of the GNX—Version 4 C Compiler optimizes for optimal speed without undue increase in code size. There are several things that can be done to improve code density:

- Optimize with the NOSPEED_OVER_SPACE flag on (-Os on UNIX systems).
- Squeeze the data area by using smaller alignment between variables, *i.e.*, /TARGET=(BUS=1) on VMS systems or -KB1 on UNIX systems.
- Do not use loop-unrolling optimization.

The optimizer has certain heuristics based on size considerations, whether to perform loop-unrolling. These heuristics also take into account the relevant on-chip caches. Nevertheless, if code density is important, it is advisable not to use the loop-unrolling optimization.

• Squeeze all structure definitions by using the /ALIGN=1 switch (-J1 on UNIX systems). See Section 4.2.4.

6.6.10 Using /NOOPT (-Oo) option

The /OPTIMIZE=NOOPT (-Oo) compiler option is used when the optimizer phase is not required and another flag needs to be turned off as well. For instance, /OPTIMIZE=(NOOPT, FIXED_FRAME) (-OoF) turns on fixed frame without running the optimizer, while /OPTIMIZE=(NOOPT, FIXED_FRAME) (-Of) turns off fixed frame but runs the optimizer.

With /OPTIMIZE=NOOPT (-00) by itself, the optimizer is not run, but the codegenerator performs all its optimizations (see Section 5.3).

6.6.11 Runtime Feedback Optimization

In the runtime feedback optimization the optimizer makes use of runtime information collected during a previous runs of the program in order to better predict the program's run-time behavior. As a result, optimizations done by the compiler will be based on assumptions closer to the real behavior of the program. Runtime feedback optimization effects mainly the following optimizations:

- Loop Unrolling
- Register Allocation

Runtime optimization is invoked by using the <code>-OB</code> (/OPT=RUNTIME_FEEDBACK) compiler option. The PIT file in the current directory (or the one specified in the PITFILE environment variable global symbol) is assumed to contain the data for the runtime feedback option.

NOTE: The runtime information used for the runtime feedback optimization should represent the real run-time behavior of the program. Otherwise, the optimizer relies on false assumptions. This may actually cause degradation in code quality.

For more details on runtime feedback optimization see Section 7.4.

6.7 COMPILATION TIME REQUIREMENTS

Using the optimizer slows down the compilation process. It is therefore recommended to use the optimizer only on final production versions of a program. The amounts of resources (time and memory) vary strongly from program to program and actually depend on the size of the routines in the compiled program file. The larger a routine, the more time and memory needed to optimize it. This behavior is more or less quadratic, the optimizer needs about four times the resources to optimize a routine of 1000 lines than to optimize a routine of 500 lines.

If time or memory requirements are unacceptable and routines cannot be reduced to reasonable (500 lines) size, it is possible to turn off some optimizations, using the NOCODE_MOTION (-Om on UNIX systems) and/or the NOREGISTER_ALLOCATION (-Or on UNIX systems) flags.

On UNIX host systems, an optimization flag (-0number) is available to set an upper limit on the memory requirements of the optimizer to a certain number of megabytes. This can be useful on host systems without virtual memory or with a limited swap-space configuration. If necessary, the optimizer then skips certain optimizations on huge routines only, in order to stay under the chosen limit. In such cases, an appropriate message is given. This flag is only necessary when compiling modules with extremely large procedures (over 500 lines in a single procedure), a case when the optimizer may need a larger swap space than the one currently available. For instance,

-02

limits the optimizer to 2 Mbytes of swap space.

An alternate method for setting an upper limit on memory requirements, on native systems, is to use the environment variable AVAIL_SWAP, which sets the maximum swap space requirement of the optimizer in megabyte units. This environment variable should be set to the number of megabytes to be used. The user can choose from 1 Mbyte to 16 Mbytes. If the user's choice is outside of these parameters, the default value of 4 Mbytes is chosen. For instance,

setenv AVAIL SWAP 2

makes 2 Mbytes of swap space the default. This can be overridden using the previously described -0number option.

PROFILE INFORMATION

7.1 INTRODUCTION

Profile information is statistical data about the run-time behavior of a program. Such information can be gathered by compiling the program using the -B option on UNIX (/GATHER on VMS), and executing the compiled program with typical inputs. Each execution of the compiled program results in the accumulation of profile information in a special file. Profile information is used by the optimizer and the tool sprof.

The optimizer can use profile information to achieve better code optimization. Code can be recompiled using the compiler option -OB on UNIX (/OPT=RUNTIME FEEDBACK on VMS).

sprof processes profile information and displays it as an annotated source file. You can use sprof's output to:

- Pinpoint the most often executed sections of program code in order to determine areas for concentrated hand optimization.
- Test the expected relative frequency of execution of different code sections.
- Provide indication of test coverage.
- Discover bugs by spotting unexpected execution of code lines.

7.2 GATHERING PROFILE INFORMATION

7.2.1 The Profile Information

Profile information is gathered for each basic block (see Section 5.2). A precise trace of every basic block execution is recorded. From this information the execution rate of each line can be deduced.

7.2.2 Code Compilation

When compiling a program on which we want to gather profile information additional code is generated by the compiler. Also a profile information tables (PIT) file is created. When the program is executed, the additional code that was generated accumulates profile information internally and adds it to the PIT file before the program exits.

The extra code is generated by the compiler and the PIT file is created by the tool pgen after the linking phase.

The following operations are performed:

- Allocation of buffer space in the .bss section (uninitialized data). The buffer space
 is used for basic block execution counters, which keep track of the number of
 times each basic block is executed.
- 2. Insertion of extra code at the beginning of each basic block. This code increments the proper basic block execution counter each time a block is executed.
- 3. Generation of additional symbolics. Additional symbolics are generated in order to map the basic blocks in the executable file to source lines.
- 4. Linking of the program with the object file pfb_exit.o (pfb_exit.obj on VMS and db_pfb_exit.o for cross compilation on a Series 32000/UNIX system). This file includes a customized version of the standard library _exit routine, which accumulates the internally accumulated profiling information into the PIT file at the end of each execution.
- Creation of the PIT file. The tool pgen is invoked to create and initialize the PIT file.

See section 7.2.4 for more information on the PIT file and section 7.2.3 for more information on pgen.

NOTE: Profile information cannot be gathered on an optimized program.

The optimizer is not to be invoked together with the profile information gathering option.

7.2.3 Pgen

The tool pgen reads the executable file and generates the profile information table (PIT) file that is used by sprof and the compiler (see Section 7.2.4).

On UNIX systems pgen is automatically invoked by the driver when called with the -B flag (Section 7.2.6). On VMS systems pgen must be invoked separately after linking the program compiled with the /GATHER qualifier (Section 7.2.7).

7.2.4 The PIT File

The PIT (Profile Information Tables) file contains accumulated profile information.

The PIT file is created by pgen (see Section 7.2.3) and modified by the additional code in the pfb exit object file (Section 7.2.5) just before execution is completed.

The PIT file is used by sprof and by the profile feedback option. PIT is the default name for the file in which profile information accumulates. In order to override this default name, the environment variable (logical name on VMS) PITFILE is used. If PITFILE is set during profile gathering, information is accumulated to a file bearing this name.

7.2.5 The pfb exit.o(pfb exit.obj) File

This file must be linked with the profiled program in order to enable profileinformation accumulation. The pfb exit.o (pfb exit.obj or pfb exit.obj on VMS) file includes a customized version of the exit routine of the C runtime-library, which accumulates the profile information into the PIT file at the end of program execution.

In the UNIX environment, linkage with pfb exit.o is performed automatically by the compiler when compiling with the -B option (see Section 7.2.6). In the VMS environment, linking with pfb exit.obj must be done by the user (see Section 7.2.7).

- NOTES: 1. For cross compilation on a Series 32000/UNIX system, the file is named db pfb exit.o.
 - In native environment, pfb_exit.o uses the standard I/O library of libc for writing the PIT file.
 In the cross environment virtual I/O facilities of the cross C library, which are based on debugger and monitor support, are used.
 - 3. For modular compilation (-X option on UNIX, /MODULAR on VMS) a special version of the pfb_exit.o exists. On native UNIX environment it is called Xdb_pfb_exit.o, on cross configuration UNIX environment Xpfb_exit.o, on VMS environment Xpfb exit.obj.

7.2.6 Compilation in the UNIX Environment

The syntax for compilation to gather profile information on the UNIX environment is:

```
nmcc -B[<pitfile>] filename (cross-support configuration)
cc -B[<pitfile>] filename (native configuration)
```

<pitfile> is the name of the PIT file to be created. If <pitfile> is omitted, the
default name PIT is given. Note that there should be no space between -B and
<pitfile>.

The compiler driver automatically calls all the necessary subprograms when invoked with the -B option. This includes linking with the special pfb_exit.o file (Section 7.2.5) and calling pgen (Section 7.2.3) after linking.

7.2.7 Compilation in the VMS Environment

The syntax for compilation to gather profile information on the VMS environment consists of three steps:

1. Compilation. Use the /GATHER compiler qualifier in the syntax:

```
nmcc /GATHER my module.c
```

2. Link the program with pfb exit.obj. Use the syntax:

```
nmeld gnxdir:crt0.obj, gnxdir:pfb_exit.obj,-
my_module.obj, ...,gnxdir:libc.a
```

The pfb exit.obj file must appear before the standard libraries.

For modular compilation (/MODULAR) linkage is done with an object file named Xpfb exit.obj.

3. Run the tool pgen to create the PIT file. Use the syntax

```
pgen <executable name> [<PIT file name>]
```

If <PIT file name> is not specified, it will be named PIT by default.

7.2.8 Code Execution

The profiled program can be executed repeatedly with any desired inputs. Profile information from each execution accumulates in the PIT file.

7.2.9 Disabling Profile Information Accumulation

Profiling information will not be accumulated under any of the following conditions:

- No PIT file exists
- Read or create permission for the PIT file is denied
- The PIT file is not in the expected format
- The PIT file and executable file are incompatible
- The cumulative number of executions of a certain basic block in the profiled program exceeds the maximum count limit, which equals the maximum unsigned long integer minus one (decimal 4294967295).

These conditions cause only the accumulation of profiling information to be disabled, and do not affect the normal operation or semantics of the profiled program.

The PIT file can be removed or renamed temporarily in order to disable the accumulation of profiling information.

7.2.10 Redefining Standard libc Symbols

The following standard libc symbols (routines and variables) are used by the profile gathering code:

```
cleanup
close
errno
etext
exit
```

```
fclose
fgets
fopen
fprintf
fputs
getenv
 iob
mktemp
           (link and unlink in native SYS-V systems) (strrchr in native SYS-V systems)
rename
rindex
sprintf
sscanf
sys_errlist
sys_nerr
unlink
```

Redefining any of these symbols can cause unexpected results.

7.2.11 Execution Time Considerations

The additional code produced in each basic block for gathering profile information may slow down a CPU-bound program by a factor of 20%-30% (without taking I/O into account).

On cross systems where the loading of a program and I/O operations are on slow serial lines, use of sprof may slow down execution significantly. This is because I/O operation will be performed during the accumulation of the profiling information (just before exiting).

7.2.12 Space Considerations

The additional code and space needed the PIT file adds approximately 20%-30% to the original code and uninitialized data size.

Additional symbolic information is also produced. However this symbolic information occupies only disk space and is not loaded into memory (since it is not a component of real code or data).

7.3 SPROF - THE GNX SOURCE PROFILER

The sprof profiler provides high-level language information about the runtime behavior of a program. The profile consists of an annotated listing of the source file. A number is printed at the beginning of each line to indicate the number of times that line was executed. sprof is supplied as part of the GNX—Version 4 C compiler package, and runs on all cross and native GNX supported environments.

Unlike other profilers, sprof's provides information on the basic block level (see Section 5.2). This means that sprof does not sample the program periodically, but instead gathers a precise trace of every basic block execution. From this information the execution rate of each line can be deduced.

sprof does not provide either timing information about a program or function caller/callee statistics. Rather, sprof gives an exact count of source-line executions. A standard profiler, such as the UNIX profiler prof (supported by the compiler in native environment), can be used to collect timing or caller/callee information.

7.3.1 Example

```
main()
        int i;
        for (i = -6; i \le 7; i++)
if (i \ge 0)
                                                                    line
14
               printf("factorial(%d) = %d0,i,fact(i));
   int fact(n)
int n;
                                                                   *line
        if (n == 0)
            return 1;
                                                                   *line
                                                                     ine
             return (n * fact(n - 1));
```

Figure 7-1. Example of sprof Output

The number "1" annotating the first basic block on line 5 indicates that the main program was executed once. The "14" and "8" annotations of lines 6 and 7 indicate that the body of the loop in main was executed 14 times, of which only 8 resulted in a call to the function fact.

It can be deduced that the function fact did not returned implicitly (i.e. without using an explicit return), as shown by the "0" annotation of line 17. In contrast, the function main did return implicitly, as can be seen from the "1" annotation besides the closing brace of function main (line 8).

The runtime behavior of the function fact is also illustrated. fact was called 36 times. Only 8 of these invocations were from main, therefore the remaining 28 calls were recursive. This is further shown by the annotation of "28" on line 16.

7.3.2 Running SPROF

After program execution and accumulation of profile information in the PIT file, sprof can be called to process and present profile information. Various options are available to control the output (see Section 7.3.3).

In order to process profile information the following files must exist:

- The source file(s)
- · The executable file
- The PIT file

Figure 7-2 illustrates the flow of data through sprof. nnm is the standard GNX utility for displaying the symbol table of a COFF object file (see the GNX Commands and Operations Manual).

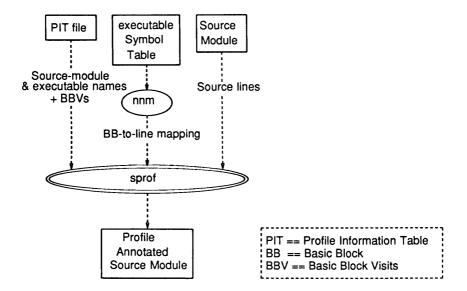


Figure 7-2. sprof Data Flow Description

7.3.3 SPROF Invocation

The tool sprof is supplied as a standard part of the GNX-Version 4 C compiler package both on UNIX and VMS.

SYNTAX (UNIX):

```
sprof [-d source_dir] [-e exec_file ]
        [-p pit_file] [-o output_file]
[-f[fmc<margin width>]] [source file ...]
```

SYNTAX (VMS):

```
SPROF [/DIRECTORY=source dir] [/EXECUTABLE=iexec file ]
       [/PITFILE=pit file] [/OUTPUT=output file]
       [/FORMAT=([format_option [, ...] ]
       source file ...]
```

Where:

The directory where the source file is located. The default is the source dir current directory.

The name of the executable file. The default is the executable exec file name as found in the PIT file.

pit file The name of the PIT file to be used. The default name is PIT or the value of the environment variable PITFILE.

The name of the output file to be generated. The default is stanoutput file dard output (SYS\$OUTPUT on VMS).

The name of the source file to be profiled. The default is all source file source files in the executable file that were compiled with the -B option.

The -f options (format option on VMS) are:

- f (FORMFEED on VMS)
 - output a FORM-FEED character between output source files
- m (MARK MARGIN on VMS)
 - mark the MARGIN of the source with vertical bar "!" characters.
- c (COMPACT on VMS)

print counts COMPACT (the count for sequential basic blocks is only printed if different from the previous basic block). The default is to print counts for every basic block.

<margin-width> (WIDTH=margin width on VMS)

The width reserved for printing profiling counts. The default is 8. A negative value will left-justify the counts.

CAUTION

Compatibility of the PIT file and the program source file is determined only by the creation date of the PIT file. Therefore, a source file with the same name, and an older date than the PIT file, but with contents different from the program source file, will not be recognized by sprof as being incompatible with the PIT file. Such a situation will result in an incorrect sprof output.

7.3.4 Counts and Basic Blocks

sprof prints basic block counts according to the following rules:

- 1. Print a count only for lines which start a basic block. (No count will be printed for lines which consist of declarations or calls to a cpp macro defined to nothing.)
- 2. Print a count only for the first basic block of a group that is mapped to the same line.

Example:

```
100 if (a < 5) j = 3;
```

The count in the profiled line is the number of times the first basic block in this line (i.e if (a<5)) was executed. However, the number of times the condition was true can not be deduced from this output. This information is provided if the code is written in the form

100 if
$$(a < 5)$$

27 $\dot{j} = 3;$

7.4 RUNTIME FEEDBACK OPTIMIZATION

The runtime feedback optimization option is used to enable the compiler to tune optimizations based on statistical information on typical run-time behavior.

