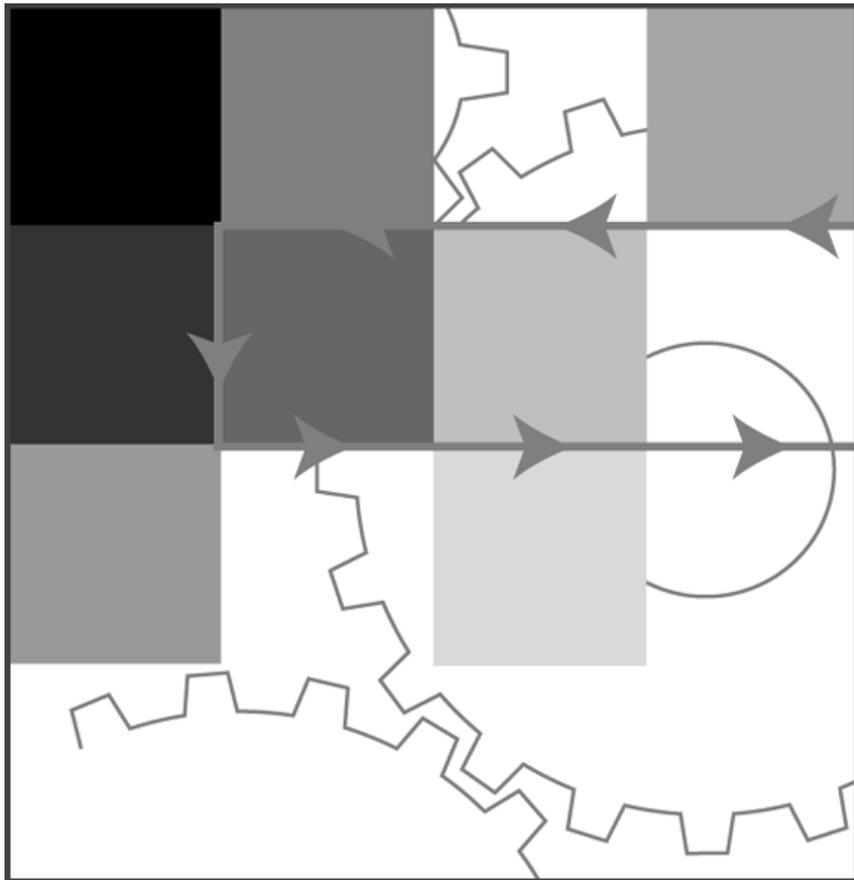


GNUPro® Toolkit

GNUPro Compiler Tools

- *Using GNU CC*
- *The C Preprocessor*



GNUPro 2001

Copyright © 1991-2001 Red Hat[®], Inc. All rights reserved.

Red Hat[®], GNUPro[®], the Red Hat Shadow Man logo[®], Source-Navigator[™], Insight[™], Cygwin[™], eCos[™], and Red Hat Embedded DevKit[™] are all trademarks or registered trademarks of Red Hat, Inc. ARM[®], Thumb[®], and ARM Powered[®] are registered trademarks of ARM Limited. SA[™], SA-110[™], SA-1100[™], SA-1110[™], SA-1500[™], SA-1510[™] are trademarks of ARM Limited. All other brands or product names are the property of their respective owners. “ARM” is used to represent any or all of ARM Holdings plc (LSE; ARM: NASDAQ; ARMHY), its operating company, ARM Limited, and the regional subsidiaries ARM INC., ARM KK, and ARM Korea Ltd.

AT&T[®] is a registered trademark of AT&T, Inc.

Hitachi[®], SuperH[®], and H8[®] are registered trademarks of Hitachi, Ltd.

IBM[®], PowerPC[®], and RS/6000[®] are registered trademarks of IBM Corporation.

Intel[®], Pentium[®], Pentium II[®], and StrongARM[®] are registered trademarks of Intel Corporation.

Linux[®] is a registered trademark of Linus Torvalds.

Matsushita[®], Panasonic[®], PanaX[®], and PanaXSeries[®] are registered trademarks of Matsushita, Inc.

Microsoft[®] Windows[®] CE, Microsoft[®] Windows NT[®], Microsoft[®] Windows[®] 98, and Win32[®] are registered trademarks of Microsoft Corporation.

MIPS[®] is a registered trademark and MIPS I[™], MIPS II[™], MIPS III[™], MIPS IV[™], and MIPS16[™] are all trademarks or registered trademarks of MIPS Technologies, Inc.

Mitsubishi[®] is a registered trademark of Mitsubishi Electric Corporation.

Motorola[®] is a registered trademark of Motorola, Inc.

Sun[®], SPARC[®], SunOS[™], Solaris[™], and Java[™], are trademarks or registered trademarks of Sun Microsystems, Inc..

UNIX[®] is a registered trademark of The Open Group.

NEC[®], VR5000[™], VRC5074[™], VR5400[™], VR5432[™], VR5464[™], VRC5475[™], VRC5476[™], VRC5477[™], VRC5484[™] are trademarks or registered trademarks of NEC Corporation.

All other brand and product names, services names, trademarks and copyrights are the property of their respective owners.

Permission is granted to make and distribute verbatim copies of this documentation, provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this documentation under the conditions for verbatim copying, provided also that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this documentation into another language, under the above conditions for modified versions.

While every precaution has been taken in the preparation of this documentation, the publisher assumes no responsibility for errors or omissions, or for damages resulting from the use of the information within the documentation. For licenses and use information, see “General Licenses and Terms for Using GNUPro Toolkit” in the GNUPro Toolkit *Getting Started Guide*.

How to Contact Red Hat

Red Hat Corporate Headquarters

2600 Meridian Parkway

Durham, NC 27713 USA

Telephone (toll free): +1 888 REDHAT 1 (+1 888 733 4281)

Telephone (main line): +1 919 547 0012

Telephone (FAX line): +1 919 547 0024

Website: <http://www.redhat.com>

Contents

Overview of GNUPro Compiler Tools	1
Using GNU CC Contents	2
The C Preprocessor Contents	3
GNUPro Compiler Tools Appendixes Content.....	4

Using GNU CC

Compile C, C++, Objective C, FORTRAN, Java, or CHILL	9
GNU CC Command Options	11
Options Controlling the Kind of Output	15
Options Controlling C Dialect	19
Options Controlling C++ Dialect	25
Compiling C++ Programs	32
Options Requesting or Suppressing Warnings	35
Options Controlling Debugging	45
Options Controlling Optimization	53
Fine-tuning Optimizations	58
Frequently Used Optimization Options	60
Options Controlling Preprocessing	61

Options Controlling the Assembler	65
Options for Linking	67
Options for Searching Directories	71
Options for Specifying Targets and Compiler Versions	73
Hardware Models and Configurations	75
AMD 29K Options	76
ARC Options	78
ARM/StrongARM Options	78
ARM THUMB Options	83
Clipper Options	85
DEC Alpha Options	85
Hitachi H8/300 Options	89
Hitachi SH Options	90
HPPA Options	90
IBM RS/6000 and PowerPC Options	92
IBM RT Options	101
Intel x86 Options	102
Intel 960 Options	105
Matsushita MN10200 Options	107
Matsushita MN10300/AM33 Options	107
MIPS Options	107
Mitsubishi D10V Options	112
Mitsubishi M32R/D/X Options	113
Motorola 68000 Options	114
Motorola 88000 Options	116
NEC V850 Options	119
SPARC Options	120
System V Options	124
Options Controlling Code Generation Conventions	127
The <code>offset-info</code> Option	135
Environment Variables Affecting GCC	137
Running the <code>protoize</code> Program	141
Extensions to the C Language Family	145
Statements and Declarations in Expressions	146
Locally Declared Labels	147
Labels as Values	148
Nested Functions	149
Constructing Function Calls	151
Naming an Expression's Type	152
Referring to a Type with the <code>typeof</code> Keyword	152
Generalized Lvalues	153
Conditionals with Omitted Operands	154

Double-word Integers.....	154
Complex Numbers	155
Arrays of Length Zero.....	156
Arrays of Variable Length	156
Macros with Variable Numbers of Arguments	157
Non-lvalue Arrays May Have Subscripts	158
Arithmetic on <code>void</code> Pointers and Function pointers.....	158
Non-constant Initializers	159
Constructor Expressions.....	159
Labeled Elements in Initializers.....	160
Case Ranges.....	161
Cast to a Union Type	161
Declaring Attributes of Functions.....	161
Prototypes and Old-style Function Definitions.....	167
Compiling Functions for Interrupt Calls	168
C++ style Comments	168
Dollar Signs in Identifier Names	168
The <code>ESC</code> Character in Constants	169
Inquiring on Alignment of Types or Variables	169
Specifying Attributes of Variables.....	169
Specifying Attributes of Types	173
An inline Function Is as Fast as a Macro.....	176
Assembler Instructions with C Expression Operands.....	177
Constraints for asm Operands.....	181
Simple Constraints.....	181
Multiple Alternative Constraints	184
Constraint Modifier Characters	184
Constraints for Particular Machines	185
Controlling Names Used in Assembler Code	193
Variables in Specified Registers	193
Defining Global Register Variables.....	194
Specifying Registers for Local Variables	195
Alternate Keywords	196
Incomplete enum Types	196
Function Names as Strings.....	197
Getting the Return or Frame Address of a Function.....	197
Extensions to the C++ Language Family	199
Named Return Values in C++	199
Minimum and Maximum Operators in C++	201
The <code>goto</code> and Destructors in GNU C++	202
Declarations and Definitions in One Header	202
Where's the Template?	204

Type Abstraction Using Signatures	207
gcov, a Test Coverage Program	209
Introduction to <code>gcov</code> Test Coverage	209
Invoking the <code>gcov</code> Program	210
Using <code>gcov</code> with GCC Optimization	213
Brief Description of <code>gcov</code> Data Files	213

The C Preprocessor

Overview of the C Preprocessor	217
What the C Preprocessor Provides	218
Transformations Made Globally	219
Preprocessing Directives	221
Header Files	223
Uses of Header Files	223
The <code>#include</code> Directive	224
How <code>#include</code> Works	225
Once-only Include Files	226
Inheritance and Header Files	227
Macros	229
Simple Macros	229
Macros with Arguments	231
Predefined Macros	233
Standard Predefined Macros	233
Non-standard Predefined Macros	235
Stringification	237
Concatenation	238
Undefining Macros	239
Redefining Macros	240
Pitfalls and Subtleties of Macros	240
Improperly Nested Constructs	241
Unintended Grouping of Arithmetic	241
Swallowing the Semicolon	242
Duplication of Side Effects	243
Self-referential Macros	243
Separate Expansion of Macro Arguments	244
Cascaded Use of Macros	246
Newlines in Macro Arguments	247
Conditionals	249
Why Conditionals are Useful	250
Syntax of Conditionals	250

The <code>#if</code> Directive	250
The <code>#else</code> Directive.....	251
The <code>#elif</code> Directive.....	251
Keeping Deleted Code for Future Reference.....	252
Conditionals and Macros.....	252
Assertions.....	254
The <code>#error</code> and <code>#warning</code> Directives.....	255
Combining Source Files	257
Other Preprocessing Directives	259
C Preprocessor Output	261
Invoking the C Preprocessor	263

Appendices

GNU General Public License	271
Preamble.....	271
Terms and Conditions for Copying, Distribution and Modification.....	272
How to Apply These Terms to Your New Programs.....	277
Contributors to GNU CC	279
Funding Free Software	287
Protect Your Freedom; Fight “Look and Feel”	289
Installing GCC	293
Installing GCC on Systems When It Exists	294
Installing GCC on UNIX Systems	294
Configurations That GCC Supports.....	303
Compilation in a Separate Directory.....	318
Building and Installing a Cross-compiler	319
Steps of Cross-compilation	319
Configuring a Cross-compiler.....	320
Tools and Libraries for a Cross-compiler	320
libgcc.a and Cross-compilers	321
Cross-compilers and Header Files	323
Standard Header File Directories	324
Actually Building the Cross-compiler	324
collect2 and Cross-compiling.....	325
Installing GCC on Sun	326
Installing GCC on VMS.....	327
Using GCC on VMS	330
Known Problems with GCC	337
Actual Bugs Not Fixed Yet.....	338
Installation Problems.....	338

Cross-compiler Problems	343
Interoperation	344
Problems Compiling Certain Programs	349
Incompatibilities of GCC	350
Fixed Header Files	353
Standard Libraries	353
Disappointments and Misunderstandings	354
Common Misunderstandings with GNU C++	356
Declare and Define Static Members	356
Temporaries May Vanish.....	356
protoize and unprotoize Warnings	357
Certain Changes GCC Will Not Use.....	358
Warning Messages and Error Messages	361
Reporting Bugs	363
Have You Found a Bug?	363
Where to Report Bugs.....	364
How to Report Bugs.....	365
Sending Patches for GCC	370
How to Get Help with GCC	373
Index	375

Overview of GNUPro Compiler Tools

The following documentation details the two parts of the *GNUPro Compiler Tools*.

- For the first part detailing the GNUPro compiler, see “Using GNU CC Contents” on page 2.
- For the second part detailing the C preprocessor, see “The C Preprocessor Contents” on page 3.
- For the appendixes, see “GNUPro Compiler Tools Appendixes Content” on page 4.

Using GNU CC Contents

The following documentation discusses the location and contents of *Using GNU CC* for working with GCC, the GNU Compiler Collection, which is for integrating C, C++, Objective C, FORTRAN, Java and CHILL resources.

- “Compile C, C++, Objective C, FORTRAN, Java, or CHILL” on page 9
This section provides a general introduction to the GNU Compiler Collection.
- “GNU CC Command Options” on page 11
This section provides a summary of all options for the GNU Compiler Collection.
- “Options Controlling the Kind of Output” on page 15
This section provides a discussion of the output formats for compiling.
- “Options Controlling C Dialect” on page 19
This section provides a discussion of compiling with C languages.
- “Options Controlling C++ Dialect” on page 25
This section provides a discussion of compiling with C++ languages.
- “Options Requesting or Suppressing Warnings” on page 35
This section provides a discussion of warning options for compiling.
- “Options Controlling Debugging” on page 45
This section provides a discussion of debugging options when compiling.
- “Options Controlling Optimization” on page 53
This section provides a discussion of options for compiling with optimizations.
- “Options Controlling Preprocessing” on page 61
This section provides a discussion of options for compiling with preprocessor work.
- “Options Controlling the Assembler” on page 65
This section provides a discussion of options for compiling with assembly code.
- “Options for Linking” on page 67
This section provides a discussion of options linking code when compiling.
- “Options for Searching Directories” on page 71
This section provides a discussion of directory search options when compiling.
- “Options for Specifying Targets and Compiler Versions” on page 73
This section provides a discussion of working with specific targets when compiling.
- “Hardware Models and Configurations” on page 75
This section provides a discussion of specific configurations when compiling.
- “The offset-info Option” on page 135
This section provides a discussion of the `offset-info` option, which simplifies access to C structs from the assembler.

- “Environment Variables Affecting GCC” on page 137
This section provides a discussion of environment variables for compiling.
- “Running the `protoize` Program” on page 141
This section provides a discussion of the `protoize` program for compiling, which is for adding *prototypes* to a program, thus converting the program to ANSI C in one respect..
- “Extensions to the C Language Family” on page 145
This section provides a discussion of extensions to C, which are features not found in ANSI standard C, when compiling.
- “Extensions to the C++ Language Family” on page 199
This section provides a discussion of extensions to C++, including some C extensions which are compatible with C++ programs when compiling.
- “`gcov`, a Test Coverage Program” on page 209
This section provides a discussion of `gcov`, which is a test coverage program to analyze your programs for creating more efficient, faster running code, to use as a profiling tool, to discover where your optimization efforts will best affect your code, and to use with the other profiling tool, `gprof`, for assessing which parts of your code use the greatest amount of computing time.

The C Preprocessor Contents

The following documentation discusses the location and contents of *The C Preprocessor* for working with the C preprocessor, a *macro processor* that is used automatically by the C compiler to transform your program before actual compilation. The C preprocessor allows you to define macros, which are abbreviations for longer constructs.

- “Overview of the C Preprocessor” on page 217
This section provides a summary of `cpp`, the C preprocessor.
- “Transformations Made Globally” on page 219
This section provides a discussion of transformations that the preprocessor always makes on all the input it receives, even in the absence of directives.
- “Preprocessing Directives” on page 221
This section provides a discussion of preprocessing directives, lines in your program that start with `#`, followed by an *identifier*, which is the directive name.
- “Header Files” on page 223
This section provides a discussion of a file containing C declarations and macro definitions to be shared between several source files.
- “Macros” on page 229
This section provides a discussion of macros, a sort of abbreviation that you can define once and later use.

- “Conditionals” on page 249
This section provides a discussion of conditionals, directives that allow a part of a program to be ignored during compilation, on some conditions.
- “Combining Source Files” on page 257
This section provides a discussion of using source files when compiling, with the C preprocessor informing the C compiler where each line of C code has as its origin, which source file and which line number.
- “Other Preprocessing Directives” on page 259
This section provides a discussion of other preprocessing directives, which are the *null directive*, the `#pragma` directive, and the `#ident` directive
- “C Preprocessor Output” on page 261
This section provides a discussion of C preprocessor output when compiling.
- “Invoking the C Preprocessor” on page 263
This section provides a discussion of invoking the C preprocessor when compiling.

GNUPro Compiler Tools Appendixes **Content**

The following documentation discusses content in the appendixes.

- “GNU General Public License” on page 271
This appendix details the agreements for using the compiler tools.
- “Contributors to GNU CC” on page 279
This appendix lists the contributors to GCC and what they have accomplished.
- “Funding Free Software” on page 287
This appendix details some of the Free Software Foundation’s approach with the free tools of the GNU collection, as well as the philosophy behind the approach.
- “Protect Your Freedom; Fight “Look and Feel”” on page 289
This appendix details some of the Free Software Foundation’s tenets and how to protect them.
- “Installing GCC” on page 293
This appendix provides the general approach for installation of GCC from net distributions; it is unnecessary if you have already installed GNUPro Toolkit, and its general purpose is for rebuilding the compiler tools from net distribution sources.
- “Known Problems with GCC” on page 337
This appendix provides a discussion of compiler problems.

- “Reporting Bugs” on page 363
This appendix provides a discussion of reporting problems; it is not necessary if you have installed GNUPro Toolkit as a Red Hat customer.
- “How to Get Help with GCC” on page 373
This appendix provides a discussion of getting help with installing, using or changing the GNU Compiler Collection; it is not necessary if you have installed GNUPro Toolkit as a Red Hat customer.

Using GNU CC

1

Compile C, C++, Objective C, FORTRAN, Java, or CHILL

Several versions of the compiler integrate C, C++, Objective C, FORTRAN, Java, and CHILL resources, using the name, GNU Compiler Collection (GCC). GCC can compile programs written in any of these languages. In the GNUPro documentation, there is only discussion of the options for the C, C++, and Objective C compilers. Consult the documentation of the other front ends for the options to use when compiling programs or when using tools that are written in other languages. The following documentation discusses some of the essentials within GCC.

- “GNU CC Command Options” on page 11 gives a summary of all the types of options that the GCC compiler provides.
- “Extensions to the C Language Family” on page 145 discusses extensions available in C and Objective C, even though most of them are available in C++.
- “Extensions to the C++ Language Family” on page 199 discusses extensions that apply *only* to C++.
- “`gcov`, a Test Coverage Program” on page 209 discusses using a GCC tool that analyzes your programs for creating more efficient, providing faster running code, and that works as a profiling tool, to discover where your optimization efforts will best affect your code, and that works with another profiling tool, `gprof`, for assessing which parts of your code use the greatest amount of computing time.

GCC refers to the compilation system as a whole, and more specifically to the

language independent part of the compiler; for example, the optimization options affect the behavior of GCC, or sometimes just the *compiler*, so that, when referring to C++ compilation, call the compiler G++. Since there is only one compiler, it is also accurate to call it GCC, no matter what the language context; however, the term, G++, is more useful when the emphasis is on compiling C++ programs.

Front ends for other languages, such as Ada 9X, CHILL, FORTRAN, Java, Modula-3, and Pascal, are still in development. These front ends are built in subdirectories of GCC and link to it. The result is an integrated compiler that can compile programs written in C, C++, Objective C, or any of the languages for which you have installed front ends.

G++ is a compiler, not only a preprocessor. G++ builds object code directly from your C++ program source. There is no intermediate C version of the program. By contrast, for example, some other implementations use a program that generates a C program from your C++ source. Avoiding an intermediate C representation of the program means that you get better object code and better debugging information. The GNU debugger (GDB) works with this information in the object code to give you comprehensive C++ source-level editing capabilities (for more information, see “Using GDB with Different Languages” on page 95 and “C and C++” on page 101 in *Debugging with GDB* in **GNUPro Debugging Tools**).

2

GNU CC Command Options

When you invoke the GNU compiler, GCC, it normally does preprocessing, compilation, assembly and linking. Some options allow you to control the output, in order to stop the compiling process at intermediate stages. Some options control the compiler itself, while still other options control the preprocessor, assembler and linker. For example, the `-c` option provides for compiling or assembling the source files without running the linker (see `-c` on page 16 in “Options Controlling the Kind of Output”), with the output then consisting of object files output by the assembler, while passing other options on to another stage of processing.

Most of the command line options with GCC are useful for C programs; when an option is only useful with another language (usually C++), the explanation explicitly states such usage. If the description for a particular option does not mention a source language, you can use that option with all supported languages. GCC accepts options and filenames as operands. Many options have multi-letter names; therefore, multiple single-letter options may *not* be grouped. You can mix options and other arguments. Order does matter when you use several options of the same kind; for instance, if you specify `-L` more than once, the directories are searched in the order specified. Many options have long names starting with `-f` or with `-w` (for example, `-fstrength-reduce` or `-wformat`). Names often have positive and negative forms (the negative form of `-ffoo` would be `-fno-foo`); this documentation generally discusses only one of

the forms, whichever one not being default. The following list shows where to find the corresponding documentation for specific types of options.

- “Options Controlling the Kind of Output” on page 15
- “Options Controlling C Dialect” on page 19
- “Options Controlling C++ Dialect” on page 25
- “Options Requesting or Suppressing Warnings” on page 35
- “Options Controlling Debugging” on page 45
- “Options Controlling Optimization” on page 53
- “Options Controlling Preprocessing” on page 61
- “Options Controlling the Assembler” on page 65
- “Options for Linking” on page 67
- “Options for Searching Directories” on page 71
- “Options for Specifying Targets and Compiler Versions” on page 73
- “Hardware Models and Configurations” on page 75 contains descriptions of compiler options for the following specific processor configurations.
 - “AMD 29K Options” on page 76
 - “ARC Options” on page 78
 - “ARM/StrongARM Options” on page 78
 - “ARM THUMB Options” on page 83
 - “Clipper Options” on page 85
 - “DEC Alpha Options” on page 85
 - “Hitachi H8/300 Options” on page 89
 - “Hitachi SH Options” on page 90
 - “HPPA Options” on page 90
 - “IBM RS/6000 and PowerPC Options” on page 92
 - “IBM RT Options” on page 101
 - “Intel x86 Options” on page 102
 - “Intel 960 Options” on page 105
 - “Matsushita MN10200 Options” on page 107
 - “Matsushita MN10300/AM33 Options” on page 107
 - “MIPS Options” on page 107
 - “Mitsubishi D10V Options” on page 112
 - “Mitsubishi M32R/D/X Options” on page 113
 - “Motorola 68000 Options” on page 114
 - “Motorola 88000 Options” on page 116

- “NEC V850 Options” on page 119
- “SPARC Options” on page 120
- “System V Options” on page 124
- “Options Controlling Code Generation Conventions” on page 127

3

Options Controlling the Kind of Output

The following documentation discusses types of output and source files. Compilation can involve up to four of the following stages, always in the following order.

- preprocessing
- compiling
- assembling
- linking

The first three stages apply to an individual source file: *preprocessing* establishes the type of source code to process; *compiling* produces an object file; *assembling* establishes the syntax that the compiler expects for symbols, constants, expressions and the general directives; *linking* completes the compilation process by combining all the object files (those newly compiled, and those specified as input) into an executable file.

For any given input file, the filename suffix determines what kind of compilation is done, and the output that produces as an object file (signified in the following documentation as *file*).

file.c

C source code which must be preprocessed.

file.i

C source code which should not be preprocessed.

file.i

C++ source code which should not be preprocessed.

file.m

Objective C source code.

IMPORTANT! You must link with the `libobjc.a` library to make an Objective C program work.

file.h

C header file (not to be compiled or linked).

file.cc

file.cxx

file.cpp

file.C

C++ source code which must be preprocessed. In `.cxx`, the last two letters must both be, literally, `x`; likewise, `.C` refers to a literal capital C.

file.s

Assembler code.

file.S

Assembler code which must be preprocessed.

other

An object file to be fed straight into linking. Any filename without a recognized suffix is treated this way.

Specify the input language with the `-x` option:

`-x language`

Specify the *language* for the following input files (rather than letting the compiler choose a default based on the filename suffix): `c`, `objective-c`, `c++`, `c-header`, `cpp-output`, `c++-cpp-output`, `assembler`, and `assembler-with-cpp`; this option applies to all subsequently specified input files until the next `-x` option.

`-x none`

Turn off any specification of a language, so that subsequent files are handled according to their filename suffixes, as if `-x` has not been used at all.

If you only want some of the stages of compilation, you can use `-x`, or filename suffixes, to tell GCC where to start, and one of the options `-c`, `-S`, or `-E` to say where GCC is to stop.

IMPORTANT! Some combinations (for example, `-x cpp-output -E`) instruct GCC to do nothing.

`-c`

Compile or assemble the source files, but do not link. The ultimate output is in the form of an object file for each source file. By default, the object filename for a source file is made by replacing the suffix, such as `.c`, `.i`, or `.s`, with `.o`. Unrecognized input files, not requiring compilation or assembly, are ignored.

- S Stop after the stage of compilation proper; do not assemble. The output is in the form of an assembler code file for each non-assembler input file specified. By default, the assembler filename for a source file is made by replacing the suffix, such as `.c` or `.i`, with `.s`. Input files that don't require compilation are ignored.
- E Stop after the preprocessing stage, but do not run the compiler. The output is in the form of preprocessed source code, which is sent to the standard output. Input files which don't require preprocessing are ignored.
- o *file* Place output in file, *file*. This applies regardless of output being produced, whether it be an executable file, an object file, an assembler file, or preprocessed C code. Since only one output file can be specified at a time, do not use `-o` when compiling more than one input file, unless you are producing an executable file as output. If `-o` is not specified, the default is to put an executable file in `a.out`, the object file for *source.suffix* in *source.o*, its assembler file in *source.s*, and all preprocessed C source on standard output.
- v Specifies printing (with standard error output) the commands executed to run the stages of compilation. Also specifies printing the version number of the compiler driver program and of the preprocessor and the compiler proper.
- pipe Use pipes rather than temporary files for communication between the various stages of compilation. This fails to work on some systems where the assembler is unable to read from a pipe, although the GNU assembler has no trouble.

4

Options Controlling C Dialect

The following options control the dialect of C (or languages derived from C, such as C++ and Objective C) that the compiler accepts.

`-ansi`

Support all ANSI standard C programs.

This turns off certain features of GNU C that are incompatible with ANSI C, such as the `asm`, `inline` and `typeof` keywords, and predefined macros such as `unix` and `vax` that identify the type of system you are using. It also enables the undesirable and rarely used ANSI `trigraph` feature, disallows the dollar sign symbol (\$) as part of identifiers, and disables recognition of C++ style comments with double forward-slash (`//`).

The alternate keywords, `__asm__`, `__extension__`, `__inline__`, and `__typeof__`, continue to work despite `-ansi`. You would not want to use them in an ANSI C program, of course, but it is useful to put them in header files that might be included in compilations done with `-ansi`. Alternate predefined macros, such as `__unix__` and `__vax__`, are also available, with or without `-ansi`.

The `-ansi` option does not cause non-ANSI programs to be rejected gratuitously. For that, `-pedantic` is required in addition to `-ansi`. See “Options Requesting or Suppressing Warnings” on page 35.

The `__STRICT_ANSI__` macro is predefined when the `-ansi` option is used. Some header files may notice this macro and refrain from declaring certain functions or

defining certain macros that the ANSI standard doesn't call for; this is to avoid interfering with any programs that might use these names for other things.

The functions, `alloca`, `abort`, `exit`, and `_exit`, are not built-in functions when `-ansi` is used.

`-fstd`

Determine the language standard. A value for this option must be provided; there are the following possible values.

- `iso9899:1990`
Same as `-ansi`.
- `iso9899:199409`
ISO C as modified in `amend.1`.
- `iso9899:199x`
ISO C 9x.
- `c89`
Same as `-std=iso9899:1990`.
- `c9x`
Same as `-std=iso9899:199x`.
- `gnu89`
Default, `iso9899:1990` with GNU extensions.
- `gnu9x`
`iso9899:199x` with GNU extensions.

Even when this option is not specified, you can still use some of the features of newer standards in so far as they do not conflict with previous C standards. For example, use `__restrict__` even when `-fstd=c9x` is not specified.

`-fallow-single-precision`

Do not promote single precision math operations to double precision, even when compiling with `-traditional`.

Traditional K&R C promotes all floating point operations to double precision, regardless of the sizes of the operands. On the architecture for which you are compiling, single precision may be faster than double precision. If you must use `-traditional`, but want to use single precision operations when the operands are single precision, use this option. This option has no effect when compiling with ANSI or GNU C conventions (the default).

`-fcond-mismatch`

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

`-flang-isoc9x`

Enable support for features found in the C9X standard. In particular, enable support for the C9X `restrict` keyword.

Even when this option is not specified, you can still use some C9X features in so far as they do not conflict with previous C standards. For example, you may use

`__restrict__` even when `-ang-isoc9x` is not specified.

`-fno-asm`

Do not recognize `asm`, `inline` or `typeof` as a keyword, so that code can use these words as identifiers. You can use, instead, the `__asm__`, `__inline__` or `__typeof__` keywords. `-ansi` implies `-fno-asm`.

In C++, this switch only affects the `typeof` keyword, since `asm` and `inline` are standard keywords.

You may want to use the `-fno-gnu-keywords` flag instead, as it also disables the other, C++-specific, extension keywords such as `headof`.

`-fno-builtin`

Don't recognize built-in functions that do not begin with two leading underscores. Currently, the functions affected include `abort`, `abs`, `alloca`, `cos`, `exit`, `fabs`, `ffs`, `labs`, `memcpy`, `strcpy`, `sin`, `sqrt`, `strcmp`, `strncpy`, and `strlen`. GCC normally generates special code to handle certain built-in functions more efficiently; for instance, calls to `alloca` may become single instructions that adjust the stack directly, and calls to `memcpy` may become inline copy loops. The resulting code is often both smaller and faster, but since the function calls no longer appear as such, you cannot set a breakpoint on those calls, nor can you change the behavior of the functions by linking with a different library.

The `-ansi` option prevents `alloca` and `ffs` from being builtin functions, since these functions do not have an ANSI standard meaning.

`-ffreestanding`

Assert that compilation takes place in a freestanding environment. This implies `-fno-builtin`. A freestanding environment is one in which the standard library may not exist, and program startup may not necessarily be at `main`. The most obvious example is an OS kernel. This is equivalent to `-fno-hosted`.

The `-ansi` option prevents `alloca` and `ffs` from being built-in functions, since these functions do not have an ANSI standard meaning.

`-fhosted`

Assert that compilation takes place in a hosted environment. This implies `-fbuiltin`. A hosted environment is one in which the entire standard library is available, and in which `main` has a return type of `int`. Examples are nearly everything except a kernel.

This is equivalent to `-fno-freestanding`.

`-fsigned-bitfields`

`-funsigned-bitfields`

`-fno-signed-bitfields`

`-fno-unsigned-bitfields`

These options control whether a bitfield is signed or unsigned, when the declaration does not use either `signed` or `unsigned`. By default, such a bitfield is

signed, because this is consistent: the basic integer types such as `int` are signed types.

However, when `-traditional` is used, bitfields are all, no matter what, unsigned.

`-fsigned-char`

Let the type, `char`, be signed, like `signed char`.

IMPORTANT! `-fno-unsigned-char` is the negative form of `-funsigned-char`. Likewise, `-fno-signed-char` is equivalent to `-funsigned-char`.

You may wish to use `-fno-builtin` as well as `-traditional` if your program uses names that are normally GNU C builtin functions for other purposes of its own.

You cannot use `-traditional` if you include any header files that rely on ANSI C features. Some vendors are starting to ship systems with ANSI C header files and you cannot use `-traditional` on such systems to compile files that include any system headers.

`-funsigned-char`

Let the type, `char`, be unsigned, like `unsigned char`.

Each kind of machine has a default for what `char` should be. It is either like `unsigned char`, by default, or like `signed char`, by default.

Ideally, a portable program should always use `signed char` or `unsigned char` when it depends on the *signedness* of an object. But many programs have been written to use plain `char` and expect it to be signed, or expect it to be unsigned, depending on the machines for which they were written. This option, and its inverse, let you make such a program work with the opposite default.

The type, `char`, is always a distinct type from each of `signed char` or `unsigned char`, even though its behavior is always just like one of those two dialect options.

`-fwritable-strings`

Store string constants in the writable data segment and don't uniquize them. This is for compatibility with old programs which assume they can write into string constants. The `-traditional` option also has this effect. Writing into string constants is a very bad idea; *constants* should be constant.

`-traditional`

Attempt to support some aspects of traditional C compilers. Specifically:

- All `extern` declarations take effect globally even if they are written inside of a function definition. This includes implicit declarations of functions.
- The `typeof`, `inline`, `signed`, `const`, and `volatile` keywords are not recognized. (You can still use the alternative keywords, such as `__typeof__`, `__inline__`, and so on.)
- Comparisons between pointers and integers are always allowed.
- Integer types, `unsigned short` and `unsigned char`, promote to `unsigned`

int.

- Out-of-range floating point literals are not an error.
- Certain constructs which ANSI regards as a single invalid preprocessing number, such as `0xe-0xd`, are treated as expressions instead.
- String constants are not necessarily constant; they are stored in writable space, and identical looking constants are allocated separately. (This is the same as the effect of `-fwritable-strings`.)
- All automatic variables not declared `register` are preserved by `longjmp`. Ordinarily, GNU C follows ANSI C: automatic variables not declared `volatile` may be clobbered.
- The character escape sequences, `\x` and `\a`, evaluate as the literal characters, `x` and `a`, respectively. Without `-traditional`, `\x` is a prefix for the hexadecimal representation of a character, and `\a` produces a ping to a terminal (like the **007** ASCII code).

You may wish to use `-fno-builtin` as well as `-traditional` if your program uses names that are normally GNU C builtin functions for other purposes of its own.

You cannot use `-traditional` if you include any header files that rely on ANSI C features. Some vendors are starting to ship systems with ANSI C header files and you cannot use `-traditional` on such systems to compile files that include any system headers.

The `-traditional` option also enables the `-traditional-cpp` option which is described in the following discussion.

`-traditional-cpp`

Attempts to support some aspects of traditional C preprocessors. Specifically:

- Comments convert to nothing at all rather than to a space. This allows traditional token concatenation.
- In preprocessing directive, the pound/number symbol (`#`) must appear as the first character of a line.
- Macro arguments are recognized within string constants in a macro definition (and their values are stringified, though without additional quote marks, when they appear in such a context). The preprocessor always considers a string constant to end at a newline.
- The predefined macro, `__STDC__`, is not defined when you use `-traditional`, but `__GNUC__` is (since the GNU extensions which `__GNUC__` indicates are not affected by `-traditional`). If you need to write header files that work differently depending on whether `-traditional` is in use, by testing both of these predefined macros you can distinguish four situations: GNU C, traditional GNU C, other ANSI C compilers, and other old C compilers. The predefined macro, `__STDC_VERSION__`, is also not defined

when you use `-traditional`. For more discussion of these and other predefined macros, see the “Standard Predefined Macros” on page 233 in *The C Preprocessor*.

- The preprocessor considers a string constant to end at a newline (unless the newline is escaped with a backslash; that is, `\n`). Without `-traditional`, string constants can contain the newline backslash character, `\n`.

`-trigraphs`

Support ANSI C trigraphs. The `-ansi` option implies `-trigraphs`.