The optimization algorithms used by the optimizer are based on assumptions and heuristics. However, run-time behavior may be different. In such a case, the compiler can achieve better optimization by operating under a different set of assumptions as suggested by the profile information.

For example, the following optimizations can be improved by using the profile feedback option:

- Register allocation Usually the optimizer heuristics used to determine register allocation are based on loop nesting and conditional execution. By using the profile information, register allocation is based on an better estimate of the use of variables
- · Loop unrolling Loop unrolling optimization enlarges code size. Therefore it is worthwhile to optimize only those loops which are entered many times. The profile information provides improved estimation of where this optimization should be performed

For more details on these optimizations see Section 5.2.

The runtime feedback mechanism is divided into two phases, described in the following two sections

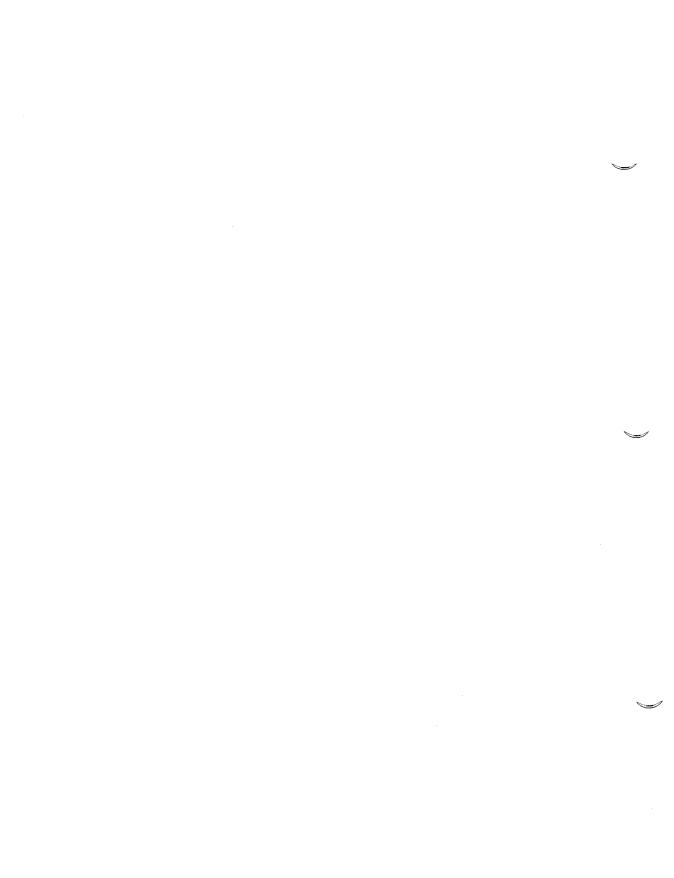
7.4.1 Profile Information Gathering

This phase involves compilation and program execution. It is described in detail in Section 7.2. The profile information is collected in the PIT file (see Section 7.2.4).

NOTE: The profile information gathered in the PIT file must represent the true run-time behavior of the program. Otherwise false assumptions are made by the optimizer and recompilation can cause degradation in program performance.

7.4.2 Runtime Feedback Compilation

The runtime feedback compilation is invoked using the -OB (/RUNTIME FEEDBACK on VMS) compiler option. The PIT file in the current directory, or the one specified in the PITFILE environment variable (logical name on VMS) is assumed to contain the data for the runtime feedback option.



INTRINSIC FUNCTIONS

8.1 INTRODUCTION

The GNX C compiler generally uses in its code selection the most efficient instructions from the Series 32000 instruction set. There are, however, certain instructions which have no natural matching C construct or that are not fully utilized by the C language. Such instructions are in particular the Application Specific Instruction Set (ASIS) of the NS32CG16, NS32CG160, NS32FX16 and NS32GX320 microprocessors (see the Series 32000/EP Embedded Processor Databook) and certain standard Series 32000 instructions. The GNX C compiler provides access to these instructions by means of intrinsic functions. The NS32CG160, NS32FX16 and NS32CG16 CPUs share the same core instructions. These microprocessors will be referred to as the CG-Core throughout this chapter. A familiarity with ASIS instructions is assumed.

In order to use intrinsic functions, include the appropriate GNX header file cg16.h, qx320.h or ns32000.h, prior to any call to the function. The function call must contain a parameter list compatible with the formal parameter list of the prototype. Otherwise it is considered a redeclaration of the function, and a proper warning message is issued by the compiler.

Redefining an intrinsic function, i.e., defining a different function with the same name as an intrinsic function, results in an error. However, it is possible to use intrinsic function names for different functions by specifying the -Fl or -Ol option (NO_STANDARD_LIBRARIES on VMS). (See Section 6.3.5).

Special compilation options generate various run-time checks for flagging improper values of parameters in calls to intrinsic functions. The compilation options and the run-time checks performed are described in Chapter 2, Section 2.5.

8.1.1 Using Intrinsic Functions

Intrinsic functions are known internally to the compiler. The syntax used for invoking intrinsic functions is the same as for regular C functions. However, no function call will appear in the code generated by the compiler. Instead, an instruction sequence containing the application specific instruction will be produced. The optimizer will attempt to allocate parameters in the registers as needed by the instruction.

NOTE: Unlike regular functions, taking the address of an intrinsic routine is not permitted. Any attempt to do so will results in an error message.

There is one intrinsic function for each supported special instruction. Generally the function's name is the instruction assembly mnemonic, with a leading underscore. This conforms to the ANSI C convention of global identifiers with leading underscores being reserved for implementation. The parameters and the result type of each function are described in full ANSI C prototype format in special GNX header files. Currently there are three such header files:

- ns32000.h
 For the general Series 32000 instructions.
- cg16.h
 For the CG-Core Application Specific instructions.
- gx320.h
 For the NS32GX320 Application Specific instructions.

A complete description of each function is given in Sections 8.3, 8.4 and 8.2, respectively. Examples of using the functions are provided in Appendix E.

8.2 General Series 32000 Intrinsic Functions

This section describes the special *Series 32000* instructions accessible by use of intrinsic functions. Supported instructions are divided into:

Single bits

Instructions that refer to a single bit in memory. They enable efficient setting (sbit), clearing (cbit), inverting (ibit) and testing (tbit) of a single bit. Single bit instructions find the first set bit in a given byte, word or double-word (ffs), and calculate an absolute bit address (cvtp).

Bit-Fields

Instructions that manipulate a consecutive group of bits in memory. They include ext, which extracts a bit-field, and ins, which inserts a bit-field.

Absolute value

Integer absolute value (abs) and floating-point absolute value (fabs and ffabs).

All definitions given in this section are supplied as part of the GNX C compiler package in the file ns32000.h.

8.2.1 Single Bit Instructions

PROTOTYPE

```
typedef int _xbit(int offset,
                int *base);
xbit cbit;
            /* clear bit */
_xbit _ibit; /* invert bit */
xbit _sbit; /* set bit
                           */
xbit tbit;
            /* test bit
```

DESCRIPTION

The cbit, ibit and sbit functions operate on the bit at offset bits from base. The return value is the original value of the specified bit.

```
_cbit
           clears the bit to zero.
           inverts the bit — zero becomes one and vice versa.
ibit
_{\mathtt{sbit}}
           sets the bit to one.
```

The tbit function returns the value of the bit residing at offset from base.

The appropriate instruction version (byte word or double word) is selected by the compiler according to the type of parameter passed for offset. Calling the functions as a procedure, such as

```
(void)_ibit(...
```

will prevent the compiler from producing the code reading the original value of the bit.

8.2.2 _ffs (Find First Set)

PROTOTYPE

DESCRIPTION

The ffs routines search for the first set bit in base. The search starts at the offset specified in the int pointed by offset, and proceeds in ascending order to the first set bit or to the last bit in base.

The routines return the value of the PSR F flag. If a set bit is found, the value of the int pointed by offset is changed to the bit number of the bit found, and the F flag in the PSR is cleared. If no set bit is found, the int pointed by offset is set to zero and the F flag in the PSR is set to 1.

Calling the functions as a procedure, such as

```
(void) ffsw(...
```

will prevent the compiler from producing the code reading the PSR F flag.

NOTE: The value of the int pointed by offset must be within tha range 0 to 7 (in ffsb routine), 0 to 15 (in ffsw routine) and 0 to 31 (in ffsd routine). Compilation with the -aa option (/CHECK=PARAMETER on VMS) generates code to check in run-time for this limitation.

8.2.3 _exti (Extract bit-field)

PROTOTYPE

```
void extb(int offset,
           void *base,
           char *dest,
           int length);
void extw(int offset,
           void *base,
           short *dest,
           int length);
void extd(int offset,
           void *base,
           int *dest,
           int length);
```

DESCRIPTION

The ext routines copy the bit-field specified by base, offset and length to the dest operand. The field is right justified in dest. High-order bits are zero-filled if the field is shorter than dest or discarded if the field is longer than dest.

The field starts at bit position

```
offset % 8
within the memory byte
   base + (offset / 8)
```

- NOTES: 1. length must be a constant. Otherwise the routine is not inlined and an emulation function is called.
 - 2. length must be in the range 1 through 32. Compilation with the -aa option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for this limitation.

CAUTION

Although a bit-field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits. A bit-field must be composed of bits from no more than four contiguous bytes.

If the offset operand is a constant expression with a value in the range 0 to 7, and length is a constant expression, the compiler will use the short version of the instruction (exts).

8.2.4 _ins (Insert Bit-field)

PROTOTYPE

```
void ins(int offset,
          unsigned int src,
          int *base,
          int length);
```

DESCRIPTION

The ins routine inserts the src operand into the bit-field specified by base, offset and length. The src operand is right-justified in the field. Highorder bits are zero-filled if src is shorter than the field or discarded if src is longer than the field.

The field starts at bit position

```
offset % 8
```

within the memory byte

```
base + (offset / 8)
```

length specifies the number of bits in the field, and must be in the range 1 through 32. Compilation with the -aa option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for this limitation.

- NOTES: 1. length must be a constant. Otherwise the routine is not inlined and an emulation function is called.
 - 2. length must be in the range 1 through 32. Compilation with the -aa option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for this limitation.

CAUTION

Although a bit-field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits. A bit-field must be composed of bits from no more than four contiguous bytes.

If the offset operand is a constant expression with a value in the range 0 to 7, and length is a constant expression, the compiler will use the short version of the instruction (inss).

The appropriate version for ins (insb, insw or insd, for byte, word or double-word respectively) is selected by the compiler according to the type of the parameter passed for src.

8.2.5 _cvtp (Convert to Bit Pointer)

PROTOTYPE

```
int cvtp(int offset,
         void *base);
```

DESCRIPTION

The cvtp function returns the absolute bit address of the memory bit specified by base and offset. The bit address specifies the number of bits from the first bit in the memory space (bit 0 of the byte at address 0) to the specified bit.

8.2.6 abs (Absolute Value)

PROTOTYPE

DESCRIPTION

The abs routines return the absolute value of their parameter val. abs compiles into the integer ABS instruction, fabs compiles into the floating-point ABSL instruction and ffabs compiles into the floating-point ABSF instruction. These routines differ from the other intrinsic routines: Their names do not begin with an underscore, and the floating-point versions' names are not the same as the assembly mnemonic.

The reason for these differences is that the routines exist in the mathematical library libm.a, and are declared in the header file math.h.

8.3 CG-Core Intrinsic Functions

The CG-Core microprocessors complements the full instruction set of the *Series 32000* processor with special graphics-oriented instructions. These include Bit-aligned Block Transfer (BITBLT) functions, line drawing, pattern replications and data compression-expansion. In addition, an interface to an external BITBLT processing unit (BPU) enables very fast BITBLT operations.

The CG-Core graphic instructions supported by intrinsic functions are:

bban	Bit-aligned Block Transfer, 4-direction bitwise AND operation for CRT applications.		
bbfo	r Bit-aligned Block Transfer, 2-direction FAST OR operation optimized for printers.		
bbor	Bit-aligned Block Transfer, general 4-direction OR operation.		
bbst	od Bit-aligned Block Transfer, 4-direction replace (Source-to-Destination) operation.		
bbxo	r Bit-aligned Block Transfer, 4-direction XOR operation.		
bitw	Bit-aligned Word Transfer, for small BITBLT OR operations.		
extb	Drives a 4-direction, 16-function external BITBLT processing unit. The BPUs supported are the DP8510 or DP8511 for the NS32CG16 CPU, or the on-chip BPU for the NS32CG160.		
movm	Move Multiple Pattern, for pattern fill, horizontal line drawing, memory clear.		
	Cat. B't. Brown links Cation. Continued to the continued to		

sbitps Set Bit Perpendicular String, for image expansion and horizontal/vertical/diagonal line drawing.

sbits Set Bit String, for image expansion and horizontal line drawing.

tbits Test Bit String, for image data compression.

In addition, the general *Series 32000* instruction set includes instructions that manipulate single bits or bit-fields in memory. These are useful for graphic operations and are supported by intrinsic functions, as described in Section 8.4.

TERMS AND CONVENTIONS

bit alignment	The ability to access any bit in memory. The Series 32000 instruction set and the CG-Core instructions enable efficient direct access to bit-aligned data.	
bit-block	A rectangular sub-area of an image. A bit-block consists of y lines of x bits, where the spacing between lines is warp bits.	
bit ordering	Bit ordering is from the least significant to the most significant bit. Bit positions increase left-to-right in the image.	

image Images are defined as binary bit-maps, where a set bit (one)

represents a black dot, and a clear bit (zero) represents a white

dot.

left-to-right A graphic operation that traverses memory in increasing

address order (see memory ordering below).

memory ordering The displayed image's memory increases in a left-to-right, top-

to-bottom raster direction. The first byte represents the image's top-left corner. The *Series 32000* architecture is "little-endian" in referring to a word (16 bits) or double-word (32 bits) — the least significant byte is stored at the lowest

address.

right-to-left A graphic operation that traverses memory in decreasing

address order (see memory ordering above).

shift The shifts in this chapter are Series 32000 logical left bit shifts.

With the imaging convention this results in image bits moving

from left to right.

warp The horizontal width of the image. Also known as raster or

pitch.

All definitions given in this section are supplied as part of the GNX C compiler package in file cq16.h.

8.3.1 _extblt (External Bit Aligned Block Transfer)

PROTOTYPE

```
void extblt(char *src addr, /* Extblt without preloading */
              char *dest addr,
              int adj width,
              int height,
              int horiz incr,
              int adj src_wrap,
              int adj dest wrap);
void extbltp(char *src addr, /* Extblt with preloading */
              char *dest addr,
              int adj width,
              int height,
              int horiz incr,
              int adj src wrap,
              int adj dest wrap);
```

PARAMETERS

src_addr	The base byte-address of the source data, the value should be even.
dest_addr	The base byte-address of the destination, the value should be even.
adj_width	Adjusted width of the image on which the operation is performed. The adjustment is the width in words of destination data multiplied by horiz_incr: (width * horiz_incr)
height	The number of lines on which the operation is performed.
horiz_incr	The horizontal step in bytes for copying. Its value should be $+2$ or -2 .
adj_src_warp	The adjusted wrap of the source.
	The adjustment is to the actual source warp in bytes minus the width in bytes (not the adjusted width) less two: (source warp - (width in bytes - 2))

adj dest warp The adjusted wrap of the destination.

The adjustment is the actual destination warp in bytes less the width in bytes (*not* the adjusted width) less two:

(destination warp - (width in bytes - 2))

DESCRIPTION

The two extblt functions drive an external BITBLT processing unit (BPU). The BPUs supported are the DP8510 or DP8511 for the NS32CG16 CPU, or the on-chip BPU for the NS32CG160. The CPU supplies addresses and bus cycles while the BPU operates on the data. For more details on the EXTBLT instruction refer to the NS32CG16 or the NS32CG160 Printer/Display Processor Programmer's Reference Supplement.

NOTE: Compilation with the -aa option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for the following:

- "src_addr" and "dest_addr" values are even.
- "horiz_incr" is +2 or -2.
- "width" value is a multiple of "horiz_incr" and has the same sign.