5

Options Controlling C++ Dialect

The following documentation describes the command line options that are only meaningful for C++ programs; you can use most of the GNU compiler options regardless of what language your program uses. For instance, you might compile a file, `firstClass.C`, like the following example's input shows. In the example, only `-frepo` is an option meant only for C++ programs; you can use the other options with any language supported by GCC. See also "Compiling C++ Programs" on page 32.

```
g++ -g -frepo -O -c firstClass.C
```

The following discussion lists options that are *only* for compiling C++ programs.

`-fno-access-control`

Turn off access checking when working around bugs in the access control code.

`-fcheck-new`

Check that the pointer returned by `operator new` is non-null before attempting to modify the storage allocated. The current ANSI C++ Working Paper* requires that `operator new` never return a null pointer, so this check is normally unnecessary.

* The C++ programming language as described in this Working Paper is based on the language as described in Chapter R (Reference Manual) of Dale Soustroup's *The C++ Programming Language* (second edition, Addison-Wesley Publishing Company, ISBN 0-201-53992-6, copyright © 1991 AT&T). That, in turn, is based on the C programming language as described in Appendix A of Kernighan and Ritchie's *The C Programming Language* (Prentice-Hall, 1978, ISBN 0-13-110163-3, copyright © 1978 AT&T). In addition, portions of this Working Paper are based on work by P.J. Plauger, which was published as The Draft Standard C++ Library; Prentice-Hall, ISBN 0-13-117003-1, copyright © 1995 P.J. Plauger). All rights in these originals are reserved.

`-fconserve-space`

Put uninitialized or runtime-initialized global variables into the common segment, compliant with what C code does. This saves space in the executable at the cost of not diagnosing duplicate definitions. If you compile with this option and your program crashes after `main()` has completed, you may have an object that is being destroyed twice because two definitions were merged. This option is no longer useful on most targets, now that support has been added for putting variables into BSS without making them common.

`-fdollars-in-identifiers`

Accept the dollar sign (\$) in identifiers. You can also explicitly prohibit use of \$ with the `-fno-dollars-in-identifiers` option. GNU C++ allows \$ by default on some target systems but not others. Traditional C allowed the character, \$, to form part of identifiers. However, ANSI C and C++ forbids \$ in identifiers.

`-fembedded-cxx`

In compliance with the Embedded C++ specification, makes the use of templates, exception handling, multiple inheritance, or RTTI illegal. This makes the use of these keywords result in warnings by default: `template`, `typename`, `catch`, `throw`, `try`, `using`, `namespace`, `dynamic_cast`, `static_cast`, `reinterpret_cast`, `const_cast`, and `typeid`. Add the `-pedantic-errors` option to make the warnings for these keywords be given as errors.

`-fno-elide-constructors`

The C++ standard allows an implementation to omit creating a temporary object, which is only used to initialize another object of the same type. Specifying this option disables that optimization, forcing G++ to call the copy constructor in all cases.

`-fexternal-templates`

Cause template instantiations to obey `#pragma interface` and `implementation`; template instances are emitted or not according to the location of the template definition. See “Where’s the Template?” on page 204 for more explanation of templates. This option is deprecated (meaning that it will not generate error or warning messages, and that it is still in compliant with the C++ standards).

`-falt-external-templates`

Similar to `-fexternal-templates`, but template instances are emitted or not according to the place where they are first instantiated. See “Where’s the Template?” on page 204 for more explanation on templates. This option is deprecated (meaning that it will not generate error or warning messages, and that it is still in compliant with the C++ standards).

`-ffor-scope``-fno-for-scope`

If `-ffor-scope` is specified, the scope of variables declared in a *for-init-statement* is limited to the `for` loop itself, as specified by the draft C++ standard. If `-fno-for-scope` is specified, the scope of variables declared in a

for-init-statement extends to the end of the enclosing scope, as was the case in old versions of GCC, and other (traditional) implementations of C++. The default, if neither option is given, is to follow the standard, but to allow and give a warning for old-style code that would otherwise be invalid, or have different behavior.

`-fhandle-signatures`

Recognize the signature and `sigof` keywords for specifying abstract types. The default (`-fno-handle-signatures`) is not to recognize them. See “Type Abstraction Using Signatures” on page 207.

`-fhonor-std`

Treat namespace `std` as a namespace, instead of ignoring it. For compatibility with earlier versions of G++, the compiler will, by default, ignore namespace-declarations, using-declarations, using-directives, and namespace-names, if they involve `std`.

`-fno-implicit-templates`

Never emit code for templates that are instantiated implicitly (that is, by use); only emit code for explicit instantiations. See “Where’s the Template?” on page 204 for more explanation on templates.

`-finit-priority`

Support `__attribute__((init_priority (n)))` for controlling the order of initialization of file-scope objects. On ELF targets, this requires the GNU linker (version 2.10 or later). `n` represents a numerical priority from 0 to 65535. Lower is higher priority.

`-fno-implement-inlines`

To save space, do not emit out-of-line copies of inline functions controlled by `#pragma implementation`. This causes linker errors if these functions are not inlined everywhere they are called.

`-fno-optional-diags`

Disable diagnostics that the standard says a compiler doesn’t need to issue. This means the diagnostic for a name having multiple meanings within a class.

`-fpermissive`

Downgrade messages about nonconformant code from errors to warnings. By default, G++ effectively sets `-pedantic-errors` without `-pedantic`; this option reverses that behavior. This behavior and this option are superseded by `-pedantic`.

`-frepo`

Enable automatic template instantiation. This option also implies `-fno-implicit-templates`. See “Where’s the Template?” on page 204 for more explanation of templates.

`-fsquangle`

`-fno-squangle`

`-fsquangle` will enable a compressed form of name mangling for identifiers. In particular, it helps to shorten very long names by recognizing types and class

names which occur more than once, replacing them with special short ID codes. `-fsquangle` also requires any C++ libraries being used to be compiled with this option as well. The compiler has this functionality disabled (the equivalent of using `-fno-squangle`) by default.

IMPORTANT! As with all options that change the ABI, all C++ code, including `libgcc.a`, must be built with the same setting of this option.

`-ftemplate-depth-n`

Set the maximum instantiation depth for template classes to *n*. A limit on the template instantiation depth is needed to detect endless recursions during template class instantiation. ANSI/ISO C++ conforming programs must not rely on a maximum depth greater than 17.

`-fuse-cxa-atexit`

Register destructors for objects with static storage duration with the `__cxa_atexit` function rather than the `atexit` function. This option is required for fully standards-compliant handling of static destructor; it will only work if your C library supports what it implements.

`-fvtable-thunks`

Use `thunks` to implement the virtual function dispatch table (`vtable`). The traditional (`cfront`-style) approach to implementing a `vtable` was to store a pointer to the function and two offsets for adjusting the `this` pointer at the call site. Newer implementations store a single pointer to a `thunk` function, which does any necessary adjustment and then calls the target function. This option also enables a heuristic for controlling emission of a `vtable`; if a class has any non-inline virtual functions, the `vtable` will be emitted in the translation unit containing the first one of those non-inline virtual functions.

`-fname-mangling-version-n`

Control the way in which names are mangled. Version 0 is compatible with versions of G++ before 2.8. Version 1 is the default. Version 1 allows correct mangling of function templates. For example, version 0 mangling does not mangle `foo int`, `double` and `foo int, char` given the following declaration.

```
template class T, class U void foo(T t);
```

`-fno-default-inline`

Do not assume `inline` for functions defined inside a class scope. See “Options Controlling Optimization” on page 53.

IMPORTANT! Such class scope defined functions will have linkage like inline functions; they just won’t be inlined by default.

`-fno-gnu-keywords`

Do not recognize `classof`, `headof`, `signature`, `sigof`, or `typeof` as a keyword, so that code can use these words as identifiers. Instead, use the `__classof__`, `__`

`headof__`, `__signature__`, `__sigof__`, and `__typeof__` keywords. `-ansi` implies `-fno-gnu-keywords`.

`-fno-nonnul-objects`

Don't assume that a reference is initialized to refer to a valid object. Although the current C++ Working Paper prohibits null references, some old code may rely on them; you can use `-fno-nonnul-objects` to turn on checking. Until there is a standard with which everyone agrees, the compiler only does this checking for conversions to virtual base classes.

`-fguiding-decls`

Treat a function declaration with the same type as a potential function template instantiation as though it declares that instantiation, not a normal function. If a definition is given for the function later in the translation unit (or another translation unit if the target supports weak symbols), that definition will be used; otherwise the template will be instantiated. This behavior reflects the C++ language prior to September 1996, when guiding declarations were removed. This option implies `-fname-mangling-version-0` and will not work with other name mangling versions. Like all options that change the ABI, all C++ code, including `libgcc.a` must be built with the same setting of this option.

`-foperator-names`

Recognize the `and`, `bitand`, `bitor`, `compl`, `not`, `or`, and `xor` operator name keywords as synonyms for the symbols they refer to. `-ansi` implies `-foperator-names`.

`-fstrict-prototype`

Within an `extern "C"` linkage specification, treat a function declaration with no arguments, such as `int foo();`, as declaring the function to take no arguments. Normally, such a declaration means that the function `foo` can take any combination of arguments, as in C. `-pedantic` implies `-fstrict-prototype` unless overridden with `-fno-strict-prototype`. Specifying this option will also suppress implicit declarations of functions. This option no longer affects declarations with C++ linkage.

`-fthis-is-variable`

Permit assignment to `this`. The incorporation of user-defined free store management into C++ has made assignment to `this` an anachronism. Therefore, by default it is invalid to assign to `this` within a class member function; that is, GNU C++ treats `this` in a member function of class `X` as a non-lvalue of type `X*`. For backwards compatibility, make it valid with `-fthis-is-variable`.

`-nostdinc`

Do not search for header files in the standard directories specific to C++, but do still search the other standard directories when building the C++ library.

IMPORTANT! As with all options that change the ABI, all C++ code, including `libgcc.a`, must be built with the same setting of this option.

`-Wctor-dtor-privacy`

Warn when a class seems unusable, because all the constructors or destructors in a class are private and the class has no friends or public static member functions.

`-Wno-deprecated`

Do not warn about usage of deprecated features.

`-Weffc++`

Warn about violation of some style rules from *Effective C++* by Scott Myers.

`-Wno-non-template-friend`

See “Options Requesting or Suppressing Warnings” on page 35.

`-Wnon-virtual-dtor`

Warn when a class declares a non-virtual destructor that should probably be virtual, because it looks like the class will be used polymorphically.

`-Wold-style-cast`

Warn if an old-style (C-style) cast is used within a C++ program. The new-style casts (`static_cast`, `reinterpret_cast`, and `const_cast`) are less vulnerable to unintended effects.

`-Woverloaded-virtual`

Warnings that apply only to C++ programs. See “Options Requesting or Suppressing Warnings” on page 35.

`-Wno-pmf-conversions`

Disable the diagnostic for converting a bound pointer to member function to a plain pointer.

`-Wreorder`

Warn when the order of member initializers given in the code does not match the order in which they must be executed, such as the following initializers.

- `-Wunknown-pragmas`

Warn when a `#pragma` directive is encountered which is not understood by GCC. If this command line option is used, warnings will even be issued for unknown pragmas in system header files. This is not the case if the warnings were only enabled by the `-Wall` command line option.

- `-Wall`

All of the previous `-w` options combined. This enables all the warnings about constructions that some users consider questionable, and that are easy to avoid (or modify to prevent the warning), even in conjunction with macros.

The following options are not implied by `-Wall`; some of them warn about constructions that users generally do not consider questionable, but for which occasionally you might wish to check; others warn about constructions that are necessary or hard to avoid in some cases, and there is no simple way to modify the code to suppress the warning.

- `-W`

Print extra warning messages for the following events:

- A function can return either with or without a value. Falling off the end

of the function body is considered returning without a value; for example, this function would evoke the following warning:

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

- An expression-statement or the left-hand side of a comma expression contains no side effects. To suppress the warning, cast the unused expression to void. For example, an expression such as `x[i, j]` will cause a warning, but `x[(void)i, j]` will not.
- An unsigned value is compared against zero with `<` or `<=` input.
- A comparison like `x<=y<=z` appears; this is equivalent to `(x<=y ? 1 : 0) <= z`, which is a different interpretation from that of ordinary mathematical notation.
- Storage-class specifiers like `static` are not the first things in a declaration. According to the C Standard, this usage is obsolete.
- If `-Wall` or `-Wunused` is also specified, warn about unused arguments.
- A comparison between signed and unsigned values could produce an incorrect result when the signed value is converted to unsigned. But don't warn if you specify the `-Wno-sign-compare` option.
- An aggregate has a partly bracketed initializer. For example, the following code would evoke such a warning, because braces are missing around the initializer for `x.h`:

```
struct s ( int f, g; );
struct t ( struct s h; int i; );
struct t x = ( 1, 2, 3 );
```

- An aggregate has an initializer which does not initialize all members. For example, the following code would cause such a warning, because `x.h` would be implicitly initialized to zero:

```
struct s ( int f, g, h; );
struct s x = ( 3, 4 );
```

`-Wsign-promo`

Warn when overload resolution chooses a promotion from unsigned or enumeral type to a signed type over a conversion to an unsigned type of the same size. Previous versions of the GNU C++ compiler would try to preserve unsignedness, but the standard mandates the current behavior.

`-Wsynth`

Warn when the GNU C++ compiler's synthesis behavior does not match that of `cfront`; for instance, the following example shows the input to use for managing such behavior.

```
struct A ( operator int ();
A& operator = (int));
main () ( A a,b; a = b;
```

In the previous example, the GNU C++ compiler will synthesize a default `A& operator = (const A&);`, while `cfront` will use the user-defined `operator =`.

`-fhuge-objects`

Support virtual function calls for objects that exceed the size representable by a `short int` call. You should not use this option by default; if you need to use it, the compiler will tell you so. This option is not useful when compiling with the `-fvtable-thunks` option

IMPORTANT! As with all options that change the ABI, all C++ code, including `libgcc`, must be built with the same setting of this option.

Compiling C++ Programs

C++ source files conventionally use one of the following suffixes.

- `.C`
- `.cc`
- `.cpp`
- `.cxx`

Preprocessed C++ files use the `.ii` suffix. GCC recognizes files with these names and compiles them as C++ programs even if you call the compiler the same way as for compiling C programs (usually with the call, `gcc`). However, C++ programs often require class libraries as well as a compiler that understands the C++ language and, under some circumstances, you might want to compile programs from standard input, or, otherwise, without a suffix that options them as C++ programs.

`G++` is a program that calls GCC with the default language set to C++, and automatically specifies linking against the GNU class library, `libg++`[†]. On many systems, the script, `g++`, is also installed as `c++`. When you compile C++ programs, you may specify many of the same command line options that you use for compiling programs in any language, or command line options meaningful for C and related languages, or options that are meaningful only for C++ programs.

See “Options Controlling C Dialect” on page 19 for documentation with explanations of options for languages related to C. See “Options Controlling C++

[†] Prior to release 2 of the compiler, there was a separate `G++` compiler. That version was based on GNU CC, but not integrated with it. Versions of `G++` with a 1.xx version number (for example, `G++` version 1.37 or 1.42) are much less reliable than the versions integrated with GCC 2. Moreover, combining `G++` 1.xx with a GCC version 2 will simply *not* work.

Dialect” on page 25 for documentation with explanations of options that are meaningful only for C++ programs.

6

Options Requesting or Suppressing Warnings

Warnings are diagnostic messages reporting constructions which are not inherently erroneous; they may warn of risky constructions or constructions, actually, in error.

You can request many specific warnings with options beginning with `-w`; for instance, use `-Wimplicit` to request warnings on implicit declarations. Some warning options also have a negative form beginning `-Wno-`, which turn off warnings; for instance, `-Wno-implicit` turns off a warning whenever a function is used before being declared.

The following options control the amount and kinds of warnings produced by GNU CC.

`-fsyntax-only`

Check the code for syntax errors, but do not do anything beyond checking for syntax errors.

`-pedantic`

Issue all the warnings demanded by strict ANSI C and ISO C++; reject all programs that use forbidden extensions. Valid ANSI C programs should compile properly with or without this option (though a rare few will require `-ansi`). Without this option, some GNU extensions and traditional C features are supported; with this option, they are rejected. `-pedantic` does not cause warning messages using alternate keywords whose names begin and end with `__`. Pedantic warnings are also disabled in the expression that follows `__extension__`. However,

only system header files should use these escapes; application programs should avoid them. See “Alternate Keywords” on page 196. This option has no support.

`-pedantic-errors`

Like `-pedantic`, except that errors are produced rather than warnings.

`-w`

Inhibit all warning messages.

`-W`

Print extra warning messages for the following events.

- A nonvolatile automatic variable might be changed by a call to `longjmp`. These warnings are possible only for optimizing before compiling. The compiler sees only the calls to `setjmp`. It cannot know where `longjmp` will be called; a signal handler could call it at any point in the code. As a result, you may get a warning even when there is no problem because `longjmp` cannot be called at a place where it would cause a problem.
- A function can return either with or without a value. (Falling off the end of the function body is considered returning without a value.) The following example shows a function call that would evoke such a warning.

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

- An expression that does not do anything generates a warning; such expressions include an expression-statement or the left-hand side of a comma expression containing no side effects. To suppress the warning, cast the unused expression to `void`. For example, an expression such as `x[i, j]` causes a warning, but one such as `x[(void)i, j]` does not.
- An unsigned value is compared against zero with `<` or `<=`.
- A comparison like `x<=y<=z` appears; this is equivalent to `(x<=y ? 1 : 0) <= z`, which is a different interpretation from that of ordinary mathematical notation.
- Storage-class specifiers like `static` are not the first things in a declaration. According to the C Standard, this usage is obsolete.
- If `-Wall` or `-Wunused` is also specified, warn about unused arguments.
- An aggregate has an initializer which does not initialize all members. For example, the following code would cause such a warning, because `x.h` would be implicitly initialized to zero:

```
struct s { int f, g; };
struct s x = { 3, 4 };
```

`-Wall`

All of the `-w` options combined.

- Waggregate-return
Warn if any functions that return structures or unions are defined or called. In languages where you can return an array, this also elicits a warning.
- Wbad-function-cast
Warn whenever a function call is cast to a non-matching type. For example, warn if `int malloc()` is cast to `anything*`.
- Wcast-align
Warn whenever a pointer is cast such that the required alignment of the target is increased. For example, warn if a `char*` is cast to an `int*` on machines where integers can only be accessed at two- or four-byte boundaries.
- Wcast-qual
Warn whenever a pointer is cast that removes a type qualifier from the target type. For example, warn if a `const char*` is cast to an ordinary `char*`.
- Wchar-subscripts
Warn if an array subscript has type `char`. This is a common cause of error, especially when type is `signed` on some machines.
- Wcomment
Warn whenever a `/*` comment's start characters appears in a `/*` comment without closing comment characters, or whenever a backslash/newline appears in a `//` comment.
- Wconversion
Warn if a prototype causes a type conversion that is different from what would happen to the same argument in the absence of a prototype. This includes conversions of fixed point to floating and vice versa, and conversions changing the width or signedness of a fixed point argument except when the same as the default promotion. Also, warn if a negative integer constant expression is implicitly converted to an unsigned type. For example, warn about the assignment `x = -1` if `x` is unsigned. But do not warn about explicit casts like `(unsigned) -1`.
- Werror
Make all warnings into errors.
- Wformat
Check calls to `printf`, `scanf`, and other functions, to make sure that the arguments supplied have types appropriate to the format string specified.
- Wid-clash-len
Warn whenever two distinct identifiers match in the first `len` characters. This may help you prepare a program that will compile with some obsolete compilers.
- Wimplicit
Same as the `-Wimplicit-int` and `-Wimplicit-function-declaration` flags.
- Wimplicit-int
Warn when a declaration does not specify a type.
- Wimplicit-function-declaration
Give a warning or error message when a function is used before being declared.

`-Wno-import`

Inhibit warning messages about the use of `#import`.

`-Werror-implicit-function-declaration`

Give a warning or error when a function is used before being declared.

`-Wfloat-equal`

Warn if floating point values are used in equality comparisons. The idea behind this is that sometimes it is convenient to consider floating point values as approximations to infinitely precise real numbers. If you are doing this, then you need to compute (by analyzing the code, or in some other way) the maximum or likely maximum error that the computation introduces. Then, allow for the error when performing comparisons (and when producing output, although that is a different problem). In particular, instead of testing for equality, you would check to see whether the two values have ranges that overlap, using relational operators, so equality comparisons are probably mistaken.

`-Winline`

Warn if a function can not be inlined, even if the function is declared as inline, or if `-finline-functions` was given as an option.

`-Wlarger-than-len`

Warn whenever an object of larger than `len` bytes is defined.

`-Wlong-long`

`-Wno-long-long`

Warn if `long long` type is used. `-Wlong-long` is default. To inhibit the warning messages, use `-Wno-long-long`. The `-Wlong-long` and `-Wno-long-long` options are taken into account only when the `-pedantic` option is used.

`-Wmain`

Warn if the type of `main` is suspicious. `main` should be a function with external linkage, returning `int`, taking either zero, two, or three arguments of appropriate types.

`-Wmissing-declarations`

Warn if a global function is defined without a previous declaration, even if the definition itself provides a prototype. Use this option to detect global functions that are not declared in header files.

`-Wmissing-noreturn`

Warn about functions which might be candidates for the `noreturn` attribute. Note that these are only possible candidates, not absolute ones. Care should be taken to verify manually which functions actually do not ever return before adding the `noreturn` attribute; otherwise, subtle code generation bugs could be introduced.

`-Wmultichar`

Warn if a multicharacter constant (`FOOF`) is used. Usually they indicate a typo in the code, as they have implementation defined values, and should not be used in portable code.

`-Wno-import`

Inhibit warning messages about the use of `#import`.

`-Wpacked`

Warn if a structure is given the packed attribute, but the packed attribute has no effect on the layout or size of the structure. Such structures may be misaligned. In the following code example, the `f.x` variable in `struct bar` will be misaligned even though `struct bar` does not itself have the packed attribute:

```
struct foo ( int x;  char a, b, c, d;
__attribute__((packed));
struct bar (  char z;  struct foo f;);
```

`-Wpadded`

Warn if padding is included in a structure, either to align an element of the structure or to align the whole structure. Sometimes when this happens, it is possible to rearrange the fields of the structure to reduce the padding and so make the structure smaller.

`-Wparentheses`

Warn if parentheses are omitted in certain contexts, for instance, when there is an assignment in a context where a *truth value* is expected, or when operators are *nested* whose precedence may be confusing.

Also warn about constructions where there may be confusion to which `if` statement an `else` branch belongs. The following is an example of such a case.

```
{
  if (a)
    if (b)
      foo ();
  else
    bar ();
}
```

In C, every `else` branch belongs to the innermost possible `if` statement, which in the previous example is `if (b)`; this is often not what is expected, as the previous example shows with its indentation for the code. When there is the potential for this confusion, GNU C will issue a warning when this flag is specified. To eliminate the warning, add explicit braces around the innermost `if` statement so that there is no way the `else` could belong to the enclosing `if`. The resulting code would look like the following declaration.

```
{
  if (a)
  {
    if (b)
      foo ();
  }
  else
    bar ();
}
```

`-Wpointer-arith`

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

`-Wredundant-decls`

Warn if anything is declared more than once in the same scope, even in cases where multiple declaration is valid and changes nothing.

`-Wreturn-type`

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

`-Wshadow`

Warn whenever a local variable shadows another local variable.

`-Wsign-compare`

Warn when a comparison between signed and unsigned values could produce an incorrect result when the signed value is converted to unsigned. This warning is also enabled by `-w`. To get the other warnings of `-w` without this warning, use `-Wno-sign-compare`.

`-Wswitch`

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

`-Wtrigraphs`

Warn if any trigraphs are encountered (assuming they are enabled).

`-Wundef`

Warn if an undefined identifier is evaluated in an `#if` directive.

`-Wuninitialized`

An automatic variable is used without first being initialized.

These warnings are possible only in compiling for optimization, because they require data flow information that is computed only when optimizing. If you don't specify `-O`, you simply won't get these warnings.

These warnings occur only for variables that are candidates for register allocation. Therefore, they do not occur for a variable that is declared `volatile`, or whose address is taken, or whose size is other than one, two, four or eight bytes. Also, they do not occur for structures, unions, or arrays, even when they are in registers.

IMPORTANT! There may be no warning about a variable that is used only to compute a value that itself is never used. Such computations may be deleted by data flow analysis before the warnings display.

These warnings are made optional because the code might be correct despite

appearing to have an error. The following is one example of how this optimization warning occurs.

```
{
    int x;
    switch (y)
    {
        case 1: x = 1;
            break;
        case 2: x = 4;
            break;
        case 3: x = 5;
            }
    foo (x);
}
```

If the value of `y` is always 1, 2 or 3, then `x` is always initialized, but GNU CC does not make this determination.

-Wunknown-pragmas

Warn when a `#pragma` directive is encountered which is not understood by GCC. If used, warnings will be issued for unknown pragmas in system header files, which is not the case if the warnings were enabled by `-Wall`.

-Wunreachable-code

Warn if the compiler detects that code will never be executed, especially when the compiler detects that at least a whole line of source code will never be executed, because some condition is never satisfied, or because it is after a procedure that never returns.

This option may produce a warning even though there are circumstances under which part of the affected line can be executed. Care should be taken when removing apparently unreachable code when, for instance, a function is inlined, where a warning may mean that the line is unreachable in only one inlined copy of the function.

IMPORTANT! `-Wunreachable-code` is not made part of `-Wall` because, in a debugging version of a program, there is often substantial code that checks correct functioning of the program and is unreachable because the program does work, or its code is selectable at compile time.

-Wunused

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

To suppress this warning for an expression, simply cast it to `void`. For unused variables, parameters and labels, use the `unused` attribute (for explanation of specifying this attribute, see “Specifying Attributes of Variables” on page 169).

The following code would have no warning message because `save_y` is used only

if it is set. Avoid such warnings by declaring all the functions you use that never return as `noreturn` (see “Declaring Attributes of Functions” on page 161).

```
{
    int save_y; if (change_y) save_y = y, y = new_y;
    ...
    if (change_y) y = save_y;
}
```

`-Wunused-function`

Warn whenever a static function is declared but not defined or a non-inline static function is unused.

`-Wunused-label`

Warn whenever a label is declared but not used. To suppress this warning use the `unused` attribute (see “Specifying Attributes of Variables” on page 169).

`-Wunused-parameter`

Warn whenever a function parameter is not used aside from its declaration. To suppress this warning use the `unused` attribute (see “Specifying Attributes of Variables” on page 169).

`-Wunused-variable`

Warn whenever a local variable or non-constant static variable is not used aside from its declaration. To suppress this warning use the `unused` attribute (see “Specifying Attributes of Variables” on page 169).

`-Wunused-value`

Warn whenever a statement computes a result that is explicitly not used. To suppress this warning cast the expression to `void`.

`-Wunknown-pragmas`

Warn when a `#pragma` directive is encountered which is not understood by GCC. If this command line option is used, warnings will even be issued for unknown pragmas in system header files. This is not the case if the warnings were only enabled by the `-Wall` command line option.

`-Wbad-function-cast`

(C only) Warn whenever a function call is cast to a non-matching type. For example, warn if `int malloc()` is cast to `anything*` in the code.

`-Wmissing-prototypes`

(C only) Warn if a global function is defined without a previous prototype declaration. This warning is issued even if the definition itself provides a prototype. The aim is to detect global functions that fail to be declared in header files.

`-Wnested-externs`

(C only) Warn if an `extern` declaration is encountered within a function.

`-Wstrict-prototypes`

(C only) Warn if a function is declared or defined without specifying the argument types. (An old-style function definition is permitted without a warning if preceded by a declaration which specifies the argument types.)

`-Wtraditional`

Warn about certain constructs that behave differently in traditional and ANSI C.

- Macro arguments occurring within string constants in the macro body. These would substitute the argument in traditional C, but are part of the constant in ANSI C.
- A function declared external in one block and then used after the end of the block.
- A `switch` statement that has an operand of type `long`.

`-Wreorder`

(C++ only) Warn when the order of member initializers given in the code does not match the order in which they must be executed. For instance, in the following example, the compiler will warn that the member initializers for `i` and `j` will be rearranged to match the declaration order of the members.

```
struct A {
    int i;
    int j;
    A(): j (0), i (1) { }
};
```

`-Wtemplate-debugging`

(C++ only) When using templates, warn if debugging is not yet fully available.

`-Wwrite-strings`

Give string constants the type `const char[length]` so that copying the address of one into a non-`const char*` pointer will get a warning.

The following warnings will help you find code, at compile time, that can write into a string constant, but only if you have been using `const` in declarations and prototypes.

`-Wall` does not request these warnings.

`-Wno-non-template-friend`

Disable warnings when non-templated friend functions are declared within a template.

With the advent of explicit template specification support in G++, if the name of the friend is an unqualified identifier (such as `friend foo(int)`), the C++ language specification demands that the friend declare or define an ordinary, non-template function. Before G++ implemented explicit specification, unqualified identifiers could be interpreted as a particular specialization of a templated function. Because this non-conforming behavior is no longer the default behavior for G++, `-Wnon-template-friend` allows the compiler to check existing code for potential trouble spots, and is on by default. This compiler behavior can also be disabled with `-fguiding-decls`, which activates the older,

non-specification compiler code, or disabled with `-Wno-non-template-friend`, which keeps the conformant compiler code but disables the warning.

`-Wold-style-cast`

(C only) Warn if an old-style cast is used within a program.

`-Woverloaded-virtual`

(C++ only) Warns when a derived class function declaration may be an error in defining a virtual function. In a derived class, the definitions of virtual functions must match the type signature of a virtual function declared in the base class. With this option, the compiler warns when you define a function with the same name as a virtual function, but with a type signature that does not match any declarations from the base class.

`-Wsynth`

(C++ only) When G++'s synthesis behavior does not match that of `cfront`. For instance, see the following declaration:

```
struct A {
    operator int ();
    A& operator = (int);
};
main ()
{
    A a,b;
    a = b;
}
```

In this previous example, G++ will synthesize a default:

```
A& operator = (const A&);,
```

`cfront` will use the user-defined expression, `operator=`.

7

Options Controlling Debugging

When using GCC, some of what the debugger can help with is profiling, which allows you to learn where your program spent its time and which functions called which other functions while it was executing. This information can show you which pieces of your program are slower than you expected, and which might be candidates for rewriting to make your program execute faster. It can also tell you which functions are being called more or less often than you expected. This may help you spot bugs that had otherwise been unnoticed. Since the profiler uses information collected during the actual execution of your program, it can be used on programs that are too large or too complex to analyze by reading the source. However, how your program is run will affect the information that shows up in the profile data. If you don't use some feature of your program while it is being profiled, no profile information will be generated for that feature.

WARNING! Profiling for even a few seconds on some systems produces a very large file.

The following documentation discusses some special options for debugging either your program or debugging GCC for special needs for your program.

-a

Generate extra code to write profile information for basic blocks, which records the number of times each basic block is executed, the basic block start address, and the function name containing the basic block. If `-g` is used, the line number and filename of the start of the basic block is recorded. If not overridden by the

machine description, the default action is to append to the text file, `bb.out`. This data could be analyzed by a program like `tcov`, a program which produces a test coverage analysis of a compiled program, taking source files as arguments and producing an annotated source listing.

However, the format of the data is not what `tcov` expects. `gprof`, the GNU profiler tool, can help to process this data. For documentation on `gprof`, see <http://sources.redhat.com/binutils/docs-2.10/gprof.html>.

`-ax`

Generate extra code to profile basic blocks. The executable produces output that is a superset of what `-a` produces. Additional output is the source and target address of the basic blocks where a jump takes place, the number of times a jump is executed, and (optionally) the complete sequence of basic blocks being executed. The output is appended to file, `bb.out`.

You can examine different profiling aspects without recompilation. Your executable will read a list of function names from file `bb.in`. Profiling starts when a function on the list is entered and stops when that invocation is exited. To exclude a function from profiling, use a hyphen (-) to prefix its name.

If a function name is not unique, specify its location.

Use the following example's declaration (where *functionname* designates the unique function to specify for the location).

```
/path/file-name.d: functionname
```

Your executable will write the available paths and filenames in file `bb.out`.

The following function names have a special meaning.

- `__bb_jumps__`
Write source, target, and frequency of jumps to file `bb.out`.
- `__bb_hidecall__`
Exclude function calls from frequency count.
- `__bb_showret__`
Include function returns in frequency count.
- `__bb_trace__`
Write the sequence of basic blocks to the `bbtrace.gz` file. The file will be compressed using the program, `gzip`, which must exist in your `PATH`. On systems *without* the `popen` function, the file will be named `bbtrace` and will not be compressed.

IMPORTANT! `__bb_hidecall__` and `__bb_showret__` will not affect the sequence written to `bbtrace.gz`.

The following example shows the use of different profiling parameters in a `bb.in` file. Assume a `foo` function consists of basic blocks 1 and 2, called twice from block 3 of the `main` function. After the calls, block 3 transfers control to block 4 of `main`. With `__bb_trace__` and `main` contained in a `bb.in` file, a sequence of

blocks is written to a `bbtrace.gz` file: 0 3 1 2 1 2 4. The return from block 2 to block 3 is not shown, because the return is to a point inside the block and not to the top. The block address 0 always indicates that control is transferred to the trace from somewhere outside the observed functions. With `-foo` added to `bb.in`, the blocks of the `foo` function are removed from the trace, so only 0 3 4 remains. With `__bb_jumps__` and `main` contained in a `bb.in` file, jump frequencies will be written to a `b.out` file. The frequencies are obtained by constructing a trace of blocks and incrementing a counter for every neighboring pair of blocks in the trace. The 0 3 1 2 1 2 4 trace displays the following frequencies.

```
Jump from block 0x0 to block 0x3 executed 1 time(s)
Jump from block 0x3 to block 0x1 executed 1 time(s)
Jump from block 0x1 to block 0x2 executed 2 time(s)
Jump from block 0x2 to block 0x1 executed 1 time(s)
Jump from block 0x2 to block 0x4 executed 1 time(s)
```

With `__bb_hidecall__`, due to call instructions, control transfer is removed from the trace; that is, the trace is cut into three parts: 0 3 4, 0 1 2, and 0 1 2. With `__bb_showret__`, control transfer is added to the trace.

The trace becomes: 0 3 1 2 3 1 2 3 4.

IMPORTANT! The previous trace is not the same as the sequence written to `bbtrace.gz`. It is solely used for counting jump frequencies.

`-dletters`

Make debugging dumps during compilation at times specified by *letters*. This is only used for debugging the compiler. The filenames for most of the dumps are made by appending a word to the source filename (for example, `file.c.rtl` or `file.c.jump`). What follows are the possible letters for use in *letters*, and their meanings for use in a filename, *file*, with their appropriate extensions.

- a Produce all the dumps previously listed.
- A Annotate the assembler output with miscellaneous debugging information.
- b Dump after computing branch probabilities to `file.bp`.
- c Dump after instruction combination to `file.combine`.
- d Dump after delayed branch scheduling to `file.dbr`.
- D Dump after purging `ADDRESSOF` to `file.addressof`.
- f Dump after flow analysis to `file.flow`.
- g Dump after global register allocation to `file.greg`.

- G Dump after the Global Common Sub-expression Elimination (GCSE) pass to *file.gcse*.
 - j Dump after first jump optimization to *file.jump*.
 - J Dump after last jump optimization to *file.jump2*.
 - k Dump after conversion from registers to stack to *file.stack*.
 - l Dump after local register allocation to *file.lreg*.
 - L Dump after loop optimization to *file.loop*.
 - m Print statistics on memory usage, at the end of the run, to standard error.
 - M Dump after performing the machine dependent reorganization pass to *file.mach*.
 - N Dump after the register move pass to *file.regmove*.
 - P Annotate the assembler output with a comment indicating which pattern and alternative was used.
 - r Dump after register transfer language (RTL) generation to *file.rtl*.
 - R Dump after the second instruction scheduling pass to *file.sched2*.
 - s Dump after the Common Sub-expression Elimination (CSE) pass, including the jump optimization that sometimes follows the CSE to *file.cse*.
 - S Dump after the first instruction scheduling pass to *file.sched*.
 - t Dump after the second CSE pass (including the jump optimization that sometimes follows CSE) to *file.cse2*.
 - x Generate RTL for a function instead of compiling it. Usually used with *r*.
 - Y Dump debugging information during parsing to standard error.
- `-fdump-unnumbered`
When performing debugging dumps (see previous descriptions for `-d` option),

suppress instruction numbers and line number note output. This allows using `diff` on debugging dumps for compiler invocations with different options, in particular with and without `-g` (see the following descriptions for the `-g` option).