8.3.2 BITBLT instructions

PROTOTYPE

```
typedef void _bbfunc(char *src_addr,
                     char *dest addr,
                     unsigned char shift val,
                     unsigned int height,
                     unsigned int mask1,
                     unsigned int mask2,
                     int adj src warp,
                     int adj dest warp,
                     unsigned short width);
__bbor_da, /* bbor with decreasing addresses */
_bbor_sda, /* bbor with inverted source and
                   * decreasing addresses */
       _bband, /* plain bband */
       bband s,
                  /* bband with inverted source */
       _bband_da, /* bband with decreasing addresses */
       bband sda, /* bband with inverted source and
                    * decreasing addresses */
       _bbxor, /* plain bbxor */
       bbxor s,
                  /* bbxor with inverted source */
       _bbxor_da, /* bbxor with decreasing addresses */
       bbxor sda, /* bbxor with inverted source and
                   * decreasing addresses */
       _bbstod, /* plain bbstod */
       bbstod s, /* bbstod with inverted source */
       bbstod da, /* bbstod with decreasing addresses */
       bbstod sda; /* bbstod with inverted source and
                    * decreasing addresses */
```

DESCRIPTION

The BITBLT instructions perform a full two-operand, bit-aligned block transfer. Sixteen intrinsic BITBLT functions are supplied as an interface to these instructions. They are divided into four groups, according to the operator between the source and destination bits: _bbor for an OR operator, _bband for an AND operator, _bbxor for a XOR operator, and _bbstod for a source to destination copy, which overwrites the destination bits.

Within each group there are four variants, resulting from the combination of:

- The direction of the transfer increasing or decreasing addresses. Increasing addresses indicate a left-to-right operation; decreasing addresses indicate a right-to-left operation.
- The reading of the source data with or without inversion. The inversion is a logical negation of each source bit.

A separate intrinsic function is supplied for each combination of variants. These are coded as suffixes to the function names. The suffixes are _s for inverted source, _da for decreasing addresses, and _sda for inverted source and decreasing addresses. No suffix signifies a simple version of the function (i.e. no source inversion, increasing addresses). There are therefore 16 different BITBLT functions, as shown in the table below:

Operator	plain	inverted source	decreasing addresses	inverted source and decreasing addresses
AND	_bband	_bband_s	_bband_da	_bband_sda
OR	_bbfor	_bbor_s	_bbor_da	_bbor_sda
XOR	_bbxor	_bbxor_s	_bbxor_da	_bbxor_sda
STOD	_bbstod	_bbstod_s	_bbstod_da	_bbstod_sda

NOTE: No plain bbor function is supplied since the CG-Core bbfor instruction has the same functionality with faster performance.

A bottom-to-top BITBLT operation can be performed by giving negative source and destination warp values, and beginning from the bottom line of the image. Thus, combined with the _da option, the BITBLT functions can manipulate blocks of data beginning in any of its four corners.

All functions have the prototype given above. The parameters are the same as the instruction operands with the following meanings:

- src addr is the base byte-address of the source bit-block
- dest addr is the base byte-address of the destination.
- height is the vertical size of the bit-block, which specifies the number of lines to be transferred.
- width is the horizontal size of the transferred bit block. It is the number of whole words on one line containing the bit-block.
- mask1 and mask2 are bit masks "protecting" those bits at the left and right
 of the source word block that do not belong to the bit-block from affecting the
 bits in the destination block. A bit set to one means that this bit should not
 be protected. A clear bit means that this bit should not affect the destination
 bit.

The upper 16 bits of the mask parameters must be clear, otherwise behavior is undefined. Compilation with the -aa compiler option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for this condition.

• shift_val contains the difference between the bit offsets of the source and destination bit-blocks, relative to the word block:

```
shift_val = destination bit offset - source bit offset
```

shift_val must be positive. src_addr and dest_addr may need to be adjusted to ensure this.

• adj_src_warp describes the adjusted source warp. For left-to-right operations the source warp must be adjusted to:

```
adj src warp = source warp - 2*(width - 1)
```

For right-to-left operations the source warp must be adjusted to:

```
adj src warp = source warp + 2*(width - 1)
```

 adj_dest_warp describes the adjusted destination warp. For left-to-right operations the destination warp must be adjusted to:

```
adj dest warp = destination warp - 2*(width - 1)
```

For right-to-left operations the destination warp must be adjusted to:

```
adj dest warp = destination warp + 2* (width - 1)
```

For more details about the BITBLT instructions refer to the appropriate CG-core Processor Programmer's Reference Supplement.

8.3.3 _bitwt (Bit Aligned Word Transfer)

PROTOTYPE

```
void bitwt (unsigned short **src addr,
             unsigned **dest addr,
             int shift val);
```

DESCRIPTION

The bitwt function performs a bit-aligned transfer of a short int. The 16 bits of data at **src addr are read, zero extended to an unsigned int and shifted to the left by the number of bit positions specified in shift val. The 32 bits of data at **dest addr are read, ORed with the shifted source and the result is written into **dest addr. Then *src addr and *dest addr are incremented by two to point to the next shorts to be operated on ((byte *) incrementing).

shift val must be in the range of 0 to 15. Compilation with the -aa option (/CHECK=PARAMETER on VMS) will generate code to check in run-time for this limitation.

The bitwt function is useful for implementing a high-speed "source OR destination" BITBLT operation, when the source data is aligned such that it does not need masking The implementation consists of a simple loop containing the bitwt function and add instructions that adjust the source and destination warps.

8.3.4 _movmp (Move Multiple Pattern)

PROTOTYPE

DESCRIPTION

The _movmpb function copies count times the byte specified by the pattern function to dest_addr. Each count is spaced incr bytes from the previous one. Only the low order byte of pattern will be copied, if a pattern parameter larger than a byte is passed.

Similarly, the _movmpw function copies the word pattern. The incr is in units of bytes. If a double-word pattern parameter is passed, only its low order word will be copied. If a byte pattern parameter is passed, it will be zero-extended.

_movmpd copies the double-word pattern. If a byte or word pattern parameter is passed, it will be zero-extended.

This function is useful for quickly clearing large memory blocks. For example, in printer applications a page image can be cleared prior to drawing the next page. The _movmp function can also be used for drawing horizontal lines and for creating simple patterns.

8.3.5 _sbits (Set Bit String)

PROTOTYPE

```
int sbits (void *dest addr,
           int bit offset,
           unsigned run length,
           unsigned *lookup table);
```

DESCRIPTION

The sbits function operates on run length consecutive bits, starting with the bit at bit offset from the byte at *dest addr. run length must be in the range 0 to 25. An OR operation is performed between these bits and a double-word from lookup table. The double-word used is

```
lookup table[run length + 32*(bit offset % 8)]
```

The function returns zero if run length is 25 or less, and returns one if run length is 26 or greater. Calling the functions as a procedure, such as

```
(void) sbits(...
```

will prevent the compiler from producing the code for the return value.

The lookup table parameter is used to lookup any table. The Sbits macro is an interface to the most common "black" lookup table (all bits in the run will be set). This lookup table is provided in the libcgl6.a archive. The prototype for the Sbits macro is:

```
#define Sbits(addr,offset,length) \
       sbits((addr),(offset),(length), sbits tbl)
extern char _sbits_tbl[];
```

8.3.6 _sbitps (Set Bit Perpendicular String)

PROTOTYPE

DESCRIPTION

The _sbitps function sets run_length bits, starting at the bit at *bit_offset from the byte at *src_addr. The set bits are separated by dest warp bits.

The function can be used to draw vertical lines by passing dest_warp equal to the image warp. Forty-five degree lines are drawn by dest_warp equal to the image warp plus or minus one. Other applications include expansion and/or rotation of images (in conjunction with the tbits function) and filling.

8.3.7 _tbits (Test Bit String)

PROTOTYPE

```
int tbits0(void *src addr,
                                    /* count a series of zeros */
            int *bit offset,
            int *run length,
            int max run length,
            int max bit offset,
            unsigned int *Fflag);
int tbits1(void *src addr,
                                    /* count a series of ones */
            int *bit offset,
            int *run length,
            int max_run_length,
            int max bit offset,
            unsigned int *Fflag);
```

DESCRIPTION

The tbits0 function counts the number of consecutive clear bits, starting from the bit at *bit offset from the byte at src addr. Counting will terminate either at the first set bit, or if max run length or max bit offset bits were tested before a set bit was found. The number of bits scanned will be placed in *run length, and *bit offset will be the offset upon termination.

Similarly, thits1 counts the number of consecutive set bits.

Both functions return the value of the PSR L flag. Calling the functions as a procedure, such as

```
(void) tbits1(...
```

will prevent the compiler from producing the code for the return value.

The value of the PSR F flag is placed in the Fflag value. Calling the functions with the macro IGNORE PARAM (defined in the cg16.h) as the actual Fflag parameter will prevent the compiler from producing the code for assigning the value of the PSR F flag.

Table 8-1. Effect of thits on PSR L and F flags

CONDITION	L	F	NOTES
*run_length < max_run_length bit found	1	F	F reflects last bit tested
*run_length >= max_run_length bit not found	1	F	F reflects last bit tested
*bit_offset >= max_bit_offset bit not found	0	F	F reflects last bit tested
*bit_offset >= max_bit_offset on entry	0	0/1	0 for _tbits0 1 for _tbits1

8.4 NS32GX320 Intrinsic Functions

The NS32GX320 high-performance 32-bit microprocessor combines the full instruction set of the Series 32000 family with the following on-chip integrated features: instruction and data caches, a 2-channel DMA controller, a 15-level interrupt control unit (ICU) and three 16-bit timers. In addition, Digital Signal Processing is supported by four special instructions:

mulwd Multiply Word to Double cmuld Complex Multiply Double Complex Multiply and Accumulate Double cmacd mactd Multiply and Accumulate Twice Double

These instructions are accessible from the GNX C compiler using the special NS32GX320 intrinsic routines.

All prototype definitions given in this section are supplied as part of the GNX C package in the file gx320.h.

8.4.1 NS32GX320 typedefs

WCOMPLEX and DCOMPLEX are typedefs, defined for cmuld and cmacd. They designate data types for fixed-point complex arithmetic.

8.4.2 _mulwd (Multiply Word to Double)

PROTOTYPE

```
int mulwd(short src1, int* src2);
```

DESCRIPTION

The _mulwd function returns the integer multiplication of src1 operand, by the lower 16 bits of *src2.

8.4.3 _cmuld (Complex Multiply Double)

PROTOTYPE

DESCRIPTION

The _cmuld function assigns the complex multiplication of the source parameters (*src1 and src2) to *result. For example

```
result->re = src1.re*src2.re - src1.im*src2.im
result->re = src1.re*src2.im + src1.im*src2.re
```

8.4.4 _cmacd (Complex Multiply and Accumulate Double)

PROTOTYPE

```
void cmacd(DCOMPLEX *accum,
            WCOMPLEX src1,
            WCOMPLEX src2);
```

DESCRIPTION

The _cmacd function accumulates the complex multiplication of the source parameters (src1 and src2) into *result. For example

```
accum->re = accum->re + src1.re*src2.re - src1.im*src2.im
accum->re = accum->im + src1.re*src2.im + src1.im*src2.re
```

8.4.5 _mactd (Multiply and Accumulate Twice Double)

PROTOTYPE

DESCRIPTION

The _mactd function accumulates the result of the two multiplications of src1.s1*src2.s1 and src1.s2*src2.s2 into *result. For example

```
accum = accum + src1.s1*src2.s1 + src1.s2*src2.s2
```

SERIES 32000 STANDARD CALLING CONVENTIONS

A.1 INTRODUCTION

The main goal of standard calling conventions is to enable the routines of one program to communicate with different modules, even when written in multiple programming languages. The Series 32000 standard calling conventions support various special language features (such as the ability to pass a variable number of arguments, which is allowed in C), by using the different calling mechanisms of the Series 32000 architecture. These conventions are employed only to call "externally visible" routines. Calls to internal routines may employ even faster calling sequences by passing arguments in registers, for instance.

Basically, the calling sequence pushes arguments on top of the stack, executes a call instruction, and then pops the stack, using the fewest possible instructions to execute at the maximum speed. The following sections discuss the various aspects of the Series 32000 standard calling conventions.

A.2 CALLING CONVENTION ELEMENTS

Elements of the standard calling sequence are as follows:

The Argument Stack

Arguments are pushed on the stack from right to left; therefore, the leftmost argument is pushed last. Consequently, the leftmost arguments are always at the same offset from the frame pointer, regardless of how many arguments are actually passed. This allows functions with a variable number of arguments to be used.

NOTE: This does not imply that the actual parameters are always evaluated from right to left. Programs cannot rely on the order of parameter evaluation.

The run-time stack must be aligned to a full double-word boundary. Argument lists always use a whole number of double-words; pointer and integer values use a double-word (by extension, if necessary), floating-point values use eight bytes and are represented as long values;

structures/unions use a multiple of double-words.

NOTE: Stack alignment is maintained by all GNX — Version 4 compilers through aligned allocation and de-allocation of local variables. Interrupt routines and other assembly-written interface routines are advised to maintain this double-word alignment.

The caller routine must pop the arguments off the stack upon return from the called routine.

NOTE: The compiler uses a more efficient organization of the stack frame if the -OF (FIXED_FRAME in VMS) optimization is enabled. In that case, programs should not rely on the organization of the stack frame.

Saving Registers

General registers R0, R1, and R2 and floating registers F0, F1, F2, and F3 are temporary or scratch registers whose values may be changed by a called routine. Also included in this list of scratch registers is the long register L1 of the NS32181/NS32381/NS32580 FPU. It is not necessary to save these registers on procedure entry or restore them before exit. If the other registers (R3 through R7, F4 through F7, and L3 through L7 of the NS32181/NS32381/NS32580) are used, their values should be saved (onto the stack or in other memory locations) by the called routine immediately upon procedure entry and restored just before executing the return instruction. This should be performed because the caller routine may rely on the values in these registers not changing.

NOTE: Interrupt and trap service routines are required to save/restore all registers that they use. If the service routine calls another routine it must save scratch registers as well.

Returned Value

An integer or a pointer value that returns from a function, returns in (part of) register R0.

Floating-point values return in floating point registers: A float value is returned in register F0. A double value is returned in register pair F0-F1.

If a function returns a structure or union, the calling function passes an additional argument at the beginning of the argument list. This argument points to where the called function returns the structure. The called function copies the structure into the specified location just before returning from the function. Note that functions that return such types must be correctly declared as such, even if the return value is ignored. For details see Chapter 4.

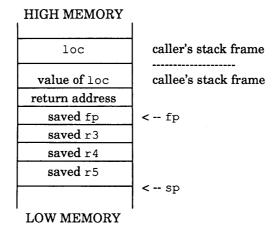
Example:

```
int iglob;
m()
{
     int loc;
     a = ifunc(loc);
ifunc(p1)
int p1;
{
     int i, j, k;
     j = 0;
     for (i = 1; i \le p1; i++)
         j = j + f(i);
     return(j);
```

The compiler may generate the following code:

```
_m:
         enter
                  [], 4
                                 # Allocate local variable
         movd
                  -4(fp), tos
                                 # Push argument
         bsr
                   ifunc
                  $(-4)
                                 # Pop argument off stack
         adjspb
                  r0,_iglob
         movd
                                 # Save return value
         exit
                  []
                  $(0)
         ret
_ifunc:
                  [r3,r4,r5],0
                                 # Save safe registers
         enter
         movd
                  8(fp),r5
                                 # Load argument to temp register
         movqd
                  $(0),r4
                                 # Initialize j
         cmpqd
                  $(1),r5
                  .LL1
         bqt
                  $(1),r3
                                 # Initialize i
         movad
.LL2:
         movd
                  r3,tos
                                 # Push argument
         bsr
                   f
                  $ (-4)
         adjspb
                                 # Pop argument off stack
         addd
                  r0,r4
                                 # Add return value to j
                                 # Increment i
         addqd
                  $(1),r3
         cmpd
                  r3, r5
         ble
                  .LL2
.LL1:
         movd
                  r4,r0
                                 # Return value
         exit
                  [r3,r4,r5]
                                 # Restore safe registers
                  $(0)
         ret
```

After the enter instruction is executed by ifunc(), the stack will look like this:



Appendix B

MIXED-LANGUAGE PROGRAMMING

B.1 INTRODUCTION

Mixed-language programs are frequently used for a couple of reasons. First, one language may be more convenient than another for certain tasks. Second, code sections, already written in another language (e.g., an already existing library function), can be reused by simply calling them.

A programmer who wishes to mix several programming languages needs to be aware of subtle differences between the compilation of the various languages. The following sections describe the issues the user needs to be aware of when writing mixed-language programs and then compiling and linking such programs successfully.

B.1.1 Writing Mixed-Language Programs

The mixed-language programmer should be aware of the following topics:

- Name Sharing Potential conflicts including permitted name-lengths, legal characters in identifiers, compiler case sensitivity, and high-level to assemblylevel name transformations.
- Calling Conventions The way parameters are passed to functions, which registers must be saved, and how values are returned from functions. See Appendix A for details.
- Declaration Conventions The demands that different languages impose when referring to an outside symbol (be it a function or a variable) that is not defined locally in the referring source file. Note that this is also true of references to an outside symbol that is not in the same language as that of the referring source file.