`-fdump-translation-unit-file`

(C++ only) Dump a representation of the tree structure for the entire translation unit to a *file*.

`-fpretend-float`

When running a cross-compiler, pretend that the target machine uses the same floating point format as the host machine. This causes incorrect output of the actual floating constants, but the actual instruction sequence will probably be the same as GCC would make when running on the target machine.

`-fprofile-arcs`

Instrument *arcs* during compilation. For each function of your program, GCC creates a program flow graph, then finds a spanning tree for the graph. Only arcs that are not on the spanning tree have to be instrumented: the compiler adds code to count the number of times that these arcs are executed. When an arc is the only exit or only entrance to a block, the instrumentation code can be added to the block; otherwise, a new basic block must be created to hold the instrumentation code. Since not every arc in the program must be instrumented, programs compiled with this option run faster than programs compiled with `-a`, which adds instrumentation code to every basic block in the program. The tradeoff is that, since `gcov` does not have execution counts for all branches, it must start with the execution counts for the instrumented branches, and then iterate over the program flow graph until the entire graph has been solved. Thus, `gcov` runs a little more slowly than a program which uses information from `-a`; for more information on `gcov`, see “`gcov`, a Test Coverage Program” on page 209.

`-fprofile-arcs` also makes it possible to estimate branch probabilities, and to calculate basic block execution counts. In general, basic block execution counts do not give enough information to estimate all branch probabilities. When the compiled program exits, it saves the arc execution counts to a file called *sourcename.da*. Use the `-fbranch-probabilities` option when recompiling to optimize using estimated branch probabilities (see also “Options Controlling Optimization” on page 53).

`-ftest-coverage`

Create data files for the `gcov` code-coverage utility (see “`gcov`, a Test Coverage Program” on page 209). The data filenames begin with the name of your source file:

- *sourcename.bb*

A mapping from basic blocks to line numbers, which `gcov` uses to associate basic block execution counts with line numbers.

- *sourcename.bbg*

A list of all arcs in the program flow graph. This allows `gcov` to reconstruct

the program flow graph, so that it can compute all basic block and arc execution counts from the information in the *sourcename.da* file (this last file is the output from `-fprofile-arcs`).

`-g`

Produce debugging information in the operating system's native format (such as *stabs*, COFF, XCOFF, or DWARF), so that GDB works with the format. On most systems using *stabs* format, `-g` enables use of extra debugging information that only GDB can use; this extra information makes debugging work better in GDB, but may make other debuggers crash or refuse to read the program.

If you want to control for certain whether to generate the extra information, use `-gstabs+`, `-gstabs`, `-gxcoff+`, `-gxcoff`, `-gdwarf+1`, or `-gdwarf-1` (see the explanations for each option in the following discussions).

Unlike most other C compilers, GCC allows you to use `-g` with `-O`. The shortcuts using optimized code may produce unexpected results. Some variables you declared may not exist at all; flow of control may briefly move where you did not expect it; some statements may not be executed because they compute constant results or their values were already at hand; some statements may execute in different places because they were moved out of loops. Nevertheless, optimized output is possible to debug, making it reasonable to use optimization for programs that might have bugs.

The following options are useful when GCC is generated with the capability for more than one debugging format.

`-glevel`

`-ggdblevel`

`-gstabslevel`

`-gcofflevel`

`-gxcofflevel`

`-gdwarflevel`

`-gdwarf-2level`

Request debugging information, using *level* to specify how much information; there are three default levels, the default being 2. Level 1 produces minimal information, enough for making backtraces in parts of the program that you don't plan to debug; this includes descriptions of functions and external variables, but no information about either local variables or line numbers. Level 3 includes extra information, such as all the macro definitions present in the program. Some debuggers support macro expansion when you use the `-g3` option.

`-ggdb`

Produce debugging information in the native format (if that is supported), including GDB extensions if at all possible.

`-gstabs`

Produce debugging information in *stabs* format (if that is supported), without GDB extensions. This is the format used by DBX on most BSD systems. On

MIPS, Alpha, and System V Release 4 systems, this option produces `stabs` debugging output which is not understood by either DBX or SDB. On System V Release 4 systems, this option requires the GNU assembler.

`-gstabs+`

Produce debugging information in `stabs` format (if that is supported), using GNU extensions understood only by GDB. The use of these extensions is likely to make other debuggers crash or refuse to read the program.

`-gcoff`

Produce debugging information in COFF format (if that is supported). This is the format used by SDB on most System V systems prior to System V Release 4.

`-gxcoff`

Produce debugging information in XCOFF format (if that is supported). This is the format used by the DBX debugger on IBM RS/6000 systems.

`-gxcoff+`

Produce debugging information in XCOFF format, using GNU extensions understood only by GDB. The use of these extensions is likely to make other debuggers crash or refuse to read the program, and may cause assemblers other than the GNU assembler to fail with an error.

`-gdwarf`

Produce debugging information in DWARF version 1 format (if that is supported). This is the format used by SDB on most System V Release 4 systems.

`-gdwarf+`

Produce debugging information in DWARF version 1 format (if that is supported), using GNU extensions understood only by GDB. The use of these extensions is likely to make other debuggers crash or refuse to read the program.

`-gdwarf-2`

Produce debugging information in DWARF version 2 format (if that is supported). This is the format used by DBX on IRIX 6.

`-p`

Generate extra code to write profile information suitable for the analysis program `prof`. You must use this option when compiling the source files for which you want debugging data, and you must also use it when linking.

`-pg`

Generate extra code to write profile information suitable for the analysis program `gprof`. You must use this option when compiling the source files for which you want debugging data, and you must also use it when linking.

`-print-file-name=library`

Print the full absolute name of the library file, *library*, used when linking, without doing anything else. With this option, GCC does not compile or link anything; it just prints the filename.

`-print-prog-name=program`

Like `-print-file-name`, searches for a *program* such as `cpp`.

`-print-libgcc-file-name`

Same as `-print-file-name=`. Useful when you use `-nostdlib` or `-nodefaultlibs` when you want to link with `libgcc.a`, for instance, using the following example's command, where *files* signifies the files needed to link.

```
gcc -nostdlib files
```

`-print-search-dirs`

Only print the name of the configured installation directory and a list of program and library directories GCC will search. This is useful when GCC prints the following error message:

```
installation problem, cannot exec cpp: No such file or directory
```

To resolve this you either need to put `cpp` and the other compiler components where GCC expects to find them, or set the environment variable,

`GCC_EXEC_PREFIX`, to the directory where you installed them. Don't forget the trailing slash; see "Environment Variables Affecting GCC" on page 137.

`-Q`

Prints each function name as it is compiled, with some status about each pass when it finishes.

`-save-temps`

Store the usual temporary intermediate files permanently; place them in the current directory and name them based on the source file. Compiling `foo.c` with `-c-save-temps` would produce `foo.i` and `foo.s` files, as well as a `foo.o` file.

8

Options Controlling Optimization

The following documentation discusses options controlling compiler optimizations.

Options of the form, `-fflag`, specify machine-independent options. Most options have both positive and negative forms; the negative form of `-ffoo` is `-fno-foo`.

`-falign-functions`

`-fno-align-functions`

`-falign-functions=n`

`-falign-functions` aligns the start of functions to the next power-of-two greater than *n*, skipping up to *n* bytes. For instance, `-falign-functions=32` aligns functions to the next 32-byte boundary, but `-falign-functions=24` would align to the next 32-byte boundary only if this can be done by skipping 23 bytes or less.

`-fno-align-functions` and `-falign-functions=1` are equivalent and mean that functions will not be aligned. Some assemblers only support this option when *n* is a power of two; in that case, it is rounded up. If *n* is not specified, use a machine-dependent default.

`-falign-labels`

`-falign-labels=n`

Align all branch targets to a power-of-two boundary, skipping up to *n* bytes like `-falign-functions`. This option may make code slower, because it must insert dummy operations for when the branch target is reached in the usual flow of the code. If `-falign-loops` or `-falign-jumps` are applicable and are greater than

this value, then the value is used instead. If n is not specified, use a machine-dependent default which is likely to be 1, meaning no alignment.

`-falign-loops`

`-falign-loops= n`

Align loops to a power-of-two boundary, skipping up to n bytes like `-falign-functions`. The loop will be executed many times, which will make up for any execution of the dummy operations. If n is not specified, use a machine-dependent default.

`-falign-jumps`

`-falign-jumps= n`

Align branch targets to a power-of-two boundary, for branch targets where the targets can only be reached by jumping, skipping up to n bytes like `-align-functions`. In this case, no dummy operations must be executed. If n is not specified, use a machine-dependent default.

`-fbranch-probabilities`

After running a program compiled with `-fprofile-arcs` (see “Options Controlling Debugging” on page 45), you can compile it again using `-fbranch-probabilities`, to improve optimizations based on guessing the path a branch might take.

`-fcaller-saves`

Enable values to be allocated in registers that will be clobbered by function calls, by emitting extra instructions to save and restore the registers around such calls. Such allocation is done only when it seems to result in better code than would otherwise be produced. This option is enabled by default on certain machines, usually those which have no call-preserved registers to use instead.

`-fcse-follow-jumps`

In common subexpression elimination, scan through jump instructions when the target of the jump is not reached by any other path. For example, when common subexpression elimination (CSE) encounters an `if` statement with an `else` clause, CSE follows the jump when the condition tested is false.

`-fcse-skip-blocks`

This is similar to `-fcse-follow-jumps`, but causes CSE to follow jumps which conditionally skip over blocks. When CSE encounters an `if` statement with no `else` clause, `-fcse-skip-blocks` causes CSE to follow the jump around the body of the `if`.

`-flive-range`

Perform live range splitting of variables at loop boundaries. This option is enabled by default at `-O2` optimization and higher for targets using stabs debug symbols.

`-fdelayed-branch`

If supported for the target machine, attempt to reorder instructions to exploit instruction slots available after delayed branch instructions.

- `-fdelete-null-pointer-checks`

Use global dataflow analysis to identify and eliminate useless null pointer checks. Programs which rely on null pointer dereferences not halting the program may not work properly with this option. Use `-fno-delete-null-pointer-checks` to disable optimizing for programs depending on that behavior.
- `-fexpensive-optimizations`

Perform a number of minor optimizations that are relatively expensive.
- `-ffast-math`

This option allows GCC to violate some ANSI or IEEE rules and/or specifications in the interest of optimizing code for speed. For example, it allows the compiler to assume that arguments to the `sqrt` function are non-negative numbers and that no floating point values are NaNs. The `-ffast-math` option sets `-fno-math-errno`. This option should never be turned on by any `-O` option since it can result in incorrect output for programs that depend on an exact implementation of IEEE or ANSI rules/specifications for math functions. A program that relies on IEEE exceptions for math error handling may want to use this option for speed while maintaining IEEE arithmetic compatibility.
- `-ffloat-store`

Do not store floating point variables in registers, and inhibit other options that might change whether a floating point value is taken from a register or memory. This option prevents undesirable excess precision on machines such as the 68000 where the floating registers (of the 68881) keep more precision than a `double` is supposed to have. For most programs, the excess precision does only good, but a few programs rely on the precise definition of IEEE floating point, so using `-ffloat-store` is better for such programs.
- `-fforce-addr`

Force memory address constants to be copied into registers before doing arithmetic on them. This may produce better code just as `-fforce-mem` may.
- `-fforce-mem`

Force memory operands to be copied into registers before doing arithmetic on them. This produces better code by making all memory references potential common subexpressions. When they are not common subexpressions, instruction combination should eliminate the separate register-load. The `-O2` option turns on this option.
- `-fno-math-errno`

Do not set `ERRNO` after calling math functions that are executed with a single instruction, such as `sqrt`. A program that relies on IEEE exceptions for math error handling may want to use this option for speed while maintaining IEEE arithmetic compatibility. The default is `-fmath-errno`. The `-ffast-math` option sets `-fno-math-errno`.

`-ffunction-sections`

`-fdata-sections`

Place functions or data into their own section in the output file, if the target supports arbitrary sections. The function's name determines the section's name in the output file. Use these options on systems where the linker can perform optimizations to improve locality of reference in the instruction space. HPPA processors running HP-UX and SPARC processors running Solaris 2 have linkers with such optimizations. Other systems using the ELF object format as well as AIX may have these optimizations in the future. Only use these options when there are significant benefits from doing so. When you specify these options, the assembler and linker will create larger object and executable files and they will also run slower. You will not be able to use `gprof` on all systems when specifying these options. You may have problems debugging when specifying these options and also using the `-g` option.

`-fgcse`

Perform a global common subexpression elimination (GCSE) pass. This pass also performs global constant and copy propagation.

`-finline-limit=n`

By default, GCC limits the size of functions that can be inlined. This option controls the size of functions that are explicitly marked as `inline` (marked with the `inline` keyword or defined within the class definition in C++). *n* is the specification of the size of functions that can be inlined in the number of pseudo instructions, not counting parameter handling. The default value of *n* is 10000. Increasing this value can result in more inlined code at the cost of compilation time and memory consumption. Decreasing the value usually makes the compilation faster and less code will be inlined, which may create slower running programs. This option is particularly useful for programs that use heavy inlining, such as those based on recursive templates with C++. Pseudo instruction represents an abstract measurement of function's size; it does not represent a count of assembly instructions and, as such, its exact meaning might change.

`-finline-functions`

Integrate all simple functions into their callers. The compiler heuristically decides which functions are simple enough to be worth integrating in this way. If all calls to a given function are integrated, and the function is declared `static`, then the function is normally not output as assembler code in its own right.

`-fkeep-inline-functions`

Even if all calls to a given function are integrated, and the function is declared `static`, output a separate run-time callable version of the function. This switch does not affect `extern inline` functions.

`-fmove-all-movables`

Forces all invariant computations in loops to be moved outside the loop.

`-fno-default-inline`

C++ only. Do not make member functions inline by default merely because they are defined inside the class scope. Otherwise, when you specify `-O`, member functions defined inside class scope are compiled inline by default; that is, you don't need to add `inline` in front of the member function name.

`-fno-defer-pop`

Always pop the arguments to each function call as soon as that function returns. For machines which must pop arguments after a function call, the compiler normally lets arguments accumulate on the stack for several function calls and pops them all at once.

`-fno-function-cse`

Do not put function addresses in registers; make each instruction that calls a constant function contain the function's address explicitly. `fno-function-cse` results in less efficient code, but some strange hacks that alter the assembler output may be confused by the optimizations performed when this option is not used.

`-fno-inline`

Don't pay attention to the `inline` keyword. This option is useful for keeping the compiler from expanding any functions inline.

IMPORTANT! If you are not optimizing, no functions can be expanded inline.

`-fomit-frame-pointer`

Don't keep the frame pointer in a register for functions that don't need one. This avoids the instructions to save, set up, and restore frame pointers; it also makes an extra register available in many functions.

`-fkeep-static-consts`

Emit variables declared `static const` when optimization isn't turned on, even if the variables weren't referenced. Disabled by default, `-fno-keep-static-consts` will force the compiler to check if the variable was referenced, regardless of whether or not optimization is turned on.

WARNING! `-fno-keep-static-consts` may make debugging impossible on some machines. On some machines, such as VAX machine, this option has no effect because the standard calling sequence automatically handles the frame pointer and nothing is saved by pretending it doesn't exist. The `FRAME_POINTER_REQUIRED` machine-description macro controls whether a target machine supports this option. See "Constraints for Particular Machines" on page 185 to determine register usage with your target machine.

`-fno-peephole`

Disable any machine-specific peephole optimizations.

`-foptimize-sibling-calls`

Optimize sibling and tail recursive calls.

`-foptimize-register-moves`

`-fregmove`

Some machines only support two operands per instruction, and on such machines, GCC might perform extra copying; `-fregmove` or `-foptimize-register-moves` override the default for the machine to do the copy before register allocation.

`-frerun-cse-after-loop`

Re-run common subexpression elimination after loop optimizations has been performed.

`-frerun-loop-opt`

Run the loop optimizer twice.

`-freduce-all-givs`

Forces all general-induction variables in loops to be strength-reduced.

IMPORTANT! When compiling programs written in FORTRAN, `-fmove-all-moveables` and `-freduce-all-givs` are enabled by default when you use the optimizer.

Some options may generate better or worse code; results are highly dependent on the structure of loops within the source code.

`-fschedule-insns`

If supported on the target machine, attempt to reorder instructions to eliminate execution stalls due to required data being unavailable. This helps machines that have slow floating point or memory load instructions, allowing other instructions to be issued until the result of the load or floating point instruction is required.

`-fschedule-insns2`

Similar to `-fschedule-insns`, but requests an additional pass of instruction scheduling after register allocation has been done. This is useful on machines with a relatively small number of registers and where memory load instructions take more than one cycle.

`-fssa`

Perform optimizations in static single assignment form. Each function's flow graph is translated into Static Single Assignment (SSA) form, optimizations are performed, and the flow graph is translated back from SSA form. Currently, no SSA-based optimizations are implemented, but converting into and out of SSA form is not an invariant operation, and generated code may differ.

Fine-tuning Optimizations

Use the following options in the rare cases when you want to fine-tune optimizations.

`-fstrength-reduce`

Perform optimizations of loop strength reduction and elimination of iteration variables.

`-fstrict-aliasing`

Allows the compiler to assume the strictest aliasing rules applicable to the

language being compiled. For C and C++, this activates optimizations based on the type of expressions. In particular, an object of one type is assumed never to reside at the same address as an object of a different type, unless the types are almost the same. For example, an `unsigned int` can alias an `int`, but an `unsigned int` cannot alias a `void*` or a `double`. A character type may alias any other type. Pay special attention to code like the following example shows:

```
union a_union {
    int i;
    double d;
};
int f() {
    a_union t;
    t.d = 3.0;
    return t.i;
}
```

There is a common practice of reading from a different union member than the one to which was most recently written (a practice known as *type-punning*). Even with `-fstrict-aliasing`, type-punning is allowed, provided the memory is accessed through the union type. So, the previous example's code will work as expected. However, the following example's code might not:

```
int f() {
    a_union t;
    int* ip;
    t.d = 3.0;
    ip = &t.i;
    return *ip;
}
```

`-fthread-jumps`

Perform optimizations when checking to see if a jump branches to a location where another comparison subsumed by the first is found. If so, the first branch is redirected to either the destination of the second branch or a point immediately following it, depending on whether the condition is known to be true or false.

`-funroll-all-loops`

Perform the optimization of loop unrolling. This is done for all loops and usually makes programs run more slowly. `-funroll-all-loops` implies `-fstrength-reduce` as well as `-frerun-cse-after-loop`.

`-funroll-loops`

Perform the optimization of loop unrolling. This is only done for loops whose number of iterations can be determined at compile time or run time.

`-funroll-loop` implies both `-fstrength-reduce` and `-frerun-cse-after-loop`.

Frequently Used Optimization Options

The following options control the most frequently used optimizations. On most machines, the `-O` option turns on the `-fthread-jumps` and `-fdelayed-branch` options, but specific machines may handle it differently.

IMPORTANT! If you use multiple `-O` options, with or without level numbers, the last such option is the one that is effective.

`-O`
`-O0`
`-O1`

Optimize. Optimizing compilation takes somewhat more time, and a lot more memory for a large function. Without `-O`, the compiler's goal is to reduce the cost of compilation and to make debugging produce the expected results. Statements are independent: if you stop the program with a breakpoint between statements, you can then assign a new value to any variable or change the program counter to any other statement in the function and get exactly the results you would expect from the source code. Without `-O`, the compiler only allocates variables declared `register` in registers. The resulting compiled code is a little worse than produced by PCC without `-O`. With `-O`, the compiler tries to reduce code size and execution time. When you specify `-O`, the compiler turns on `-fthread-jumps` and `-fdefer-pop` on all machines. The compiler turns on `-fdelayed-branch` on machines that have delay slots, and `-fomit-frame-pointer` on machines that can support debugging even without a frame pointer. On some machines the compiler also turns on other options. Do not optimize with the `-O0` option.

`-O2`

Optimize even more. GNU CC performs nearly all supported optimizations that do not involve a space-speed tradeoff. The compiler does not perform loop unrolling or function inlining when you specify `-O2`. As compared to `-O`, this option increases both compilation time and the performance of the generated code. `-O2` turns on all of these optimizations except `-funroll-loops` and `-funroll-all-loops`, optimizations for loop unrolling, function inlining, life shortening, and static variable optimizations. It also turns on frame pointer elimination on machines where doing so does not interfere with debugging.

`-O3`

Optimize yet more. `-O3` turns on all optimizations specified by `-O2` and also turns on the option, `-finline-functions`.

`-Os`

Optimize for size. `-Os` enables all `-O2` optimizations that do not typically increase code size. It also performs further optimizations designed to reduce code size.

9

Options Controlling Preprocessing

The following documentation discusses options controlling the C preprocessor, when it is run on each C source file before actual compilation.

If you use the `-E` option, nothing is done except preprocessing. Some of these options make sense only together with `-E` because they cause the preprocessor output to be unsuitable for actual compilation. All instances of `-D` on the command line are processed before using `-U` option calls.

- `-E`
Run only the C preprocessor. Preprocess all the C source files specified and output the results to standard output or to a specified output file.
- `-Aquestion(answer)`
Assert the answer *answer* for *question*, in case it is tested with a preprocessing conditional such as `#if question(answer)`. `-A-` disables the standard assertions that normally describe the target machine.
- `-C`
Tell the preprocessor not to discard comments. Used with the `-E` option.
- `-dD`
Tell the preprocessor to pass all macro definitions into the output, in their proper sequence in the rest of the output.

- `-dM`
Tell the preprocessor to output only a list of the macro definitions that are in effect at the end of preprocessing. Used with the `-E` option.
- `-Dmacro`
Define *macro* with the string *1* as its defined value.
- `-Dmacro=defn`
Define a *macro* with *defn*. All instances of `-D` on the command line are processed before using `-U` option calls.
- `-dN`
Like `-dD` except that the macro arguments and contents are omitted. Only `#define name` is included in the output.
- `-H`
Print the name of each header file used, in addition to other normal activities.
- `-idirafter dir`
Add the *dir* directory to the second include path. The directories on the second include path are searched when a header file is not found in any of the directories in the main include path (the one that `-I` adds to; see also “Options for Searching Directories” on page 71).
- `-include file`
Process *file* as input before processing the regular input file. However, any `-D` and `-U` options on the command line are always processed before `-include file`, regardless of the order in which they are written. All the `-include` and `-imacros` options are processed in the order in which they are written.
- `-imacros file`
Process *file* as input, discarding the resulting output, before processing the regular input file. Because the output generated from *file* is discarded, the only effect of `-imacros file` is to make the macros defined in *file* available for use in the main input. Any `-D` and `-U` options on the command line are always processed before the `-imacros file`, regardless of the order in which they are written. All the `-include` and `-imacros` options are processed in the order in which they are written.
- `-iprefix prefix`
Specify *prefix* as the prefix for subsequent `-iwithprefix` options.
- `-iwithprefix dir`
Add a *dir* directory to the second include path. The directory’s name is made by concatenating *prefix* and *dir*, where *prefix* was specified previously with `-iprefix`. If *prefix* is not specified, the directory containing the installed passes of the compiler is used as the default.
- `-iwithprefixbefore dir`
Add a *dir* directory to the main include path. The directory’s name is made by concatenating *prefix* and *dir* like `-iwithprefix`.

- `-isystem dir`
Add a *dir* directory to the beginning of the second include path, marking it as a system directory, so that it gets the same treatment as is applied to the standard system directories.
- `-isystem-c++ dir`
Same behavior as with `-isystem`, but does not make headers in *dir* be implicitly evaluated as if they include the `extern "C"` linkage specification.
- `-M`
Tell the preprocessor to output a rule suitable for `make` describing the dependencies of each object file. For each source file, the preprocessor outputs one `make` rule whose target is the object filename for that source file and whose dependencies are all the `#include` header files it uses. This rule may be a single line or may be continued with `\` newline if it is long. The list of rules is printed on standard output instead of the preprocessed C program. `-M` is only used with `-E`. Another way to specify output of a `make` rule is by setting the environment variable, `DEPENDENCIES_OUTPUT` (see “Environment Variables Affecting GCC” on page 137).
- `-MD`
Like `-M` but the dependency information is written to a file made by replacing `.c` with `.d` at the end of the input filenames. This is in addition to compiling the file as specified. `-MD` does not inhibit ordinary compilation the way `-M` does. In Mach, you can use the utility `md` to merge multiple dependency files into a single dependency file suitable for using with the `make` command.
- `-MM`
Like `-M` but the output mentions only the user header files included with `include "file"`. System header files included with `#include file` are omitted.
- `-MMD`
Like `-MD` except output mentions only user header files, not system header files.
- `-MG`
Treat missing header files as generated files and assume they live in the same directory as the source file. If you specify `-MG`, you must also specify either `-M` or `-MM`. `-MG` is not supported with `-MD` or `-MMD`.
- `-nostdinc`
Do not search the standard system directories for header files. Only the directories you have specified with `-I` options (and the current directory, if appropriate) are searched. See “Options for Searching Directories” on page 71 for information on `-I`. By using both `-nostdinc` and `-I-`, you can limit the include-file search path to only those directories you specify explicitly.
- `-P`
Tell the preprocessor not to generate `#line` directives. Used with the `-E` option.
- `-trigraphs`
Support ANSI C trigraphs. The `-ansi` option also has this effect.

`-undef`

Do not predefine any nonstandard macros, including architecture flags.

`-Umacro`

Undefine a macro, *macro*. The `-U` options are evaluated after all `-D` options, but before any `-include` and `-imacros` options.

`-Wp, option`

Pass *option* as an option to the preprocessor. If *option* contains commas, it is split into multiple options at the commas.

10

Options Controlling the Assembler

You can pass options to the assembler. The following option is the only one regularly in use by the assembler.

`-wa, option`

Pass *option* as an option to the assembler. If *option* is used multiple times, use commas.

11

Options for Linking

The following documentation discusses options that are available for the compiler to link object files into an executable output file. They are meaningless if the compiler is not doing any link operations. For more information on the GNU linker, see “ld, the GNU Linker” on page 24 of *Getting Started Guide*; see also *Using ld* in *GNUPro Development Tools*.

object-file-name

A filename that does not end in a special recognized suffix is considered to name an object file or library. Object files are distinguished from libraries by the linker according to the file contents. If linking is done, these object files are used as input to the linker.

-c
-S
-E

If any of these options are used, then the linker is not run, and object filenames should not be used as arguments; for more information, see “Options Controlling the Kind of Output” on page 15.

-llibrary

Search the library named *library* when linking. It makes a difference where in the command you write this option; the linker searches processes, libraries, and object files, in the order that you specify them; a `foo.o -lz bar.o` command

searches a `z` library after a `foo.o` file but before a `bar.o` file, and if `bar.o` refers to functions in `z`, those functions may not be loaded.

If you specify asking the linker to search a standard list of directories for a `liblibrary.a` library, the linker then uses that file as if it had been specified by name. The directories searched include several standard system directories plus any that you specify with `-L`; for more information on `-L`, see “Options for Searching Directories” on page 71. Normally the files found this way are library files, which are archive files whose members are object files. The linker handles an archive file by scanning through it for members defining symbols that have been referenced but not yet defined; if a file is found that is an ordinary object file, it is linked in the usual fashion.

The only difference between using an `-l` option and specifying a filename is that `-l` surrounds library with `lib` and `.a`, searching several directories.

`-nostartfiles`

Do not use the standard system startup files when linking. The standard system libraries are used normally, unless you use either `-nodefaultlibs` or `-nostdlib`.

`-nodefaultlibs`

Do not use the standard system libraries when linking. Only the libraries you specify will be passed to the linker. The standard startup files are used normally, unless you use `-nostartfiles`.

`-nostdlib`

Do not use the standard system startup files or libraries when linking. Neither the startup files nor the unspecified libraries will be passed to the linker. One of the standard libraries bypassed by `-nostdlib` and `-nodefaultlibs` is `libgcc.a`, a library of internal subroutines that GCC uses to overcome shortcomings of particular machines, or that GCC uses for special needs of some languages. See “Building and Installing a Cross-compiler” on page 319 and “`libgcc.a` and Cross-compilers” on page 321 for more discussion of `libgcc.a`.

In most cases, you need `libgcc.a` even when you want to avoid other standard libraries. In other words, when you specify `-nostdlib` or `-nodefaultlibs` you should usually specify `-lgcc` as well. This ensures that there will be no unresolved references to internal GCC library subroutines. For example, `__main` is used to ensure C++ constructors will be called; see “`collect2` and Cross-compiling” on page 325 for more discussion.

`-lobjc`

You need this special case of the `-l` option in order to link an Objective C program.

`-s`

Remove all symbol table and relocation information from the executable.

`-static`

On systems supporting dynamic linking, this prevents linking with the shared libraries. On other systems, this option has no effect.

`-shared`

Produce a shared object which can then be linked with other objects to form an executable. Not all systems support this option. You must also specify `-fpic` or `-fPIC` on some systems when you specify this option; for more information, see these options discussed with “Options Controlling Code Generation Conventions” on page 127.

`-symbolic`

Bind references to global symbols when building a shared object. Warn about any unresolved references, unless overridden by the `-xlinker -z -xlinker defs` link editor option. Only a few systems support this option.

`-wl, option`

Pass *option* as an option to the linker. If *option* contains commas, it is split into multiple options at the commas.

`-xlinker option`

Pass *option* as an option to the linker. You can use this to supply system-specific linker options which GCC does not know how to recognize. If you want to pass an option that takes an argument, you must use `-xlinker` twice, once for the option and once for the argument. To pass `-assert` definitions, you must use the command, `-xlinker -assert -xlinker definitions`; it does not work to write `-xlinker "-assert definitions"` because this passes the entire string as a single argument, which is not what the linker expects.

`-u symbol`

Make the symbol, *symbol*, undefined, to force linking of library modules to define it. You can use `-u` multiple times with different symbols to force loading of additional library modules.

12

Options for Searching Directories

The following options specify directories to search for header files, for libraries and for parts of the compiler.

-Bprefix

This option specifies where to find the executables, libraries, include files, and data files of the compiler itself. The compiler driver program runs one or more of the subprograms, `cpp`, `cc1`, `as` and `ld`. It tries *prefix* as a prefix for each program it tries to run, both with and without specifying the *machine/version/* (see “Options for Specifying Targets and Compiler Versions” on page 73).

For each subprogram to be run, the compiler driver first tries, if any, the `-Bprefix`. If that name is not found, or if `-B` was not specified, the driver tries two standard prefixes, which are `/usr/lib/gcc/` and `/usr/local/lib/gcc-lib/`. If neither of those results in a filename that is found, a search transpires for the unmodified program name using the directories that you specified in your `PATH` environment variable.

`-B` prefixes that effectively specify directory names also apply to libraries in the linker, because the compiler translates these options into `-L` options for the linker. They also apply to includes files in the preprocessor, because the compiler translates these options into `-isystem` options for the preprocessor. In this case, the compiler appends `include` to the prefix.

The run-time support file, `libgcc.a`, can also be searched for by using the `-B`

prefix, if needed. If it is not found there, the two standard prefixes discussed in the previous discussion are tried, and that is all. The file is left out of the link if it is not found by those means.

Another way to specify a prefix much like the `-B prefix`, is to use the environment variable, `GCC_EXEC_PREFIX`. See “The `offset-info` Option” on page 135.

`-Idir`

Add the directory, *dir*, to the head of the list of directories to be searched for header files. This can be used to override a system header file, substituting your own version, since these directories are searched before the system header file directories. If you use more than one `-I` option, the directories are scanned in left-to-right order; the standard system directories come after.

`-I-`

Any directories you specify with `-I` options before the `-I-` option are searched only for the case of `#include file`; they are not searched for `#include file..`. If additional directories are specified with `-I` options after the `-I-` options, these directories are searched for all `#include` directives. (Ordinarily all `-I` directories are used this way.) In addition, the `-I-` option inhibits the use of the current directory (where the current input file came from) as the first search directory for `#include file`. There is no way to override this `-I-` effect. With `-I`, you can specify searching the directory which was current when the compiler was invoked. That is not exactly the same as what the preprocessor does by default, but it is often satisfactory.

`-I-` does not inhibit the use of the standard system directories for header files. Thus, `-I-` and `-nostdinc` are independent.

`-Ldir`

Add directory, *dir*, to the list of directories to be searched for `-l`.

`-specs=file`

Process *file* after the compiler reads in the standard *specs* file, in order to override the defaults that the `gcc` driver program uses when determining what switches to pass to `cc1`, `cc1plus`, `as`, `ld`, and other GNU tools. More than one `-specs=file` can be specified on the command line, and they are processed in order, from left to right.

Options for Specifying Targets and Compiler Versions

By default, GCC compiles code for the same type of machine that you are using. However, it can also be installed as a cross-compiler, to compile for some other type of machine. In fact, several different configurations of GCC, for different target machines, can be installed side by side. Then you specify which one to use with the `-b` option. In addition, older and newer versions of GCC can be installed side by side. One of them (probably the newest) will be the default, but you may sometimes wish to use another version, using the following arguments.

`-b machine`

The argument, *machine*, specifies the target machine for compilation. This is useful when you have installed GCC as a cross-compiler. The value to use for *machine* is the same as was specified as the machine type when configuring GCC as a cross-compiler. For example, if a cross-compiler was configured with `configure i386v`, meaning to compile for an 80386 running System V, then you would specify `-b i386v` to run that cross compiler. When you do not specify `-b`, it normally means to compile for the same type of machine that you are using.

`-v version`

The argument, *version*, specifies which version of GCC to run. This is useful when multiple versions are installed. For example, *version* might be `2.0`, meaning to run GCC version 2.0. The default, *version*, when you do not specify `-v`, is the last version of GCC that you installed.

The `-b` and `-v` options actually work by controlling part of the filename used for the

executable files and libraries used for compilation. A given version of GCC, for a given target machine, is normally kept in a standard directory, similarly named as in the following example, where *machine* and *version* signify the machine and its compliant version that you need to specify.

```
/usr/local/lib/gcc-lib/machine/version.
```

Thus, sites can customize the effect of `-b` or `-v` either by changing the names of these directories or adding alternate names (or symbolic links). If in a directory,

```
/usr/local/lib/gcc-lib/, where the 80386 file is a link to the i386v file, then -b 80386 becomes an alias for -b i386v.
```

In one respect, the `-b` or `-v` do not completely change to a different compiler: the top-level driver for the compiler that you originally invoked continues to run and invoke the other executables (preprocessor, compiler per se, assembler and linker) which do the real work. However, since no real work is done in the driver program, it usually does not matter that the driver program in use is not the one for the specified target and version.

The only way that the driver program depends on the target machine is in the parsing and handling of special machine-specific options. However, this is controlled by a file which is found, along with the other executables, in the directory for the specified version and target machine. As a result, a single installed driver program adapts to any specified target machine and compiler version.

The driver program executable does control one significant thing, however: the default version and target machine. Therefore, you can install different instances of the driver program, compiled for different targets or versions, under different names.

For example, if the driver for version 2.0 is installed as `other-gcc` and that for version 2.1 is installed as `gcc`, then the command `gcc` will use version 2.1 by default, while `ogcc` will use 2.0 by default. However, you can choose either version with either command with the `-v` option.

Hardware Models and Configurations

The following documentation discusses the options for the GNU compiler, GCC, using the following hardware processors and their configurations. For specific options, look up the specific option in this documentation's index which starts on page 375.

- “AMD 29K Options” on page 76
- “ARC Options” on page 78
- “ARM/StrongARM Options” on page 78
- “ARM THUMB Options” on page 83
- “Clipper Options” on page 85
- “DEC Alpha Options” on page 85
- “Hitachi H8/300 Options” on page 89
- “Hitachi SH Options” on page 90
- “HPPA Options” on page 90
- “IBM RS/6000 and PowerPC Options” on page 92
- “IBM RT Options” on page 101
- “Intel x86 Options” on page 102
- “Intel 960 Options” on page 105
- “Matsushita MN10200 Options” on page 107

- “Matsushita MN10300/AM33 Options” on page 107
- “MIPS Options” on page 107
- “Mitsubishi D10V Options” on page 112
- “Mitsubishi M32R/D/X Options” on page 113
- “Motorola 68000 Options” on page 114
- “Motorola 88000 Options” on page 116
- “NEC V850 Options” on page 119
- “SPARC Options” on page 120
- “System V Options” on page 124

In “Options for Specifying Targets and Compiler Versions” on page 73, there is discussion of the standard option, `-b`, which chooses among different installed compilers for completely different target machines, such as Motorola 68K or Hitachi SH series.