To help the programmer avoid these potential problems, a set of rules for writing mixed-language programs has been devised. Each rule consists of a short mnemonic name (for easy reference), the audience of interest for the rule, and a brief description of the rule.

Figure B-1 summarizes all of the rules in the context of each possible cross-language pair.

	С	Pascal	FORTRAN 77	Series 32000 Assembly
Series 32000 Assembly	'_' prefix	'_' prefix include ext case sensitivity	'_' prefix '_' suffix ref args case sensitivity	
FORTRAN 77	'_' suffix ref args case sensitivity	'_' suffix include ext ref args		'_' prefix '_' suffix ref args case sensitivity
Pascal	include ext case sensitivity		'_' suffix include ext ref args	'_' prefix include ext case sensitivity
C		include ext case sensitivity	'_' suffix ref args case sensitivity	'_' prefix

Figure B-1. Cross-Language Pairs

RULE 1 case sensitivity

This rule is of interest to every programmer who mixes programming languages.

C and Series 32000 assembly are case sensitive while FORTRAN 77 and Pascal are not. Programmers who share identifiers between these two groups of languages must take this into account. To avoid problems with case sensitivity, the programmer can:

- 1. Take care to use case-identical identifiers in all sources and compile FORTRAN 77 and Pascal sources using the case-sensitive option (/CASE SENSITIVE on VMS, -d on UNIX).
- 2. Use only lower-case letters for identifiers which are shared with FORTRAN 77 or Pascal since the FORTRAN 77 and Pascal compilers fold all identifiers to lower-case if not given the case-sensitive option.

RULE 2 '_' prefix

This rule is of interest to those who mix high-level languages with assembly code.

All compilers map high-level identifier names into assembly symbols by prepending these names with an underscore. This ensures that user-defined names are never identical to assembly reserved words. For example, a high-level symbol NAME, which can be a function name, a procedure name, or a global variable name, generates the assembly symbol NAME.

Assembly written code which refers to a name defined in any high-level language should, therefore, prepend an underscore to the high-level name. Stated from a high-level language user viewpoint, assembly symbols are not accessible from high-level code unless they start with an underscore.

RULE 3 ''suffix

This rule is of interest to those who mix FORTRAN 77 with C, Pascal, or assembly code.

The FORTRAN 77 compiler appends an underscore to each high-level identifier name (in addition to the action described in RULE 2). The reason for an appended underscore is to avoid clashes with standard-library functions that are considered part of the language, e.g., the FORTRAN 77 WRITE instruction. For example, a FORTRAN 77 identifier NAME is mapped into the assembly symbol NAME.

Therefore, a C, Pascal, or assembly program that refers to an FORTRAN 77 identifier name should append an underscore to that name. Stated from an FORTRAN 77 user viewpoint, it is impossible to refer to an existing C, Pascal, or assembly symbol from FORTRAN 77 unless the symbol terminates with an underscore.

RULE 4 ref args

This rule is of interest to those who mix FORTRAN 77 with other languages.

Any language which passes an argument to a FORTRAN 77 routine must pass its address. This is because a FORTRAN 77 argument is always passed by reference, *i.e.*, a routine written in FORTRAN 77 always expects addresses as arguments.

Routines not written in FORTRAN 77 cannot be called from an FORTRAN 77 program if the called routines expect any of their arguments to be passed by value. Only routines which expect all their arguments to be passed by reference can be called from FORTRAN 77.

The Pascal program must declare all FORTRAN 77 routine arguments using var. C programs must prepend the address operator & to FORTRAN 77 routine arguments in the call. The C or Pascal programmer

who wants to pass an unaddressable expression (such as a constant) to a FORTRAN 77 routine, must assign the expression to a variable and pass the variable, by reference, as the argument.

RULE 5 include ext

This rule is of interest to Pascal programmers who want to share variables between different source files which may or may not be written in Pascal.

Pascal sources which share global variables or routines must make these variables known to separately compiled modules. This is done by the import and export attributes, or by inclusion of a .h file which contains the variables or routines. ERROR in line number 290 incorrect number of fields line is: .if!C'Pascal'See The Series 32000 GNX-Version 4 Pascal Optimizing Compiler Reference Manual. ERROR line292contains a . No matching

In addition to these rules, a few points should be noted. First, GNX—Version 4 FORTRAN 77 allows identifiers longer than the six character maximum of traditional FORTRAN compilers. Second, the family of GNX—Version 4 Compilers allows the use of underscores in identifiers. Both of these enhancements simplify name sharing.

Importing Routines and Variables

The general conventions of all languages must be kept in mixed-language programs. In particular, externals must be declared in those program sections which import them. The following are examples of declarations of external (imported) functions/procedures and external (imported) variables in each language. The examples are in the form:

```
caller language: external (imported) functions/procedures
or
external (imported) variables
```

Example:

```
C: extern int func_();
          or
          extern int var name_;
```

Note that the strict reference C model (ANSI C standard) is assumed. If the model is relaxed, then the external declarations are not mandatory.

```
FORTRAN 77: INTEGER func
or
COMMON /var_name/ local_name

Pascal.h file: function func_: integer; external;
```

```
procedure proc ; external;
              #include "var def.h"
              where the file var def.h contains:
                var var name : integer;
Pascal .h file:
              import function func : integer;
import/export
              var export i: integer;
Series 32000:
              .globl func
   assembly
              .globl _var_name_
```

B.1.2 Compiling Mixed-Language Programs

After writing different program parts in different languages, keeping in mind the rules previously mentioned, the mixed-language programmer must still link and load these parts to make them run successfully. The following three points should be mentioned in conjunction with the successful linking and loading of programs.

- External library (standard or nonstandard) routines must be bound with the user-written code that calls them.
- Initialization code which arranges to pass program parameters to the main program and then calls the main program, sometimes has to be bound with userwritten code.
- The entry point of the code, i.e., the location where the program starts executing, should be determined.

Libraries:

The following table (Table B-1) lists libraries associated with each compiler. When programming with mixed languages, the libraries associated with the languages used must be bound with the program during the link phase of compilation.

Initialization code and Entry-points:

Normally, the entry point of the final executable file is called start. The code that follows this entry-point is an initialization code that prepares the run-time environment and arranges parameters to be passed to the user-written main program. The initialization object file which is linked by default is called crt0.o. The crt0.o file always calls main.

The assembly-symbol that starts the user main program is main (the underscore is prepended by the C compiler) in the case of C programs and MAIN in Pascal or FORTRAN 77 programs.

Table B-1. Compilers and their Associated Libraries

COMPILER (DRIVER) NAME	LIBRARIES
cc (cross nmcc)	libc
f77 (cross nf77)	libF77, libI77, libm, libc
pc (cross nmpc)	libpas, libm, libc

Note that the last two compilers completely ignore the user's main program name. Therefore, in C, the user-written code is called directly from crt0.o. In Pascal and FORTRAN 77, _main is located in the respective standard library which performs additional initializations before calling the user entry-point MAIN .

B.1.3 Compilation on UNIX Operating Systems

National Semiconductor's GNX tools (assembler, linker, etc.) on UNIX systems relieve a user's concern about external libraries, initialization code, and entry-points. This is due to the coherence and consistency of the GNX — Version 4 Compilers and their integration through the use of a common driver.

When using a GNX — Version 4 Compiler on a UNIX system, the user does not directly call the compiler front end, optimizer, code generator, assembler or linker. Instead, the calls are indirectly made through the driver program.

The driver program accepts a variable number of filename arguments and options and knows how to identify language-specific options. The driver also identifies the languages in which its filename arguments are written by the names of these arguments. Therefore, the driver can arrange to compile and bind the programs with the needed libraries in order to run the program successfully.

The driver program used by C, Pascal and FORTRAN 77 programmers is exactly the same program on UNIX systems, named differently for each language. The respective driver names are cc, pc and f77 (nmcc, nmpc and nf77 for cross-support).

The driver program looks at its own name in order to determine the libraries that are to be bound with the program. In addition, the driver links additional libraries according to the name extensions of any of its filename arguments. For instance, cc also links libm and libpas when one of the filename arguments is a Pascal source (recognized by the .p, .pas, .P or .PAS extensions).

The -v (/VERBOSE on VMS) option of the driver verbosely outputs all driver actions. With this option the interested user can track problems that might arise (such as undefined symbols from the linker).

As mentioned in the previous section, different languages use different initialization codes that reside in language-specific standard libraries. It is necessary that the correct language initialization code be linked with a mixed-language program. The driver program helps do this, but it needs to know in which language the main program is written.

To ensure that the correct initialization code is linked with a mixed-language program, the user should call the driver that corresponds to the language of the main program module within the mixed-language program.

For example, suppose there are four source modules written in four different languages (c_utils.c written in C, f_utils.f written in FORTRAN 77, p_utils.p written in Pascal, and s_utils.s written in assembly language), and there is a fifth module that has already been compiled separately (obj.o, an object module). Assuming there is a main program written in FORTRAN 77, the f77 driver should be used.

```
f77 main.f c_utils.c f_utils.f p_utils.p s_utils.s obj.o
```

If the main program is written in C, cc is used, and so on.

B.1.4 Compilation on VMS Operating Systems

Under the GNX tools on VMS systems, the linking phase is separate from the compilation phase; therefore, it demands separate actions from the user.

The interested user should refer to the language tools manuals (assembler, linker, etc.) for a complete description of how to use them on VMS systems.

B.2 COMPILING THE MIXED-LANGUAGE EXAMPLE

The example listed in Section B.3 consists of a number of program modules written in languages different from the main program which is written in C.

B.2.1 Compiling the Example on a UNIX System

To compile the program modules on a UNIX system, type the command:

```
nmcc c_main.c\
c fun.c f77_fun.f pas fun.p asm_fun.s
```

This assumes that all the program modules are in the same directory. If the program compiles and links successfully, the result is an executable file that, when run, prints

B.2.2 Compiling the Example on a VMS System

To compile the modules on a VMS system, type the following commands:

```
nmcc c_main.c
nmcc c_fun.c
nf77 f77_fun.f
nmpc pas_fun.p
nasm asm fun.s
```

After successful linking, the result is an executable file that, when run, prints the line "Passed OK !!!".

B.3 PROGRAM MODULE LISTINGS

The different program modules are listed in this section.

c_main.c

```
/*-----
    Example of a C program which communicates with C, Pascal,
    Fortran 77, and Assembly external functions, via direct
    calls as well as via a global variable.
    Parameter passing by reference is accomplished by passing the
    addresses of the characters variables 'a', 'b', 'c', 'd' and 'e'.
char str_[] = "Passed OK !!!0; global ('exported') string
main() {
   char a, b, c, d, e;
   int three = 3;
                       FORTRAN must get its parameters by reference
                         So we put this constant into a variable ...
                       & &
& &
   if (c_func (&a, 0)
                                 in C arrays start with 0
                                 in Pascal they start at 1 in f77, at 1
       pas_func (&b, 2)
       f77_func_(&c, &three) &&
       asm_func (&e, 4))
                                 in assembly, at 0
           printf("%c%c%c%c%c%s", a, b, c, d, e, str_+5);
             Should print "Passed OK !!!"
}
dummy()
{
}
```

```
/*
    * Declaration of the public character string 'str[]' and definition
    * of the C function 'c_func()'.
    * Note the appending of an underscore to the external symbol 'str_'
    * which is shared with FORTRAN 77.
    */
    extern char str_[];
    int c_func(c_ptr, index)
    char    *c_ptr;
    int index;
{
        *c_ptr = str_[index];
        return 1;
}
```

f77_fun.f

```
C The FORTRAN 77 function:
C All parameters are passed by reference
C The COMMON statement aliases the external array 'str' as 'text'
C LOGICAL FUNCTION f77_func(c, index)
    CHARACTER c
    INTEGER index
    COMMON /str/ text
    CHARACTER text(15)
    c = text(index)
    f77_func = .TRUE.
    RETURN
    END
```

pas_fun.p

```
The Pascal function 'pas func()'
(* 'str[]' character-array declaration *)
#include "str pas.h";
(* make this function visible to outsiders ('export') *)
function pas func (var c: char; index: integer): boolean; external;
function pas_func();
begin
   c := str [index];
   pas_func := TRUE;
end;
```

str_pas.h

```
(* 'str[]' character-array declaration for Pascal *)
var
   str_ : packed array [1..15] of char;
```

asm_fun.s

```
The 32000 Assembly Language Function 'asm_func'
       The function includes an artificial use of r7, to demonstrate the
       need to save it upon entry and restore upon exit, as opposed to
       r0, r1 and r2; f0, f1, f2 and f3 which can be used freely without
       saving or restoring. This is according to the Series 32000
       standard calling convention.
       The function return value is placed in r0, also according to the
       standard calling convention.
      .globl
                _str_
                           # Import the global str[] array.
                _asm_func # Export (make visible) the assembly function.
      .globl
      .align 4
_asm_func:
      enter
                [r7],0
                                            # Set frame, show saving of r7
      movb
                _str_+0(12(fp)),0(8(fp))
                                            # argument 1 ← str[argument 2]
                $(1),r7
                                            # artificial use of r7
      movqd
      movd
                r7, r0
                                            # return value ← TRUE
                                            # Unwind frame, restore r7
      exit
                [r7]
      ret
                $(0)
                                            # Return to caller
```

ERROR DIAGNOSTICS

C.1 INTRODUCTION

The GNX C compiler has a superior error handling mechanism. In most cases, the compiler continues to compile when an error is found. An error message is displayed, providing information on the type of error, the source filename, the line number location of the error. Generally, the compiler attempts to minimize the effects of errors on compilation.

C.2 ERROR MESSAGES

Errors are divided into six categories:

- 1. Limitation Errors
- 2. System Errors
- 3. Severe Errors
- 4. Syntax Errors
- 5. Caution Errors
- 6. Warnings

C.2.1 Error Messages Format

The general syntax of an error message is

filename, line m:c) [category]: message

Where:

is the source file name. filename

is the line number location of the error.

is a lower case letter used to mark the error position on the source line.

is the error category. category

The error message is followed by an echo of the source line. The errors are marked with the appropriate lower case letter corresponding to the error message as displayed in the syntax.

Example:

C.2.2 System Errors

System errors are related to the operating system or the environment in which the compiler runs. For example:

```
[system]: Can't open file filename
[system]: Ran out of memory
```

C.2.3 Limitation Errors

Limitation errors are caused by exceeding compiler limitations. Generally a limitation error causes the suspension of code generation. In these cases the limitation message includes "no object file produced". However, in some cases limitation messages are warnings, and code generation continues.

The following is a complete list of limitation errors:

1. [Limitation]: array size too large; no object file produced

This error message is produced if an array size exceeds the maximal number which can be represented in 29 bits (536870911).

2. [Limitation]: structure too large; no object file produced

This error message is produced if a structure size exceeds the maximal number which can be represented in 29 bits (536870911).

[Limitation]: array dimension too large; no object file produced

This error message is produced if a number greater than the maximal number which can be represented in 29 bits (536870911) is used for an array's dimension.

4. [Limitation]: cumulative size of structure members is too large; no object file produced

C-2 ERROR DIAGNOSTICS

This error message is produced for structures for which the cumulative size of the structure's members exceeds the maximal number which can be represented in 29 bits (536870911).

5. [Limitation]: not enough space on frame for compiler-produced temporaries; no object file produced

This error message is produced when the cumulative size of the local variables and the temporary variables created by the compiler for computations exceeds the maximal number which can be represented in 29 bits (536870911).

6. [Limitation]: cumulative size of local variables is too large; no object file produced

This error message is produced when the cumulative size of the local variables and the temporary variables created by the compiler for computations exceeds the maximal number which can be represented in 29 bits (536870911).

7. [Limitation]: size too large for symbolic information, confuse the debugger

This error message is produced for objects (structures/unions or arrays) who's size is greater than the maximal number which can be represented in 16 bits (65535). Code continues to be generated, however the operation of the debugger may be affected.

8. [Limitation]: array dimension too large for symbolic information; may confuse the debugger

This error message is produced if a number greater than the maximal number which can be represented in 16 bits (65535) is used for an array's dimension. Code continues to be generated, however the operation of the debugger may be affected.

C.2.4 Syntax Errors

When the compiler detects a syntactic error, an attempt is made to fix the error by internally changing the input to a syntactically legal phrase. This is done in order to continue compilation and produce a maximum number of useful diagnostics; it should not be used as a means to correct your program.

An error is fixed by the deletion, insertion or replacement of a token, or by the skipping of a language phrase. The appropriate action is selected by the compiler based on the context of the source code, semantic information and a set of heuristics. Once a successful change in the source code is introduced, a syntax error is reported to indicate the change. If the change does not result in a legal syntactic phrase, the compiler skips to a point where the text is synchronized with the language syntax.