Targets can have special options, starting with `-m`, to specify various hardware *models*, or configurations—for example, Motorola’s 68010 or 68020 processors, floating coprocessor or none. A single installed version of the compiler can compile for any model or configuration, according to the specified options.

Some configurations of the compiler also support additional special options, usually for compatibility with other compilers on the same platform.

AMD 29K Options

The following `-m` options are defined for the AMD Am29000 processor.

`-m29000`

Generate code for the Am29000. This is the default specification for the AMD 29K series.

`-m29050`

Generate code for the Am29050.

`-mbw`

`-mnbw`

With `-mbw`, generate code that assumes the system supports byte and halfword write operations. This is the default, not `-mnbw`. With `-mnbw`, generate code that assumes the systems does not support byte and halfword write operations. `-mnbw` implies `-mndw`.

`-mdw`

`-mndw`

With `-mdw`, generate code that assumes the `DW` bit is set, for example, that byte and halfword operations are directly supported by the hardware. This is the default, not `-mndw`. With `-mndw`, generate code that assumes the `DW` bit is not set.

`-mkernel-registers`

`-muser-registers`

Generate references to registers, `gr64` through `gr95`, instead of to the `gr96` through `gr127` registers. Use `-mkernel-registers` when compiling kernel code that wants a set of global registers disjoint from that used by user-mode code.

With `-muser-register`, use the normal set of global registers, `gr96` through `gr127`; this is the default, not `-mkernel-registers`.

IMPORTANT! With `-mkernel-registers`, register names in `-f` flags must use the normal user-mode names.

`-mlarge`

`-msmall`

`-mnormal`

With `-mlarge`, always use `calli` instructions; specify `-mlarge` if you expect a single file to compile into more than 256 KB of code.

With `-mnormal`, use the normal memory model; generate `call` instructions only when calling functions in the same file and `calli` instructions otherwise. This works if each file occupies less than 256 KB but allows the entire executable to be larger than 256 KB. This is the default, not `-msmall` or `-mlarge`.

With `-msmall`, use a small memory model that assumes that all function addresses are either within a single 256 KB segment or at an absolute address of less than 256k. This allows the `call` instruction to be used instead of a `const`, `consth`, `calli` sequence.

`-mno-reuse-arg-regs`

`-mreuse-arg-regs`

`-mno-reuse-arg-regs` tells the compiler to use (or, with `-mreuse-arg-regs`, not to use) only incoming argument registers for copying out arguments. This helps detect calling a function with fewer arguments than those with which it was declared.

`-mstack-check`

`-mno-stack-check`

Inserts (or, with `-mno-stack-check`, does not insert) a call to `__msp_check` after each stack adjustment. `-mstack-check` is often used for kernel code.

`-mstorem-bug`

`-mno-storem-bug`

`-mstorem-bug` handles (or, with `-mno-storem-bug`, does not handle) 29K processors which cannot handle the separation of a `mtsrin` insn and a `storem` instruction; this option works for most 29000 series chips to date, excepting the 29050 series).

ARC Options

The following options are for the ARC processor.

`-EB`

Compile code for big endian mode.

`-EL`

Compile code for little endian mode. This is the default.

`-mmangle-cpu`

Prepend the name of the CPU to all public symbol names. In multiple-processor systems, there are many ARC variants with different instruction and register set characteristics. This flag prevents code compiled for one CPU to be linked with code compiled for another CPU. No facility exists for handling variants that are almost identical. This is an all or nothing option.

`-mcpu=cpu`

Compile code for ARC variant, *cpu*. Which variants are supported depend on the configuration. All variants support the default, `-mcpu=base`.

`-mtext=text section`

`-mdata=data section`

`-mrodata=readonly data section`

Put functions, data, and read-only data in *text section*, *data section*, and *readonly data section*, respectively, by default. This can be overridden with the `section` attribute; see “Specifying Attributes of Variables” on page 169.

ARM/StrongARM Options

The following `-m` options are defined for Advanced RISC Machines (ARM) architectures.

`-mapcs`

`-mapcs-frame`

`-mno-apcs-frame`

`-mapcs` is a synonym for `-mapcs-frame`. With `-mapcs-frame`, generate a stack frame (or, with `-mno-apcs-frame`, do not generate a stack frame) to be compliant with the ARM Procedure Call Standard (APCS) for all functions, even if this is not strictly necessary for correct execution of the code.

`-mapcs-26`

Generate code for a processor running with a 26-bit program counter, and conforming to the function calling standards for the APCS 26-bit option. This option replaces the `-m2` and `-m3` options of previous releases of the compiler.

`-mapcs-32`

Generate code for a processor running with a 32-bit program counter, and

conforming to the function calling standards for the APCS 32-bit option. This option replaces the `-m6` option of previous releases of the compiler.

`-mapcs-stack-check`

`-mno-apcs-stack-check`

Generate code to check the amount of stack space available upon entry to every function (that actually uses some stack space). If there is insufficient space available then either the `__rt_stkovf_split_small` or the `__rt_stkovf_split_big` function will be called, depending upon the amount of stack space required. The run time system is required to provide these functions. The alternative, `-mno-apcs-stack-check`, produces smaller code.

`-mapcs-float`

`-mno-apcs-float`

Pass floating point arguments using the float point registers. This is one of the variants of the APCS. This option is recommended if the target hardware has a floating point unit or if a lot of floating point arithmetic is going to be performed by the code. The default is `-mno-apcs-float`, since integer only code is slightly increased in size if `-mapcs-float` is used.

`-mapcs-reentrant`

`-mno-apcs-reentrant`

Generate reentrant, position independent code with `-mapcs-reentrant`; this is the equivalent to specifying the `-fpic` option. `-mno-apcs-reentrant` is the default.

`-matpcs`

`-mno-atpcs`

Generate code that conforms to the ATPCS (ARM THUMB Procedure Call Standard). This ABI standard is different from the APCS standard and is incompatible in that small structures are returned from a function. For the APCS, a structure like the following example would be returned in memory, whereas, for the ATPCS, it is returned in a register. `-mno-atpcs` disables such conformance.

```
struct { char a; char b; }
```

`-mbuggy-return-in-memory`

`-mno-buggy-return-in-memory`

This options enables compatibility with older versions of the GNU compiler for ARM that incorrectly interpreted the APCS. These older versions would incorrectly return some small structures in memory instead of in a register. This only affects functions which return a structure containing a single float value like the following example shows.

```
struct { float a; }
```

Functions that return a structure with the following features have the same values.

- Less than or equal to 32 bits in size
- Contains exactly one field which is not a bitfield.
- This field is the first field in the structure.
- This field is itself a structure type.

- The contents of the following example's sub-structure are at least two non-bitfield types:

```
typedef struct { char a; char b }
sub_structure;
typedef struct { sub_structure s; }
structure;
```

By default, `-mno-buggy-return-in-memory` is not enabled.

`-msched-prolog`

`-mno-sched-prolog`

Prevent the reordering of instructions in the function prolog, or the merging of those instruction with the instructions in the function's body. This means that all functions will start with a recognisable set of instructions (or in fact one of a choice from a small set of different function prologues), and this information can be used to locate the start of functions inside an executable piece of code.

`-msched-prolog`, the default of these two options, disables the functionality.

`-mlittle-endian`

Generate code for a processor running in little-endian mode. This is the default for all standard configurations.

`-mbig-endian`

Generate code for a processor running in big-endian mode; the default is to compile code for a little-endian processor.

`-mwords-little-endian`

This option only applies when generating code for big-endian processors.

Generate code for a little-endian word order but a big-endian byte order. That is, a byte order of the form 32107654.

IMPORTANT! Use this option only you require compatibility with code for big-endian ARM processors generated by versions of the GNU compiler prior to version 2.8.

`-malignment-traps`

`-mno-alignment-traps`

`-malignment-traps` generates code that will not trap if the MMU has alignment traps enabled. On ARM architectures prior to ARMv4, there were no instructions to access half-word objects stored in memory. However, when reading from memory, a feature of the ARM architecture allows a word load to be used, even if the address is unaligned, and the processor core will rotate the data as it is being loaded. This option tells the compiler that such misaligned accesses will cause a MMU trap and that it should instead synthesise the access as a series of byte accesses. The compiler can still use word accesses to load half-word data if it knows that the address is aligned to a word boundary. `-malignment-traps` is ignored when compiling for ARM architecture 4 or later, since these processors have instructions to directly access half-word objects in memory.

`-mno-alignment-traps` generates code that assumes that the MMU will not trap unaligned accesses. This produces better code when the target instruction set does

not have half-word memory operations (implementations prior to ARMv4).

Do not use `-mno-alignment-traps` to access unaligned word objects, since the processor will only fetch one 32-bit aligned object from memory.

The default setting for most targets is `-mno-alignment-traps`, since this produces better code when there are no half-word memory instructions available.

`-mhard-float`

Generate output containing floating point instructions. This is the default.

`-msoft-float`

Generate output containing library calls for floating point.

WARNING! The requisite libraries are not available for all ARM targets. Normally the facilities of the machine's usual C compiler are used, but this cannot be done directly in cross-compilation. You must make your own arrangements to provide suitable library functions for cross-compilation.

`-msoft-float` changes the calling convention in the output file; therefore, it is only useful if you compile all of a program with this option. In particular, you need to compile `libgcc.a`, the library that comes with GCC, with `-msoft-float` in order for this to work.

`-mthumb-interwork`

`-mno-thumb-interwork`

`-mthumb-interwork` generates code supporting calls between the ARM and THUMB instruction sets. Without this option, the two instruction sets cannot be reliably used inside one program. The default is `-mno-thumb-interwork`, since slightly larger code is generated when `-mthumb-interwork` is specified.

`-mcpu=name`

This specifies the *name* of the target ARM processor. GCC uses this name to determine what kind of instructions it can use when generating assembly code.

Permissible names are: `arm2`, `arm250`, `arm3`, `arm6`, `arm60`, `arm600`, `arm610`, `arm620`, `arm7`, `arm7m`, `arm7d`, `arm7dm`, `arm7di`, `arm7dmi`, `arm70`, `arm700`, `arm700i`, `arm710`, `arm710c`, `arm7100`, `arm7500`, `arm7500fe`, `arm7tdmi`, `arm8`, `strongarm`, `strongarm110`.

`-march=architecture`

This specifies the *architecture* name of the target ARM architecture so the compiler can determine what kind of instructions it can use when generating assembly code. This option can be used in conjunction with or instead of the `-mcpu=` option. Permissible names are: `armv2`, `armv2a`, `armv3`, `armv3m`, `armv4`, `armv4t`.

`-mfpe=number`

This specifies the version (*number*) of the floating point emulation available on the target. Permissible values are 2 and 3.

`-mstructure-size-boundary=n`

The size of all structures and unions will be rounded up to a multiple of the

number of bits (n) set by this option . Permissible values are 8 and 32. The default value varies for different toolchains. For the COFF targeted toolchain the default value is 8. Specifying the larger number can produced faster, more efficient code, but can also increase the size of the program. The two values are potentially incompatible. Code compiled with one value cannot necessarily expect to work with code or libraries compiled with the other value, if they exchange information using structures or unions. Programmers are encouraged to use the 32 value as future versions of the toolchain may default to this value.

`-mbsd`

This option only applies to RISC iX. Emulate the native BSD-mode compiler. This is the default if `-ansi` is not specified.

`-mxopen`

This option only applies to RISC iX. Emulate the native X/Open-mode compiler.

`-mshort-load-bytes`

`-mno-short-load-bytes`

With `-mshort-load-bytes`, do not try to load half-words (for example, `short`) by loading a word from an unaligned address. For some targets the MMU is configured to trap unaligned loads; use this option to generate code that is safe in these environments. With `-mno-short-load-bytes`, use unaligned word loads to load half-words (for example., `short`). This option produces more efficient code, but the MMU is sometimes configured to trap these instructions.

`-mshort-load-words`

`-mno-short-load-words`

`-mshort-load-words` is a synonym for `-mno-short-load-bytes`.

`-mno-short-load-words` is a synonym for `-mshort-load-bytes`.

`-mabort-on-noreturn`

`-mnoabort-on-noreturn`

Generate a call to the function `abort` at the end of a `noreturn` function. It will be executed if the function tries to return.

`-mlong-calls`

`-mno-long-calls`

`-mlong-calls` tells the compiler to perform function calls by first loading the address of the function into a register and then performing a subroutine call on this register, necessary if the target function will lie outside of the 64 megabyte addressing range of the offset based version of subroutine call instruction.

Even if `-mlong-calls` is enabled, not all function calls will be turned into `long` calls. The heuristic is that static functions, functions having the `short-call` attribute, functions inside the scope of a `#pragma no_long_calls` directive, and functions whose definitions have already been compiled within the current compilation unit, will not be turned into `long` calls. The exception to this rule is that weak function definitions, functions with the `long-call` attribute or the `section` attribute, and functions within the scope of a `#pragma long_calls`

directive, will always be turned into `long` calls.

This feature is not enabled by default. Specifying `--no-long-calls` will restore the default behaviour, as will placing the function calls within the scope of a `#pragma long_calls_off` directive. These switches have no effect on how the compiler generates code to handle function calls using function pointers.

`-mno-symrename`

This option only applies to RISC iX. Do not run the assembler post-processor, `symrename`, after code has been assembled. Normally it is necessary to modify some of the standard symbols in preparation for linking with the RISC iX C library; this option suppresses this pass. The post-processor is never run when the compiler is built for cross-compilation.

`-mnop-fun-dllimport`

`-mnop-nop-fun-dllimport`

Disable the support for the `dllimport` attribute.

`-msingle-pic-base`

`-mno-single-pic-base`

Treat the register used for PIC addressing as read-only, rather than loading it in the prologue for each function. The run-time system is responsible for initialising this register with an appropriate value before execution begins.

`-mpic-register=reg`

`-mno-pic-register=reg`

Specify the register (`reg`) to be used for PIC addressing. The default is `R10` unless stack-checking is enabled, when `R9` is used.

`-msa1110-fix`

Due to a bug in the StrongARM 1110 hardware the `LDRSH` and `LDRSB` instructions can fail to execute correctly under certain circumstances. Specifying `-msa1110-fix` enables a workaround in the compiler which causes each and every `LDRSH` and `LDRSB` instruction to be issued twice. The second instruction is guaranteed to work.

ARM THUMB Options

The following `m` options are available for ARM THUMB architectures.

`-mtpcs-frame`

`-mno-tpcs-frame`

`-mtpcs-frame` generates a stack frame that is compliant with the THUMB

Procedure Call Standard (TPCS) for all non-*leaf* functions (functions that do not call any other functions). Its alternative, the default, is `-mno-apcs-frame`.

`-mtpcs-leaf-frame`

`-mno-tpcs-leaf-frame`

`-mtpcs-leaf-frame` generates a stack frame that is compliant with the THUMB

- Procedure Call Standard (TPCS) for all leaf functions (functions that do not call any other functions). Its alternative, the default, is `-mno-apcs-leaf-frame`.
- `-mlittle-endian`
Generate code for a processor running in little-endian mode. This is the default for all standard configurations.
- `-mbig-endian`
Generate code for a processor running in big-endian mode.
- `-mthumb-interwork`
`-mno-thumb-interwork`
`-mthumb-interwork` generates code that allows calling between the THUMB and ARM instruction sets. Without `-mthumb-interwork`, the two instruction sets cannot be reliably used inside one program. Its alternative, the default, is `-mno-thumb-interwork`, since slightly smaller code is generated with it.
- `-matpcs`
`-mno-atpcs`
Generate code that conforms to the ATPCS (ARM THUMB Procedure Call Standard). This ABI standard is different from the APCS standard and is incompatible in that small structures are returned from a function. For the APCS, a structure like the following example would be returned in memory, whereas, for the ATPCS, it is returned in a register. `-mno-atpcs` disables such conformance.
- ```
struct { char a; char b; }
```
- `-mstructure-size-boundary=n`  
The size of all structures and unions will be rounded up to a multiple of the number of bits (*n*) set by this option. Permissible values are 8 and 32. The default value varies for different toolchains. For the COFF targeted toolchain, the default value is 8. Specifying the larger number can produce faster, more efficient code, but can also increase the size of the program. The two values are potentially incompatible. Code compiled with one value cannot necessarily expect to work with code or libraries compiled with the other value, if they exchange information using structures or unions. Use the 32 value as future versions of the toolchain may default to this value.
- `-mnop-fun-dllimport`  
`-mnop-nop-fun-dllimport`  
Disable the support for the *dllimport* attribute.
- `-mcallee-super-interworking`  
`-mno-callee-super-interworking`  
`-mcallee-super-interworking` gives all externally visible functions in the file that is being compiled an ARM instruction set header so that it can switch to THUMB mode before executing the rest of the function. This allows these functions to be called from non-interworking code.
- `-mcaller-super-interworking`  
`-mno-caller-super-interworking`  
`-mcaller-super-interworking` allows calls using function pointers (including virtual functions) to execute correctly regardless of whether or not the target code

has been compiled for interworking. There is a small overhead in the cost of executing a function pointer if this option is enabled.

`-msingle-pic-base`

`-mno-single-pic-base`

Treat the register used for PIC addressing as read-only, rather than loading it in the prologue for each function. The run-time system is responsible for initialising this register with an appropriate value before execution begins.

`-mpic-register=reg`

`-mno-pic-register=reg`

Specify the register (*reg*) to be used for PIC addressing. The default is `R10` unless stack-checking is enabled, when `R9` is used.

## Clipper Options

The following `-m` options are defined for the Clipper processor's implementations.

`-mc300`

Produce code for a C300 Clipper processor. This is the default.

`-mc400`

Produce code for a C400 Clipper processor; for example, use floating point registers, `f8` through `f15`.

## DEC Alpha Options

The following `-m` options are defined for the DEC Alpha family of processor's implementations.

`-mfp-reg`

`-mno-fp-regs`

Generate code that uses (does not use) the floating point register set.

`-mno-fp-regs` implies `-msoft-float`. If the floating point register set is not used, floating point operands are passed in integer registers as if they were integers and floating point results are passed in `$0` instead of `$f0`.

This is a non-standard calling sequence, so any function with a floating point argument or return value called by code compiled with `-mno-fp-regs` must also be compiled with that option. A typical use of this option is building a kernel that does not use, and hence need not save and restore, any floating point registers.

`-mno-soft-float`

`-msoft-float`

Use (do not use) the hardware floating point instructions for floating point operations. When `-msoft-float` is specified, functions in `libgcc1.c` will be used to perform floating point operations. Unless they are replaced by routines that emulate the floating point operations, or compiled in such a way as to call such

emulation routines, these routines will issue floating point operations. If you are compiling for an Alpha without floating-point operations, you must ensure that the library is built so as not to call them.

**IMPORTANT!** Alpha implementations without floating point operations are required to have floating-point registers.

`-malpha-as`

`-mgas`

Select whether to generate code to be assembled by the vendor-supplied assembler (using the option, `-malpha-as`) or by the GNU assembler, using `-mgas`.

`-mieee`

The Alpha architecture implements floating point hardware optimized for maximum performance. It is mostly compliant with the IEEE floating point standard. However, for full compliance, software assistance is required. This option generates code fully IEEE compliant code *except* that the `inexact` flag is not maintained (compare following description for `-mieee-with-inexact`). If this option is turned on, the CPP macro, `_IEEE_FP`, is defined during compilation.

The option is a shorthand for `-D_IEEE_FP -D_IEEE_FP_INEXACT` plus `-mieee-conformant`, and `-mfp-trap-mode=sui`, and `-mtrap-precision=i`.

The resulting code is less efficient but is able to correctly support denormalized numbers and exceptional IEEE values such as not-a-number and plus/minus infinity.

Other Alpha compilers call this option `-ieee_with_no_inexact`.

`-mieee-with-inexact`

This is like `-mieee` except the generated code also maintains the IEEE `inexact` flag. Turning on this option causes the generated code to implement fully-compliant IEEE math. The option is a shorthand; it signifies the declaration of `-D_IEEE_FP -D_IEEE_FP_INEXACT` plus `-mieee-conformant`, and `-mfp-trap-mode=sui`, and `-mtrap-precision=i`. On some Alpha implementations the resulting code may execute significantly slower than the code generated by default. Since there is very little code that depends on the `inexact` flag, you should normally not specify this option. Other Alpha compilers call this option `-ieee_with_inexact`.

`-mieee-conformant`

This option marks the generated code as IEEE conformant. You must not use this option unless you also specify `-mtrap-precision=i` and either `-mfp-trap-mode=su` or `-mfp-trap-mode=sui`. Its only effect is to emit the `.eflag 48` line in the function prologue of the generated assembly file. Under DEC UNIX, this has the effect that IEEE-conformant math library routines will be linked in.

`-mfp-trap-mode=trap mode`

This option controls what floating point related traps are enabled. Other Alpha

compilers call this option `-fptm trap mode`, where *trap mode* is the definition of what can be set by one of the following four values.

- `n`  
This is the default (normal) setting. The only traps that are enabled are the ones that cannot be disabled in software (such as division by zero trap).
- `u`  
In addition to the traps enabled by `n`, underflow traps are enabled as well.
- `su`  
Like `u`, but the instructions are marked to be safe for software completion (see Alpha architecture manuals for details).
- `sui`  
Like `su`, but inexact traps are enabled as well.

`-mfp-rounding-mode=rounding mode`

Selects the IEEE rounding mode. Other Alpha compilers call this option `-fprm rounding mode`. The rounding mode can be one of the following four values.

- `n`  
Normal IEEE rounding mode. Floating point numbers are rounded towards the nearest machine number or towards the even machine number in case of a tie.
- `m`  
Round towards minus infinity.
- `c`  
Chopped rounding mode. Floating point numbers are rounded towards zero.
- `d`  
Dynamic rounding mode. A field in the floating point control register (`fpcr`; see Alpha architecture reference manuals for details) controls the rounding mode in effect. The C library initializes this register for rounding towards plus infinity. Thus, unless your program modifies the `fpcr`, `d` corresponds to round towards plus infinity.

`-mtrap-precision=trap precision`

In the Alpha architecture, floating point traps are imprecise. This means without software assistance it is impossible to recover from a floating trap and program execution normally needs to be terminated. GCC can generate code that can assist operating system trap handlers in determining the exact location that caused a floating point trap. Depending on the requirements of an application, different levels of precision can be selected, such as with the following options.

- `p`  
Program precision. This option is the default and means a trap handler can only identify which program caused a floating point exception.
- `f`

Function precision. The trap handler can determine the function that caused a floating point exception.

- `i`  
Instruction precision. The trap handler can determine the exact instruction that caused a floating point exception.

Other Alpha compilers provide the equivalent options, `-scope_safe` and `-resumption_safe`.

`-mbuild-constants`

Normally GCC examines a 32- or 64-bit integer constant to see if it can construct it from smaller constants in two or three instructions. If it cannot, it will output the constant as a literal and generate code to load it from the data segment at runtime.

Use this option to require GCC to construct all integer constants using code, even if it takes more instructions (the maximum is six).

You would typically use this option to build a shared library dynamic loader. If itself a shared library, it must relocate itself in memory before it can find the variables and constants in its own data segment.

`-mcpu=cpu type`

Set the instruction set, register set, and instruction scheduling parameters for machine type, *cpu type*. You can specify either the EV style name or the corresponding chip number. GCC supports scheduling parameters for the EV4 and EV5 family of processors and will choose the default values for the instruction set from the processor you specify. If you do not specify a processor type, GCC will default to the processor on which the compiler was built. Supported replacements for *cpu type* use the following values.

- `ev4`  
`21064`  
Schedules as an EV4 and has no instruction set extensions.
- `ev5`  
`21164`  
Schedules as an EV5 and has no instruction set extensions.
- `ev56`  
`21164a`  
Schedules as an EV5 and supports the BWX extension.
- `pca56`  
`21164pc`  
`21164PC`  
Schedules as an EV4 and supports the BWX and MAX extensions.
- `ev6`  
`21264`  
Schedules as an EV5 (until Digital releases the scheduling parameters for the EV6) and supports the BWX, CIX, and MAX extensions.

-mbwx  
 -mno-bwx  
 -mcix  
 -mno-cix  
 -mmax  
 -mno-max

Indicate whether GCC should generate code to use the optional BWX, CIX, and MAX instruction sets. The default is to use the instruction sets supported by the CPU type specified, using the `-mcpu=option` or that of the CPU on which GCC was built if none was specified.

-mmemory-latency=*time*

Sets the latency the scheduler should assume for typical memory references as seen by the application.

This number is highly dependant on the memory access patterns used by the application and the size of the external cache on the machine. Valid options for *time* use the following values.

- number  
A decimal number representing clock cycles.

- L1  
L2  
L3

main

The compiler contains estimates of the number of clock cycles for typical EV4 and EV5 hardware for the Level 1, 2 and 3 caches (also called *Dcache*, *Scache*, and *Bcache*), as well as to main memory.

**IMPORTANT!** L3 is only valid for EV5.

## Hitachi H8/300 Options

The following `-m` options are defined for the H8/300 family of processors' implementations.

-mrelax

Shorten some address references at link time, when possible; this option uses the linker option, `-relax`. See “ld and the H8/300 Processors” on page 63 in *Using ld of GNUPro Development Tools* for a fuller description.

-mh

Generate code for the H8/300H.

-ms

Generate code for the H8/S.

-mint32

Make `int` data 32 bits by default.

`-malign-300`

On the H8/300H, use the same alignment rules as for the H8/300. The default for the H8/300H is to align longs and floats on 4 byte boundaries. `-malign-300` causes them to be aligned on 2 byte boundaries. This option has no effect on the H8/300.

## Hitachi SH Options

The following `-m` options are defined for the Hitachi SH processors family implementations.

`-m1`

Generate code for the SH1.

`-m2`

Generate code for the SH2.

`-m3`

Generate code for the SH3.

`-m3e`

Generate code for the SH3e.

`-m4`

Generate code for the SH4.

`-mb`

Compile code for the processor in big endian mode.

`-ml`

Compile code for the processor in little endian mode.

`-mdalign`

Align doubles at 64 bit boundaries. This changes the calling conventions, and some functions from the standard C library will not work unless you recompile it first with `-mdalign`.

`-mrelax`

Shorten some addresses at link time, when possible; uses the linker option, `-relax`.

## HPPA Options

The following `-m` options are defined for the HPPA family.

`-march=architecture`

This specifies the *architecture* name of the target HPPA architecture so the compiler can determine what kind of instructions it can use when generating assembly code. This option can be used in conjunction with or instead of the `-mcpu=` option. Permissible names are `1.0` for PA 1.0, `1.1` for PA 1.1, and `2.0` for PA 2.0 processors. See `/usr/lib/sched.models` on an HP-UX system to

determine the proper architecture option for your machine. Code compiled for lower numbered architectures will run on higher numbered architectures, but higher numbered architectures will run on lower numbered architectures.

`-mbig-switch`

Generate code suitable for big switch tables. Use this option only if the assembler/linker complain about out of range branches within a switch table.

`-mdisable-fpregs`

Prevent floating point registers from being used in any manner. This is necessary for compiling kernels which perform lazy context switching of floating point registers. If you use this option and attempt to perform floating point operations, the compiler will abort.

`-mdisable-indexing`

Prevent the compiler from using indexing address modes. This avoids some rather obscure problems when compiling MIG generated code under MACH.

`-mfast-indirect-calls`

Generate code that assumes calls never cross space boundaries. This allows GCC to emit code which performs faster indirect calls.

This option will not work in the presence of shared libraries or nested functions.

`-mgas`

Enable the use of assembler directives only GAS understands.

`-mjump-in-delay`

Fill delay slots of function calls with unconditional jump instructions by modifying the return pointer for the function call to be the target of the conditional jump.

`-mlong-load-store`

Generate 3-instruction load and store sequences as some-times required by the HP/UX 10 linker. This is equivalent to the `+k` option to the HP compilers.

`-mno-big-switch`

Disables `-mbig-switch` functionality.

`-mno-disable-fpregs`

Disables `-mdisable-fpregs` functionality.

`-mno-disable-indexing`

Disables `-mdisable-indexing` functionality.

`-mno-gas`

Disables `-mgas` functionality.

`-mno-jump-in-delay`

Disables `-mjump-in-delay` functionality.

`-mno-long-load-store`

Disables `-mlong-load-store` functionality.

`-mno-portable-runtime`

Disables `-mportable-runtime` functionality.

`-mno-soft-float`

Disables `-msoft-float` functionality.

`-mno-space-regs`

Generate code that assumes the target has no space registers. This allows GCC to generate faster indirect calls and use unscaled index address modes. Such code is suitable for level 0 PA systems and kernels.

`-msoft-float`

Generate output containing library calls for floating point.

**WARNING!** The requisite libraries are not available for all HPPA targets. Normally the facilities of the machine's usual C compiler are used, but this cannot be done directly in cross-compilation. You must make your own arrangements to provide suitable library functions for cross-compilation. The embedded target `hppa1.1-*-pro` does provide software floating point support.

`-msoft-float` changes the calling convention in the output file; therefore, it is only useful if you compile *all* of a program with this option. In particular, you need to compile `libgcc.a`, the library that comes with GCC, with `-msoft-float` in order for this to work.

`-mpa-risc-1-0`

Generate code for a PA 1.0 processor.

`-mpa-risc-1-1`

Generate code for a PA 1.1 processor.

`-mpa-risc-2-0`

Generate code for a PA 2.0 processor.

`-mportable-runtime`

Use the portable calling conventions proposed by HP for ELF systems.

`-mschedule=cpu type`

Schedule code according to the constraints for the machine type (signified by the *cpu type*). The choices for *cpu type* are 700 for 7n0 machines, 7100 for 7n5 machines, and 7100 for 7n2 machines. 7100 is the default for *cpu type*.

**IMPORTANT!** The 7100LC scheduling information is incomplete and using 7100LC often leads to bad schedules. It's probably best to use 7100 instead of 7100LC for the 7n2 machines.

`-mspace-regs`

Disables `-mno-space-regs` functionality.

## IBM RS/6000 and PowerPC Options

These `-m` options are defined for the IBM RS/6000 and PowerPC processor's implementations.

`-mcpu=cpu_type`

Set architecture type, register usage, choice of mnemonics, and instruction

scheduling parameters for machine type *cpu\_type*. Supported values for *cpu\_type* are *rs6000*, *rios1*, *rios2*, *rsc*, *601*, *602*, *603*, *603e*, *604*, *604e*, *620*, *power*, *power2*, *powerpc*, *403*, *505*, *801*, *821*, *823*, *860* and *common*.

The `-mcpu=power`, `-mcpu=power2`, and `-mcpu=powerpc` specify generic POWER, POWER2 and pure PowerPC (that is, not MPC601) architecture machine types, with an appropriate, generic processor model assumed for scheduling purposes.

Specifying `-mcpu=rios1`, `-mcpu=rios2`, `-mcpu=rsc`, `-mcpu=power`, or `-mcpu=power2` enables the `-mpower` option and disables the `-mpowerpc` option; `-mcpu=601` enables both the `-mpower` and `-mpowerpc` options; `-mcpu=602`, `-mcpu=603`, `-mcpu=603e`, `-mcpu=604`, `-mcpu=620`; `-mcpu=403`, `-mcpu=505`, `-mcpu=821`, `-mcpu=860` and `-mcpu=powerpc` enable the `-mpowerpc` option and disable the `-mpower` option; `-mcpu=common` disables both the `-mpower` and `-mpowerpc` options.

AIX versions 4 or greater selects `-mcpu=common` by default, so that code will operate on all members of the RS/6000 and PowerPC families. In that case, GCC will use only the instructions in the common subset of both architectures plus some special AIX common-mode calls, and will not use the MQ register. GCC assumes a generic processor model for scheduling purposes.

Specifying `-mcpu=rios1`, `-mcpu=rios2`, `-mcpu=rsc`, `-mcpu=power`, or `-mcpu=power2` also disables the `-mnew-mnemonics` option.

Specifying `-mcpu=601`, `-mcpu=602`, `-mcpu=603`, `-mcpu=603e`, `-mcpu=604`, `-mcpu=620`, `-mcpu=403`, or `-mcpu=powerpc` also enables the `-mnew-mnemonics` option.

Specifying `-mcpu=403`, `-mcpu=821`, or `-mcpu=860` also enables the `-msoft-float` option.

`-mtune=cpu_type`

Set the instruction scheduling parameters for machine type, *cpu\_type*, but do not set the architecture type, register usage, choice of mnemonics like `-mcpu=cpu_type` would. The same values for *cpu\_type* are used for `-mtune=cpu_type` as for `-mcpu=cpu_type`. The `-mtune=cpu_type` option overrides the `-mcpu=cpu_type` option in terms of instruction scheduling parameters.

- mpower
- mno-power
- mpower2
- mno-power2
- mpowerpc
- mpowerpc64
- mno-powerpc
- mpowerpc-gpopt
- mno-powerpc-gpopt
- mpowerpc-gfxopt
- mno-powerpc-gfxopt

Provides support for the two related instruction set architectures for the RS/6000 and PowerPC. The POWER instruction set includes those instructions supported by the `rios` chip set used in the original RS/6000 systems. The PowerPC instruction set is the architecture of the Motorola MPC5<sub>xx</sub>, MPC6<sub>xx</sub>, MCP8<sub>xx</sub> and the IBM 4<sub>xx</sub> microprocessors. The PowerPC architecture defines 64-bit instructions, but they are not supported by any current processors. Neither architecture is a subset of the other; however, there is a large common subset of instructions supported by both. An MQ register is included in processors supporting the POWER architecture.

Use these options to specify which instructions are available on the processor you are using. The default value of these options is determined when configuring GCC. Specifying the `-mcpu=cpu_type` overrides the specification of these options. A general recommendation is that you use the `-mcpu=cpu_type` option rather than any of these options.

The `-mpower` option allows GCC to generate instructions that are found only in the POWER architecture, using the MQ register. Specifying `-mpower2` implies `-mpower` and also allows GCC to generate instructions that are present in the POWER2 architecture but not the original POWER architecture.

The `-mpowerpc` option allows GCC to generate instructions that are found only in the 32-bit subset of the PowerPC architecture. Specifying `-mpowerpc-gpopt` implies `-mpowerpc` and also allows GCC to use the optional PowerPC architecture instructions in the General Purpose group, including floating point square root. Specifying `-mpowerpc-gfxopt` implies `-mpowerpc` and also allows GCC to use the optional PowerPC architecture instructions in the Graphics group, including floating point select. If you specify both `-mno-power` and `-mno-powerpc` options, GCC will use only the instructions in the common subset of both architectures plus some special AIX common-mode calls, and will not use the MQ register. Specifying both `-mpower` and `-mpowerpc` permits GCC to use any instruction from either architecture and to allow use of the MQ register; specify this for the Motorola MPC601.

- mnew-mnemonics
- mold-mnemonics

Selects mnemonics to use in the generated assembler code. `-mnew-mnemonics`

requests output that uses the assembler mnemonics defined for the PowerPC architecture, while `-mold-mnemonics` requests the assembler mnemonics defined for the POWER architecture. Instructions defined in only one architecture have only one mnemonic; GCC uses that mnemonic irrespective of which of these options is specified.

PowerPC assemblers support both the old and new mnemonics, as will later POWER assemblers. Current POWER assemblers only support the old mnemonics. Specify `-mnew-mnemonics` if you have an assembler that supports them, otherwise specify `-mold-mnemonics`.

The default value of these options depends on how GCC was configured.

Specifying `-mcpu=cpu_type` sometimes overrides the value of these option.

Unless you are building a cross-compiler, you should normally not specify either `-mnew-mnemonics` or `-mold-mnemonics`, but should instead accept the default.

`-mfull-toc`

`-mminimal-toc`

`-mno-sum-in-toc`

`-mno-fp-in-toc`

Modifies the generation of the TOC (Table Of Contents), which is created for every executable file. The `-mfull-toc` option is selected by default. In that case, GCC will allocate at least one TOC entry for each unique non-automatic variable reference in your program. GCC will also place floating point constants in the TOC. However, only 16,384 entries are available in the TOC.

If you receive a linker error message that saying you have overflowed the available TOC space, you can reduce the amount of TOC space used with the `-mno-fp-in-toc` and `-mno-sum-in-toc` options.