NOTE: In some cases the change introduced by the compiler may not be appropriate for your source code.

Examples:

1. Token deletion

```
"filename", line 1: a) [syntax]: Unexpected ',' int i,,j;
```

The unexpected token ',' was deleted internally in attempt to continue compilation.

2. Token insertion

```
"filename", line 3: a) [syntax]: ';' may be missing before '}' i= i }
```

The token ';' was inserted before the token '}' in an attempt to continue compilation.

3. Token replacement

```
"filename", line 4: a) [syntax]: ']' unexpected, ';' may be more appropriate
i ++]
```

The unexpected ']' token was replaced with a ';' in an attempt to continue compilation.

4. Skipping a language phrase

```
"stam.c", line 6: a) [syntax]: Unexpected "p", more errors may follow... skipping until ';' on line 13 xxxx ;p
```

Indicates the beginning of the skipped phrase.

```
"stam.c", line 13: a) [syntax]: ... skipped until here (from
               line 6)
foo(){};
-----a----a-------
```

Indicates the end of the skipped phrase.

Additional errors occasionally result from the fact that part of the program was skipped by the compiler.

C.2.5 Severe Errors

Severe errors are caused when the compiler detects semantic violations of the language rules, where the programmer's intention is not clear to the compiler. In this case, an error message is displayed and code generation is terminated. For example:

```
"stam.c", line 10: a) [severe]: traditional and prototype
             parameters cannot be mixed
foo(int,a){};
```

C.2.6 Caution Errors

Caution errors are issued for erroneous language constructs which the compiler either "thinks" may be on purpose or "guesses" the programmers intention. An error message is displayed and generation of code continues. If the programmer is satisfied with the compiler's action, the program produced may be run.

For example, if i was declared as volatile int i;:

```
"stam.c", line 2: a) [caution]: volatile pointer mismatch
int *p = &i;
----a----a-----
```

C.2.7 Warnings

Warning messages are issued for input which conforms to the language, but is deemed to be inappropriate by the compiler in the context found. An error message is displayed and generation of code continues.

For example:

Warning messages can be disabled by the -w compiler option on UNIX systems (/NOWARNING on VMS).

Appendix D

COMPILER OPTIONS

D.1 INTRODUCTION

This appendix contains tables for quick reference to the GNX-Version 4 C compiler options. These tables list:

- Options to the compiler on UNIX systems
- Options to the compiler on VMS systems
- Options to the compiler that pass to the C preprocessor on UNIX systems
- Options to the compiler that pass to the C preprocessor on VMS systems
- Options to the compiler that pass to the linker

(Options that pass to the linker are relevant only for UNIX systems.)

Table D-1. UNIX Operating System Options Sheet 1 of 2

OPTION	FUNCTION
-A	Allocate variables as standard.
-aflags	Generate runtime checks.
-В	Add code for profile information gathering.
-c	Suppress loading, force production of object file in file.o.
-d	This option is useful only when compiling Pascal and FORTRAN 77 programs.
-Fflags	Set optimization flags but do not call optimizer.
-f	Use floating-point emulation library.
-g	Prepare symbolic debug information for debugger.
-Jwidth	Force alignment boundary within structs to width.
-KCcpu -KFfpu -KBbus	Set target CPU. Set target FPU. Set target buswidth.
-llib	Use lib as a program library.
-m	Use m4 as the preprocessor for FORTRAN 77 and assembly files.
-n	Put C source lines as comments into assembly output file.
-N [flags]nnn	This option is useful when compiling FORTRAN 77 programs.

Table D-1. UNIX Operating System Options Sheet 2 of 2

OPTION	FUNCTION
-Oflags	Perform optimizations according to flags.
-X	Generate code that conforms to the Series 32000 architectural feature of modularity.
-p	Prepare profiling information for profiling.
-Q	Compile only, verify for syntax errors.
-R	Put all literal strings in read-only memory.
-S	Do not assemble, leave assembly in file.s.
Т	This option is only useful when compiling FORTRAN 77 programs.
-v	Verbose: list the subprograms as actually called by the driver.
-vn	List the subprograms to be called, but do not actually execute them.
-Wx,options	Pass options to compiler phase x. x can be p (C preprocessor), a (assembler), or 1 (linker).
-w	Suppress warnings.
-w66	This option is only useful when compiling FORTRAN programs.
-Zc	Use an alternate library.

Table D-2. VMS Operating System Options Sheet 1 of 2

OPTION	FUNCTION
/[NO]OBJECT[=filename]	[do not] Generate an object file during the compila- tion process.
/[NO]OPTIMIZE[=(flags[])]	[do not] Perform optimizations [according to flags].
/CHECK[=(option[,])]	Generate run-time checks.
/[NO]DEBUG	[do not] Prepare symbolic debug information for debugger.
/[NO]GATHER	[do not] Add code for profile information gathering.
/[NO]PROFILE	[do not] Prepare profiling information for profiling.
/[NO]ASM[=filename]	[do not] Generate an assembler file during the compilation process.
/[NO]ANNOTATE	[do not] Put C source lines as comments into assembly output file.
/[NO]ROM_STRINGS	Put all literal strings in read-only memory.
/ALIGN[=width]	Force alignment boundary within structs to width.
/[NO]WARNING	[do not] Output warning diagnostics.
/[NO]STANDARD	Allocate variables as standard.

Table D-2. VMS Operating System Options Sheet 2 of 2

OPTION	FUNCTION
/[NO]PRE_PROCESSOR	[do not]Run the source code through the cpp preprocessor.
/[NO]VERBOSE	[do not] List the compiler subprograms called by the driver.
/[NO]VN	[do not] List the subprograms to be called, but do not actually call them.
/TARGET=(CPU=cpu) /TARGET=(FPU=fpu) /TARGET=(BUSWIDTH=bus)	Set target CPU. Set target FPU. Set target buswidth.
/[NO]MODULAR	Generate code that conforms to the Series 32000 architectural feature of modularity.
/[NO]ERROR[=filename]	[do not] Generate an error log file during the compilation process.

Table D-3. Options Passed to the Preprocessor — UNIX Systems

OPTION	FUNCTION
-C	Prevent the macro preprocessor from removing comments.
-Dname=def -Dname	Define <i>name</i> to have the value <i>def</i> . Define <i>name</i> to have the value 1.
-E	Run only the preprocessor, send the result to stdout.
-Idir	Look for include files in dir after looking in the current directory.
-M	Generate makefile dependencies (cpp option).
-P	Run only the preprocessor, send the result to a preprocessed source. file.
-Uname	Undefine name.

Table D-4. Options Passed to the Preprocessor — VMS Systems

OPTION	FUNCTION		
/[NO]COMMENT	[do not] Prevent the preprocessor from removing comments.		
/DEFINE=(name [=def] [,])	Define name to the preprocessor.		
/[NO]EXPAND[=filename]	[do not] Generate a source file after preprocessing.		
/INCLUDE=(include_dir[,])	Look for include files in <i>include_dir</i> after looking for them in the current directory.		
/UNDEFINE=(name[,])	Undefine name to the preprocessor.		

Table D-5. Options Recognized and Passed to the Linker

OPTION	FUNCTION
-e epname	Define epname as entry point.
-o out	Name the compilation output file out.
-r	Retain relocation.
-s	Strip.
-u symname	Undefine symname in symbol table.
-v	Print linker version information.
-x	Do not preserve local symbols in the symbol table.
-i	Intialize variables in runtime.

EMBEDDED PROGRAMMING HINTS

E.1 INTRODUCTION

The GNX C compiler provides certain features which allow for programming of embedded applications in C. These features help solve the following issues:

- full control over memory allocation including RAM, ROM, stack space, trap and interrupt vectors, peripheral memory-mapped control registers.
- startup actions performed at system reset including initializing stack pointers, configuration registers, peripheral control registers, and timers.
- initialization of RAM data variables usually by copying from ROM or by zeroing.
- interrupt/trap handling

This appendix provides suggestions and examples for using the C compiler in embedded applications.

E.2 VOLATILE AND CONST

The const and volatile type qualifiers can be used in embedded applications to indicate ROM entities and memory mapped entities, respectively. A general overview of the semantics and use of these qualifiers is explained below. For further detail see Section E.2.5 and the ANSI C standard.

E.2.1 Const Type Qualifier

The value of an object (any lvalue expression) whose type includes the const qualifier cannot be modified. The const qualifier can be used for several purposes:

- 1. Constant strings can be made a part of the program code and placed into
- 2. Protecting variables from being changed. If during run-time an attempt is made to change a const variable, a trap will occur.

A non-volatile global or static object declared as const, will be allocated in read-only memory (the .text area) if it is initialized.

For example:

The const syntax allows for the declaration of both 'constant pointers' and 'pointers to constants'. For example:

The types pointer to const object and const pointer to object, as in the above example, have different meanings. The value of a pointer to const object can be modified; however the value of the pointed object can not be modified. In contrast, the value of a const pointer to an object can not be modified; however the value of the pointed object can be modified.

For example:

E.2.2 Volatile Type qualifier

The value of an object (any lvalue expression) whose type includes the volatile qualifier can be used or changed by asynchronous events (such as I/O or interrupts). Such an object should not be subject to any optimization that will change or delay references to it.

By using the volatile qualifier, you can specify volatile objects. Therefore, full optimization is carried out on all other objects, including global variables and pointer dereferences.

For example, in the following code

```
volatile int i;
int j;
foo() {
    for (i=1 ; i < j; i++) {
    }
}
```

the compiler can put j in a register. But for i this optimization is not permitted.

The volatile syntax allows for the declaration of both 'volatile pointers' and 'pointers to volatiles'.

For example:

```
/* pc is defined as pointer to char */
char * pc;
volatile char * pvc; /* pvc is defined as pointer to */
                   /* volatile char
char * volatile vpc; /* vpc is defined volatile pointer */
                    /* to char
volatile char * volatile vpvc;
                    /* vpvc is defined as volatile */
                    /* pointer to volatile char
```

The types pointer to volatile object and volatile pointer to object, as in the above example, have different meanings. References to a pointer to volatile object can be optimized; however references to the pointed object can not be optimized. In contrast, references to a volatile pointer to an object can not be optimized; however references to the pointed object can be optimized.

Assignment of pointer to volatile object to pointer to object is permitted only if an explicit cast is used. Not using an explicit cast causes an error message.

r not del

For example:

E.2.3 Memory Allocation

Memory allocation is performed by the operating system in native programming environments such as UNIX. However, embedded applications require the ability to control memory allocation. This is achieved by specifying in the linker directive file:

- · the memory ranges of various program sections.
- the division of program sections into ROM and RAM.
- the sections to be copied from ROM to RAM at program startup.

A complete description of the linker directive file is provided in Chapter 3 of the GNX Linker User's Guide. Figure E-1 is an example of a simple linker definition file for defining two areas of memory.

```
MEMORY {
    ROM : origin=0x1000 length=0x2000
    RAM : origin=0x10000 length=0x80000
}

SECTIONS {
    .text INTO(ROM) : { *(.text) }
    .data INTO(RAM) : { *(.data) }
    . . .
}
```

Figure E-1. Example of Linker Directive File

E.2.4 Initialized C Variables

The C programming language allows compile-time initialization of global and static variables. In addition, uninitialized global and static variable are defined by the C language to have a zero value at program startup.

In native environment, initialization is handled by the compiler and the operating system. In cross environment, when loading the program with the GNX debugger, the debugger performs these initializations. However in embedded applications, all initialized data resides in ROM and must be explicitly copied to RAM at program startup. The GNX linker directive file and the GNX run-time library are used to automatically

initialize RAM variables.

Refer to the GNX Linker User's Guide for further details.

E.2.5 Programming Memory Mapped Devices

When writing code for the registers of memory mapped peripheral, correctly and efficient accessing these entities can be problematic. However, the GNX C compiler allows optimization of such code.

The volatile qualifier should be used to specify the memory mapped entities. This allows the optimizer to perform optimizations without changing or delaying references to these entities.

An example of the correct way to code memory mapped entities is:

```
#define ctrl reg *((volatile short *)0xffe8)
foo()
{
   return ctrl reg;
```

This will result in:

```
movxwd @(65512),r0
ret
```

NOTE: Do not define a global pointer variable, such as volatile short *ctrl reg = (short *) CTRL REG; for memory mapped entities. Dereferencing such a pointer, as in *ctrl reg, will result in less efficient code.

E.3 ASM STATEMENTS

The asm keyword (see Section 3.3.2) provides for the unlimited insertion of assembly language statements into any position in the code. It is recommended to use this feature only for actions not codeable in standard C, such as processor register manipulation and svc calls on Series 32000 development boards.

Extreme care should be taken especially when asm is used in conjunction with the optimizer. See Section 6.6.6 for details.

The following example is a routine that will change the CPU status to supervisor mode on a NS32CG16ED board:

```
change to supervisor()
{
                               /* svc call number 9 */
    asm("movd $9,r0");
    asm("movd $0x55555555,r1"); /* security code */
    asm("svc");
}
```

The following example assumes that the stack pointer symbol is defined in the linker directive file (see the GNX Linker User's Guide for more details) as stack pointer = 0x003ffff0.

The following code will load the stack pointer register with the above value:

```
asm("addr
           stack pointer, r0");
asm("lprd
           sp, r0");
```

E.4 EXAMPLES OF PROGRAMMING WITH INTRINSIC FUNCTIONS

This section describes programming with intrinsic functions. More details can be found in Chapter 8.

E.4.1 NS32CG16 bit instructions

An example of a graphic application based on certain special NS32CG16 core bit operations is illustrated in this section. The image is represented by a bit-map 80 bits wide and 21 lines high. The picture is drawn by printing the bit-map in an ascii format in the following way:

- 1. A set bit in the bit-map is represented by the character '*'.
- 2. A clear bit is represented by a space.

It is important to include the proper header file, with intrinsic routines declarations, in your application. In this example, cg16.h is included.

The following definitions are used throughout the example:

```
#define PAGE WIDTH IN BYTES
                              10
#define PAGE WIDTH IN BITS
                              (PAGE WIDTH IN BYTES
#define PAGE HEIGHT
#define Page(y,x) (page + (y) *PAGE_WIDTH_IN_BYTES + (x))
```

The bitmap is kept in the following char array:

```
char page[PAGE WIDTH IN BYTES * PAGE HEIGHT];
```

A 'V' is drawn in the upper-left part of the image using the spitps intrinsic function:

```
draw a v() {
  int offset = 3;
   sbitps(Page(2,0),&offset,7,PAGE WIDTH IN BITS + 1);
   _sbitps(Page(2,0),&offset,8,-PAGE WIDTH IN BITS + 1);
}
```

The resulting image is shown in Figure E-2

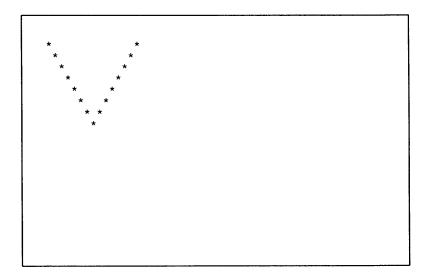


Figure E-2. The Image

A reversed image of 'V' figure is drawn by using the bbors intrinsic function:

```
/*
* Copy a bit-block (reversed) containing the V, to
* the lower right corner of itself.
* Mask off the areas to the right and left of the image,
* to protect the image of the V drawn above.
_bbor_s(Page(1,0),Page(9,2),3,10,0x0fffe,0x0000f,
        PAGE_WIDTH_IN_BYTES-2, PAGE_WIDTH IN BYTES-2, 2);
```

The resulting image is shown in Figure E-3.

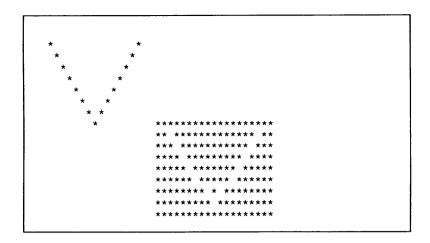


Figure E-3. The Image with the Reversed Shape

To print the image in the format of the example, the intrinsic routine _tbit() (a general 32000 instruction) is used:

```
dump as ascii()
* Displays the bitmap "page" as an ascii picture.
* Each bit is displayed as a character:
           '*' if it is ON, ' ' (space) if it is OFF.
 *----*/
#include <ns32000.h>
dump as ascii()
{
 int line, offset;
 char *pict;
 for (line = 0, pict = page;
           line < PAGE HEIGHT;
              line++, pict += PAGE WIDTH IN BYTES)
 {
/*
 * Use of thit to test each bit on the current line.
 */
     for (offset = 0; offset < PAGE WIDTH IN BITS; offset++)</pre>
         putchar( tbit(offset, pict) ? '*' : ' ');
     putchar('\n');
  }
}
```

The code is compiled using the following syntax:

UNIX environment

nmcc -KCG16 filename

VMS environment

NMCC /TARGET=(CPU=CG16) filename

E.4.2 NSGX320 specific instructions

An example of a part of an implementation of a digital FIR (Finite Impulse Response) filter is illustrated in this section.

The following definitions of the types WCOMPLEX and DCOMPLEX (used in the example) are found in gx320.h (see Section 8.4.1):

```
typedef struct WCOMPLEX {
      short re;
      short im:
} WCOMPLEX;
typedef struct DCOMPLEX {
      long re;
      long im;
} DCOMPLEX;
```

This example shows a common operation in various DSP (Digital Signal Processing) applications. The function performs a complex multiply and accumulate operation on two complex vectors of length ten, and then scales down the complex result from 32 bit to 16 bits. The C code is:

```
#include <qx320.h>
#define HALF 16384
#define SHIFT 15
WCOMPLEX b[10];
WCOMPLEX a[10];
WCOMPLEX c;
DCOMPLEX init result = {0,0};
for fir()
   DCOMPLEX result:
    int i:
    /* initialize result to 0 */
   result = init result;
    for(i = 0 ; i < 10; i++)
        _cmacd(&result, a[i], b[i]);
    /* now scale down from 32 bits back to 16 bits */
   c.re = (result.re + HALF) >> SHIFT;
   c.im = (result.im + HALF) >> SHIFT;
}
```

The following code was produced for the for_fir() function, when compiled with -KCGX320 and -O:

```
for fir:
movd
        init result,r0
        init result+(4),r1
movd
        _a, b
cmacd
        a+(4), b+(4)
cmacd
cmacd
        a+(8), b+(8)
        _a+(12),_b+(12)
cmacd
        a+(16), b+(16)
cmacd
cmacd
        _a+(20),_b+(20)
        _a+(24),_b+(24)
cmacd
        _a+(28),_b+(28)
cmacd
        _a+(32),_b+(32)
cmacd
cmacd
        a+(36), b+(36)
        $(0),tos
movqd
        $(0),tos
movqd
        r0,0(sp)
movd
movd
        r1,4(sp)
movd
        0(sp), r0
movd
        4(sp),r1
        $(16384),r0
addd
addd
        $(16384),r1
ashd
        (-15), r0
ashd
        (-15), r1
cmpd
        tos, tos
movw
        r0, c
        r1,_c+(2)
movw
ret
```

NOTE: The nop instructions were used because r0 and r1 can not be accessed for two instructions after the cmacd instruction.

E.5 PROGRAMMING TRAP/INTERRUPT ROUTINES

The example used in this section is a clock display for the time of day. The routine clock_handler handles a clock interrupt, which occurs TICKS_PER_SECOND times per second. The time display is updated every second.

Since the routine does not use floating-point registers, save_regs=int_regs is specified in the pragma directive (i.e. only integer registers are saved). The saved registers include all those used by the routine. Scratch registers are also saved because the handler calls the routine update time display.

The C code for the clock interrupt handler is:

```
#pragma interrupt(clock int routine, save regs=int regs);
void clock int routine (void)
   static int counter;
   static int hours:
   static int minutes;
   static int seconds;
   counter++;
   if (counter == TICKS PER SECOND)
        seconds++;
        counter = 0;
        if (seconds == 60)
          {
            minutes++;
            seconds = 0:
            if (minutes == 60)
                 hours++;
                 minutes = 0;
                 if (hours == 24)
                    hours = 0;
               }
         update time display (hours, minutes, seconds);
     }
}
```

Certain CPUs of the Series 32000/EP microprocessor family can be set to work in either direct or indirect exception mode.

When direct exception mode is enabled the address of the trap handler (residing in the interrupt dispatch table) is interpreted by the CPU as a pointer. The clock interrupt entry in the interrupt dispatch table should be set to the address of clock int routine. The following line is inserted to the clock interrupt entry in the initialization of the interrupt dispatch table

```
.double @ clock int routine
```

When direct exception mode is disabled (or non existent), the address of the trap handler (residing in the interrupt dispatch table) is interpreted by the CPU as an external procedure descriptor (i.e. mod + offset). The clock interrupt entry in the interrupt dispatch table should be set to the descriptor of _clock_int_routine. The following line should inserted to the clock interrupt entry in the initialization of the interrupt dispatch table

```
.xpd clock int routine
```

In addition the interrupt handler should be associated to a module table entry. It is recommended to do so by adding the following asm statement to the C source file before the interrupt handled definition

```
asm(".module");
```

For more details on modular and direct exception mode see the Series 32000 instruction set and the GNX Assembler manual.

INDEX

	c -	D O	A	2-2
A - A 2.6	_ prefix	B-2	Assembly program	2-2
A A 2-6 B -a 2-4, 2-14 abs 8-10 Absolute value instructions 8-2 Accumation of profile information in pfb_exit object file sime considerations 7-2 improving code integrations 6-12 improving code 6-11 set bit 8-3 integra variables integrate allocation 6-15 static functions 6-12 register allocation 6-15 static functions 6-12 register allocation 6-15 static functions 6-15 static functions 6-16 Address at standard 8-10 Allocate variables as standard 8-10 Allocate variables as standard 8-10 Allocate variables as standard 8-11 Allocate variables as standard 8-12 Agnment reference 8-13 Agnment stack in calling sequence 8-14 Agnment stack in calling sequence 8-15 Agnment stack in calling sequence 8-15 Agnment stack in calling sequence 8-15 Agnment stack in calling sequence 8-16 Agnment stack in calling sequence 8-17 Agnment stack in calling sequence 8-18 Agnment 8-18 Agnme	_ surnx	B-3	•	2.2
A A 2-4, 2-14 abs				
A	A		AVAIL_SWAP	0-17
Absolute value				
Absolute value	-A		${f B}$	
Absolute value Absolute value instructions Absolute value instructions Absolute value instructions Accumulation of profile information 7-2	-a	,		
Absolute value instructions 8-2 count printed by sprof 7-9, 7-10 Accumulation of profile information disabling 7-5 sprof information 7-2 gathering profile information 7-2 for profile information 7-5 sprof information 7-7 Additional code for profile information 7-6 Basic-block			_	2-4
Accumulation of profile information disabling 7-2 disabling gathering profile information sprof information 7-2 sprof information 7-7 7-7 7-7 7-7 7-7 7-7 7-7 7-7 7-7 7-7				
disabling	Absolute value instructions		count printed by sprof	7-9, 7-10
Additional code for profile information in pft_exit object file space considerations space considerations time considerations 7-6 bit aligned word transfer space considerations 8-19 Additional guidelines as as statements 6-14 Bit instructions 8-2, 8-11 Additional guidelines as as statements 6-14 find first set 8-4 floating-point computations integer variables 6-11 set bit 8-3 integer variables 6-11 bit bit 8-3 local variables 6-16 direction 8-16 optimizing for space 6-16 direction 8-16 optimizing for space 6-16 direction 8-16 pointer usage 6-12 source inversion 8-16 setimp() 6-15 Bit field 8-16 setimp() 6-15 Bit field 8-5, 8-7 Address taking 15-6 Bit field 8-2, 8-11 Allication of memory 8-2 Bit field 8-16 Allocate variables as standard 2-6, 2-10 Board 4-1 Allocation of memory 8	Accumulation of profile information	7-2	gathering profile information	
in pfb_exit object file 7-3 definition of F-1 space considerations 7-6 Bit aligned word transfer 8-19 Additional guidelines 6-14 linstructions 8-2, 8-11 Additional guidelines 6-14 find first set 8-3 asm statements 6-12 invert bit 8-3 improving code 6-11 set bit 8-3 integer variables 6-11 test bit 8-3 local variables 6-11 bitblt 8-3 optimizing for space 6-16 direction 8-16 optimizing for space 6-15 suffixes 8-16 setjmp() 6-15 suffixes 8-16 setjmp() 6-15 Bit-field 8-16 setjmp() 6-15 Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 AllIGN 2-10 bitwt 8-19 all field 8-2, 8-1 8-2 AllIGN <td< td=""><td>disabling</td><td></td><td></td><td>7-7</td></td<>	disabling			7-7
space considerations 7-6 Bit aligned word transfer 8-19 ttime considerations 7-6 Bit instructions 8-2, 8-11 Additional guidelines clear bit 8-3 asm statements 6-14 find first set 8-4 floating-point computations 6-12 invert bit 8-3 improving code 6-11 test bit 8-3 integer variables 6-11 bit bit 8-3 local variables 6-16 direction 8-16 optimizing for space 6-16 direction 8-16 optimizing for space 6-16 direction 8-16 optimizer usage 6-12 source inversion 8-16 setjimp() 6-15 BITBLT instructions 8-16 setjimp() 6-15 BITBLT instructions 8-2, 8-11 Address taking Bit-field instructions 8-2, 8-11 All LIGN 2-10 bitwt 8-19 Allocate variables as standard 2-6, 2-10 developement F-1	Additional code for profile information	7-2	Basic-block	
time considerations 7-6	in pfb_exit object file	7-3	definition of	F-1
Additional guidelines clear bit 8-3 as ms statements 6-14 find first set 8-4 floating-point computations 6-12 invert bit 8-3 improving code 6-11 set bit 8-3 integer variables 6-11 bit bit 8-3 local variables 6-11 bit bit 90 optimizing for space 6-16 direction 8-16 pointer usage 6-12 source inversion 8-16 register allocation 6-15 suffixes 8-16 setimp() 6-15 BITBLT instructions 8-16 setimp() 6-15 Bit field instructions 8-2, 8-16 Address taking Bit field instructions 8-2, 8-11 of intrinsic functions 8-2 Bitfields 8-5, 8-7 Aldress taking 8-2 Bit field instructions 8-2, 8-11 All LIGN 2-10 bit field instructions 8-2, 8-11 All Called variables as standard 2-6, 2-10 developement F-1	space considerations	7-6	Bit aligned word transfer	8-19
asm statements 6-14 find first set 8-4 floating-point computations 6-12 invert bit 8-3 improving code 6-11 set bit 8-3 improving code 6-11 set bit 8-3 improving code 6-11 set bit 8-3 integer variables 6-11 test bit 8-3 integer variables 6-11 test bit 8-3 integer variables 6-11 bit bit 8-16 direction 8-16 direction 8-16 direction 8-16 set bit 8-10 direction 8-16 set bit 8-16 direction 8-16 set bit 8-1 set bit <td>time considerations</td> <td>7-6</td> <td>Bit instructions</td> <td>8-2, 8-11</td>	time considerations	7-6	Bit instructions	8-2, 8-11
asm statements 6-14 floating-point computations improving code 6-11 set bit 8-3 integer variables 8-16 direction 8-16 direction 8-16 direction 8-16 direction 8-16 set bit 8-16 direction 8-16 set bit 8-2 source inversion 8-16 set bit 8-17 set bit 8-17 set bit 8-17 set bit 8-17 set bit	Additional guidelines		clear bit	8-3
floating-point computations 6-12 invert bit 8-3 improving code 6-11 set bit 8-3 integer variables 6-11 test bit 8-3 local variables 6-11 Bitblt optimizing for space 6-16 direction 8-16 pointer usage 6-15 suffixes 8-16 setjimp() 6-15 suffixes 8-16 setjimp() 6-15 BITBLT instructions 8-16 static functions 6-11 Bit-field 8-5, 8-7 Address taking Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 of intrinsic functions 0-2, 8-2, 8-2, 8-2, 8-2, 8-2, 8-2, 8-2, 8		6-14	find first set	8-4
improving code 6-11 set bit 8-3 integer variables 6-11 test bit 8-3 local variables 6-11 Bitblt 8-3 optimizing for space 6-16 direction 8-16 pointer usage 6-12 source inversion 8-16 register allocation 6-15 BITBLT instructions 8-16 setimp() 6-15 BITBLT instructions 8-15 static functions 6-11 Bit-field instructions 8-2, 8-11 Address taking Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 /ALIGN 2-10 _bitwt 8-19 ALIGN 2-10 _bitwt 8-19 Allocate variables as standard 2-6, 2-10 _developement F-1 Allocate variables as standard 2-6, 2-11 _developement F-1 Allocate variables as standard 2-6, 2-11 _developement _c Anotated source file listing _c _c		6-12	invert bit	8-3
integer variables 6-11 bit bit bit		6-11	set bit	8-3
local variables optimizing for space 6-16 direction 8-16 pointer usage 6-12 source inversion 8-16 register allocation 6-15 suffixes 8-16 setjmp() 6-15 BITBLT instructions 8-15 static functions 6-11 Bit-field 85-5, 8-7 Address taking Bit-field Bit-field 8-15 8-2, 8-11 of intrinsic functions 8-2 Bitfields 3-7, 4-1 ALIGN 2-10 _bitwt 8-19 = 1 for space optimization 6-16 Board Alignment 2-6, 2-10 developement F-1 Allocate variables as standard 2-6, 2-11 Allocation of memory E-4 ANNOTATE 2-10, 6-10 C for debugging optimized code 6-9 ANSI C Standard 3-1 Calling conventions B-1 Application specific instruction set 8-1 Standard A-1 Argument Reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack in calling sequence 4-8, 5-10 CG G-16		6-11	test bit	8-3
optimizing for space pointer usage 6-16 counter usage 6-12 counter usage 8-16 counter usage 8-18 counter usage 8-18 counter usage 8-2 counter usage 8-2 counter usage 8-19 counter usage 9-2 counter	•	6-11	Bitblt	
Pointer usage			direction	8-16
register allocation 6-15 setjmp() 8-15 setjmp() 8-16 setjmp() 8-15 setjmp() 8-16 setjmp() 8-16 setjmp() 8-16 setjmp() 8-16 setjmp() 8-16 setjmp() 8-16 setjmp() 8-15 setjmp() 8-16 setjmp() 8-15 setjmp() 8-15 setjmp() 8-18 setjmp() 8-18 setjmp() 8-18 setjmp() 8-19 setjmp() 8-10 setjmp() 8-11 setjmp() 8-11 setjmp() 8-11 setjmp() 8-11 setjmp() 8-11 setjmp() 8-				
setjmp() 6-15 static functions BITBLT instructions 8-15 static functions Address taking of intrinsic functions 8-2 bit-field instructions 8-2, 8-11 bitwt 8-2, 8-11 bitwt 8-19 bitwt 8-10 bitwt				
static functions 6-11 Bit-field instructions 8-5, 8-7 Address taking of intrinsic functions 8-2 Bit-field instructions 8-2, 8-11 ALIGN 2-10 _bitwt 8-19 -1 for space optimization 6-16 Board Alignment 2-6, 2-10 _developement F-1 Allocate variables as standard 2-6, 2-11 _developement F-1 Allocation of memory E-4	•			
Address taking Bit-field instructions 8-2, 8-11 of intrinsic functions 8-2 Bitfields 3-7, 4-1 /ALIGN 2-10 _bitwt 8-19 = 1 for space optimization 6-16 Board Alignment 2-6, 2-10 developement F-1 Allocation of memory E-4	• • •			
of intrinsic functions 8-2 class and processed points and processed points are in calling sequence Bitfields 3-7, 4-1 class and points are in calling sequence 3-7, 4-1 class are in calling sequence 3-7, 4-1 class are in calling sequence Board 3-7, 4-1 class are in calling sequence Board 3-7, 4-1 class are in calling sequence 3-7, 4-1 class are in call in		0 11		
/ALIGN 2-10 _bitwt 8-19 =1 for space optimization 6-16 Board Alignment 2-6, 2-10 developement F-1 Allocate variables as standard 2-6, 2-11 developement F-1 Allocation of memory E-4 F-1 E-4 /ANNOTATE 2-10, 6-10 C C for debugging optimized code 6-9 C 2-5 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C Clanguage extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 var B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 _cbit 8-2, 8-3 in calling sequence A-1		8-2		
=1 for space optimization 6-16 Board Alignment 2-6, 2-10 developement F-1 Allocate variables as standard 2-6, 2-11 developement F-1 Allocation of memory E-4 C C /ANNOTATE 2-10, 6-10 C C for debugging optimized code 6-9 C 2-5 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C Clanguage extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 var B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1				•
Alignment 2-6, 2-10 developement F-1 Allocate variables as standard 2-6, 2-11 Allocation of memory E-4 /ANNOTATE 2-10, 6-10 C C for debugging optimized code 6-9 C 2-5 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C Clanguage extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM				0-13
Allocate variables as standard 2-6, 2-11 Allocation of memory E-4 /ANNOTATE 2-10, 6-10 C for debugging optimized code 6-9 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C C language extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 <td></td> <td></td> <td></td> <td>E 1</td>				E 1
Allocation of memory	3	•	developement	r-1
/ANNOTATE 2-10, 6-10 C for debugging optimized code 6-9 -C 2-5 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C C language extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 Asm statements 6-14 Char 4-1		•		
for debugging optimized code 6-9 Annotated source file listing -C 2-5 by sprof 7-7 -c 2-5 ANSI C C language extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Case sensitivity B-1 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 Asm statements 6-14 Char Char 4-1			C	
Annotated source file listing by sprof 7-7 C C 2-5 ANSI C extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set Argument reference B-3 Case sensitivity Tererce B-3 Caution Errors C-5 Argument stack in calling sequence A-1 CG16 ASIS B-1 CG160 B-1 ASM C-1 CG160 B-1 CG160 B-1 ASM C-1 CG160 B-1 CG160 B-1 CG160 B-1 CG160 B-1 CG160 CG-Core CG-Core CG-Core CG-Core CG-Core CHanging default optimization options C-4 Asm statements C-1 Changing default optimization options C-4 Asm statements		•	C	
by sprof 7-7 -c 2-5 ANSI C C language extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1		6-9	6	
ANSI C C language extensions 1-3 extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1				
extensions 3-1 Calling conventions ANSI C standard 3-1 in mixed language programming B-1 Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1		7-7	-	_
ANSI C standard Application specific instruction set Argument reference var B-3 Case sensitivity B-1 Var B-3 Caution Errors C-5 Argument stack in calling sequence A-1 CG16 ASIS B-1 CG160 B-1 ASM 3-6, E-6 CG-Core B-1 CG-Core B-1 CG16.h B-1 CG16.h B-1 CG-Core B-1 CG-Core B-1 ASM CHASM CHAS		0.1		1-3
Application specific instruction set 8-1 standard A-1 Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1			•	ъ.
Argument Calling sequence 4-8, 5-10 reference B-3 Case sensitivity B-1 var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1				
reference B-3 var Case sensitivity B-1 var Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS ASIS 8-1 CG160 8-1 ASIA Asm 3-6, E-6 CG-Core 8-1 CG-Core /ASM 2-10 ch16.h 8-2 Changing default optimization options Asm statements 6-14 Char 4-1		8-1		
var B-3 Caution Errors C-5 Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1	•		.	,
Argument stack _cbit 8-2, 8-3 in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1			•	
in calling sequence A-1 CG16 8-1 ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1		B-3		
ASIS 8-1 CG160 8-1 Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1	9			,
Asm 3-6, E-6 CG-Core 8-1 /ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1	in calling sequence			
/ASM 2-10 ch16.h 8-2 for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1	ASIS	8-1		8-1
for debugging optimized code 6-9 Changing default optimization options 6-4 Asm statements 6-14 Char 4-1		•		
Asm statements 6-14 Char 4-1	/ASM	2-10	ch16.h	8-2
Asm statements 6-14 Char 4-1	for debugging optimized code	6-9	Changing default optimization options	6-4
Assembler B-6 /CHECK 2-10, 2-14		6-14	Char	4-1
	Assembler	B-6	/CHECK	2-10, 2-14