`-mno-fp-in-toc` prevents GCC from putting floating point constants in the TOC and `-mno-sum-in-toc` forces GCC to generate code to calculate the sum of an address and a constant at run-time instead of putting that sum into the TOC. You may specify one or both of these options. Each causes GCC to produce very slightly slower and larger code at the expense of conserving TOC space.

If you still run out of space in the TOC even when you specify both of these options, specify `-mminimal-toc` instead. This option causes GCC to make only one TOC entry for every file. When you specify this option, GCC will produce code that is slower and larger but which uses extremely little TOC space. Use this option only on files that contain less frequently executed code.

`-maix64`

`-maix32`

Enable 64-bit AIX ABI and calling convention: 64-bit pointers, 64-bit `long` type, and the infrastructure needed to support them. Specifying `-maix64` implies `-mpowerpc64` and `-mpowerpc`, while `-maix32` disables the 64-bit ABI and implies `-mno-powerpc64`; the default is `-maix32`.

`-mxl-call`

`-mno-xl-call`

On AIX, `-mxl-call` passes floating point arguments to prototyped functions beyond the register save area (RSA) on the stack in addition to floating point register arguments. The AIX calling convention was extended but not initially documented to handle an obscure K&R C case of calling a function that takes the address of its arguments with fewer arguments than declared. AIX XL compilers assume that floating point arguments which do not fit in the RSA are on the stack when they compile a subroutine without optimization. Because always storing floating point arguments on the stack is inefficient and rarely needed, this option is not enabled by default and only is necessary when calling subroutines compiled by AIX XL compilers without optimization. `-mno-xl-call` disables `-mxl-call`.

`-mthreads`

Supports *AIX threads*. Link an application written to use *pthread*s with special libraries and startup code to enable the application to run.

`-mpe`

Supports *IBM RS/6000 SP Parallel Environment* (PE). Link an application written to use message passing with special startup code to enable the application to run. PE should be installed in the standard location (`/usr/lpp/ppc/ppe.poe/`), or the `specs` file must be overridden with the `-specs=` option to specify the appropriate directory location. The Parallel Environment does not support threads, so the `-mpe` option and the `-mthreads` option are incompatible.

`-msoft-float`

`-mhard-float`

`-msoft-float` generates code that does not use the floating point register set. Software floating point emulation is provided if you use the `-msoft-float` option, and pass the option to GCC when linking. `-mhard-float` does use the floating point register set.

`-mmultiple`

`-mno-multiple`

`-mmultiple` generates code that uses the load multiple word instructions and the store multiple word instructions. These instructions are generated by default on POWER systems, and not generated on PowerPC systems. Do not use `-mmultiple` on little endian PowerPC systems, since those instructions do not work when the processor is in little endian mode. `-mno-multiple` disables use of the load and store multiple word instructions.

`-mstring`

`-mno-string`

`-mstring` generates code that uses the load string instructions and the store string word instructions, allows saving multiple registers and small block moves. These instructions are generated by default on POWER systems, and not generated on PowerPC systems. `-mno-string` disables what `-mstring` provides.

**WARNING!** Do not use `-mstring` on little endian PowerPC systems, since the instructions

do not work when the processor is in little endian mode.

`-mupdate`

`-mno-update`

`-mupdate` generates code that uses the load or store instructions that update the base register to the address of the calculated memory location. These instructions are generated by default. If you use `-mno-update`, there is a small window between the time that the stack pointer is updated and the address of the previous frame is stored, which means code that walks the stack frame across interrupts or signals may get corrupted data.

`-mfused-madd`

`-mno-fused-madd`

`-mfused-madd` generates code that uses the floating point multiply and accumulate instructions. These instructions are generated by default if hardware floating is used. `-mno-fused-madd` disables what `-mfused-madd` provides.

`-mno-bit-align`

`-mbit-align`

For System V.4 and embedded PowerPC systems, `-mno-bit-align` disables forcing structures and unions that contain bit fields, so that they align to the base type of the bit field. For example, by default a structure containing nothing but 8 unsigned bitfields of length 1 would be aligned to a 4 byte boundary and have a size of 4 bytes. By using `-mno-bit-align`, the structure would be aligned to a 1 byte boundary and be one byte in size. `-mbit-align` enables what `-mno-bit-align` disables.

`-mno-strict-align`

`-mstrict-align`

For System V.4 and embedded PowerPC systems, `-mno-strict-align` disables assuming that unaligned memory references will be handled by the system. `-mstrict-align` enables what `-mno-strict-align` disables.

`-mrelocatable`

`-mno-relocatable`

For embedded PowerPC systems, `-mrelocatable` generates code that allows (does not allow) the program to be relocated to a different address at runtime. If you use `-mrelocatable` on any module, all objects linked together must be compiled with `-mrelocatable` or `-mrelocatable-lib`. `-mno-relocatable` does not allow the program to be relocated to a different address at runtime.

`-mrelocatable-lib`

`-mno-relocatable-lib`

For embedded PowerPC systems, `-mrelocatable-lib` generates code that allows the program to be relocated to a different address at runtime. Modules compiled with `-mrelocatable-lib` can be linked with either modules compiled without `-mrelocatable` and `-mrelocatable-lib` or with modules compiled with the `-mrelocatable` options. `-mno-relocatable-lib` does not allow the program to be relocated to a different address at runtime.

`-mno-toc`

`-mtoc`

For System V.4 and embedded PowerPC systems, `-mno-toc` provides for not assuming that register 2 contains a pointer to a global area that points to the addresses used in the program. `-mtoc` disables providing for assuming that register 2 contains a pointer to a global area that points to the addresses used in the program.

`-mlittle`

`-mlittle-endian`

For System V.4 and embedded PowerPC systems, `-mlittle` allows for compiling code for the processor in little endian mode. `-mlittle-endian` works the same as `-mlittle`.

`-mbig`

`-mbig-endian`

For System V.4 and embedded PowerPC systems, `-mbig` allows for compiling code for the processor in big endian mode. `-mbig-endian` works the same as `-mbig`.

`-mcall-aix`

For System V.4 and embedded PowerPC systems, allows for compiling code using calling conventions that are similar to those used on AIX. This is the default if you configured GCC using `powerpc-*-eabiaix`.

`-mcall-sysv`

For System V.4 and embedded PowerPC systems, allows for compiling code using calling conventions that adheres to the March 1995 draft of the System V ABI, PowerPC processor supplement. This is the default unless you configured GCC using `powerpc-*-eabiaix`.

`-mcall-sysv-eabi`

Specifies both `-mcall-sysv` and `-meabi` options.

`-mcall-sysv-noeabi`

Specifies both `-mcall-sysv` and `-mnoeabi` options.

`-mcall-solaris`

For System V.4 and embedded PowerPC systems, allows for compiling code for the Solaris operating system.

`-mcall-linux`

For System V.4 and embedded PowerPC systems, allows for compiling code for Linux operating systems.

`-mprototype`

`-mno-prototype`

For System V.4 and embedded PowerPC systems, provides for assuming that all calls to variable argument functions are properly prototyped. Otherwise, the compiler must insert an instruction before every non-prototyped call to set or clear bit 6 of the condition code register (CR), indicating whether floating point values were passed in the floating point registers (in case the function takes a variable

arguments). With `-mprototype`, only calls to prototyped variable argument functions will set or clear the bit.

`-msim`

For embedded PowerPC systems, provides for assuming that the startup module is called `sim-crt0.o` and the standard C libraries are `libsim.a` and `libc.a`. This is default for `powerpc-*-eabisim` configurations.

`-mmvme`

For embedded PowerPC systems, provides for assuming that the startup module is called `mvme-crt0.o` and the standard C libraries are `libmvme.a` and `libc.a`.

`-mads`

For embedded PowerPC systems, provides for assuming that the startup module is called `crt0.o` and the standard C libraries are `libads.a` and `libc.a`.

`-myellowknife`

For embedded PowerPC systems, provides for assuming that the startup module is called `crt0.o` and `libyk.a` and `libc.a` are the standard C libraries.

`-memb`

For embedded PowerPC systems, sets the `PPC_EMB` bit in the ELF options header to indicate that EABI extended relocations are used.

`-mmpc860c0=num`

Only applicable to MPC860 chips when producing ELF executables with the GNU linker; does not cause any changes to the `.o` files, does cause the linker to perform a check for problematic conditional branches, and implements alternatives. The problem is that some chips may treat the target instruction as a no-op, given the following conditions:

- The processor is an MPC860, version C0 or earlier.
- A forward conditional branch is executed.
- The branch is predicted as not taken.
- The branch is taken.
- The branch is located in the last 5 words of a page.
- The branch target is located on a subsequent page.

The optional argument, `=num`, is the number of words that are checked at the end of each text page. It may be any value from 1 to 10, defaulting to 5.

`-meabi`

`-mno-eabi`

For System V.4 and embedded PowerPC systems, `-meabi` provides for adhering to the Embedded Applications Binary Interface (EABI) as a set of modifications to the System V.4 specifications. Selecting `-meabi` means that the stack is aligned to an 8 byte boundary, a function, `__eabi`, is called from `main` to set up the EABI environment, and the `-msdata` option can use both `r2` and `r13` to point to two separate small data areas. Selecting `-mno-eabi` means that the stack is aligned to a 16 byte boundary, without calling an initialization function from `main`; `-msdata`

will only use `r13` to point to a single small data area. `-meabi` is on by default if you configured GCC using one of the `powerpc*-*-eabi*` options.

`-msdata=eabi`

For System V.4 and embedded PowerPC systems, puts small initialized const global and static data in the `.sdata2` section, which is pointed to by register, `r2`. Puts small initialized non-const global and static data in the `.sdata` section, which is pointed to by register, `r13`. Puts small uninitialized global and static data in the `.sbss` section, which is adjacent to the `.sdata` section. `-msdata=eabi` is incompatible with `-mrelocatable`. `-msdata=eabi` also sets the `-memb` option.

`-msdata=default`

`-msdata`

For System V.4 and embedded PowerPC systems, if `-meabi` is used, `-msdata=default` or `-msdata` compile code the same as `-msdata=eabi`; otherwise, compile code the same as `-msdata=sysv`.

`-msdata=sysv`

For System V.4 and embedded PowerPC systems, puts small global and static data in the `.sdata` section, which is pointed to by register `r13`. Puts small uninitialized global and static data in the `.sbss` section, which is adjacent to the `.sdata` section. `-msdata=sysv` is incompatible with `-mrelocatable`.

`-msdata-data`

For System V.4 and embedded PowerPC systems, puts small global and static data in the `.sdata` section. Puts small uninitialized global and static data in the `.sbss` section. Do not use the `r13` register to address small data. This is the default behavior unless other `-msdata` options are used.

`-msdata=none`

`-mno-sdata`

For embedded PowerPC systems, `-mno-sdata` or `-msdata=none` puts all initialized global and static data in the `.data` section, and all uninitialized data in the `.bss` section.

`-G num`

For embedded PowerPC systems, puts global and static items less than or equal to `num` bytes into the small data or bss sections instead of the normal data or bss section. By default, `num` is 8. The `-G num` switch is also passed to the linker. All modules should be compiled with the same `-G num` value.

`-mregnames`

`-mno-regnames`

On System V.4 and embedded PowerPC systems, `-mregnames` does (and `-mno-regnames` does not) emit register names in the assembly language output using symbolic forms.

`-mno-traceback`

`-mtraceback`

For embedded PowerPC systems, `-mno-traceback` provides for not generating a trace-back tag before the start of the function. This tag can be used by the

debugger to identify where the start of a function is. `-mtraceback` provides for generating a trace-back tag before the start of the function.

## IBM RT Options

The following `-m` options are defined for the IBM RT PC architectures.

`-mcall-lib-mul`

Calls `lmul$$` for integer multiples.

`-mfp-arg-in-fpregs`

`-mfp-arg-in-gregs`

`-mfp-arg-in-fpregs` uses a calling sequence incompatible with the IBM calling convention in which floating point arguments are passed in floating point registers. `-mfp-arg-in-gregs`, the default call of this group, uses the normal calling convention for floating point arguments.

`-mfull-fp-blocks`

`-mminimum-fp-blocks`

`-mfull-fp-blocks`, the default call of this group, generates full-size floating point data blocks, including the minimum amount of scratch space recommended by IBM. `-mminimum-fp-blocks` allows for not including extra scratch space in floating point data blocks; this results in smaller code, but slower execution, since scratch space must be allocated dynamically.

`-mhc-struct-return`

`-mnohc-struct-return`

`-mhc-struct-return` returns structures of more than one word in memory, rather than in a register, providing compatibility with the MetaWare HighC (`hc`) compiler.

Use `-fpcc-struct-return` for compatibility with the Portable C Compiler (`pcc`).

`-mnohc-struct-return` returns some structures of more than one word in registers, when convenient; this is the default.

For compatibility with the IBM-supplied compilers, use `-fpcc-struct-return` or `-mhc-struct-return`.

`-min-line-mul`

Uses an in-line code sequence for integer multiplies.

**IMPORTANT!** `varargs.h` and `stdarg.h` will not work with floating point operands when using this option.

# Intel x86 Options

The following `-m` options are defined for the Intel x86 family of computers (where *x* is either the i386, i486, i586, i686, Pentium, or the Pentium Pro processors from Intel, or the K6 from AMD).

`-mcpu=cpu type`

Sets the default values for the machine type, *cpu type*, when scheduling instructions; the choices for *cpu type* are i386, i486, i586, i686, pentium, pentiumpro, or k6. While picking a specific *cpu type* will schedule things appropriately for that particular chip, the compiler will not generate any code that does not run on the i386 without the `-march=cpu type` option being used. i586 is equivalent to pentium, and i686 is equivalent to pentiumpro.

`-march=cpu type`

Sets the default values for the target architecture, *cpu type*, when scheduling instructions; the choices for *cpu type* are i386, i486, i586, i686, pentium, pentiumpro, or k6. Specifying `-march=cpu type` is the same as specifying `-mcpu=cpu type`.

`-m386`

`-m486`

`-mpentium`

`-mpentiumpro`

Synonyms for `-mcpu=i386`, `-mcpu=i486`, `-mcpu=pentium`, and `-mcpu=pentiumpro`, respectively; these synonyms are deprecated.

`-mieee-fp`

`-mno-ieee-fp`

Control whether or not the compiler uses IEEE floating point comparisons. These handle correctly the case where the result of a comparison is unordered.

`-mno-fancy-math-387`

Some 387 emulators do not support the `sin`, `cos` and `sqrt` instructions for the 387. Specify this option to avoid generating those instructions. This option is the default on FreeBSD. As of revision 2.6.1, these instructions are not generated unless you also use the `-ffast-math` switch.

`-mno-fp-ret-in-387`

Do not use the FPU registers for return values of functions.

The usual calling convention has functions return values of types `float` and `double` in an FPU register, even if there is no FPU. The idea is that the operating system should emulate an FPU.

The option `-mno-fp-ret-in-387` causes such values to be returned in ordinary CPU registers instead.

`-msoft-float`

Generate output containing library calls for floating point.

**WARNING!** The requisite libraries are not part of GCC. Normally the facilities of the

machine's usual C compiler are used, but this can't be done directly in cross-compilation. You must make your own arrangements to provide suitable library functions for cross-compilation. On machines where a function returns floating point results in the 80387 register stack, some floating point opcodes may be emitted even if `-msoft-float` is used.

`-msvr3-shlib`

`-mno-svr3-shlib`

Control whether GCC places uninitialized locals into `bss` or `data`. `-msvr3-shlib` places these locals into `bss`. These options are meaningful only on System V Release 3.

`-mno-wide-multiply`

`-mwide-multiply`

Control whether GCC uses the `mul` and `imul` that produce 64 bit results in `eax:edx` from 32 bit operands to do `long long` multiplies and 32-bit division by constants.

`-mrtcd`

Use a different function-calling convention, in which functions that take a fixed number of arguments return with the `ret num` instruction, which pops their arguments while returning. This saves one instruction in the caller since there is no need to pop the arguments there. You can specify that an individual function is called with this calling sequence with the function attribute `stdcall`. You can also override the option, `-mrtcd`, by using the function attribute, `cdecl`. See "Declaring Attributes of Functions" on page 161.

**WARNING!** This calling convention is incompatible with the one normally used on UNIX, so you cannot use it if you need to call libraries compiled with the UNIX compiler.

You must provide function prototypes for all functions that take variable numbers of arguments (including `printf`); otherwise incorrect code will be generated for calls to those functions.

In addition, seriously incorrect code will result if you call a function with too many arguments. Normally, extra arguments are harmlessly ignored.

`-malign-double`

`-mno-align-double`

Control whether GCC aligns `double`, `long double`, and `long long` variables on a two word boundary or a one word boundary. Aligning double variables on a two word boundary will produce code that runs somewhat faster on a Pentium at the expense of more memory.

**WARNING!** If you use the `-malign-double` switch, structures containing the above types will be aligned differently than the published application binary interface specifications for the 386.

`-mreg-alloc=regs`

Control the default allocation order of integer registers. The `regs` string is a

series of letters specifying a register. The supported letters are: `a` for allocating EAX; `b` for allocating EBX; `c` for allocating ECX; `d` for allocating EDX; `s` for allocating ESI; `D` for allocating EDI; `B` for allocating EBP.

`-mregparm=num`

Control how many registers are used to pass integer arguments. By default, no registers are used to pass arguments, and at most 3 registers can be used. You can control this behavior for a specific function by using the function attribute `regparm`; see also “Declaring Attributes of Functions” on page 161.

**WARNING!** If you use this switch, and `num` is nonzero, then you must build all modules with the same value, including any libraries. This includes the system libraries and startup modules.

`-malign-loops=num`

Align loops to a 2 raised to a `num` byte boundary. If `-malign-loops` is not specified, the default is 2.

`-malign-jumps=num`

Align instructions that are only jumped to a 2 raised to a `num` byte boundary. If `-malign-jumps` is not specified, the default is 2 if optimizing for a 386, and 4 if optimizing for a 486.

`-malign-functions=num`

Align the start of functions to a 2 raised to `num` byte boundary.

If `-malign-jumps` is not specified, the default is 2 if optimizing for a 386, and 4 if optimizing for a 486.

`-mpreferred-stack-boundary=num`

Attempt to keep the stack boundary aligned to a 2 raised to `num` byte boundary. If `-mpreferred-stack-boundary` is not specified, the default is 4 (16 bytes or 128 bits).

The stack is required to be aligned on a 4 byte boundary. On Pentium and PentiumPro, `double` and `long double` values should be aligned to an 8 byte boundary (see `-malign-double`) or suffer significant run time performance penalties. On Pentium III, the Streaming SIMD<sup>†</sup> Extention (SSE) data type `__m128` suffers similar penalties if it is not 16 byte aligned. To ensure proper alignment of these values on the stack, the stack boundary must be as aligned as that required by any value stored on the stack. Further, every function must be generated such that it keeps the stack aligned. Thus calling a function compiled with a higher preferred stack boundary from a function compiled with a lower preferred stack boundary will most likely misalign the stack. It is recommended that libraries that use callbacks always use the default setting. This extra alignment does consume extra stack space. Code that is sensitive to stack space

<sup>†</sup> Single Instruction Multiple Data. Usually, processors process one data element in one instruction, a processing style called Single Instruction Single Data, or SISD. In contrast, processors having the SIMD capability process more than one data element in one instruction.

usage, such as embedded systems and operating system kernels, may want to reduce the preferred alignment to `-mpreferred-stack-boundary=2`.

`-mthreads`

Support thread-safe exception handling on Mingw32. Code that relies on thread-safe exception handling must compile and link all code with the `-mthreads` option. When compiling, `-mthreads` defines `-D_MT`; when linking, it links in a special thread helper library, `-lmingwthrd`, which cleans up per thread exception handling data.

`-mno-align-stringops`

Do not align destination of inlined string operations. This switch reduces code size and improves performance in case the destination is already aligned, but GCC doesn't know about it.

`-minline-all-stringops`

By default GCC inlines string operations only when destination is known to be aligned at least to 4 byte boundary. This enables more inlining, increase code size, but may improve performance of code that depends on fast `memcpy`, `strlen` and `memset` for short lengths.

`-mpush-args`

Use `PUSH` operations to store outgoing parameters. This method is shorter and usually equally fast as method using `SUB/MOV` operations and is enabled by default. In some cases disabling it may improve performance because of improved scheduling and reduced dependencies.

`-maccumulate-outgoing-args`

If enabled, the maximum amount of space required for outgoing arguments will be computed in the function prologue. This is faster on most modern CPUs because of reduced dependencies, improved scheduling and reduced stack usage when preferred stack boundary is not equal to 2. The drawback is a notable increase in code size. This switch implies `-mno-push-args`.

## Intel 960 Options

The following `-m` options are defined for the Intel 960 implementations.

`-mcpu type`

Assume the defaults for the machine type `cpu type` for some of the other options, including instruction scheduling, floating point support, and addressing modes. The choices for `cpu type` are `ka`, `kb`, `mc`, `ca`, `cf`, `sa`, and `sb`. The default is `kb`.

`-masm-compat`

Enable compatibility with the iC960 assembler.

`-mclean-linkage`

These options are not fully implemented.

`-mcode-align`

`-mno-code-align`

`-mcode-align` aligns code to 8-byte boundaries for faster fetching; currently turned on by default for C-series implementations only; `-mno-code-align` disables the functionality of `-mcode-align`.

`-mcomplex-addr`

`-mno-complex-addr`

`-mcomplex-addr` assumes (`-mno-complex-addr` does not assume) that the use of a complex addressing mode is a win on this implementation of the i960. Complex addressing modes may not be worthwhile on the K-series, but they definitely are on the. The default is currently `-mcomplex-addr` for all processors except the CB and CC.

`-mleaf-procedures`

`-mno-leaf-procedures`

`-mleaf-procedures` (or do not) attempts to alter leaf procedures to be callable with the `bal` instruction as well as `call`. This will result in more efficient code for explicit calls when the `bal` instruction can be substituted by the assembler or linker, but less efficient code in other cases, such as calls using function pointers, or using a linker that doesn't support this optimization. `-mno-leaf-procedures` disables the functionality of `-mleaf-procedures`.

`-mic-compat`

`-mic2.0-compat`

`-mic3.0-compat`

Enable compatibility with iC960 v2.0 or v3.0.

`-mintel-asm`

Enable compatibility with the iC960 assembler.

`-mno-old-align`

`-mold-align`

`-mold-align` enables structure-alignment compatibility with Intel's GCC release version 1.3 (based on GCC 1.37). This option implies `-mstrict-align`.

`-mno-old-align` disables the functionality of `-mold-align`, the default.

`-mstrict-align`

`-mno-strict-align`

`-mstrict-align` doesn't permit (`-mno-strict-align` permits) unaligned accesses.

`-mnumerics`

`-msoft-float`

The `-mnumerics` option indicates that the processor does support floating point instructions. The `-msoft-float` option indicates that floating point support should not be assumed.

`-mtail-call`

`-mno-tail-call`

`-mtail-call` makes (`-mno-tail-call` does not make) additional attempts (beyond those of the machine-independent portions of the compiler) to optimize tail-recursive calls into branches. You may not want to do this because the

detection of cases where this is not valid is not totally complete. The default is `-mno-tail-call`.

## Matsushita MN10200 Options

The following `-m` option is defined for Matsushita MN10200 architectures.

`-mrelax`

Indicate to the linker that it should perform a relaxation optimization pass to shorten branches, calls and absolute memory addresses. This option only has an effect when used on the command line for the final link step. This option makes symbolic debugging impossible.

## Matsushita MN10300/AM33 Options

These `-m` options are defined for Matsushita MN10300 or AM33 architectures.

`-mmult-bug`

`-mmult-bug` generates code to avoid bugs in the multiply instructions for the MN10300 processors. `-mmult-bug` is the default, not `-mno-mult-bug`.

`-mno-mult-bug`

`-mno-mult-bug` does not generate code, avoiding bugs in the multiply instructions for the MN10300 processors.

`-mam33`

`-mno-am33`

`-mam33` generates code using features specific to the AM33 processor. `-mno-am33`, the default, does not generate code specific to the AM33 processor.

`-mrelax`

Indicates to the linker that the linker should perform a relaxation optimization pass to shorten branches, calls and absolute memory addresses. This option only has an effect when used on the command line for the final link step.

This option makes symbolic debugging impossible.

## MIPS Options

The following `-m` options are defined for the MIPS family of computers.

`-mabicalls`

`-mno-abicalls`

`-mabicalls` emits (`-mno-abicalls` does not emit) the pseudo operations, `.abicalls`, `.cpload`, and `.cprestore`, which some System V.4 ports use for position independent code.

`-mcpu=cpu type`

Assumes the defaults for the machine type *cpu type* when scheduling instructions. The choices for *cpu type* are `r2000`, `r3000`, `r4000`, `r4400`, `r4600`, and `r6000`. While picking a specific *cpu type* will schedule things appropriately for that particular chip, the compiler will not generate any code that does not meet level 1 of the MIPS ISA (instruction set architecture) without the `-mips2` or `-mips3` switches being used.

`-membedded-data`

`-mno-embedded-data`

Allocate variables to the read-only data section first if possible, then next in the small data section if possible, otherwise in data. This gives slightly slower code than the default, reducing the amount of RAM required when executing, and thus may be preferable for some embedded systems.

`-muninit-const-in-rodata`

When used together with `-membedded-data`, it will always store uninitialized `const` variables in the read-only data section.

`-membedded-pic`

`-mno-embedded-pic`

Generate PIC code suitable for some embedded systems. All calls are made using PC relative address, and all data is addressed using the register, `$gp`. This requires coordination with the GNU `as` assembler and the GNU `ld` linker, which do most of the work. `-mno-embedded-pic` disables the functionality of `-membedded-pic`.

`-mfp32`

Assume that 32 32-bit floating point registers are available. This is the default.

`-mfp64`

Assume that 32 64-bit floating point registers are available. This is the default when the `-mips3` option is used.

`-mgas`

Generate code for the GNU assembler. This is the default on the OSF/1 reference platform, using the `OSF/rose` object format.

`-mgp32`

Assume that 32 32-bit general purpose registers are available. This is the default.

`-mgp64`

Assume that 32 64-bit general purpose registers are available. This is the default when the `-mips3` option is used.

`-mgpopt`

`-mno-gpopt`

The `-mgpopt` switch says to write all of the data declarations before the instructions in the text section, this allows the MIPS assembler to generate one word memory references instead of using two words for short global or static data items. This is on by default if optimization is selected.

- mhalf-pic
- mno-half-pic
  - Put pointers to extern references into the data section and load them up, rather than put the references in the text section.
- mhard-float
  - Generate output containing floating point instructions. This is the default if you use the unmodified sources.
- mint64
  - Types `long`, `int`, and `pointer` are 64 bits. This works only if `-mips3` is also specified.
- mips1
  - Issues instructions from level 1 of the MIPS ISA. This is the default. `r3000` is the default `cpu type` at this ISA level.
- mips2
  - Issues instructions from level 2 of the MIPS ISA (branch likely, square root instructions). `r6000` is the default CPU type at this ISA level.
- mips3
  - Issue instructions from level 3 of the MIPS ISA (64 bit instructions). `r4000` is the default `cpu type` at this ISA level. This option does not change the sizes of any of the C data types.
- mips4
  - Issue instructions from level 4 of the MIPS ISA (conditional move, prefetch, enhanced FPU instructions). `r8000` is the default `cpu type` at this ISA level.
- mlong64
  - Force `long` types to be 64 bits wide. See `-mlong32` for an explanation of the default, and the width of pointers.
- mlong32
  - Force `long`, `int`, and `pointer` types to be 32 bits wide.
- mlong-calls
- mno-long-calls
  - Do all calls with the `JALR` instruction, which requires loading up a function's address into a register before the call. You need to use this switch, if you call outside of the current 512 megabyte segment to functions that are not through pointers.
- mmemcpy
- mno-memcpy
  - The `-mmemcpy` switch makes all block moves call the appropriate string function (`memcpy` or `bcopy`) instead of possibly generating inline code.
- mmips-as
  - Generate code for the MIPS assembler, and invoke `mips-tfile` to add normal debug information. This is the default for all platforms except for the OSF/1 reference platform, using the OSF/rose object format. If either of the `-gstabs` or

`-gstabs+` switches are used, the `mips-tfile` program will encapsulate the `stabs` within MIPS ECOFF.

`-mmips-tfile`

`-mno-mips-tfile`

The `-mno-mips-tfile` switch causes the compiler not post-process the object file with the `mips-tfile` program, after the MIPS assembler has generated it to add debug support. If `mips-tfile` is not run, then no local variables will be available to the debugger. In addition, `stage2` and `stage3` objects will have the temporary filenames passed to the assembler embedded in the object file, which means the objects will not compare the same. The `-mno-mips-tfile` switch should only be used when there are bugs in the `mips-tfile` program that prevents compilation.

`-mabiccalls`

`-mno-abiccalls`

`-mabiccalls` emits (`-mno-abiccalls` does not emit) the pseudo operations `.abiccalls`, `.cpload`, and `.cprestore`, which some System V.4 ports use for position independent code.

`-msplit-addresses`

`-mno-split-addresses`

Generate code to load the high and low parts of address constants separately. This allows the GNU compiler to optimize away redundant loads of the high order bits of addresses. This optimization is enabled by default for some embedded targets where GNU `as` and GNU `ld` are standard.

`-mrnames`

`-mno-rnames`

The `-mrnames` switch says to output code using the MIPS software names for the registers, instead of the hardware names (`a0` instead of `$4`). The only known assembler that supports this option is the Algorithmics assembler.

`-mstats`

`-mno-stats`

For each non-inline function processed, the `-mstats` switch causes the compiler to emit one line to the standard error file to print statistics about the program (number of registers saved, stack size, etc.).

`-msoft-float`

Generate output containing library calls for floating point.

**WARNING!** The requisite libraries are not part of GCC. Normally the facilities of the machine's usual C compiler are used, but this can't be done directly in cross-compilation. You must make your own arrangements to provide suitable library functions for cross-compilation.

`-m4650`

Turns on `-msingle-float`, `-mmad`, and, at least for now, `-mcpu=r4650`.

`-msingle-float`

`-mdouble-float`

The `-msingle-float` switch tells GCC to assume that the floating point

coprocessor only supports single precision operations, as on the `r4650` chip. The `-mdouble-float` switch permits GCC to use double precision operations. This is the default.

`-mmad`

`-mno-mad`

Permit use of the `mad`, `madu` and `mul` instructions, as on the `r4650` chip.

`-EL`

Compile code for the processor in little endian mode. The requisite libraries are assumed to exist.

`-EB`

Compile code for the processor in big endian mode. The requisite libraries are assumed to exist.

`-G num`

Put global and static items less than or equal to `num` bytes into the sections, `small data` or `bss`, instead of the `normal data` or `bss` section. This allows the assembler to emit one word memory reference instructions based on the global pointer (`gp` or `$28`), instead of the normal two words used. By default, `num` is 8. The `-G num` switch is also passed to the assembler and linker. All modules should be compiled with the same `-G num` value.

`-nocpp`

Tell the MIPS assembler to not run its preprocessor over user assembler files (with a suffix, `.s`) when assembling them.

`-mabi=32`

`-mabi=n32`

`-mabi=64`

`-mabi=n64`

`-mabi=eabi`

Generate code for the indicated ABI. The default instruction level is `-mips1` for 32, `-mips3` for `n32`, and `-mips4`, otherwise. Conversely, with `-mips1` or `-mips2`, the default ABI is 32; otherwise, the default ABI is 64.

`-mfix7000`

Pass an option to the GNU assembler, which will cause `nops` to be inserted if the read of the destination register of an `mfhi` or `mflo` instruction occurs in the subsequent two instructions.

`-mno-crt0`

Do not include the default `crt0`.

`-malign-loops=num`

Align loops to a 2 raised to a `num` byte boundary. If `-malign-loops` is not specified, the default is 2.

`-malign-jumps=num`

Align instructions that are only jumped to a 2 raised to a `num` byte boundary. If `-malign-jumps` is not specified, the default is 2.

`-malign-functions=num`

Align the start of functions to a 2 raised to a *num* byte boundary. If

`-malign-functions` is not specified, the default is 2. MIPS targets may override the default value.

`-mmax-skip-loops=num`

Maximum number of padding bytes allowed to satisfy a loop alignment request.

The default value is zero which specifies no limit on the number of padding bytes. Specific MIPS targets may override the default value.

`-mmax-skip-jumps=num`

Maximum number of padding bytes allowed to satisfy a loop alignment request.

The default value is zero which specifies no limit on the number of padding bytes. Specific MIPS targets may override the default value.

`-mmax-skip-functions=num`

Maximum number of padding bytes allowed to satisfy a loop alignment request.

The default value is zero which specifies no limit on the number of padding bytes.

This option may have no effect when combined with other options such as `-ffunction-sections`. Specific MIPS targets may override the default value.

## Mitsubishi D10V Options

These `-m` options are defined for the Mitsubishi D10V processor's implementations.

`-mint32`

`-mint16`

Make `int` data 32 (or 16) bits by default. The default is `-mint16`.

`-mdouble64`

`-mdouble32`

Make double data 64 (or 32) bits by default. The default is `-mdouble32`.

`-maddac3`

`-mno-addac3`

Enable (disable) the use of `addac3` and `subac3` instructions. The `-maddac3` instruction also enables the `-maccum` instruction.

`-maccum`

`-mno-accum`

Enable (disable) the use of the 32-bit accumulators in compiler generated code.

`-mno-cond-move`

`-mcond-move`

Disable (or enable) conditional move instructions, eliminating short branches.

`-mno-asm-optimize`

`-masm-optimize`

Disable (enable) passing `-O` to the assembler when optimizing. The assembler uses the `-O` option to parallelize adjacent short instructions automatically where possible.

`-mno-small-insns`

`-msmall-insns`

Disable (enable) converting some long instructions into two short instructions, which can eliminate some nops and enable more code to be conditionally executed.

`-mbranch-cost=n`

Increase the internal costs of branches to *n*. Higher costs means that the compiler will issue more instructions to avoid doing a branch. The default is 1.

`-mcond-exec=n`

Specify the maximum number of conditionally executed instructions that replace a branch. The default is 4.

## Mitsubishi M32R/D/X Options

The following `-m` options are defined for Mitsubishi M32R/D/X architectures.

`-mcode-model=small`

Assume all objects live in the lower 16MB of memory (so that their addresses can be loaded with the `ld24` instruction), and assume all subroutines are reachable with the `bl` instruction. This is the default. The addressability of a particular object can be set with the `model` attribute.

`-mcode-model=medium`

Assume objects may be anywhere in the 32 bit address space (the compiler will generate `seth/add3` instructions to load their addresses), and assume all subroutines are reachable with the `bl` instruction.

`-mcode-model=large`

Assume objects may be anywhere in the 32 bit address space (the compiler will generate `seth/add3` instructions to load their addresses), and assume subroutines may not be reachable with the `bl` instruction (the compiler will generate the much slower `seth/add3/jl` instruction sequence).

`-msdata=none`

Disable use of the small data area. Variables will be put into one of `.data`, `bss`, or `.rodata` (unless the `section` attribute has been specified). This is the default. The small data area consists of sections `.sdata` and `.sbss`. Objects may be explicitly put in the small data area with the `section` attribute using one of these sections.

`-msdata=sdata`

Put small global and static data in the small data area, but do not generate special code to reference them.

`-msdata=use`

Put small global and static data in the small data area, and generate special instructions to reference them.

`-G num`

Put global and static objects less than or equal to *num* bytes into the small data or

bss sections instead of the normal data or bss sections. The default value of *num* is 8. The `-msdata` option must be set to one of `sdata` or `use` for this option to have any effect. All modules should be compiled with the same `-G num` value.

Compiling with different values of *num* may or may not work; if it doesn't, the linker will give an "incorrect code will not be generated" error message.

`-mcond-exec=n`

Specifies the maximum number of conditionally executed instructions that replace a branch. The default is 4.

## Motorola 68000 Options

The following explains the `-m` options defined for the Motorola 68000 series. The default values for these options depends on which style of 68000 was selected when the compiler was configured; the defaults for the most common choices are also given with the following options.

`-m68000`

`-mc68000`

Generate output for a 68000. This is the default when the compiler is configured for 68000-based systems.

`-m68020`

`-mc68020`

Generate output for a 68020. This is the default when the compiler is configured for 68020-based systems.

`-m68881`

Generate output containing 68881 instructions for floating point. This is the default for most 68020 systems unless `-nfp` was specified when the compiler was configured.

`-m68030`

Generate output for a 68030. This is the default when the compiler is configured for 68030-based systems.