cleanup		-1	2-8
for profile gathering	7-5	-1 -M	2-8 2-7
clear bit	8-2, 8-3	-m	2-6
close	o -, o o	/MODULAR	2-11
for profile gathering	7-5	-n	2-5
_cmacd	8-29	-0	2-4, 6-3
CMDDIR	2-18	-0	2-5
_cmult	8-28	/OBJECT	2-9
	2-2, 5-9, B-6	/OPTIMIZE	2-9, 6-3
Code portability	4-1, 6-5	-P	2-8
Code-generator	,	- -p	2-4
definition of	F-1	PRE_PROCESSOR	2-12
CODE_MOTION optimization option	6-2, 6-4,	-Q	2-4
<u>-</u>	6-17	-Ř	2-5
Coloring algorithm	5-8	-r	2-8
Command line	2-2	/ROM_STRINGS	2-10
.comment	3-7	-S	2-5
COMMENT	2-12	-s	2-8
Common subexpression elimination	5-1, 5-6	/STANDARD	2-11
Common subexpressions	6-13	/TARGET	2-11
Compatibility	0-10	-U	2-8
PIT file and source file	7-10	-u	2-8
Compilation	7-10	/UNDEFINE	2-12
for profile information	7-2	-V	2-8
Compilation options	1-2	-v -v	2-6
UNIX	2-2, 2-4	/VERBOSE	2-11
VMS	2-2, 2-4	-vn	2-11
	2-3 2-1	/VN	2-11
Compilation process			2-11 2-8
Compilation time requirements	6-17	-W	2-6 2-6
Compilation unit	T3 1	-W	2-6 2-10
definition of	F-1	/WARNING	
Compile but do not link	2-5	-X	2-7
Compile leaving assembly files	2-5, 2-10	-X	2-8
Compiler options	0.0	-Z	2-6
- A	2-6	Compiler structure	2-1 2-2
-a	2-4, 2-14	code generator	
/ALIGN	2-10	driver	2-1
/ANNOTATE	2-10	front end	2-1
/ASM	2-10	language parser	2-1
-В	2-4	macro preprocessor	2-1
-C	2-5	optimizer	2-1
-C	2-5	Compiling mixed-language programs	B-5
/CHECK	2-10, 2-14	Compiling system code	6-6
/COMMENT	2-12	Complex multiply and accumulate double	
-D	2-7	Complex multiply double	8-28
/DEBUG	2-10	Configuration	
/DEFINE	2-12	cross	F-1
-E	2-7	native	F-2
-e	2-8	Const	3-2, E-1
/ERROR	2-11	definition of	F-1
/EXPAND	2-12	Constant folding	5-1, 5-2
- F	2-4	Conversion	
-f	2-7	definition of	F-1
-g	2-4	Convert to bit pointer	8-9
/GATHER	2-10, 7-4	Copy propagation	5-2
-I	2-7	Count of source-line executions	7-7
/INCLUDE	2-12	Cross configuration	
- J	2-6	definition of	F-1
-K	2-6, 2-13	cvtp	8-2
	-		

auto.	0.0	English shirt Cl.	
_cvtp	8-9	Executable object file definition of	F-2
		exit	F-Z
D		for profile gathering	7-5
_		_exit routine	
-D	2-7	for profile information	7-3
Data flow analysis	5-2	/EXPAND	2-12
db_pfb_exit.o	7-4	_extblt	8-13
Dbug		Extensions	
definition of	F-1	\$ sign in identifiers	3-7
Dead code removal	5-1, 5-5	ANSI C	3-1
/DEBUG	2-10	asm keyword	3-6
disabling FIXED_FRAME optimiz		bitfields	3-7
Debugging of optimized code	6-9	const	3-2
Declaration	77.4	enumerated type	3-3
definition of	F-1	floating-point constants	3-2
Declaration conventions	ъ.	for embedded programming	3-4
in mixed language programming	B-1	function prototype	3-1
Default optimization options	6-3 6-4	ident	3-7 3-4
changing Define		Interrupt/Trap Routines Support	3- 4 3-7
/DEFINE	2-7, 2-12 2-12	intrinsic routines	3-7 3-2
Define entry point	2-12 2-8	pragma signed keyword	3-2 3-2
Development board	2-0	string concatenation	3-2
definition of	F-1	structure handling	3-3
Directive file	F-1	unsigned constants	3-3
linker	E-4	volatile	3-3
Disabling profile information accumu		Extensions to structures	3-3
tion	7-5	Extensions to the C language	3-1
Driver		External bit aligned block transfer	8-13
definition of	F-1	External functions	B-4
Driver program	2-1	External procedures	B-4
		External variables	B-4
		exti	8-5
${f E}$		Extract bit-field	8-5
-E	2-7		
-е	2-8	F	
Embed source lines as comments	2-5, 2-10		
Embedded extensions	3-4	- F	2-4
Embedded programming hints	E-1	- f	2-7
Embedded support	3-1	fabs	8-10
Emulation (floating-point)		fclose	
definition of	F-1	for profile gathering	7-5
Entry point	B-5	Features	1-3
Enumerated type	3-3	ffabs	8-10
Environment variables	2-18	ffs	8-2
AVAIL_SWAP	6-17	_ffsb	8-4
PITFILE	6-16, 7-3	_ffsd	8-4
errno for profile gathering	7-5	_ffsw	8-4
Error	1-0	fgets for profile gathering	7-5
detection	C-1	Filename conventions	2-3
recovery from	C-1	Files	2-3 2-3
/ERROR	2-11	assembly	2-3 2-2
etext	2-11	executable	2-2 2-2
for profile gathering	7-5	object	2-2 2-2
Executable program	2-2	Find first set	8-2, 8-4
L. AB. AW.	22	Fixed frame	5-1, 5-10
			0 1, 0-10

FIXED_FRAME optimization option	6-2, 6-4,	I	
	6-7		
-Fl	8-1	-I	2-7
FLOAT_FOLD optimization option		_ibit	8-2, 8-3
Floating-point arithmetic	4-10	#ident	3-7
Floating-point computations	6-12	Identifiers	0.5
Floating-point constants	3-2	\$ sign	3-7
	2-7, 2-17, F-1	IEEE standards	TP 0
Cross-Configuration/UNIX system	$\begin{array}{c} 2-17 \\ 2-17 \end{array}$	definition of	F-2 4-1
native configuration VAX/VMS system	2-17 2-18	Implementation issues	B-4
Flow optimizations	5-1, 5-4	Importing routines and variables Improved annotation	6-10
fopen	0-1, 0-4	/INCLUDE	2-12
for profile gathering	7-5	INCLUDEPATH	2-12
fprintf	1-0	Incompatibilities with GNX C compiler	2-10
for profile gathering	7-5	version 3	1-5
fputs	1-0	Induction variable elimination	5-1, 5-7
for profile gathering	7-5	Initialization	0 1, 0 .
FRAME_ALLOCATION optimization		of structures	3-3
option	6-2, 6-4	Initialization code	B-5
Front end	2-1, B-6	Initialization of variables	E-4
Function call	- 1, 2 0	Initializer	
intrinsic	8-1	definition of	F-2
Function prototype		_ins	8-7
intrinsic functions	8-1	Insert bit-field	8-7
Function return value	A-2	Instructions	
Functions		application specific	8-1
Function prototypes	3-1	Integer variables	6-11
FX16	8-1	Intermediate form	2-1
		Interrupt and trap routines	
		programming	3-4
G		Interrupt handler routine	
		programming examples	E-12
-g	2-4	Intrinsic function	
disabling -OF optimization	6-9	redefinition	8-1
/GATHER	2-10, 7-4	Intrinsic functions	8-1
Gather profile information	2-10	general description	8-1
Gathering profile information	2-4, 7-2	NS32GX320	8-25
Generate an error log file	2-11	programming examples	E-7
Generate makefile dependencies	2-7	use	8-1
Generate modular code	2-7, 2-11	Intrinsic routines	3-7
getenv		Run-time parameter checks	2-15
for profile gathering	7-5	invert bit	8-2, 8-3
Global variables	B-4	Invocation syntax	
gn320.h	8-2	UNIX	2-2
GTS		VMS	2-9
target setup	2-2	_iob	
Guidelines on using the optimizer	6-1	for profile gathering	7-5
GX320	8-1		
		-	
TT		J	
Н		-J	2-6
Header files		-J -J1	∠-0
for intrinsic functions	0.0		6-16
Hints	8-2	for space optimization	0-10
for embedded programming	E-1		
Host machine	E-1		
definition of	F-2		
deninition of	r-Z		

K		Macro	
		definition of	F-2
-K	2-6	Macro preprocessor	2-1
-KB1		_mactd	8-30
for space optimization	6-16	Main program	B-5
Keyword		Memory allocation	4-9, E-4
definition of	F-2	Memory layout optimizations	5-1, 5-10
Keywords	_	Memory mapped devices	
asm	3-6, E-6	programming	E-5
const	E-1	Memory representation	4-1
Signed	3-2	Mixed-language programming	2-3, 4-8, B-1
volatile	E-2	Compilation on UNIX operating	
		Compilation on VMS operating	systems B-7
L		mktemp	7 5
L		for profile gathering	7-5
1	0.0	/MODULAR	2-11
-l	2-8	with profile information	7-4
Language parser	2-1	Monitor definition of	F-2
Leave comments in	2-5, 2-12	Move multiple pattern	
Libc symbols	7 5		8-20
used for profile gathering LIBPATH	7-5 2-18	_movmp Multiply and accumulate twice dou	8-20 1ble 8-30
Library function	2-10	Multiply and accumulate twice dot Multiply word to double	8-27
•	B-1	mulwd	8-27
reuse	6-7, B-5	_muiwa	0-41
Library routines Limitation Errors	6-7, B-5 C-2		
Linker	2-2, 2-3, B-6	N	
	2-2, 2-3, B-6 2-8	14	
compiler options passed to definition of	F-2	n	25 6 10
Linker directive file	F-2 E-4	-n for debugging optimized code	2-5, 6-10 6-9
example	E-4 E-4	Name sharing	0-3
Linker version	2-8	in mixed language programmin	g B-1
Linking phase	B-7	Native configuration	g D-1
Literal strings in read-only memory		definition of	F-2
Local variables	6-11	Nburn	r-2
Longimp()	6-15	definition of	F-2
Loop	0-10	NIL pointer checks	2-16
definition of	F-2	No local symbols in symbol table	2-10
Loop invariant code motion	5-1	NOOPT optimization option	6-2, 6-4, 6-16
Loop invariant expressions	5-6	NO_STANDARD_LIBRARIES	8-1
Loop unrolling	5-1, 5-4	ns32000.h	8-2
LOOP_UNROLLING optimization		NS32CG16	8-1
	6-4	NS32CG160	8-1
Low-level interface	6-7	NS32FX16	8-1
relying on frame structure	6-7	NS32GX320	8-1
relying on register order	6-7	intrinsic functions	8-25
using asm statements	6-7	typedefs for intrinsic functions	8-26
Lvalue			
definition of	F-2		
		О	
M		-0	2-4, 6-3
		Object	
-M	2-7	definition of	F-2
-m	2-6	/OBJECT	2-9
Machine		Object code program	2-2
host	F-2	Object file	
target	F-3	definition of	F-2
		-O1	0 1