`-m68040`

Generate output for a 68040. This is the default when the compiler is configured for 68040-based systems.

This option inhibits the use of 68881/68882 instructions that have to be emulated by software on the 68040. If your 68040 does not have code to emulate those instructions, use `-m68040`.

`-m68060`

Generate output for a 68060. This is the default when the compiler is configured for 68060-based systems. This option inhibits the use of 68020 and 68881/68882 instructions that have to be emulated by software on the 68060. If your 68060 does not have code to emulate those instructions, use `-m68060`.

`-m68020-40`

Generate output for a 68040, without using any of the new instructions. This results in code which can run relatively efficiently on either a 68020/68881 or a 68030 or a 68040. The generated code does use the 68881 instructions that are emulated on the 68040.

`-m68020-60`

Generate output for a 68060, without using any of the new instructions. This results in code which can run relatively efficiently on either a 68020/68881 or a 68030 or a 68040. The generated code does use the 68881 instructions that are emulated on the 68060.

`-mfpa`

Generate output containing SunFPA instructions for floating point.

`-mbitfield`

Do use the bit-field instructions. The `-m68020` option implies `-mbitfield`. This is the default if you use a configuration designed for a 68020.

`-mnobitfield`

Do not use the bit-field instructions. The `-m68000` option implies `-mnobitfield`.

`-mrtd`

Use a different function-calling convention, in which functions that take a fixed number of arguments return with the `rtd` instruction, which pops their arguments while returning. This saves one instruction in the caller as there is no need to pop the arguments there.

The calling convention using the `-mrtd` option is incompatible with the one normally used on UNIX, so you cannot use it if you need to call libraries compiled with the UNIX compiler.

Also, you must provide function prototypes for all functions that take variable numbers of arguments (including `printf`); otherwise incorrect code will be generated for calls to those functions.

In addition, seriously incorrect code will result if you call a function with too many arguments. Normally, extra arguments are harmlessly ignored.

The `rtd` instruction is supported by the 68010, 68020, 68030, 68040, and 68060 processors but not by the 68000 or the 5200 processors.

`-mshort`

Consider type `int` to be 16 bits wide, like `short int`.

`-msoft-float`

Generate output containing library calls for floating point.

**WARNING!** The requisite libraries are not available for all Motorola 68000 series targets. Normally, the facilities of the machine's usual C compiler are used, but this can't be done directly in cross-compilation; make your own arrangements to provide suitable library functions for cross-compilation. `m68k-*-aout` and

`m68k-* -coff` do provide software floating point support.

`-mpcrel`

Use the PC-relative addressing mode of the 68000 directly, instead of using a global offset table. This option implies `-fpic`, allowing at most a 16-bit offset for PC-relative addressing. `-fPIC` is not currently supported with `-mpcrel`, though this could be supported for 68020 and higher processors.

`-malign-int`

`-mno-align-int`

Control whether GCC aligns `int`, `long`, `long long`, `float`, `double`, and `long double` variables on a 32-bit boundary (`-malign-int`) or a 16-bit boundary (`-mno-align-int`). Aligning variables on 32-bit boundaries produces code that runs somewhat faster on processors with 32-bit busses at the expense of more memory.

**WARNING!** Using the `-malign-int` switch, GCC will align structures having the above types differently than most published ABI specifications for the `m68k`.

`-mno-strict-align`

`-mstrict-align`

`-mno-strict-align` does not (`-mstrict-align` does) assume that unaligned memory references will be handled by the system.

## Motorola 88000 Options

The following `-m` options are defined for Motorola 88000 architectures.

`-m88000`

Generate code that works well on both the `m88100` and the `m88110`.

`-m88100`

Generate code that works best for the `m88100`, but that also runs on the `m88110`.

`-m88110`

Generate code that works best for the `m88110`, and may not run on the `m88100`.

`-mbig-pic`

Obsolete option to be removed from the next revision. Use `-fPIC`.

`-mno-check-zero-division`

`-mcheck-zero-division`

`-mcheck-zero-division` does (`-mno-check-zero-division` doesn't) generate code that guarantees that integer division by zero will be detected. By default, detection is guaranteed. Some models of the MC88100 processor fail to trap upon integer division by zero under certain conditions. By default, when compiling code that might be run on such a processor, GNU C generates code that explicitly checks for zero-valued divisors and traps with exception number 503 when one is detected. Use of `-mno-check-zero-division` suppresses such checking for code generated to run on an MC88100 processor. GNU C assumes that the MC88110 processor correctly detects all instances of integer division by zero. When

`-m88110` is specified, both `-mcheck-zero-division` and `-mno-check-zero-division` are ignored, and no explicit checks for zero-valued divisors are generated.

`-mtrap-large-shift`

`-mhandle-large-shift`

`-mtrap-large-shift` includes code to detect bit-shifts of more than 31 bits; `-mtrap-large-shift` traps such shifts and `-mhandle-large-shift` emits code to handle bit-shifts properly. By default, GCC makes no special provision for large bit shifts.

`-midentify-revision`

Include an `ident` directive in the assembler output recording the source filename, compiler name and version, timestamp, and compilation flags.

`-mocs-debug-info`

`-mno-ocs-debug-info`

Include (or omit) additional debugging information (about registers used in each stack frame) as specified in the 88open Object Compatibility Standard, “OCS.” This extra information allows debugging of code that has had the frame pointer eliminated. The default for DG/UX, SVr4, and Delta 88 SVr3.2 is to include this information; other 88k configurations omit this information by default.

`-mocs-frame-position`

When emitting COFF debugging information for automatic variables and parameters stored on the stack, use the offset from the canonical frame address which is the stack pointer (register 31) on entry to the function. The DG/UX, SVr4, Delta88 SVr3.2, and BCS configurations use `-mocs-frame-position`; other 88k configurations have the default, `-mno-ocs-frame-position`.

`-mno-ocs-frame-position`

When emitting COFF debugging information for automatic variables and parameters stored on the stack, use the offset from the frame pointer register (register 30). When this option is in effect, the frame pointer is not eliminated when debugging information is selected by the `-g` switch.

`-moptimize-arg-area`

`-mno-optimize-arg-area`

Control how function arguments are stored in stack frames.

`-moptimize-arg-area` saves space by optimizing them, but this conflicts with the 88open specifications. The alternative, `-mno-optimize-arg-area`, agrees with 88open standards. By default, GCC does not optimize the argument area.

`-mserialize-volatile`

`-mno-serialize-volatile`

`-mserialize-volatile` does (`-mno-serialize-volatile` doesn't) generate code to guarantee sequential consistency of volatile memory references. By default, consistency is guaranteed.

The order of memory references made by the MC88110 processor does not always match the order of the instructions requesting those references. In

particular, a `load` instruction may execute before a preceding `store` instruction. Such reordering violates sequential consistency of volatile memory references, when there are multiple processors. When consistency must be guaranteed, GNU C generates special instructions, as needed, to force execution in the proper order.

The MC88100 processor does not reorder memory references and so always provides sequential consistency. However, by default, GNU C generates the special instructions to guarantee consistency even when you use `-m88100`, so that the code may be run on an MC88110 processor. If you intend to run your code only on the MC88100 processor, use `-mno-serialize-volatile`.

The extra code generated to guarantee consistency may affect the performance of your application. If you know that you can safely forgo this guarantee, use `-mno-serialize-volatile`.

`-mno-underscores`

In assembler output, emit symbol names without adding an underscore character at the beginning of each name. The default is to use an underscore as prefix on each name.

`-mshort-data-num`

Generate smaller data references by making them relative to `r0`, allowing loading of a value using a single instruction (rather than the usual two). You control which data references are affected by specifying `num` with this option. For example, if you specify `-mshort-data-512` (`num` being 512 bytes), then the data references affected are those involving displacements of less than 512 bytes.

`-mshort-data-num` is not effective for `num` values greater than 64K.

`-msvr4`

`-msvr3`

Turn on (`-msvr4`) or off (`-msvr3`) compiler extensions related to System V release 4 (SVr4), controlling the following implementations:

- which variant of the assembler syntax to emit,
- what `-msvr4` makes the C preprocessor recognize `#pragma weak` used on System V release 4.
- what `-msvr4` makes GCC issue additional declaration directives used in System V release 4.

`-msvr4` is the default for the `m88k-motorola-sysv4` and `m88k-dg-dgux` M88K configurations. `-msvr3` is the default for all other M88K configurations.

`-muse-div-instruction`

Use the `div` instruction for signed integer division on the MC88100 processor. By default, the `div` instruction is not used.

On the MC88100 processor the signed integer division instruction `div` traps to the operating system on a negative operand. The operating system transparently completes the operation, but at a large cost in execution time. By default, when compiling code that might be run on an MC88100 processor, GNU C emulates

signed integer division using the unsigned integer division instruction `divu`), thereby avoiding the large penalty of a trap to the operating system. Such emulation has its own, smaller, execution cost in both time and space. To the extent that your code's important signed integer division operations are performed on two nonnegative operands, it may be desirable to use the `div` instruction directly.

On the MC88110 processor the `div` instruction (also known as the `divs` instruction) processes negative operands with-out trapping to the operating system. When `-m88110` is specified, `-muse-div-instruction` is ignored, and the `div` instruction is used for signed integer division.

**IMPORTANT!** The result of dividing `INT MIN` by `-1` is undefined. In particular, the behavior of such a division with and without `-muse-div-instruction` may differ.

`-mversion-03.00`

This option is obsolete, and is ignored.

`-mwarn-passed-structs`

Warn when a function passes a struct as an argument or result. Structure-passing conventions have changed during the evolution of the C language, and are often the source of portability problems. By default, GCC issues no such warning.

## NEC V850 Options

The following `-m` options are defined for NEC V850 implementations.

`-mlong-calls`

`-mno-long-calls`

`-mlong-calls` treat alls calls as being far away (or near with `-mno-long-calls`).

If calls are assumed to be far away, the compiler will always load the functions address up into a register, and call indirect through the pointer.

`-mno-ep`

`-mep`

`-mno-ep` does not optimize (`-mep` does optimize) the basic blocks that use the same index pointer 4 or more times to copy pointer into the `ep` register, using the shorter `sld` and `sst` instructions. The `-mep` option is on by default if you optimize.

`-mno-prolog-function`

`-mprolog-function`

`-mno-prolog-function` does not use (`-mprolog-function`, does use) external functions to save and restore registers at the prolog and epilog of a function. The external functions are slower, but use less code space if more than one function saves the same number of registers. The `-mprolog-function` option is on by default if you optimize.

`-mspace`

Makes code as small as possible. At present, this just turns on the `-mep` and `-mprolog-function` options.

`-mtda=n`

Puts static or global variables whose size is *n* bytes or less into the tiny data area to which the `ep` register points. The tiny data area can hold up to 256 bytes in total (128 bytes for byte references).

`-msda=n`

Puts static or global variables whose size is *n* bytes or less into the small data area to which the `gp` register points. The small data area can hold up to 64K.

`-mzda=n`

Puts static or global variables whose size is *n* bytes or less into the first 32K of memory.

`-mv850`

Specifies that the target processor is the V850.

`-mbig-switch`

Generates code suitable for big switch tables. Use this option only if the assembler/linker complain about out of range branches within a switch table.

`-mapp-regs`

`-mapp-regs` causes the `r2` and `r5` registers to be used in the code generated by the compiler. `-mno-app-regs` causes the `r2` and `r5` registers to be treated as fixed registers. `-mapp-regs` is the default.

`-mv850e`

Specifies that the target processor is the V850e.

The preprocessor constant, `__v850e__`, will be defined if this option is used. If neither `-mv850` nor `-mv850e` are defined, then a default target processor will be chosen and the relevant `__v850__` or `__v850e__` preprocessor constant will be defined.

The preprocessor constants, `__v850` and `__v850__`, are always defined, regardless of which processor variant is the target.

`-mdisable-callt`

`-mno-disable-callt`

`-mdisable-callt` will suppress generation of the `CALLT` instruction for the v850e versions of the v850 architecture. `-mno-disable-callt`, the default, allows the `CALLT` instruction to be used.

## SPARC Options

The following `-m` switches are supported on the SPARC architectures.

`-mcpu=cpu_type`

Set the instruction set, register set, and instruction scheduling parameters for machine type, *cpu\_type*. Supported values for *cpu\_type* are `v7`, `cypress`, `v8`, `supersparc`, `sparclite`, `hypersparc`, `sparclite86x`, `f930`, `f934`, `sparclet`, `tsc701`, `v9` (only supported on true 64 bit targets), and `ultrasparc`.

Default instruction scheduling parameters are used for values that select an architecture and not an implementation. These are v7, v8, sparclite, sparcllet, and v9.

| <i>Architecture</i> | <i>Implementation</i>    |
|---------------------|--------------------------|
| v7                  | cypress                  |
| v8                  | supersparc, hypersparc   |
| sparclite           | f930, f934, sparclite86x |
| sparcllet           | tsc701                   |
| v9                  | ultrasparc               |

`-mtune=cpu_type`

Sets the instruction scheduling parameters for machine type, `cpu_type`, but does not set the instruction set or register set that the option, `-mcpu=cpu_type`, would. The same values for `-mcpu=cpu_type`, are used for `-mcpu=cpu_type`, though the only useful values are those that select a particular CPU implementation: `cypress`, `supersparc`, `hypersparc`, `f930`, `f934`, `sparclite86x`, `tsc701`, and `ultrasparc`.

`-mno-app-regs`

`-mapp-regs`

Specify `-mapp-regs` to generate output using the global registers, 2 through 4, which the SPARC SVR4 ABI reserves for applications. These options are the default. To be fully SVR4 ABI compliant at the cost of some performance loss, specify `-mno-app-regs`. You should compile libraries and system software with `-mno-app-regs`.

`-mfpu`

`-mhard-float`

Generate output containing floating point instructions. These options are the default.

`-mno-fpu`

`-msoft-float`

Generate output containing library calls for floating point.

**WARNING!** The requisite libraries are not available for all SPARC targets. Normally the facilities of the machine's usual C compiler are used, but this cannot be done directly in cross-compilation. You must make your own arrangements to provide suitable library functions for cross-compilation. The embedded targets, `sparc-*-aout` and `sparclite-*-*`, do provide software floating point support.

`-msoft-float` changes the calling convention in the output file; therefore, it is only useful if you compile *all* of a program with this option. In particular, you need to compile `libgcc.a`, the library that comes with GCC, with `-msoft-float` in order for this to work.

`-mhard-quad-float`

Generate output containing quad-word (long double) floating point instructions.

`-msoft-quad-float`

Generate output containing library calls for quad-word (long double) floating point instructions. The functions called are those specified in the SPARC ABI. This is the default.

As of this writing, there are no SPARC implementations that have hardware support for the quad-word floating point instructions. They all invoke a trap handler for one of these instructions, and then the trap handler emulates the effect of the instruction. Because of the trap handler overhead, this is much slower than calling the ABI library routines. Thus the `-msoft-quad-float` option is the default.

`-mno-epilogue`

`-mepilogue`

With `-mepilogue` (the default), the compiler always emits code for function exit at the end of each function. Any function exit in the middle of the function (such as a return statement in C) will generate a jump to the exit code at the end of the function. With `-mno-epilogue`, the compiler tries to emit exit code inline at every function exit.

`-mno-flat`

`-mflat`

With `-mflat`, the compiler does not generate save/restore instructions and will use a *flat* or single register window calling convention. This model uses `%i7` as the frame pointer and is compatible with the normal register window model. Code from either may be intermixed. The local registers and the input registers (0-5) are still treated as “*call saved*” registers and will be saved on the stack as necessary.

With `-mno-flat` (the default), the compiler emits save/restore instructions (except for leaf functions) and is the normal mode of operation.

`-mno-unaligned-doubles`

`-munaligned-doubles`

Assume that doubles have 8 byte alignment. `-mno-unaligned-doubles` is the default.

With `-munaligned-doubles`, GCC assumes that doubles have 8 byte alignment only if they are contained in another type, or if they have an absolute address.

Otherwise, `-munaligned-doubles` assumes they have 4 byte alignment.

Specifying `-munaligned-doubles` avoids some rare compatibility problems with code generated by other compilers. It is not the default because a performance loss results, especially for floating point code.

`-mv8`

`-msparclite`

These two options select variations on the SPARC architecture.

By default (unless specifically configured for the Fujitsu SPARClite), GCC generates code for the v7 variant of the SPARC architecture.

`-mv8` will give you SPARC v8 code. The only difference from v7 code is that the

compiler emits the integer multiply and integer divide instructions which exist in SPARC v8 but not in SPARC v7.

`-msparclite` will give you SPARClite code. This adds the `integer multiply`, `integer divide step` and `scan (ffs)` instructions which exist in SPARClite but not in SPARC v7.

These options are deprecated and will be deleted in GCC 2.9. They have been replaced with `-mcpu=xxx`.

`-mcypress`

`-msupersparc`

These two options select the processor for which the code is optimized.

With `-mcypress` (the default), the compiler optimizes code for the Cypress CY7C602 chip, as used in the SparcStation/ SparcServer 3xx series. This is also appropriate for the older SparcStation 1, 2, IPX, and so forth.

With `-msupersparc` the compiler optimizes code for the SuperSparc CPU, as used in the SparcStation 10, 1000 and 2000 series. This flag also enables use of the full SPARC v8 instruction set. These options are deprecated and will be deleted in GCC 2.9. They have been replaced with `-mcpu=xxx`.

`-malign-loops=num`

Align loops to a 2 raised to a *num* byte boundary. If `-malign-loops` is not specified, the default is 2.

`-malign-jumps=num`

Align instructions that are only jumped to a 2 raised to a *num* byte boundary. If `-malign-jumps` is not specified, the default is 2.

`-malign-functions=num`

Align the start of functions to a 2 raised to a *num* byte boundary. If `-malign-functions` is not specified, the default is 2 if compiling for a 32-bit SPARC processor, and 5 if compiling for a 64-bit SPARC processor.

The following `-m` switches are specific to the SPARCllet processors.

`-mlittle-endian`

Generate code for a processor running in little-endian mode for the SPARCllet processor.

`-mlive-g0`

Treat register `%g0` as a normal register for the SPARCllet processor. GCC will continue to clobber it as necessary but will not assume it always reads as 0.

`-mbroken-saverestore`

Generate code that does not use non-trivial forms of the `save` and `restore` instructions for the SPARCllet processor. Early versions of the SPARCllet processor do not correctly handle `save` and `restore` instructions used with arguments. They correctly handle them used without arguments. A `save` instruction used without arguments increments the current window pointer but

does not allocate a new stack frame. It is assumed that the window overflow trap handler will properly handle this case as will interrupt handlers.

The following `-m` switches are specific to the SPARC V9 processors in 64-bit environments.

`-m32`

`-m64`

Generate code for a 32 bit or 64 bit environment. The 32 bit environment sets `int`, `long` and pointer to 32 bits. The 64 bit environment sets `int` to 32 bits and `long` and pointer to 64 bits.

`-mlittle-endian`

Generate code for a processor running in little-endian mode.

`-mcmmodel=medlow`

Generate code for the Medium/Low code model. The program must be linked in the low 32 bits of the address space. Pointers are 64 bits. Programs can be statically or dynamically linked.

`-mcmmodel=medmid`

Generate code for the Medium/Middle code model. The program must be linked in the low 44 bits of the address space, the text segment must be less than 2G bytes, and data segment must be within 2G of the text segment. Pointers are 64 bits.

`-mcmmodel=medany`

Generate code for the Medium/Anywhere code model. The program may be linked anywhere in the address space, the text segment must be less than 2G bytes, and data segment must be within 2G of the text segment. Pointers are 64 bits.

`-mcmmodel=embmedany`

Generate code for the Medium/Anywhere code model for embedded systems. Assume a 32 bit text and a 32 bit data segment, both starting anywhere (determined at link time). The `%g4` register points to the base of the data segment. Pointers still 64 bits. Programs are statically linked; PIC is not supported.

`-mstack-bias`

`-mno-stack-bias`

With `-mstack-bias`, assume that the stack pointer, and frame pointer if present, are offset by `-2047` which must be added back when making stack frame references. Otherwise, assume no such offset is present.

## System V Options

The following additional options are available on System V Release 4 architectures for compatibility with other compilers on those systems.

`-G`

Create a shared object. It is recommended that `-symbolic` or `-shared` be used instead.

-Qy

Identify the versions of each tool used by the compiler, in a `.ident` assembler directive in the output.

-Qn

Refrain from adding `.ident` directives to the output file (this is the default).

-YP, *dirs*

Search the directories, *dirs*, and no others, for libraries specified with `-l`.

-Ym, *dir*

Look in the directory, *dir*, to find the M4 preprocessor. The assembler uses this option.



# Options Controlling Code Generation Conventions

---

These machine-independent options control the interface conventions used in code generation. Most of them have both positive and negative forms; the negative form of `-ffoo` would be `-fno-foo`. In the following options, only one of the forms is listed—the one which is not the default. You can figure out the other form by either removing ‘no-’ or adding it.

`-fcall-saved-reg`

Treat the register named *reg* as an allocatable register saved by functions. It may be allocated even for temporaries or variables that live across a call. Functions compiled this way will save and restore the register *reg* if they use it. Use of this flag for a register that has a fixed pervasive role in the machine’s execution model, such as the stack pointer or frame pointer, will produce disastrous results.

A different sort of disaster will result from the use of this flag for a register in which function values may be returned. This flag does not have a negative form, because it specifies a three-way choice.

`-fcall-used-reg`

Treat the register named *reg* as an allocatable register that is clobbered by function calls. It may be allocated for temporaries or variables that do not live across a call. Functions compiled this way will not save and restore the register *reg*.

Use of this flag for a register that has a fixed pervasive role in the machine’s

execution model, such as the stack pointer or frame pointer, will produce disastrous results. This flag does not have a negative form, because it specifies a three-way choice.

`-fexceptions`

Enable exception handling. Generates extra code needed to propagate exceptions. For some targets, this implies GNU CC will generate frame unwind information for all functions, which can produce significant data size overhead, although it does not affect execution. If you do not specify this option, GNU CC will enable it by default for languages which normally require exception handling like C++, while disabling it for languages like C that do not normally require it. However, you may need to enable this option when compiling C code that needs to interoperate properly with exception handlers written in C++. You may also wish to disable this option if you are compiling older C++ programs that don't use exception handling.

`-funwind-tables`

Similar to `-fexceptions`, except that it will just generate any needed static data, but will not affect the generated code in any other way. You will normally not enable this option; instead, a language processor that needs this handling would enable it on your behalf.

`-ffixed-reg`

Treat the register named *reg* as a fixed register; generated code should never refer to it (except perhaps as a stack pointer, frame pointer or in some other fixed role). *reg* must be the name of a register. The register names accepted are machine-specific and are defined in the `REGISTER_NAMES` macro in the machine description macro file.

This flag does not have a negative form, because it specifies a three-way choice.

`-finhibit-size-directive`

Don't output a `.size` assembler directive, or anything else that would cause trouble if the function is split in the middle, and the two halves are placed at locations far apart in memory. This option is used when compiling `'crtstuff.c'`; you should not need to use it for anything else.

`-fcheck-memory-usage`

Generate extra code to check each memory access. GCC will generate code that is suitable for a detector of bad memory accesses such as `'Checker'`. If you specify this option, you can not use the `asm` or `__asm__` keywords.

You must also specify this option when you compile functions you call that have side effects. If you do not, you may get erroneous messages from the detector. Normally, you should compile all your code with this option.

If you use functions from a library that have side-effects (such as `read`), you may not be able to recompile the library and specify this option. In that case, you can enable the `'-fprefix-function-name'` option, which requests GCC to

encapsulate your code and make other functions look as if they were compiled with `-fcheck-memory-usage`. This is done by calling “stubs” which are provided by the detector. If you cannot find or build stubs for every function you call, you may have to specify `-fcheck-memory-usage` with `-fprefix-function-name`.

`-fprefix-function-name`

Request GCC to add a prefix to the symbols generated for function names. GCC adds a prefix to the names of functions defined as well as functions called. Code compiled with this option and code compiled without the option can't be linked together, unless stubs are used.

Consider compiling the following code with `-fprefix-function-name`.

```
extern void bar (int);
void
foo (int a)
{
 return prefix_bar (a + 5);
}
```

GCC will then compile the code written like the following example.

```
extern void prefix_bar (int);
void
prefix_foo (int a)
{
 return prefix_bar (a + 5);
}
```

`-fno-common`

Allocate even uninitialized global variables in the `bss` section of the object file, rather than generating them as common blocks. This has the effect that if the same variable is declared (without `extern`) in two different compilations, you will get an error when you link them. The only reason this might be useful is if you wish to verify that the program will work on other systems which always work this way.

`-fno-ident`

Ignore the `#ident` directive.

`-fno-gnu-linker`

Do not output global initializations (such as C++ constructors and destructors) in the form used by the GNU linker (on systems where the GNU linker is the standard method of handling them). Use this option when you want to use a non-GNU linker, which also requires using the `collect2` program to make sure the system linker includes constructors and destructors. (`collect2` is included in the GCC distribution.) For systems which must use `collect2`, the compiler driver GCC is configured to do this automatically.

`-fpcc-struct-return`

Return “*short*” `struct` and `union` values in memory like longer ones, rather than

in registers. This convention is less efficient, but it has the advantage of allowing intercallability between GCC-compiled files and files compiled with other compilers. The precise convention for returning structures in memory depends on the target configuration macros. Short structures and unions are those whose size and alignment match that of some integer type.

`-fpic`

Generate position-independent code (PIC) suitable for use in a shared library, if supported for the target machine. Such code accesses all constant addresses through a global offset table (GOT). The dynamic loader resolves the GOT entries when the program starts (the dynamic loader is not part of GCC; it is part of the operating system). If the GOT size for the linked executable exceeds a machine-specific maximum size, you get an error message from the linker indicating that `-fpic` does not work; in that case, recompile with `-fPIC` instead. (These maximums are 16K on the m88k, 8K on the SPARC, and 32K on the m68k and RS/6000. The 386 has no such limit.)

Position-independent code requires special support, and therefore works only on certain machines. For the 386, GCC supports PIC for System V but not for the Sun 386i. Code generated for the IBM RS/6000 is always position-independent.

`-fPIC`

If supported for the target machine, emit position-independent code, suitable for dynamic linking and avoiding any limit on the size of the global offset table. This option makes a difference on the m68k, m88k, and the SPARC.

Position-independent code requires special support, and therefore works only on certain machines.

`-freg-struct-return`

Use the convention that `struct` and `union` values are returned in registers when possible. This is more efficient for small structures than `-fpcc-struct-return`.

If you specify neither `-fpcc-struct-return` nor its contrary `-freg-struct-return`, GCC defaults to whichever convention is standard for the target. If there is no standard convention, GCC defaults to `-fpcc-struct-return`, except on targets where GCC is the principal compiler. In those cases, we can choose the standard, and we chose the more efficient register return alternative.

`-fshared-data`

Requests that the data and non-`const` variables of this compilation be shared data rather than private data. The distinction makes sense only on certain operating systems, where shared data is shared between processes running the same program, while private data exists in one copy per process.

`-fshort-enums`

Allocate to an `enum` type only as many bytes as it needs for the declared range of possible values. Specifically, the `enum` type will be equivalent to the smallest integer type which has enough room.

`-fshort-double`

Use the same size for `double` as for `float`.

`-fvolatile`

Consider all memory references through pointers to be volatile.

`-fvolatile-global`

Consider all memory references to extern and global data items to be volatile.

`-fvolatile-static`

Consider all memory references to static data to be volatile.

`-funaligned-pointers`

Assume that all pointers contain unaligned addresses. On machines where unaligned memory accesses trap, this will result in much larger and slower code for all pointer dereferences, but the code will work even if addresses are unaligned.

`-funaligned-struct-hack`

Always access structure fields using loads and stores of the declared size. This option is useful for code that dereferences pointers to unaligned structures, but only accesses fields that are themselves aligned. Without this option, GCC may try to use a memory access larger than the field. This might give an unaligned access fault on some hardware. This option makes some invalid code work at the expense of disabling some optimizations. It is strongly recommended that this option not be used.

`-foptimize-comparisons`

Optimize multiple comparisons better within `&&` and `||` expressions. This is an experimental option. In some cases it can result in worse code. It depends on many factors. Now it is known only that the optimization works well for the PowerPC 740 and PowerPC 750 processors. This option switches on the following transformations:

```
(a != 0 || b != 0) => ((a | b) != 0)
(a == 0 && b == 0) => ((a | b) == 0)
(a != b || c != d) => (((a ^ b) | (c ^ d)) != 0)
(a == b && c == d) => (((a ^ b) | (c ^ d)) == 0)
(a != 0 && b != 0) => (((a | -a) & (b | -b)) < 0)
(a != b && c != d) => x = a ^ b; y = c ^ d; (((x | -x) & (y | -y)) < 0)
(a < 0 || b < 0) => ((a | b) < 0)
(a < 0 && b < 0) => ((a & b) < 0)
(a >= 0 || b >= 0) => ((a & b) >= 0)
(a >= 0 && b >= 0) => ((a | b) >= 0)
(a < 0 || b >= 0) => ((a | ~b) < 0)
(a < 0 && b >= 0) => ((a & ~b) < 0)
(a >= 0 || b < 0) => ((~a | b) < 0)
(a >= 0 && b < 0) => ((~a & b) < 0)
(a != 0 && b < 0) => (((a | -a) & b) < 0)
(a != 0 && b >= 0) => (((a | -a) & ~b) < 0)
(a < 0 && b != 0) => (((b | -b) & a) < 0)
(a >= 0 && b != 0) => (((b | -b) & ~a) < 0)
```

`-fverbose-asm`

Put extra commentary information in the generated assembly code to make it more readable. This option is generally only of use to those who actually need to read the generated assembly code (perhaps while debugging the compiler itself).

'`-fverbose-asm`' is the default. '`-fno-verbose-asm`' causes the extra information to be omitted and is useful when comparing two assembler files.

`-fpack-struct`

Pack all structure members together without holes. Usually you would not want to use this option, since it makes the code suboptimal, and the offsets of structure members won't agree with system libraries.

`-fstack-check`

Generate code to verify that you do not go beyond the boundary of the stack. You should specify this flag if you are running in an environment with multiple threads, but only rarely need to specify it in a single-threaded environment since stack overflow is automatically detected on nearly all systems if there is only one stack.

`-fstack-limit-register=reg`

`-fstack-limit-symbol=sym`

`-fno-stack-limit`

Generate code to ensure that the stack does not grow beyond a certain value, either the value of a register or the address of a symbol. If the stack would grow beyond the value, a signal is raised. For most targets, the signal is raised before the stack overruns the boundary, so it is possible to catch the signal without taking special precautions.

For instance, if the stack starts at a `0x80000000` address and grows downwards, you can use the `-fstack-limit-symbol=__stack_limit` and

`-Wl,--defsym,__stack_limit=0x7ffe0000` flags, which will enforce a stack limit of 128K.

`-fargument-alias`

`-fargument-noalias`

`-fargument-noalias-global`

Specifies the possible relationships among parameters and between parameters and global data.

`-fargument-alias` specifies that arguments (parameters) may alias each other and may alias global storage.

`-fargument-noalias` specifies that arguments do not alias each other, but may alias global storage.

`-fargument-noalias-global` specifies that arguments do not alias each other and do not alias global storage.

Each language will automatically use whatever option is required by the language standard. You should not need to use these options.

`-fleading-underscore`

`-fno-leading-underscore`

`-fleading-underscore` and `-fno-leading-underscore` forcibly change the way C symbols are represented in the object file. One use is to help link with legacy assembly code. Not all targets provide complete support for these switches.

**WARNING!** Know what you are doing when invoking this option.

`+e0`

`+e1`

For C++ only; controls whether virtual function definitions in classes are used to generate code, or only to define interfaces for their callers.

These options are provided for compatibility with `cffront 1.x` usage; the recommended alternative G++ usage is in flux. See “Declarations and Definitions in One Header” on page 202.

With `+e0`, virtual function definitions in classes are declared `extern`; the declaration is used only as an interface specification, not to generate code for the virtual functions (in this compilation).

With `+e1`, G++ actually generates the code implementing virtual functions defined in the code, and makes them publicly visible.



# 16

## The `offset-info` Option

---

The `-offset-info output-file` option simplifies access to C structs from the assembler. For each member of each structure, the compiler will output a `.equ` directive to associate a symbol with the member's offset (in bytes) into the structure. The symbol itself is the concatenation of the structure's tag name and the member's name, separated by an underscore.

This option will output to the specified *output-file* an assembler directive, `.equ`, for each member of each structure found in each manipulation. The `.equ` directives for the structures in the header file can be obtained by using the following input: `gcc -fsyntax-only -offset-info m.s -x c m.h`.

`m.h` is the header containing the structures, and `m.s` holds the directives.

The following is a short example of output produced by `-offset-info`.

input file (for example `m.h`):

```
struct W {
 double d;
 int I;
};

struct X {
 int a;
 int b;
```

```
struct Y {
 int a;
 int b;
};

struct Y y;
struct Y yy[10];
struct Y* p;
}
```

output file (for example `m.s`):

```
.equ W_d,0
.equ W_i,8
.equ Y_a,0
.equ Y_b,4
.equ X_a,0
.equ X_b,4
.equ X_y,8
.equ X_yy,16
.equ X_p,96
```

`-offset-info` has the following caveats.

- No directives are output for bit-field members.
- No directives are output for members whose offsets (as measured in bits) is greater than the word size of the host.
- No directives are output for members whose offsets are not constants. This can happen only in structures that use some GCC specific extensions allowing for variable sized members.

# 17

## Environment Variables Affecting GCC

---

The following documentation describes several special *environment variables* that affect how GCC operates. They work by specifying directories or prefixes to use when searching for various kinds of files. Some are used to specify other aspects of the compilation environment.

**IMPORTANT!** You can also specify places to search using options such as `-B`, `-I` and `-L` (see “Options for Searching Directories” on page 71). These take precedence over places specified using environment variables, which in turn take precedence over those specified by the GCC configuration.

`LC_CTYPE`  
`LC_COLLATE`  
`LC_MESSAGES`  
`LC_MONETARY`  
`LC_NUMERIC`  
`LC_TIME`  
`LC_ALL`

Control the way that GCC uses localization information that allow GCC to work with different national conventions. GCC inspects the locale categories, `LC_CTYPE` and `LC_MESSAGES`, if it has been configured to do so. These locale categories can be set to any value supported by your installation. A typical value is `en_UK` for English in the United Kingdom. The `LC_CTYPE` environment variable specifies character classification. GCC uses it to determine the character boundaries in a

string; this is needed for some multibyte encodings that contain quote and escape characters that would otherwise be interpreted as a string end or escape. The `LC_MESSAGES` environment variable specifies the language to use in diagnostic messages. If the `LC_ALL` environment variable is set, it overrides the value of `LC_CTYPE` and `LC_MESSAGES`; otherwise, `LC_CTYPE` and `LC_MESSAGES` default to the value of the `LANG` environment variable. If none of these variables are set, GCC defaults to traditional C English language behavior.

#### `TMPDIR`

Specifies the directory to use for temporary files. GCC uses temporary files to hold the output of one stage of compilation which is to be used as input to the next stage: for example, the output of the preprocessor, which is the input to the compiler proper.

#### `GCC_EXEC_PREFIX`

Specifies a prefix to use in the names of the subprograms executed by the compiler. No slash is added when this prefix is combined with the name of a subprogram, but you can specify a prefix that ends with a slash if you wish.

If GCC cannot find the subprogram using the specified prefix, it tries looking in the usual places for the subprogram. The default value of `GCC_EXEC_PREFIX` is *prefix/lib/gcc-lib/* where *prefix* is the value of `prefix` when you ran the `configure` script. Other prefixes specified with `-B` take precedence over this prefix.

Also finds files such as `crt0.o` that are used for linking.

The prefix is used for finding directories to search for header files.

For each of the standard directories whose name normally begins with `/usr/local/lib/gcc-lib` (more precisely, with the value of `GCC_INCLUDE_DIR`), GCC tries replacing that beginning with the specified prefix to produce an alternate directory name. So, with `-Bfoo/`, GCC will search `foo/bar` where it would normally search `/usr/local/lib/bar`. These alternate directories are searched first; the standard directories come next.

#### `COMPILER_PATH`

The value of `COMPILER_PATH` is a colon-separated list of directories, much like `PATH`. GCC tries the directories thus specified when searching for subprograms, if it can't find the subprograms using `GCC_EXEC_PREFIX`.

#### `LIBRARY_PATH`

The value of `LIBRARY_PATH` is a colon-separated list of directories, much like `PATH`. When configured as a native compiler, GCC tries the directories thus specified when searching for special linker files, if it can't find them using `GCC_EXEC_PREFIX`. Linking using GCC also uses these directories when searching for ordinary libraries for the `-l` option (but directories specified with `-L` come first).

---

C\_INCLUDE\_PATH

CPLUS\_INCLUDE\_PATH

OBJC\_INCLUDE\_PATH

These environment variables pertain to particular languages. Each variable's value is a colon-separated list of directories, much like `PATH`. When GCC searches for header files, it tries the directories listed in the variable for the language you are using, after the directories specified with `-I` but before the standard header file directories.

DEPENDENCIES\_OUTPUT

If this variable is set, its value specifies how to output dependencies for `make` based on the header files processed by the compiler. This output looks much like the output from the `-M` option (see “Options Controlling Preprocessing” on page 61), but it goes to a separate file, and is in addition to the usual results of compilation. The value of `DEPENDENCIES_OUTPUT` can be just a filename, in which case the Make rules are written to that file, guessing the target name from the source filename. Or the value can have the form, *file target*, in which case the rules are written to *file* using *target* as the target name.

LANG

Used to pass locale information to the compiler. One way in which this information is used is to determine the character set to be used when character literals, string literals and comments are parsed in C and C++. When the compiler is configured to allow multibyte characters, the following values for `LANG` are recognized:

- C-JIS  
Recognizes JIS characters.
- C-SJIS  
Recognizes SJIS characters.
- C-EUCJP  
Recognizes EUCJP characters.



# 18

## Running the `protoize` Program

---

The `protoize` program is an optional part of GCC. You can use it to add *prototypes* to a program, thus converting the program to ANSI C in one respect. The companion program, `unprotoize`, does the reverse, removing argument types from any prototypes that are found.

When you run these programs, you must specify a set of source files as command line arguments. The conversion programs start out by compiling these files to see what functions they define. The information gathered about a file, `foo`, is saved in a file named `foo.x`.

After scanning comes actual conversion. The specified files are all eligible to be converted; any files they include (whether sources or just headers) are eligible as well.

Not all the eligible files are converted. By default, `protoize` and `unprotoize` convert only source and header files in the current directory.

You can specify additional directories whose files should be converted with the `-d directory` option. You can also specify particular files to exclude with the `-x file` option. A file is converted if it is eligible, its directory name matches one of the specified directory names, and its name within the directory has not been excluded.

Basic conversion with `protoize` consists of rewriting most function definitions and function declarations to specify the types of the arguments. The only ones not rewritten are those for `varargs` functions.

`protoize` optionally inserts prototype declarations at the beginning of the source file, to make them available for any calls that precede the function's definition. Or it can insert prototype declarations with block scope in the blocks where undeclared functions are called.

Basic conversion with `unprotoize` consists of rewriting most function declarations to remove any argument types, and rewriting function definitions to the old-style pre-ANSI form.

Both conversion programs print a warning for any function declaration or definition they can't convert. You can suppress these warnings with `-q`.

The output from `protoize` or `unprotoize` replaces the original source file. The original file is renamed to a name ending with `.save`. If the `.save` file already exists, then the source file is simply discarded.

`protoize` and `unprotoize` both depend on GCC itself to scan the program and collect information about the functions it uses. So neither of these programs will work until GCC is installed.

You can use the following options with `protoize` and `unprotoize`. Each option works with both programs unless otherwise stated.

`-B directory`

Look for the file, `SYSCALLS.c.x`, in the specified directory, `directory`, instead of the usual directory (normally `/usr/local/lib`). This file contains prototype information about standard system functions. This option applies only to `protoize`.

`-c compilation-options`

Use `compilation-options` as the options when running GCC to produce the `.x` files. The special option `-aux-info` is always passed in addition, to tell GCC to write a `.x` file.

**WARNING!** The compilation options must be given as a single argument to `protoize` or `unprotoize`. If you want to specify several GCC options, *you must quote the entire set of compilation options to make them a single word in the shell*. You *cannot* use certain GCC arguments because they produce the *wrong kind of output*. These include `-g`, `-O`, `-c`, `-S`, and `-o`. If you include these in the `'compilation-options'`, they are ignored.

`-C`

Rename files to end in `.c` instead of `.C` extensions. This is convenient if you are converting a C program to C++. This option applies only to `protoize`.

`-g`

Add explicit global declarations. This means inserting explicit declarations at the beginning of each source file for each function that is called in the file and was not declared. These declarations precede the first function definition that contains a call to an undeclared function. This option applies only to `protoize`.

- `-i string`  
Indent old-style parameter declarations with the string *string*. This option applies only to `protoize`.  
`unprotoize` converts prototyped function definitions to old-style function definitions, where the arguments are declared between the argument list and the `{` initial brace. By default, `unprotoize` uses five spaces as the indentation. If you want to indent with just one space instead, use a `-i " "` command.
- `-k`  
Keep the `.x` files. Normally, they are deleted after conversion is finished.
- `-l`  
Add explicit local declarations. `protoize` with `-l` inserts a prototype declaration for each function in each block which calls the function without any declaration. This option applies only to `protoize`.
- `-n`  
Make no real changes. This mode just prints information about the conversions that would have been done without `-n`.
- `-N`  
Make no `.save` files. The original files are simply deleted. Use this option with caution.
- `-p program`  
Use the program, *program*, as the compiler. Normally, GCC, is used.
- `-q`  
Work quietly. Most warnings are suppressed.
- `-v`  
Print the version number, just like `-v` for GCC.

If you need special compiler options to compile one of your program's source files, then you should generate that file's `.x` file specially, by running GCC on that source file with the `-aux-info` option and the appropriate options. Then run `protoize` on the entire set of files. `protoize` will use the existing `.x` file because it is newer than the source file. Use the following example for `protoize`.

```
gcc -Dfoo=bar file1.c -aux-info
protoize *.c
```

You need to include the special files along with the rest in the `protoize` command, even though their `.x` files already exist, because otherwise they won't get converted.

See “`protoize` and `unprotoize` Warnings” on page 236 for more information on how to use `protoize` successfully.



# Extensions to the C Language Family

---

GCC provides several language features not found in ANSI standard C. The option, `-pedantic`, directs GCC to print a warning message if any of these features is used. To test for the availability of these features in conditional compilation, check for a predefined macro, `__GNUC__`, which is always defined under GCC.

These extensions are available in C and Objective C. Most of them are also available in C++. See “Extensions to the C++ Language Family” on page 199 for extensions that apply *only* to C++.

See the following documentation for discussion on the subject of extensions to C.

- “Statements and Declarations in Expressions” on page 146
- “Locally Declared Labels” on page 147
- “Labels as Values” on page 148
- “Nested Functions” on page 149
- “Constructing Function Calls” on page 151
- “Naming an Expression’s Type” on page 152
- “Referring to a Type with the `typeof` Keyword” on page 152
- “Generalized Lvalues” on page 153
- “Conditionals with Omitted Operands” on page 154
- “Double-word Integers” on page 154

- “Complex Numbers” on page 155
- “Arrays of Length Zero” on page 156
- “Arrays of Variable Length” on page 156
- “Macros with Variable Numbers of Arguments” on page 157
- “Non-lvalue Arrays May Have Subscripts” on page 158
- “Arithmetic on `void` Pointers and Function pointers” on page 158
- “Non-constant Initializers” on page 159
- “Constructor Expressions” on page 159
- “Labeled Elements in Initializers” on page 160
- “Declaring Attributes of Functions” on page 161
- “Prototypes and Old-style Function Definitions” on page 167
- “Compiling Functions for Interrupt Calls” on page 168
- “Inquiring on Alignment of Types or Variables” on page 169
- “Specifying Attributes of Variables” on page 169
- “Specifying Attributes of Types” on page 173
- “An inline Function Is as Fast as a Macro” on page 176
- “Assembler Instructions with C Expression Operands” on page 177
- “Constraints for `asm` Operands” on page 181
- “Controlling Names Used in Assembler Code” on page 193
- “Variables in Specified Registers” on page 193
- “Alternate Keywords” on page 196
- “Incomplete enum Types” on page 196
- “Function Names as Strings” on page 197

## Statements and Declarations in Expressions

A compound statement enclosed in parentheses may appear as an expression in GCC. This allows you to use loops, switches, and local variables within an expression.

Recall that a compound statement is a sequence of statements surrounded by braces; in the following construct, parentheses go around the braces.

```
({ int y = foo (); int z;
 if (y > 0) z = y;
 else z = - y;
 z; })
```

The previous construct is a valid (though slightly more complex than necessary) expression for the absolute value of `foo()`.

The last thing in the compound statement should be an expression followed by a semicolon; the value of this subexpression serves as the value of the entire construct. If you use some other kind of statement last within the braces, the construct has type `void`, and thus effectively no value.

This feature is especially useful in making macro definitions “safe” (so that they evaluate each operand exactly once). For example, the “*maximum*” function is commonly defined as a macro in standard C as follows.

```
#define max(a,b) ((a) > (b) ? (a) : (b))
```

But this definition computes either `a` or `b` twice, with bad results if the operand has side effects. In GCC, if you know the type of the operands (in the following example, `int`), you can define the macro safely as follows.

```
#define maxint(a,b) ({int _a = (a), _b = (b); _a > _b ? _a : _b; })
```

Embedded statements are not allowed in constant expressions, such as the value of an enumeration constant, the width of a bit field, or the initial value of a static variable.

If you don’t know the type of the operand, you can still do this, but you must use `typeof` (see “Referring to a Type with the `typeof` Keyword” on page 152) or type naming (see “Naming an Expression’s Type” on page 152).

## Locally Declared Labels

Each statement expression is a scope in which *local labels* can be declared. A local label is simply an identifier.

You can jump to it with an ordinary `goto` statement, but only from within the statement expression it belongs to; a local label declaration looks like `__label__ label` or `__label__ label1, label2, ...`

Local label declarations must come at the beginning of the statement expression, right after the parenthesis and brace, `{ ``, before any ordinary declarations.

The label declaration defines the label *name*, but does not define the label itself. You must do this in the usual way, with `label:`, within the statements of the statement expression.

The local label feature is useful because statement expressions are often used in macros. If the macro contains nested loops, a `goto` can be useful for breaking out of them. However, an ordinary label whose scope is the whole function cannot be used: if the macro can be expanded several times in one function, the label will be multiply defined in that function. A local label avoids this problem, as the following example shows.

```

#define SEARCH(array, target) \
({ \
 __label__ found; \
 typeof (target) _SEARCH_target = (target); \
 typeof (*(array)) *_SEARCH_array = (array); \
 int i, j; \ int value; \
 for (i = 0; i < max; i++) \
 for (j = 0; j < max; j++) \
 if (_SEARCH_array[i][j] == _SEARCH_target) \
 { value = i; goto found; } \
 value = -1; \
found: \
 value; \
})

```

## Labels as Values

You can get the address of a label defined in the current function (or a containing function) with the unary operator, `&&`. The value has type `void *`. This value is a constant and can be used wherever a constant of that type is valid, as the following example shows.

```

void *ptr;
...
ptr = &&foo;

```

To use these values, you need to be able to jump to one. This is done with the computed `goto` statement<sup>†</sup>, `goto * exp;` as in `goto *ptr;`

Any expression of type `void *` is allowed.

One way of using these constants is in initializing a static array that will serve as a jump table:

```

static void *array[] = { &&foo, &&bar, &&hack };

```

Then you can select a label with indexing, as in the following example.

```

goto *array[i];

```

**IMPORTANT!** This does not check whether the subscript is in bounds—array indexing in C never does that.

Such an array of label values serves a purpose much like that of the `switch` statement. The `switch` statement is cleaner, so use that rather than an array unless the problem does not fit a `switch` statement very well.

Another use of label values is in an interpreter for threaded code. The labels within the interpreter function can be stored in the threaded code for super-fast dispatching.

<sup>†</sup> The analogous feature in FORTRAN is called an assigned `goto`, but that name seems inappropriate in C, where one can do more than simply store label addresses in label variables.

You can use this mechanism to jump to code in a different function. If you do that, totally unpredictable things will happen. The best way to avoid this is to store the label address only in automatic variables and never pass it as an argument.

## Nested Functions

A *nested function* is a function defined inside another function. Nested functions are not supported for GNU C++. The nested function's name is local to the block where it is defined. For example, in the following, a nested function is named `square`, and is called twice.

```
foo (double a, double b)
{
 double square (double z) { return z * z; }

 return square (a) + square (b);
}
```

The nested function can access all the variables of the containing function that are visible at the point of its definition. This is called *lexical scoping*. For example, in the following, a nested function uses an inherited variable named `offset`.

```
bar (int *array, int offset, int size)
{
 int access (int *array, int index)
 { return array[index + offset]; }
 int i;
 ...
 for (i = 0; i < size; i++)
 ... access (array, i) ...
}
```

Nested function definitions are permitted within functions in the places where variable definitions are allowed; that is, in any block, before the first statement in the block.

It is possible to call the nested function from outside the scope of its name by storing its address or passing the address to another function, as in the following declaration.

```
hack (int *array, int size)
{
 void store (int index, int value)
 { array[index] = value; }

 intermediate (store, size);
}
```

Within this block, the function `intermediate` receives the address of `store` as an argument. If `intermediate` calls `store`, the arguments given to `store` are used to store into `array`. But this technique works only so long as the containing function (`hack`, in the previous example) does not exit.

If you try to call the nested function through its address after the containing function has exited, all hell will break loose. If you try to call it after a containing scope level has exited, and if it refers to some of the variables that are no longer in scope, you may be lucky, but it's not wise to take the risk. If, however, the nested function does not refer to anything that has gone out of scope, you should be safe.

GCC implements taking the address of a nested function using a technique called *trampolines*. A paper describing trampolines is available from [maya.idiap.ch](http://maya.idiap.ch), under `pub/tmb`, in `usenix88-lexic.ps.Z`.

A nested function can jump to a label inherited from a containing function, provided the label was explicitly declared in the containing function (see “Locally Declared Labels” on page 147). Such a jump returns instantly to the containing function, exiting the nested function which did the `goto` and any intermediate functions as well, as the following example shows.

```
bar (int *array, int offset, int size)
{
 __label__ failure;
 int access (int *array, int index)
 {
 if (index > size)
 goto failure;
 return array[index + offset];
 }
 int i;
 ...
 for (i = 0; i < size; i++)
 ...access (array, i) ...
 ...
 return 0;

 /* Control comes here from access if it detects an error. */
failure:
 return -1;
}
```

A nested function always has internal linkage. Declaring one with `extern` is erroneous. If you need to declare the nested function before its definition, use `auto` (which is otherwise meaningless for function declarations).

```
bar (int *array, int offset, int size)
{
 __label__ failure;
 auto int access (int *, int);
 ...
 int access (int *array, int index)
 {
```

```

 if (index > size)
 goto failure;
 return array[index + offset];
 }
 ...
}

```

## Constructing Function Calls

Using the built-in functions described in the following, you can record the arguments a function received, and call another function with the same arguments, without knowing the number or types of the arguments. You can also record the return value of that function call, and later return that value, without knowing what data type the function tried to return (as long as your caller expects that data type).

`__builtin_apply_args ()`

This built-in function returns a pointer of type `void *` to data describing how to perform a call with the same arguments as were passed to the current function.

The function saves the `arg` pointer register, structure value address, and all registers that might be used to pass arguments to a function into a block of memory allocated on the stack. Then it returns the address of that block.

`__builtin_apply (function, arguments, size)`

This built-in function invokes `function` (type `void (*)()`) with a copy of the parameters that `arguments` (type `void *`) and `size` (type `int`) describe.

The value of `arguments` should be the value returned by

`__builtin_apply_args`. The argument `size` specifies the size of the stack argument data, in bytes.

This function returns a pointer of type `void *` to data describing how to return whatever value was returned by `function`. The data is saved in a block of memory allocated on the stack.

It is not always simple to compute the proper value for `size`. The value is used by `__builtin_apply` to compute the amount of data that should be pushed on the stack and copied from the incoming argument area.

`__builtin_return (result)`

This built-in function returns the value described by `result` from the containing function. You should specify, for `result`, a value returned by `__builtin_apply`.

# Naming an Expression's Type

You can give a name to the type of an expression using a `typedef` declaration with an initializer. To define *name* as a type name for the type of *exp*, as the following example shows.

```
typedef name = exp;
```

This is useful in conjunction with the statements-within-expressions feature. The following shows how the two together can be used to define a safe “*maximum*” macro that operates on any arithmetic type.

```
#define max(a,b) \
 ({typedef _ta = (a), _tb = (b); _ta _a = (a); _tb _b = (b); \
 _a > _b ? _a : _b; })
```

The reason for using names that start with underscores for the local variables is to avoid conflicts with variable names that occur within the expressions that are substituted for *a* and *b*. Eventually, a new form of declaration syntax would allow declaring variables whose scopes start only after their initializers; this will be a more reliable way to prevent such conflicts.

## Referring to a Type with the `typeof` Keyword

Another way to refer to the type of an expression is with `typeof`. The syntax of using of this keyword looks like `sizeof`, but the construct acts semantically like a type name defined with `typedef`.

There are two ways of writing the argument to `typeof`: with an expression, or with a type.

The following is an example with an expression; this input assumes that *x* is an array of functions, while the type described is that of the values of the functions.

```
typeof (x[0](1))
```

The following is an example with a type name as the argument.

```
typeof (int *)
```

In the following discussion, the type described is that of pointers to `int`.

If you are writing a header file that must work when included in ANSI C programs, write `__typeof__` instead of `typeof`; for more information, see “Alternate Keywords” on page 196.

A `typeof`-construct can be used anywhere a `typedef` name could be used. For example, you can use it in a declaration, in a cast, or inside of `sizeof` or `typeof`.

```
typeof (*x) y;
```

The following example shows declaring `y` as an array of such values.

```
typeof (*x) y[4];
```

The following example shows declaring `y` as an array of pointers to characters.

```
typeof (typeof (char *)[4]) y;
```

It is equivalent to the traditional C declaration, `char *y[4];`.

To see the meaning of the declaration using `typeof`, and why it might be a useful way to write, let's rewrite it with the following macros.

```
#define pointer(T) typeof(T *)
#define array(T, N) typeof(T [N])
```

Now the declaration can be rewritten the following way.

```
array (pointer (char), 4) y;
```

Thus, `array (pointer (char), 4)` is the type of arrays of 4 pointers to `char`.

## Generalized Lvalues

Compound expressions, conditional expressions and casts are allowed as *lvalues*. An lvalue is an expression referring to an object, which is a manipulable region of storage; an obvious example of an lvalue expression is an identifier. Such expressions are allowed, provided their operands are lvalues. This means that you can take their addresses or store values into them.

Standard C++ allows compound expressions and conditional expressions as lvalues, and permits casts to reference type, so use of this extension is deprecated for C++ code. For example, a compound expression can be assigned, provided the last expression in the sequence is an lvalue.

The following two expressions are equivalent.

```
(a, b) += 5
a, (b += 5)
```

Similarly, the address of the compound expression can be taken. The following two expressions are equivalent.

```
&(a, b)
a, &b
```

A conditional expression is a valid lvalue if its type is not void and the true and false branches are both valid lvalue. For example, the following two expressions are equivalent.

```
(a ? b :c)=5
(a ? b =5 :(c=5))
```

A cast is a valid lvalue if its operand is an lvalue. A simple assignment whose left-hand side is a cast works by converting the right-hand side first to the specified type, then to the type of the inner left-hand side expression. After this is stored, the

value is converted back to the specified type to become the value of the assignment. Thus, if `a` has type `char *`, the following two expressions are equivalent.

```
(int)a = 5
(int)(a = (char *) (int)5)
```

An assignment-with-arithmetic operation such as `+=` applied to a cast performs the arithmetic using the type resulting from the cast, and then continues as in the previous case. Therefore, the following two expressions are equivalent.

```
(int)a += 5
(int)(a = (char *) (int) ((int)a + 5))
```

You cannot take the address of an lvalue cast, because the use of its address would not work out coherently. Suppose that `&(int)f` were permitted, where `f` has type `float`. Then the following statement would try to store an integer bit-pattern where a floating point number belongs.

```
*&(int)f = 1;
```

This is quite different from what `(int)f = 1` would do—that would convert 1 to floating point and store it. Rather than cause this inconsistency, it is better to prohibit use of `&` on a cast. If you really do want an `int *` pointer with the address of `f`, you can simply write `(int *)&f`

## Conditionals with Omitted Operands

The middle operand in a conditional expression may be omitted. Then if the first operand is nonzero, its value is the value of the conditional expression. Therefore, the following expression has the value of `x` if that is nonzero; otherwise, the value of `y`)

```
x ? :y
```

This example is perfectly equivalent to the following.

```
x ?x:y
```

In this simple case, the ability to omit the middle operand is not especially useful. When it becomes useful is when the first operand does, or may (if it is a macro argument), contain a side effect. Then repeating the operand in the middle would perform the side effect twice. Omitting the middle operand uses the value already computed without the undesirable effects of recomputing it.

## Double-word Integers

GCC supports data types for integers that are twice as long as `int`. Simply write `long long int` for a signed integer, or `unsigned long long int` for an unsigned integer. To make an integer constant of type, `long long int`, add the suffix `LL` to the integer. To make an integer constant of type, `unsigned long long int`, add the suffix `ULL` to the integer.

You can use these types in arithmetic like any other integer types. Addition, subtraction, and bitwise Boolean operations on these types are open-coded on all types of machines. Multiplication is open-coded if the machine supports fullword-to-doubleword (a widening multiply instruction). Division and shifts are open-coded only on machines that provide special support.

The operations that are not open-coded use special library routines that come with GCC.

There may be pitfalls when you use `long long` types for function arguments, unless you declare function prototypes. If a function expects type `int` for its argument, and you pass a value of type `long long int`, confusion will result because the caller and the subroutine will disagree about the number of bytes for the argument. Likewise, if the function expects `long long int` and you pass `int`. The best way to avoid such problems is to use prototypes.

## Complex Numbers

GCC supports complex data types. You can declare both complex integer types and complex floating types, using the keyword, `__complex__`. For example, `__complex__double x;` declares `x` as a variable whose real part and imaginary part are both of type `double`.

`__complex__short int y;` declares `y` to have real and imaginary parts of type `short int`; this is not likely to be useful, but it shows that the set of complex types is complete.

To write a constant with a complex data type, use the suffix `i` or `j` (either one; they are equivalent). For example, `2.5fi` has type, `__complex__float`, and `3i` has type, `__complex__int`. Such a constant always has a pure imaginary value, but you can form any complex value you like by adding one to a real constant.

To extract the real part of a complex-valued expression `exp`, write `__real__exp`. Likewise, use `__imag__` to extract the imaginary part.

The tilde operator, (`~`), performs complex conjugation when used on a value with a complex type.

GCC can allocate complex automatic variables in a noncontiguous fashion; it's even possible for the real part to be in a register while the imaginary part is on the stack (or vice-versa). None of the supported debugging info formats has a way to represent noncontiguous allocation like this, so GCC describes a noncontiguous complex variable as if it were two separate variables of noncomplex type. If the variable's actual name is `foo`, the two fictitious variables are named `foo$real` and `foo$imag`. You can examine and set these two fictitious variables with your debugger. A future version of GDB will recognize such pairs and treat them as a single variable with a complex type.

# Arrays of Length Zero

Zero-length arrays are allowed in GCC. They are very useful as the last element of a structure which is really a header for a variable-length object.

```
struct line {
 int length;
 char contents[0];
};

{
 struct line *thisline = (struct line *)
 malloc (sizeof (struct line) + this_length);
 thisline->length = this_length;
}
```

In standard C, you would have to give `contents` a length of 1, which means either you waste space or complicate the argument to `malloc`.

## Arrays of Variable Length

Variable-length automatic arrays are allowed in GCC. These arrays are declared like any other automatic arrays, but with a length that is not a constant expression. The storage is allocated at the point of declaration and deallocated when the brace-level is exited. Use the following for example.

```
FILE *
concat_fopen (char *s1, char *s2, char *mode)
{
 char str[strlen (s1) + strlen (s2) + 1];
 strcpy (str, s1);
 strcat (str, s2);
 return fopen (str, mode);
}
```

Jumping or breaking out of the scope of the array name deallocates the storage.

Jumping into the scope is not allowed; you get an error message for it. You can use the function `alloca` to get an effect much like variable-length arrays. The function `alloca` is available in many other C implementations (but not in all).

On the other hand, variable-length arrays are more elegant.

There are other differences between these two methods. Space allocated with `alloca` exists until the containing *function* returns.

The space for a variable-length array is deallocated as soon as the array name's scope ends. (If you use both variable-length arrays and `alloca` in the same function, deallocation of a variable-length array will also deallocate anything more recently allocated with `alloca`.)

You can also use variable-length arrays as arguments to functions, as in the following example.

```
struct entry
tester (int len, char data[len][len])
{
 ...
}
```

The length of an array is computed once when the storage is allocated and is remembered for the scope of the array in case you access it with `sizeof`. If you want to pass the array first and the length afterward, you can use a forward declaration in the parameter list—another GNU extension.

```
struct entry
tester (int len; char data[len][len], int len)
{
 ...
}
```

The `int len` before the semicolon is a *parameter forward declaration*, and it serves the purpose of making the name `len` known when the declaration of `data` is parsed. You can write any number of such parameter forward declarations in the parameter list. They can be separated by commas or semicolons, but the last one must end with a semicolon, which is followed by the “*real*” parameter declarations. Each forward declaration must match a “*real*” declaration in parameter name and data type.

## Macros with Variable Numbers of Arguments

In GCC, a macro can accept a variable number of arguments, much as a function can. The syntax for defining the macro looks much like that used for a function. The following is an example.

```
#define eprintf(format, args...) \
 fprintf (stderr, format , ## args)
```

`args` is a rest argument: it takes in zero or more arguments, as many as the call contains. All of them plus the commas between them form the value of `args`, which is substituted into the macro body where `args` is used. Thus, there is the following expansion.

```
eprintf ("%s:%d: ", input_file_name, line_number)
?
fprintf (stderr, "%s:%d: " , input_file_name, line_number)
```

**IMPORTANT!** The comma after the string constant comes from the definition of `eprintf`, whereas the last comma comes from the value of `args`.

The reason for using `##` is to handle the case when `args` matches no arguments at all. In this case, `args` has an empty value. In this case, the second comma in the definition becomes an embarrassment: if it got through to the expansion of the macro, you would get something like the following output.

```
fprintf (stderr, "success!\n" ,)
```

The previous example shows invalid C syntax. `##` gets rid of the comma, so, instead, you get the following.

```
fprintf (stderr, "success!\n")
```

This is a special feature of the GCC preprocessor: `##` before a rest argument that is empty discards the preceding sequence of non-whitespace characters from the macro definition. If another macro argument precedes, none of it is discarded. It might be better to discard the last preprocessor token instead of the last preceding sequence of non-whitespace characters; in fact, there is hope to change this feature someday to do so. For now, use the macro definition so that the preceding sequence of non-whitespace characters is just a single token, and so that the meaning will not change if there is a re-definition of this feature.

## Non-lvalue Arrays May Have Subscripts

Subscripting is allowed on arrays that are not lvalues, even though the unary `&` operator is not. For example, the following declaration is valid in GCC though not valid in other C dialects.

```
struct foo {int a[4];};
```

```
struct foo f();
```

```
bar (int index)
{
 return f().a[index];
}
```

## Arithmetic on `void` Pointers and Function pointers

In GCC, addition and subtraction operations are supported on pointers to `void` and on pointers to functions. This is done by treating the size of a `void` or of a function as 1. A consequence of this is that `sizeof` is also allowed on `void` and on function types, and returns 1. The option, `-Wpointer-arith`, requests a warning if these extensions are used.

# Non-constant Initializers

As in standard C++, the elements of an aggregate initializer for an automatic variable are not required to be constant expressions in GCC. The following is an example of an initializer with run-time varying elements:

```
foo (float f, float g)
{
 float beat_freqs[2] = { f-g, f+g };
 ...
}
```

# Constructor Expressions

GCC supports constructor expressions. A constructor looks like a cast containing an initializer. Its value is an object of the type specified in the cast, containing the elements specified in the initializer.

Usually, the specified type is a structure. Assume that `struct foo` and `structure` are declared as shown in the following declaration.

```
struct foo {int a; char b[2];} structure;
```

The following is an example of constructing a `struct foo` with a constructor.

```
structure = ((struct foo) {x + y, 'a', 0});
```

The previous declaration is equivalent to writing the following input.

```
{
 struct foo temp = {x + y, 'a', 0};
 structure = temp;
}
```

You can also construct an array. If all the elements of the constructor are (made up of) simple constant expressions, suitable for use in initializers, then the constructor is an lvalue and can be coerced to a pointer to its first element, as shown in the following.

```
char **foo = (char *[]) { "x", "y", "z" };
```

Array constructors whose elements are not simple constants are not very useful, because the constructor is not an lvalue. There are only two valid ways to use it: to subscript it, or initialize an array variable with it. The former is probably slower than a `switch` statement, while the latter does the same thing an ordinary C initializer would do. The following is an example of subscripting an array constructor.

```
output = ((int[]) { 2, x, 28 }) [input];
```

Constructor expressions for scalar types and union types are also allowed, but then the constructor expression is equivalent to a cast.

# Labeled Elements in Initializers

Standard C requires the elements of an initializer to appear in a fixed order, the same as the order of the elements in the array or structure being initialized. In GCC you can give the elements in any order, specifying the array indices or structure field names they apply to. This extension is not implemented in GCC++. To specify an array index, use `[index]` or `[index]=` before the element value. Use the following example.

```
int a[6] = { [4] 29, [2] = 15 };
```

The previous specification is equivalent to the following.

```
int a[6] = { 0, 0, 15, 0, 29, 0 };
```

The index values must be constant expressions, even if the array being initialized is automatic. To initialize a range of elements to the same value, use `[ first...last]=value`) Use the following example's input.

```
int widths[] = { [0 ... 9] = 1, [10 ... 99] = 2, [100] = 3 };
```

**IMPORTANT!** The length of the array is the highest value specified + 1.

In a structure initializer, specify the name of a field to initialize with `fieldname:` before the element value.

For example, given `struct point { int x, y; };`, the initialization, `struct point p = { y: yvalue, x: xvalue };` is equivalent to the following input:

```
point p = { xvalue, yvalue struct };
```

Another syntax which has the same meaning is `.fieldname=` as in the following input:

```
struct point p = { .y = yvalue, .x = xvalue };
```

You can also use an element label (with either the colon syntax or the period-equal syntax) when initializing a union, to specify which element of the union should be used. For example, the following will convert 4 to a double to store it in the union using the second line's element.

```
union foo { int i; double d; };
```

```
union foo f = { d: 4 };
```

By contrast, casting 4 to type `union foo` would store it into the union as the integer `i`, since it is an integer. (See “Cast to a Union Type” on page 161.)

You can combine this technique of naming elements with ordinary C initialization of successive elements. Each initializer element that does not have a label applies to the next consecutive element of the array or structure. For example, `int a[6] = { [1] = v1, v2, [4] = v4 };` is equivalent to the following input.

```
int a[6] = { 0, v1, v2, 0, v4, 0 };
```

Labeling the elements of an array initializer is especially useful when the indices are characters or belong to an `enum` type, as in the following example.

```
int whitespace[256] = { [' '] = 1, ['\t'] = 1, ['\h'] = 1, \
 ['\f'] = 1, ['\n'] = 1, ['\r'] = 1 };
```

## Case Ranges

You can specify a range of consecutive values in a single case label, like `case low ... high:`. This has the same effect as the proper number of individual `case` labels, one for each integer value from `low` to `high`, inclusive. This feature is especially useful for ranges of ASCII character codes, as in a `case A ... Z:` command.

**WARNING!** Use spaces around the `...`, for otherwise it may be parsed wrong when you use it with integer values. For example, use a `case 1 ... 5:` command rather than a `case 1...5:` command.

## Cast to a Union Type

A cast to union type is similar to other casts, except that the type specified is a union type. You can specify the type either with `union tag` or with a typedef name. A cast to union is actually a constructor though, not a cast, and hence does not yield an lvalue like normal casts. (See “Constructor Expressions” on page 159.)

The types that may be cast to the union type are those of the members of the union. Thus, given the following union and variables, both `x` and `y` can be cast to type `union foo`.

```
union foo { int i; double d; };
int x;
double y;
```

Using the cast as the right-hand side of an assignment to a variable of union type is equivalent to storing in a member of the union, like the following example shows.

```
union foo u;
...
u = (union foo) x ≡ u.i = x
u = (union foo) y ≡ u.d = y
```

You can also use the union cast as a function argument:

```
void hack (union foo);
...
hack ((union foo) x);
```

# Declaring Attributes of Functions

In GCC, you declare certain things about functions called in your program which help the compiler optimize function calls and check your code more carefully. The keyword, `__attribute__`, allows you to specify special attributes when making a

declaration. This keyword is followed by an attribute specification inside double parentheses. Eight attributes, `noreturn`, `const`, `format`, `section`, `constructor`, `destructor`, `unused` and `weak` are currently defined for functions. Other attributes, including `section` are supported for variables declarations (see “Specifying Attributes of Variables” on page 169) and for types (see “Specifying Attributes of Types” on page 173).

You may also specify attributes with `__` double underscores, preceding and following each keyword. This allows you to use them in header files without being concerned about a possible macro of the same name. For example, you may use `__noreturn__` instead of `noreturn`.

#### `noreturn`

A few standard library functions, such as `abort` and `exit`, cannot return. GCC knows this automatically. Some programs define their own functions that never return. You can declare them `noreturn` to tell the compiler this fact. For example,

```
void fatal () __attribute__ ((noreturn));

void
fatal (...)
{
 ... /* Print error message.*/ ...
 exit (1);
}
```

The `noreturn` keyword tells the compiler to assume that `fatal` cannot return. It can then optimize without regard to what would happen if `fatal` ever did return. This makes slightly better code. More importantly, it helps avoid spurious warnings of uninitialized variables.

Do not assume that registers saved by the calling function are restored before calling the `noreturn` function. It does not make sense for a `noreturn` function to have a return type other than `void`. The attribute `noreturn` is not implemented in GCC versions earlier than 2.5.

An alternative way to declare that a function does not return, which works in the current version and in some older versions, is as follows:

```
typedef void voidfn ();
volatile voidfn fatal;
```

#### `const`

Many functions do not examine any values except their arguments, and have no effects except the return value. Such a function can be subject to common subexpression elimination and loop optimization just as an arithmetic operator would be. These functions should be declared with the attribute, `const`. For example, the following says that the hypothetical function, `square`, is safe to call fewer times than the program says.

```
int square (int) __attribute__ ((const));
```

The attribute `const` is not implemented in GCC versions earlier than 2.5. An alternative way to declare that a function has no side effects, which works in the current version and in some older versions, is as follows:

```
typedef int intfn ();

extern const intfn square;
```

This approach does not work in GCC++ from 2.6.0 on, since the language specifies that `const` must be attached to the return value.

**IMPORTANT!** A function that has pointer arguments and examines the data pointed to must not be declared `const`. Likewise, a function that calls a non-`const` function usually must not be `const`. It does not make sense for a `const` function to return `void`.

`format (archetype, string-index, first-to-check)`

The `format` attribute specifies that a function takes `printf` or `scanf` style arguments which should be type-checked against a format string.

For example, the following declaration causes the compiler to check the arguments in calls to `my_printf` for consistency with the `printf` style format string argument `my_format`.

```
extern int
my_printf (void *my_object, const char *my_format, ...)
 __attribute__((format (printf, 2, 3)));
```

The parameter `archetype` determines how the format string is interpreted, and should be either `printf` or `scanf`.

The parameter `string-index` specifies which argument is the format string argument (starting from 1), while `first-to-check` is the number of the first argument to check against the format string. For functions where the arguments are not available to be checked (such as `vprintf`), specify the third parameter as zero. In this case the compiler only checks the format string for consistency.

In the previous example, the format string (`my_format`) is the second argument of the function, `my_print`, and the arguments to check start with the third argument, so the correct parameters for the `format` attribute are 2 and 3.

The `format` attribute allows you to identify your own functions which take format strings as arguments, so that GCC can check the calls to these functions for errors. The compiler always checks formats for the ANSI library functions, `printf`, `fprintf`, `sprintf`, `scanf`, `fscanf`, `sscanf`, `vprintf`, `vfprintf` and `vsprintf` whenever such warnings are re-requested (using `-Wformat`), so there is no need to modify the header file, `stdio.h`.