Old fashioned compound assignmen	nt 3-4	dofino ontre noint	9.0
Old fashioned initialization	3-4	define entry point embed source lines as comments	2-8 s 2-5, 2-10
Optimization	0-4	floating-point emulation	2-3, 2-10 2-7
definition of	F-3	gather profile information	2-4, 2-10
	6-1	<u> </u>	2-4, 2-10
Optimization flags	0-1	generate error log file generate makefile dependencies	
Optimization options	C 4		
changing default	6-4	generate modular code	2-7, 2-11
default on	6-3	leave comments in	2-5, 2-12
default on VMS	6-3	linker version	2-8
Optimization options on the comma		no local symbols in symbol table	
UNIX systems	6-3	optimize	2-4, 2-9
VMS systems	6-3	pass options	2-8
Optimization techniques	5-1	pass to C preprocessor	2-12
Optimizations		profile information	2-4
common subexpression eliminat		quick compilation	2-4
constant folding	5-1	read-only memory	2-5, 2-10
dead code removal	5-1	redirect output to .i file	2-8
fixed frame	5-1	rename output file	2-5
flow optimizations	5-1	retain relocation	2-8
induction variable elimination	5-1	run cpp only	2-7, 2-12
loop invariant code motion	5-1	run-time checks	2-4, 2-10
loop unrolling	5-1	set target	2-6, 2-11, 2-13
memory layout optimizations	5-1	show do not execute	2-6, 2-11
partial redundancy elimination	5-1	specify include file directory	2-7, 2-12
peephole optimizations	5-1	specify program library	2-8
redundant assignment eliminat		strip	2-8
register allocation	5-1	undefine	2-8, 2-12
	5-1, 5-11, 6-16	undefine symbol in symbol table	•
Runtime feedback	7-11	use alternative library	2-6
strength reduction	5-1	use the m4 preprocessor	2-6
value propagation	5-1	verbose	2-6, 2-11
Optimize	2-4, 2-9	warning diagnostics	2-6, 2-10
/OPTIMIZE		Order of evaluation	2-0, 2-10 4-9
	2-9, 6-3		
CODE_MOTION	6-2, 6-4, 6-17	Overview	1-1
FIXED_FRAME	6-2, 6-4, 6-7		
FLOAT_FOLD	6-2, 6-4, 6-8		
FRAME_ALLOCATION	6-2, 6-4	P	
LOOP_UNROLLING	6-2, 6-4	_	
NOOPT	6-2, 6-4, 6-16	-P	2-8
REGISTERS_ALLOCATION	6-2, 6-4, 6-17	-p	2-4
RUNTIME_FEEDBACK	6-2, 6-4, 6-16	Partial redundancy	5-6
SPEED_OVER_SPACE	6-2, 6-4, 6-16	Partial redundancy elimination	5-1
STANDARD_LIBRARIES	6-2, 6-4, 6-7	Pass options to compilation phase	2-8
USER_REGISTERS 6-2	, 6-4, 6-7, 6-15	Pass source file to the C preprocess	or 2-12
VOLATILE	6-6	Pcc	3-1
VOLATILE optimization	6-2, 6-4	Peephole optimizations	5-1, 5-9
Optimizer	2-1, 5-2, B-6	pfb_exit.o and pfb_exit.obj	7-3
definition of	F-3	Pgen	7-3
Optimizing for space	6-16	running on VMS	7-5
Option		PIT file	7-2, 7-3
definition of	F-3	PITFILE	2-18, 6-16, 7-3
Options	2-4	Pointer usage	6-12
alignment	2-6	Pointers	
Alignment	2-10	to void	3-2
allocate variables as standard	2-6, 2-11	Portability	4-1, 6-5
compile but do not link	2-0, 2-11	Portable C compiler	3-1
compile leaving assembly files	2-5, 2-10	_	3-1 3-2
debug information		Pragma	
<u> </u>	2-4, 2-10	Prepare debug information	2-4, 2-10
define	2-7, 2-12	Prepare profile information	2-4

n	0.1	MON OMBINIOS	0.10
Preprocessor	2-1	/ROM_STRINGS	2-10
compiler options passed to	2-7	Run cpp only	2-7, 2-12
definition of	F-3	Run-time checks	2-4, 2-10, 2-14
m4	2-6	array index	2-15
macro	2-1, 2-7	intrinsic function parameters	8-1
/PRE_PROCESSOR	2-12	Intrinsic routines parameters	2-15
Preprocessor Directive		Run-Time checks	
pragma	3-2	NIL pointer	2-16
Profile information	7-1	Runtime feedback optimization	5-1, 5-11, 6-16,
customized _exit routine	7-3		7-11
gathering	7-2	Run-time library	6-7
Profilers		RUNTIME_FEEDBACK optimiza	tion
sprof	7-7	option	6-2, 6-4, 6-16
Programming hints		Run-time library	
for embedded programming	E-1	definition of	F-3
Programming in other languages	4-8		
Programming memory mapped devices			
Programming trap and interrupt routi		s	
PSR	1105 0-4	5	
L and F flags after _tbits	8-24	-S	2-5
L and r mags after _tofts	0-24		
		for debugging optimized code	6-9
		-s	2-8
Q		Safe register	5-8
		definition of	F-3
-Q	2-4	Saving registers	A-2
Qualifie r		_sbit	8-2, 8-3
definition of	F-3	_sbitps	8-22
Quick compilation	2-4	_Sbits	8-21
		Scratch register	5-8
		definition of	F-3
R			1 0
R		Series 32000 microprocessors	
	2-5	Series 32000 microprocessors NS32CG16	8-1
-R	2-5 2-5-2-8	Series 32000 microprocessors NS32CG16 NS32CG160	8-1 8-1
-R -r	2-5, 2-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16	8-1 8-1 8-1
-R -r Recommended reference book		Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320	8-1 8-1 8-1 8-1
-R -r Recommended reference book Redefinition	2-5, 2-8 1-2	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit	8-1 8-1 8-1 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions	2-5, 2-8 1-2 8-1	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string	8-1 8-1 8-1 8-2, 8-3 8-22
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file	2-5, 2-8 1-2 8-1 2-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string	8-1 8-1 8-1 8-2, 8-3 8-22 8-21
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string	8-1 8-1 8-1 8-2, 8-3 8-22 8-21
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp()	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed	8-1 8-1 8-1 8-2, 8-3 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizatioption 6-	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimization option 6- Registers	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11, 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option 6- Registers safe	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimization option 6- Registers safe saving	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 8-2, 8-3 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimization option Registers safe saving scratch	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 8-2, 8-3, 8-11 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimization option 6- Registers safe saving scratch Reliance on naive algebraic relations	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option Registers safe saving scratch Reliance on naive algebraic relations rename	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 2-8 2-7, 2-12 5-10
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization	8-1 8-1 8-2, 8-3 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization option	8-1 8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 2-8 2-7, 2-12 5-10
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option 6- Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file Retain relocation	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization	8-1 8-1 8-2, 8-3 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization option	8-1 8-1 8-2, 8-3 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option 6- Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file Retain relocation	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization sprintf	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 6-15 10-15
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option 6- Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file Retain relocation Return value	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8 7-5 2-5 2-8 4-10, 6-5	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimization sprintf for profile gathering Sprof	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 6-15 6-10 ion 6-2, 6-4, 6-16
-R -r Recommended reference book Redefinition of intrinsic functions Redirect output to .i file Redundant assignment elimination Register allocation for intrinsic functions Register allocation by coloring Register parameters Register variables REGISTER_ALLOCATION optimizati option 6- Registers safe saving scratch Reliance on naive algebraic relations rename for profile gathering Rename the output file Retain relocation Return value Returned value	2-5, 2-8 1-2 8-1 2-8 5-1, 5-2 5-8, 6-15 8-1 5-1 5-9 4-9 on 2, 6-4, 6-17 5-8, F-3 A-2 5-8, F-3 6-8 7-5 2-5 2-8 4-10, 6-5	Series 32000 microprocessors NS32CG16 NS32CG160 NS32FX16 NS32GX320 set bit Set bit perpendicular string Set bit string Set target configuration Setjmp() Severe Errors Show, but do not execute Signed Single bit instructions clear bit find first set invert bit set bit test bit Specify a program library Specify directory for included files Speed over space SPEED_OVER_SPACE optimizati option sprintf for profile gathering	8-1 8-1 8-2, 8-3 8-22 8-21 2-6, 2-11, 2-13 6-15 C-5 2-6, 2-11 3-2 8-2, 8-3, 8-11 8-2, 8-3 8-4 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 8-2, 8-3 6-15 6-10 ion 6-2, 6-4, 6-16

/DIRECTORY	7-9	Types and conversions	4-2
-е	7-9		
/EXECUTABLE	7-9		
- f	7-9	U	
/FORMAT	7-9		
-0	7-9	-U	2-8
/OUTPUT	7-9	-u	2-8
-p	7-9	Undefine	2-8, 2-12
sscanf	• •	/UNDEFINE	2-12
for profile gathering	7-5	Undefine symbol in symbol table	2-8
Stack	1.0	Undefined behavior	4-10
in calling sequence	A-1	Undetected program errors	6-5
/STANDARD	2-11	failing to declare a function	6-5
			6-5
Standard calling convention	A-1	relying on memory allocation uninitialized local variables	
STANDARD_LIBRARIES optimizat			6-5
option	6-2, 6-4, 6-7	UNIX	0.0
Statement		invocation syntax	2-2
definition of	F-3	unlink	
Static functions	6-11	for	7-5
Strength reduction	5-1, 5-6, 5-10	Unsigned constants	3-3
Strings		Use alternative library	2-6
string concatenation	3-3	Use the m4 preprocessor	2-6
Strip	2-8	USER_REGISTERS optimization o	ption 6-2,
Structure returning function	4-8, 6-5	-	6-4, 6-7, 6-15
Structures	3-3		,,
sys_errlist	0.0		
for profile gathering	7-5	v	
	1-0	•	
sys_nerr		-V	2-8
for	7-5	·	2-6 2-6
System code	6-6	-V	5-1, 5-2
System Errors	C-2	Value propagation	- ,
		Variable and structure alignment	4-2
		Variable initialization	E-4
T		Verbose	2-6, 2-11
		/VERBOSE	2-11
/TARGET	2-11	VMS	
BUS=1 for space optimization	6-16	invocation syntax	2-9
Target setup	2-2	-vn	2-6
Target machine		/VN	2-11
definition of	F-3	Void	3-2
_tbit	8-2, 8-3	Volatile	3-2, E-1, E-2
_tbits	8-23	definition of	F-3
effect on PSR L and F flags	8-24	VOLATILE optimization option	6-2, 6-4, 6-6
test bit	8-2, 8-3	Volatile variables	6-6
	,	volatile variables	0-0
Test bit string	8-23		
Timing assumptions	6-7		
TMPDIR	2-18	\mathbf{w}	
Tokens			
definition of	F-3	-W	2-8
Trap and interrupt routines		-W	2-6, C-6
programming	3-4	/WARNING	2-10, C-6
Trap handler routine		Warning diagnostics	2-6, 2-10
programming examples	E-12	Warnings	C-6
Type qualifiers		Writing Mixed-Language Programs	B-1
const	3-2		
volatile	3-2		
Type representations	4-1		
typedefs	4-1		
for NS32GX320 intrinsic functio	ns 8-26		

-X	2-7
with profile information	7-4
-x	2-8
Xdb_pfb_exit.o	7-4
Xpfb_exit.o and Xpfb_exit.obj	7-4
${f z}$	

2-6

-**Z**

INDEX 9

National Semiconductor

READER'S COMMENT FORM

In the interest of improving our documentation, National Semiconductor invites your comments on this manual.

Please restrict your comments to the documentation. Technical Support may be contacted at:

(800) 759-0105 - US and Canada ((0)8141) 103-330 - Germany only

Please rate this document according to the following categories. Include your comments below.

	EXCELLENT	GOOD	ADEQUATE	FAIR	POOR	
Readability (style)		D				
Technical Accuracy	0		0	–		
Fulfills Needs		0		0	0	
Organization	0					
Presentation (format)	0		0			
Depth of Coverage		0		0		
Overall Quality	0		•	0	0	
NAME				DATE		
ritle						
COMPANY NAME/DEPA	RTMENT					
ADDRESS						
CITY			STATE	ZI	IP	
Do you require a response?	o Yes □ No PH	ONE				
Comments:						
				·		

GNX — Version 4.0 C Optimizing Compiler Reference Manual

FOLD, STAPLE, AND MAIL

424010516-004



BUSINESS REPLY MAIL

FIRST CLASS

PERMIT NO. 409

SANTA CLARA, CA

Halandalallaan lalandallaladalaa Hadadaa H

POSTAGE WILL BE PAID BY ADDRESSEE

National Semiconductor Corporation Technical Publications Dept., M/S E295 2900 Semiconductor Drive P.O. Box 58090 Santa Clara, CA 95052 – 9968 NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



National Semiconductor

SALES OFFICES

ALABAMA Huntsville (205) 837-8960 (205) 721-9367

ARIZONA Tempe (602) 966-4563

Burnaby (604) 435-8107

CALIFORNIA Encino (818) 888-2602 Inglewood

(213) 645-4226 Roseville (916) 786-5577

San Diego (619) 587-0666 Santa Clara (408) 562-5900

Tuetin (714) 259-8880 Woodland Hills (818) 888-2602

COLORADO Boulder

> (303) 440-3400 Colorado Springs (303) 578-3319 Englewood (303) 790-8090

CONNECTICUT Fairfield (203) 371-0181 Hamden (203) 288-1560

INTERNATIONAL **OFFICES**

Electronica NSC de Mexico SA

Juventino Rosas No. 118-2 Col Guadalune inn Mexico, 01020 D.F. Mexico Tel: 52-5-524-9402

National Semicondutores Do Brasil Ltds. Av. Brig. Faria Lima, 1409

6 Andor Salas 62/64 01451 Sao Paulo, SP, Brasil Tel: (55/11) 212-5066 Telex: 391-1131931 NSBR BR

National Semiconductor GmbH Industriestrasse 10

D-8080 Furstenfeldbruck West Germany Tel: 49-08141-103-0 Telex: 527 649

National Semiconductor (UK) Ltd.

301 Harpur Centre Horne Lane Redford MK40 ITR United Kingdom Tel: (02 34) 27 00 27 Telex: 826 209

National Semiconductor Benefux Vorstlaan 100

B-1170 Brussels Belaium Tel: (02) 6725360 Telex: 61007

FLORIDA Boce Raton (305) 997-8133 Orlando

(305) 629-1720 St. Petersburg (813) 577-1380 GEORGIA

(404) 396-4048 Norcross

(404) 441-2740 ILLINOIS Schaumburg

(312) 397-8777 INDIANA

Carmel (317) 843-7160 Fort Wayne (219) 484-0722

AWO Cedar Rapids (319) 395-0090

KANSAS Overland Park (913) 451-8374 MARYLAND

Hanover (301) 796-8900 MASSACHUSETTS

Burdinaton (617) 273-3170 Waltham (617) 890-4000 MICHIGAN

W. Bloomfield (313) 855-0166

National Semiconductor (UK) Ltd. 1 Bianco Lunos Alle

DK-1868 Fredriksberg C Denmark

Tel: (01) 213211 Telex: 15179 **National Semiconductor**

Expansion 10000 28, rue de la Redoute F-92260 Fontenay-aux-Roses France

Tel: (01) 46 60 81 40 Telex: 250956

National Semiconductor S.p.A. Strada 7, Palazzo R/3

20089 Bozzano Milanofiori Italy

Telex: 10731

Tel: (02) 8242046/7/8/9

National Semiconductor AR Box 2016 Stensatravagen 13

S-12702 Skarholmen Sweden Tel: (08) 970190

National Semiconductor Calle Agustin de Foxa, 27 28036 Madrid

Spain Tel: (01) 733-2958 Telex: 46133

MINNESOTA

Bloomington (612) 835-3322 (612) 854-8200

NEW JERSEY (201) 599-0955

MEW MEXICO Albuquerque (505) 884-5601

NEW YORK **Endicott**

(607) 757-0200 Fairport (716) 425-1358 (716) 223-7700 Malvilla

(516) 351-1000 Wappinger Falls (914) 298-0680

NORTH CAROLINA Carv

(919) 481-4311 оню Dayton

> (513) 435-6886 Highland Heights (216) 442-1555 (216) 461-0191

ONTARIO Mississauna

(416) 678-2920 Nenean (404) 441-2740 (613) 596-0411 Woodbridge (416) 746-7120 PENNSYLVANIA Horsham

OREGON

(215) 675-6111 Willow Grove (215) 657-2711

(503) 639-5442

PLIERTO RICO Rio Piedias

(809) 758-9211 QUEBEC

Dollard Des Ormeaux (514) 683-0683 Lachine

(514) 636-8525

TEXAS

Austin (512) 346-3990

Houston (713) 771-3547 Richardson (214) 234-3811

UTAH Salt Lake City (801) 322-4747

WASHINGTON Bellevue (206) 453-9944

WISCONSIN Brookfield (414) 782-1818 Milwaukee (414) 527-3800

National Semiconductor Switzerland

Alte Winterthurerstrasse 53 Postfach 567 Ch-8304 Wallisellen-Zurich Switzerland Tel: (01) 830-2727

Telex: 59000 National Semiconductor Kauppakartanonkatu 7

SF-00930 Helsinki Finland Tel: (0) 33 80 33 Telev: 126116

Ltd

National Semiconductor Japan

Sanseido Bldg. 5F 4-15 Nishi Shinjuku Shiniuku-ku Tokyo 160 Japan Tel: 3-299-7001 Fax: 3-299-7000

National Semiconductor Hong Kong Ltd. Southeast Asia Marketing Austin Tower, 4th Floor

Teley: 52996 NSSEA HX

22-26A Austin Avenue Tsimshatsui, Kowloon, H.K. Tel: 852 3-7243645 Cable: NSSEAMKTG

National Semiconductor (Australia) PTY, Ltd. 1st Floor, 441 St. Kilda Rd. Melbourne, 3004

Victory, Australia Tel: (03) 267-5000 Fax: 61-3-2677458

National Semiconductor (PTE). Ltd.

200 Cantonment Road 13-01 Southpoint Singapore 0208 Tel: 2252226 Telev: BS 33877

National Semiconductor (Far East) L td

Talwan Branch P.O. Box 68-332 Tainei 7th Floor, Nan Shan Life Bldg 302 Min Chuan East Road. Taipei, Taiwan R.O.C. Tel: (86) 02-501-7227 Telex: 22837 NSTW

Cable: NSTW TAIPEI

National Semiconductor (Far East)

Ltd. Korea Office Room 612, Korea Fed. of Small Bus. Bidg. 16-2. Yoido-Dong. Youngdeungpo-Ku Secul Korea

Tel: (02) 784-8051/3 = 785-0696-8 Telex: K24942 NSRKI O