`format_arg` (*string-index*)

The `format_arg` attribute specifies that a function takes `printf` or `scanf` style arguments, modifies it (for example, to translate it into another language), and passes it to a `printf` or `scanf` style function. For example, the following declaration causes the compiler to check the arguments in calls to `my_dgettext`, whose result is passed to a `printf` or `scanf` type function for consistency with the `printf` style format string argument, `my_format`.

```
extern char *
my_dgettext (char *my_domain, const char *my_format)
 __attribute__ ((format_arg (2)));
```

The parameter, *string-index*, specifies which argument is the format string argument (starting from 1).

The `format-arg` attribute allows you to identify your own functions which modify format strings, so that GCC can check the calls to `printf` and `scanf` function whose operands are a call to one of your own function. The compiler always treats `gettext`, `dgettext`, and `dcgettext` in this manner.

`section` ("section-name")

Normally, the compiler places the code it generates in the `text` section. Sometimes, however, you need additional sections, or you need certain particular functions to appear in special sections. The `section` attribute specifies that a function lives in a particular section. For example, the following declaration puts the function, `foobar`, in the `bar` section.

```
extern void foobar (void) __attribute__ \
 ((section ("bar")));
```

Some file formats do not support arbitrary sections so the `section` attribute is not available on all platforms. If you need to map the entire contents of a module to a particular section, consider using the facilities of the linker instead.

`constructor`

`destructor`

The `constructor` attribute causes the function to be called automatically before execution enters `main()`. Similarly, the `destructor` attribute causes the function to be called automatically after `main()` has completed or `exit()` has been called. Functions with these attributes are useful for initializing data that will be used implicitly during the execution of the program. These attributes are not currently implemented for Objective C.

`unused`

This attribute, attached to a function, means that the function is meant to be possibly unused. GCC will not produce a warning for this function. GCC++ does not currently support this attribute as definitions without parameters are valid in C++.

**weak**

The `weak` attribute causes the declaration to be emitted as a weak symbol rather than a global. This is primarily useful in defining library functions which can be overridden in user code, though it can also be used with non-function declarations. Weak symbols are supported for ELF targets, and also for a.out targets when using the GNU assembler and linker.

**alias ("target")**

The `alias` attribute causes the declaration to be emitted as an alias for another symbol, which must be specified. For instance, the following declares `f` to be a weak alias for `__f`. In C++, the mangled name for the target must be used.

```
void __f () { /* do something */; }
void f () __attribute__ ((weak, alias ("__f")));
```

Not all target machines support this attribute.

**regparm (number)**

On the Intel 386, the `regparm` attribute causes the compiler to pass up to *number* integer arguments in registers *EAX*, *EDX*, and *ECX* instead of on the stack.

Functions that take a variable number of arguments will continue to be passed all of their arguments on the stack.

**stdcall**

On the Intel 386, the `stdcall` attribute causes the compiler to assume that the called function will pop off the stack space used to pass arguments, unless it takes a variable number of arguments. The PowerPC compiler for Windows NT currently ignores the `stdcall` attribute.

**cdecl**

On the Intel 386, the `cdecl` attribute causes the compiler to assume that the called function will pop off the stack space used to pass arguments. This is useful to override the effects of the `switch`, `-mrtD`. The PowerPC compiler for Windows NT currently ignores the `cdecl` attribute.

**longcall**

On the RS/6000 and PowerPC, the `longcall` attribute causes the compiler to always call the function via a pointer, so that functions which reside further than 64 megabytes (67,108,864 bytes) from the current location can be called.

**dllimport**

On the PowerPC running Windows NT, the `dllimport` attribute causes the compiler to call the function via a global pointer to the function pointer that is set up by the Windows NT dll library. The pointer name is formed by combining `__imp_` and the function name.

**dllexport**

On the PowerPC running Windows NT, the `dllexport` attribute causes the compiler to provide a global pointer to the function pointer, so that it can be called with the `dllimport` attribute. The pointer name is formed by combining `__imp_` and the function name.

`exception (except-func[, except-arg])`

On the PowerPC running Windows NT, the `exception` attribute causes the compiler to modify the structured exception table entry it emits for the declared function. The string or identifier, `except-func`, is placed in the third entry of the structured exception table. It represents a function which is called by the exception handling mechanism if an exception occurs. If it was specified, the string or identifier, `except-arg`, is placed in the fourth entry of the structured exception table.

`function_vector`

Use this option on the H8/300 and H8/300H to indicate that the specified function should be called through the function vector. Calling a function through the function vector will reduce code size, however; the function vector has a limited size (maximum 128 entries on the H8/300 and 64 entries on the H8/300H) and shares space with the interrupt vector.

`interrupt_handler`

Use this option on the H8/300 and H8/300H to indicate that the specified function is an interrupt handler. The compiler will generate function entry and exit sequences suitable for use in an interrupt handler when this attribute is present.

`eightbit_data`

Use this option on the H8/300 and H8/300H to indicate that the specified variable should be placed into the eight bit data section. The compiler will generate more efficient code for certain operations on data in the eight bit data area. Note the eight bit data area is limited to 256 bytes of data.

`tiny_data`

Use this option on the H8/300H to indicate that the specified variable should be placed into the tiny data section. The compiler will generate more efficient code for loads and stores on data in the tiny data section. Note the tiny data area is limited to slightly under 32kbytes of data.

`interrupt`

Use this option on the M32R/D to indicate that the specified function is an interrupt handler. The compiler will generate a function entry and exit sequences suitable for use in an interrupt handler when this attribute is present.

`model (model-name)`

Use this attribute on the M32R/D to set the addressability of an object, and the code generated for a function. The identifier, `model-name`, is one of `small`, `medium`, or `large`, representing each of the code models.

Small model objects live in the lower 16MB of memory (so that their addresses can be loaded with the `ld24` instruction), and are callable with the `bl` instruction.

Medium model objects may live anywhere in the 32 bit address space (the compiler will generate `seth/add3` instructions to load their addresses), and are

callable with the `bl` instruction.

You can specify multiple attributes in a declaration by separating them by commas within the double parentheses or by immediately following an attribute declaration with another attribute declaration.

Some people object to the `__attribute__` feature, suggesting that ANSI C's `#pragma` should be used instead; there are two reasons for not doing this:

- It is impossible to generate `#pragma` commands from a macro.
- There is no telling what the same `#pragma` might mean in another compiler.

These two reasons apply to almost any application that might be proposed for `#pragma`. It is basically a mistake to use `#pragma` for anything.

## Prototypes and Old-style Function Definitions

GCC extends ANSI C to allow a function prototype to override a later old-style non-prototype definition. Consider the following example.

```

/* Use prototypes unless the compiler is old-fashioned. */
#if __STDC__
#define P(x) x
#else
#define P(x) ()
#endif

/* Prototype function declaration. */
int isroot P((uid_t));

/* Old-style function definition. */
int isroot (x) /* ??? lossage here ??? */
 uid_t x;
{
 return x == 0;
}

```

Suppose the type `uid_t` happens to be `short`. ANSI C does not allow this example, because subword arguments in old-style non-prototype definitions are promoted. Therefore in this example the function definition's argument is really an `int`, which does not match the prototype argument type of `short`.

This restriction of ANSI C makes it hard to write code that is portable to traditional C compilers, because the programmer does not know whether the `uid_t` type is `short`, `int`, or `long`.

Therefore, in cases like these GCC allows a prototype to override a later old-style definition. More precisely, in GCC, a function prototype argument type overrides the

argument type specified by a later old-style definition if the former type is the same as the latter type before promotion. Thus, in GCC, the previous example is equivalent to the following declaration.

```
int isroot (uid_t);
int
isroot (uid_t x)
{
 return x == 0;
}
```

GNU C++ does not support old-style function definitions, so the previous example's extension is irrelevant.

## Compiling Functions for Interrupt Calls

When compiling code for certain platforms (currently the Hitachi H8/300 and the Tandem ST-2000), you can instruct {No value for "GCC"} that certain functions are meant to be called from hardware interrupts.

To mark a function as callable from interrupt, include the line `#pragma interrupt` somewhere before the beginning of the function's definition. (For maximum readability, you might place it immediately before the definition of the appropriate function.) `#pragma interrupt` will affect only the next function defined; if you want to define more than one function with this property, include `#pragma interrupt` before each of them.

When you define a function with `#pragma interrupt`, {No value for "GCC"} alters its usual calling convention, to provide the right environment when the function is called from an interrupt. *Such functions cannot be called in the usual way from your program.*

You must use other facilities to actually associate these functions with particular interrupts; {No value for "GCC"} can only compile them in the appropriate way.

### C++ style Comments

In GCC, you may use C++ style comments, which start with `//` and continue until the end of the line. Many other C implementations allow such comments, and they are likely to be in a future C standard. However, C++ style comments are not recognized if you specify `-ansi` or `-traditional`, since they are incompatible with traditional constructs like `dividend/*comment*/divisor`.

### Dollar Signs in Identifier Names

In GCC, you may use dollar signs in identifier names. This is because many traditional C implementations allow such identifiers. However, dollar signs are not

supported on a few target machines, typically because the target assembler does not allow them.

## The `\ESC` Character in Constants

You can use the sequence, `\e`, in a string or character constant to stand for the ASCII character, `ESC`.

# Inquiring on Alignment of Types or Variables

The keyword `__alignof__` allows you to inquire about how an object is aligned, or the minimum alignment usually required by a type. Its syntax is just like `sizeof`.

For example, if the target machine requires a double value to be aligned on an 8-byte boundary, then `__alignof__ (double)` is 8. This is true on many RISC machines.

On more traditional machine designs, `__alignof__ (double)` is 4 or even 2.

Some machines never actually require alignment; they allow reference to any data type even at an odd addresses. For these machines, `__alignof__` reports the *recommended* alignment of a type.

When the operand of `__alignof__` is an lvalue rather than a type, the value is the largest alignment that the lvalue is known to have. It may have this alignment as a result of its data type, or because it is part of a structure and inherits alignment from that structure. For example, after this declaration:

```
struct foo { int x; char y; } fool;
```

The value of `__alignof__ (fool.y)` is probably 2 or 4, the same as `__alignof__ (int)`, even though the data type of `fool.y` does not itself demand any alignment. A related feature which lets you specify the alignment of an object is `__attribute__((aligned ( alignment)))`; see “Specifying Attributes of Variables” on page 169.

# Specifying Attributes of Variables

The keyword, `__attribute__`, allows you to specify special attributes of variables or structure fields. This keyword is followed by an at-tribute specification inside double parentheses. Eight attributes are currently defined for variables: `aligned`, `mode`, `nocommon`, `packed`, `section`, `transparent_union`, `unused`, and `weak`. Other attributes are available for functions (see “Declaring Attributes of Functions” on page 161) and for types (see “Specifying Attributes of Types” on page 173). You may also specify attributes with `__` preceding and following each keyword. This allows you to use them

in header files without being concerned about a possible macro of the same name. For example, you may use `__aligned__` instead of `aligned`.

`aligned (alignment)`

This attribute specifies a minimum alignment for the variable or structure field, measured in bytes. For example, the following declaration causes the compiler to allocate the global variable `x` on a 16-byte boundary.

```
int x __attribute__ ((aligned (16))) = 0;
```

On a 68040, this could be used in conjunction with an `asm` expression to access the `move16` instruction which requires 16-byte aligned operands. You can also specify the alignment of structure fields.

For example, to create a double-word aligned `int` pair, you could write:

```
struct foo { int x[2] __attribute__ ((aligned (8))); };
```

This is an alternative to creating a union with a `double` member that forces the union to be double-word aligned. It is not possible to specify the alignment of functions; the alignment of functions is determined by the machine's requirements and cannot be changed.

You cannot specify alignment for a `typedef` name because such a name is just an alias, not a distinct type.

As in the preceding examples, you can explicitly specify the alignment (in bytes) that you wish the compiler to use for a given variable or structure field.

Alternatively, you can leave out the alignment factor and just ask the compiler to align a variable or field to the maximum useful alignment for the target machine you are compiling for. For example, you could write:

```
short array[3] __attribute__ ((aligned));
```

Whenever you leave out the alignment factor in an `aligned` attribute specification, the compiler automatically sets the alignment for the declared variable or field to the largest alignment which is ever used for any data type on the target machine you are compiling for. Doing this can often make copy operations more efficient, because the compiler can use whatever instructions copy the biggest chunks of memory when performing copies to or from the variables or fields that you have aligned this way.

The `aligned` attribute can only increase the alignment; but you can decrease it by specifying `packed` as well. (See attribute specifications for `packed`.)

**IMPORTANT!** The effectiveness of `aligned` attributes may be limited by inherent limitations in your linker. On many systems, the linker is only able to arrange for variables to be aligned up to a certain maximum alignment. (For some linkers, the maximum supported alignment may be very small.)

If your linker is only able to align variables up to a maximum of 8 byte alignment, then specifying `aligned(16)` in an `__attribute__` will still only provide you with 8 byte alignment.

See *Using LD in GNUPro Utilities* for further information.

`mode (mode)`

This attribute specifies the data type for the declaration— whichever type corresponds to the mode *mode*. This in effect lets you request an integer or floating point type according to its width.

You may also specify a mode of `byte` or `__byte__` to indicate the mode corresponding to a one-byte integer, `word` or `__word__` for the mode of a one-word integer, and `pointer` or `__pointer__` for the mode used to represent pointers.

`nocommon`

This attribute specifies requests GCC not to place a variable “*common*” but instead to allocate space for it directly. If you specify the `-fno-common` option, GCC will do this for all variables.

Specifying the `nocommon` attribute for a variable provides an initialization of zeros. A variable may only be initialized in one source file.

`packed`

The `packed` attribute specifies that a variable or structure field should have the smallest possible alignment—one byte for a variable, and one bit for a field, unless you specify a larger value with the `aligned` attribute. The following example is a structure in which the field `x` is packed, so that it immediately follows `a`:

```
struct foo
{
 char a;
 int x[2] __attribute__((packed));
};
```

`section ("section-name")`

Normally, the compiler places the objects it generates in sections like `data` and `bss`. Sometimes, however, you need additional sections, or you need certain particular variables to appear in special sections, for example to map to special hardware.

The `section` attribute specifies that a variable (or function) lives in a particular section. For example, the following small program uses several specific `section` names.

```
struct duart a __attribute__((section("DUART_A"))) = { 0
};
struct duart b __attribute__((section("DUART_B"))) = { 0
};
char stack[10000] __attribute__((section("STACK"))) = {0 };
int init_data_copy __attribute__((section("INITDATACOPY")))
= 0;
main()
{
```

```
/* Initialize stack pointer */
init_sp (stack + sizeof (stack));

/* Initialize initialized data */
memcpy (&init_data_copy, &data, &edata - &data);

/* Turn on the serial ports */
init_duart (&a); init_duart (&b);
}
```

Use the `section` attribute with an *initialized* definition of a *global* variable, as shown in the previous example. GCC issues a warning and otherwise ignores the `section` attribute in uninitialized variable declarations.

You may only use the `section` attribute with a fully initialized global definition because of the way linkers work. The linker requires each object be defined once, with the exception that uninitialized variables tentatively go in the `common` (or `bss`) section and can be multiply-defined. You can force a variable to be initialized with the `-fno-common` option or the `nocommon` attribute. Some file formats do not support arbitrary sections so the `section` attribute is not available on all platforms. If you need to map the entire contents of a module to a particular section, consider using the facilities of the linker instead.

#### `transparent_union`

This attribute, attached to a function parameter which is a union, means that the corresponding argument may have the type of any union member, but the argument is passed as if its type were that of the first union member. For more details, see “Specifying Attributes of Types” on page 173. You can also use this attribute on a `typedef` for a union data type; then it applies to all function parameters with that type.

#### `unused`

This attribute, attached to a variable, means that the variable is meant to be possibly unused. GCC will not produce a warning for this variable.

#### `weak`

See the descriptions for the `weak` attribute with “Declaring Attributes of Functions” on page 161.

#### `model ( model-name )`

Use this attribute on the M32R/D to set the addressability of an object. The identifier `model-name` is one of `small`, `medium`, or `large`, representing each of the code models. Small model objects live in the lower 16MB of memory (so that their addresses can be loaded with the `ld24` instruction). Medium and large model objects may live anywhere in the 32 bit address space (the compiler will generate `seth/add3` instructions to load their addresses). To specify multiple attributes, separate them by commas within double parentheses; use `__attribute__((aligned (16), packed))`, for example.

# Specifying Attributes of Types

The keyword `__attribute__` allows you to specify special attributes of `struct` and `union` types when you define such types. This keyword is followed by an attribute specification inside double parentheses. Three attributes are currently defined for types: `aligned`, `packed`, and `transparent_union`. Other attributes are defined for functions (see “Declaring Attributes of Functions” on page 161) and for variables (see “Specifying Attributes of Variables” on page 169).

You may also specify any one of these attributes with `__` preceding and following its keyword. This allows you to use these attributes in header files without being concerned about a possible macro of the same name. For example, you may use `__aligned__` instead of `aligned`.

You may specify the `aligned` and `transparent_union` attributes either in a `typedef` declaration or just past the closing curly brace of a complete `enum`, `struct` or `union` type definition and the `packed` attribute only past the closing brace of a definition.

`aligned` (*alignment*)

This attribute specifies a minimum alignment (in bytes) for variables of the specified type. For example, the following declarations force the compiler to insure (as fast as it can) that each variable whose type is `struct S` or `more_aligned_int` will be allocated and aligned *at least* on an 8-byte boundary.

```
struct S { short f[3]; } __attribute__((aligned (8)));
typedef int more_aligned_int __attribute__((aligned (8)));
```

On a Sparc, having all variables of type `struct S` aligned to 8-byte boundaries allows the compiler to use the `ldd` and `std` (doubleword load and store) instructions when copying one variable of type `struct S` to another, thus improving run-time efficiency.

**IMPORTANT!** The alignment of any given `struct` or `union` type is required by the ANSI C standard to be at least a perfect multiple of the lowest common multiple of the alignments of all of the members of the `struct` or `union` in question. This means that you can effectively adjust the alignment of a `struct` or `union` type by attaching an `aligned` attribute to any one of the members of such a type, but the notation illustrated in the last example is a more obvious, intuitive, and readable way to request the compiler to adjust the alignment of an entire `struct` or `union` type.

As in the preceding example, you can explicitly specify the alignment (in bytes) that you wish the compiler to use for a given `struct` or `union` type. Alternatively, you can leave out the alignment factor and just ask the compiler to align a type to the maximum alignment for the target machine for which you’re compiling:

```
struct S { short f[3]; } __attribute__((aligned));
```

Whenever you leave out the alignment factor in an `aligned` attribute

specification, the compiler automatically sets the alignment for the type to the largest alignment which is ever used for any data type on the target machine you are compiling for. Doing this can often make copy operations more efficient, because the compiler can use whatever instructions copy the biggest chunks of memory when performing copies to or from the variables which have types that you have aligned this way. In the example above, if the size of each `short` is 2 bytes, then the size of the entire `struct S` type is 6 bytes. The smallest power of two which is greater than or equal to that is 8, so the compiler sets the alignment for the entire `struct S` type to 8 bytes.

**IMPORTANT!** Although you can ask the compiler to select a time-efficient alignment for a given type and then declare only individual stand-alone objects of that type, the compiler's ability to select a time-efficient alignment is primarily useful only when you plan to create arrays of variables having the relevant (efficiently aligned) type. If you declare or use arrays of variables of an efficiently-aligned type, then it is likely that your program will also be doing pointer arithmetic (or subscripting, which amounts to the same thing) on pointers to the relevant type, and the code that the compiler generates for these pointer arithmetic operations will often be more efficient for efficiently-aligned types than for other types.

The `aligned` attribute can only increase the alignment; but you can decrease it by specifying `packed` as well. See the following discussion for `packed`.

**IMPORTANT!** The effectiveness of `aligned` attributes may be limited by inherent limitations in your linker. On many systems, the linker is only able to arrange for variables to be aligned up to a certain maximum alignment. (For some linkers, the maximum supported alignment may be very very small.) If your linker is only able to align variables up to a maximum of 8 byte alignment, then specifying `aligned(16)` in an `__attribute__` will still only provide you with 8 byte alignment. See *Using LD in GNUPro Utilities* for more information on alignment.

#### `packed`

This attribute, attached to an `enum`, `struct`, or `union` type definition, specified that the minimum required memory be used to represent the type. Specifying this attribute for `struct` and `union` types is equivalent to specifying the `packed` attribute on each of the structure or union members.

Specifying the `-fshort-enums` option on the line is equivalent to specifying the `packed` attribute on all `enum` definitions.

You may only specify this attribute after a closing curly brace on an `enum` definition, not in a `typedef` declaration, unless that declaration also contains the definition of the `enum`.

`transparent_union`

This attribute, attached to a `union` type definition, indicates that any function parameter having that `union` type causes calls to that function to be treated in a special way.

First, the argument corresponding to a `transparent union` type can be of any type in the union; no cast is required. Also, if the union contains a pointer type, the corresponding argument can be a null pointer constant or a void pointer expression; and if the union contains a void pointer type, the corresponding argument can be any pointer expression. If the union member type is a pointer, qualifiers like `const` on the referenced type must be respected, just as with normal pointer conversions.

Second, the argument is passed to the function using the calling conventions of first member of the transparent union, not the calling conventions of the union itself. All members of the union must have the same machine representation; this is necessary for this argument passing to work properly.

Transparent unions are designed for library functions that have multiple interfaces for compatibility reasons. For example, suppose the `wait` function must accept either a value of type `int *` to comply with Posix, or a value of type `union wait *` to comply with the 4.1 BSD interface. If the `wait` function's parameter were `void *`, `wait` would accept both kinds of arguments, but it would also accept any other pointer type and this would make argument type checking less useful. Instead, `<sys/wait.h>` might define the interface as follows.

```
typedef union
{
 int *__ip;
 union wait *__up;
} wait_status_pointer_t __attribute__ \
 ((__transparent_union__));
pid_t wait (wait_status_pointer_t);
```

This interface allows either `int *` or `union wait *` arguments to be passed, using the `'int *` calling convention. The program can call `wait` with arguments of either of the following types.

```
int w1 () { int w; return wait (&w); }
int w2 () { union wait w; return wait (&w); }
```

With this interface, the `wait` implementation might look like the following example's declaration.

```
pid_t wait (wait_status_pointer_t p)
{
 return waitpid (-1, p.__ip, 0);
}
```

To specify multiple attributes, separate them by commas within the double parentheses, as in the following example.

```
__attribute__ ((aligned (16), packed))
```

# An inline Function Is as Fast as a Macro

By declaring a function `inline`, you can direct GCC to integrate that function's code into the code for its callers. This makes execution faster by eliminating the function-call overhead; in addition, if any of the actual argument values are constant, their known values may permit simplifications at compile time so that not all of the inline function's code needs to be included.

The effect on code size is less predictable; object code may be larger or smaller with function inlining, depending on the particular case. Inlining of functions is an optimization and it really “works” only in optimizing compilation. If you don't use `-O`, no function is really inline.

To declare a function inline, use the `inline` keyword in its declaration, like the following example shows.

```
inline int
inc (int *a)
{
 (*a)++;
}
```

(If you are writing a header file to be included in ANSI C programs, write `__inline__` instead of `inline`. See “Alternate Keywords” on page 196.)

You can also make all “*simple enough*” functions inline with the option `-finline-functions`. Certain usage in a function definition can make it unsuitable for inline substitution.

**IMPORTANT!** In C and Objective C, unlike C++, the `inline` keyword does not affect the linkage of the function.

GCC automatically inlines member functions defined within the class body of C++ programs even if they are not explicitly declared `inline`. (You can override this with `-fno-default-inline`; see “Options Controlling C++ Dialect” on page 25.)

When a function is both `inline` and `static`, if all calls to the function are integrated into the caller, and the function's address is never used, then the function's own assembler code is never referenced. In this case, GCC does not actually output assembler code for the function, unless you specify the option, `-fkeep-inline-functions`. Some calls cannot be integrated for various reasons (in particular, calls that precede the function's definition cannot be integrated, and neither can recursive calls within the definition). If there is a nonintegrated call, then the function is compiled to assembler code as usual. The function must also be compiled as usual if the program refers to its address, because that can't be inlined.

When an `inline` function is not static, then the compiler must assume that there may be calls from other source files; since a global symbol can be defined only once in any program, the function must not be defined in the other source files, so the calls therein cannot be integrated. Therefore, a non-`static` inline function is always compiled on its own in the usual fashion.

If you specify both `inline` and `extern` in the function definition, then the definition is used only for inlining. In no case is the function compiled on its own, not even if you refer to its address explicitly. Such an address becomes an external reference, as if you had only declared the function, and had not defined it.

This combination of `inline` and `extern` has almost the effect of a macro. The way to use it is to put a function definition in a header file with these keywords, and put another copy of the definition (lacking `inline` and `extern`) in a library file. The definition in the header file will cause most calls to the function to be inlined. If any uses of the function remain, they will refer to the single copy in the library.

GCC does not `inline` any functions when not optimizing. It is not clear whether it is better to `inline` or not, in this case, but a correct implementation when not optimizing is difficult. So it is turned off.

## Assembler Instructions with C Expression Operands

In an assembler instruction, using `asm`, you can now specify the operands of the instruction using C expressions. This means no more guessing which registers or memory locations will contain the data you want to use.

You must specify an assembler instruction template much like what appears in a machine description, plus an operand constraint string for each operand.

For example, here is how to use the 68881's `fsinx` instruction:

```
asm ("fsinx %l,%0" : "=f" (result) : "f" (angle));
```

`angle` is the C expression for the input operand while `result` is that of the output operand. Each has `"f"` as its operand constraint, saying that a floating point register is required. The `=` in `=f` indicates that the operand is an output; all output operands constraints must use `=`. The constraints use the same language used in the machine description (see "Constraints for `asm` Operands" on page 181).

Each operand is described by an operand-constraint string followed by the C expression in parentheses. A colon separates the assembler template from the first output operand, and another separates the last output operand from the first input, if any. Commas separate output operands and separate inputs. The total number of operands is limited to ten or to the maximum number of operands in any instruction pattern in the machine description, whichever is greater.

If there are no output operands, and there are input operands, then there must be two consecutive colons surrounding the place where the output operands would go.

Output operand expressions must be lvalues; the compiler can check this. The input operands need not be lvalues. The compiler cannot check whether the operands have data types that are reasonable for the instruction being executed. It does not parse the assembler instruction template and does not know what it means, or whether it is valid assembler input. The extended `asm` feature is most often used for machine instructions that the compiler itself does not know exist. If the output expression cannot be directly addressed (for example, it is a bit field), your constraint must allow a register. In that case, GCC will use the register as the output of the `asm`, and then store that register into the output.

The output operands must be write-only; GCC will assume that the values in these operands before the instruction are dead and need not be generated. Extended `asm` supports input-output or read-write operands.

Use the constraint character, `+`, to indicate such an operand, and list it with the output operands.

When the constraints for the read-write operand (or an operand in which only some of the bits are to be changed) allows a register, you may, as an alternative logically split its function into two separate operands, one input operand and one write-only output operand. The connection between them is expressed by constraints which say they need to be in the same location when the instruction executes. You can use the same C expression for both operands, or different expressions. For example, in the following declaration, the (fictitious) `combine` instruction uses `bar` as its read-only source operand and `foo` as its read-write destination.

```
asm ("combine %2,%0" : "=r" (foo) : "0" (foo), "g" (bar));
```

The constraint `"0"` for operand 1 says that it must occupy the same location as operand 0. A digit in constraint is allowed only in an input operand, and it must refer to an output operand. Only a digit in the constraint can guarantee that one operand will be in the same place as another. The mere fact that `foo` is the value of both operands is not enough to guarantee that they will be in the same place in the generated assembler code. The following declaration would not work.

```
asm ("combine %2,%0" : "=r" (foo) : "r" (foo), "g" (bar));
```

Various optimizations or reloading could cause operands 0 and 1 to be in different registers; GCC knows no reason not to do so. For example, the compiler might find a copy of the value of `foo` in one register and use it for operand 1, but generate the output operand 0 in a different register (copying it afterward to `foo`'s own address). Of course, since the register for operand 1 is not even mentioned in the assembler code, the result will not work, but GCC can't tell that. Some instructions clobber specific hard registers. To describe this, write a third colon after the input operands, followed by the names of the clobbered hard registers (given as strings). The following is a realistic example for the Vax.

```
asm volatile ("movc3 %0,%1,%2"
 : /* no outputs */
 : "g" (from), "g" (to), "g" (count)
 : "r0", "r1", "r2", "r3", "r4", "r5");
```

If you refer to a particular hardware register from the assembler code, then you will probably have to list the register after the third colon to tell the compiler that the register's value is modified.

In many assemblers, the register names begin with %'; to produce one % in the assembler code, you must write %% in the input. If your assembler instruction can alter the condition code register, add `cc` to the list of clobbered registers.

GCC on some machines represents the condition codes as a specific hardware register; `cc` serves to name this register. On other machines, the condition code is handled differently, and specifying `cc` has no effect. But it is valid no matter what the machine. If your assembler instruction modifies memory in an unpredictable fashion, add `memory` to the list of clobbered registers. This will cause GCC to not keep memory values cached in registers across the assembler instruction. You can put multiple assembler instructions together in a single `asm` template, separated either with newlines (written as `\n`) or with semi-colons if the assembler allows such semicolons. The GNU assembler allows semicolons and all UNIX assemblers seem to do so. The input operands are guaranteed not to use any of the clobbered registers, and neither will the output operands addresses, so you can read and write the clobbered registers as many times as you like. The following is an example of multiple instructions in a template, assuming that the subroutine, `_foo`, accepts arguments in registers 9 and 10.

```
asm ("movl %0,r9;movl %1,r10;call _foo"
 : /* no outputs */
 : "g" (from), "g" (to)
 : "r9", "r10");
```

Unless an output operand has the `&` constraint modifier, GCC may allocate it in the same register as an unrelated input operand, on the assumption that the inputs are consumed before the outputs are produced. This assumption may be false if the assembler code actually consists of more than one instruction. In such a case, use `&` for each output operand that may not overlap an input. See “Constraint Modifier Characters” on page 184.

If you want to test the condition code produced by an assembler instruction, you must include a branch and a label in the `asm` construct, as follows.

This assumes your assembler supports local labels, as the GNU assembler and most UNIX assemblers do.

```
asm ("clr %0;frob %1;beq 0f;mov #1,%0;0:"
 : "g" (result)
 : "g" (input));
```

Speaking of labels, jumps from one `asm` to another are not supported. The compiler's optimizers do not know about these jumps, and therefore they cannot take account of them when deciding how to optimize.

Usually the most convenient way to use these `asm` instructions is to encapsulate them in macros that look like functions, as in the following example.

```
#define sin(x) ({ double __value, __arg = (x); \
 asm ("fsinx %1,%0": "=f" (__value): "f" (__arg)); \
 __value; })
```

In the previous example, the variable, `__arg`, is used to make sure that the instruction operates on a proper `double` value, and to accept only those arguments `x` which can convert automatically into a `double`.

Another way to make sure the instruction operates on the correct data type is to use a cast in the `asm`. This is different from using a variable `__arg` in that it converts more different types. For example, if the desired type were `int`, casting the argument to `int` would accept a pointer with no complaint, while assigning the argument to an `int` variable named `__arg` would warn about using a pointer unless the caller explicitly casts it.

If an `asm` has output operands, GCC assumes for optimization purposes that the instruction has no side effects except to change the output operands. This does not mean that instructions with a side effect cannot be used, but you must be careful, because the compiler may eliminate them if the output operands aren't used, or move them out of loops, or replace two with one if they constitute a common subexpression. Also, if your instruction does have a side effect on a variable that otherwise appears not to change, the old value of the variable may be reused later if it happens to be found in a register.

You can prevent an `asm` instruction from being deleted, moved significantly, or combined, by writing the keyword `volatile` after the `asm`.

For example:

```
#define set_priority(x) \
 asm volatile ("set_priority %0": /* no outputs */ : "g" (x))
```

An instruction without output operands will not be deleted or moved significantly, regardless, unless it is unreachable.

**IMPORTANT!** Even a `volatile asm` instruction can be moved in ways that appear insignificant to the compiler, such as across jump instructions. You can't expect a sequence of `volatile asm` instructions to remain perfectly consecutive. If you want consecutive output, use a single `asm`.

It is a natural idea to look for a way to give access to the condition code left by the assembler instruction. However, to implement this, there is no way to make it work reliably. The problem is that output operands might need reloading, which would result in additional following "*store*" instructions. On most machines, these

instructions would alter the condition code before there was time to test it. This problem doesn't arise for ordinary “*test*” and “*compare*” instructions because they don't have any output operands.

If you are writing a header file that should be included in ANSI C programs, write `__asm__` instead of `asm`. See “Alternate Keywords” on page 196.

## Constraints for `asm` Operands

The following details discuss what constraint letters you can use with `asm` operands. Constraints can say whether an operand may be in a register, and which kinds of register; whether the operand can be a memory reference, and which kinds of address; whether the operand may be an immediate constant, and which possible values it may have. Constraints can also require two operands to match.

### Simple Constraints

The simplest kind of constraint is a string full of letters, each of which describes one kind of operand that is permitted. Here are the letters that are allowed:

`m`

A memory operand is allowed, with any kind of address that the machine supports in general.

`o`

A memory operand is allowed, but only if the address is *offsettable*. This means that adding a small integer (actually, the width in bytes of the operand, as determined by its machine mode) may be added to the address and the result is also a valid memory address.

For example, an address which is constant is offsettable; so is an address that is the sum of a register and a constant (as long as a slightly larger constant is also within the range of address-offsets supported by the machine); but an autoincrement or autodecrement address is not offsettable. More complicated indirect/indexed addresses may or may not be offsettable depending on the other addressing modes that the machine supports.

**IMPORTANT!** In an output operand which can be matched by another operand, the constraint letter `o` is valid only when accompanied by both `<` (if the target machine has predecrement addressing) and `>` (if the target machine has preincrement addressing).

`v`

A memory operand that is not offsettable. In other words, anything that would fit the `m` constraint but not the `o` constraint.

- < A memory operand with autodecrement addressing (either predecrement or postdecrement) is allowed.
- > A memory operand with autoincrement addressing (either preincrement or postincrement) is allowed.
- r A register operand is allowed provided that it is in a general register.
- d, a, f, ... Other letters can be defined in machine-dependent fashion to stand for particular classes of registers. `d`, `a`, and `f` are defined on the 68000/68020 to stand for data, address and floating point registers.
- i An immediate integer operand (one with constant value) is allowed. This includes symbolic constants whose values will be known only at assembly time.
- n An immediate integer operand with a known numeric value is allowed. Many systems cannot support assembly-time constants for operands less than a word wide. Constraints for these operands should use `n` rather than `i`.
- I, J, K, ... P Other letters in the range `I` through `P` may be defined in a machine-dependent fashion to permit immediate integer operands with explicit integer values in specified ranges. For example, on the 68000, `I` is defined to stand for the range of values 1 to 8. This is the range permitted as a shift count in the shift instructions.
- E An immediate floating operand (expression code `const_ double`) is allowed, but only if the target floating point format is the same as that of the host machine (on which the compiler is running).
- F An immediate floating operand (expression code `const_ double`) is allowed.
- G, H `G` and `H` are defined in a machine-dependent fashion to permit immediate floating operands in particular ranges of values.

- s** An immediate integer operand whose value is not an explicit integer is allowed. This might appear strange; if an `insn` allows a constant operand with a value not known at compile time, it certainly must allow any known value. So why use `s` instead of `i`? Sometimes it allows better code to be generated. For example, on the 68000 in a `fullword` instruction it is possible to use an immediate operand; but if the immediate value is between -128 and 127, better code results from loading the value into a register and using the register. This is because the load into the register can be done with a `moveq` instruction. This happens by defining the letter `κ` to mean any integer outside the range -128 to 127, and then specifying `κs` in the operand constraints.
- g** Any register, memory or immediate integer operand is allowed, except for registers that are not general registers.
- x** Any operand whatsoever is allowed.
- 0, 1, 2, ... 9** An operand that matches the specified operand number is allowed. If a digit is used together with letters within the same alternative, the digit should come last. This is called a matching constraint and what it really means is that the assembler has only a single operand that fills two roles which `asm` distinguishes. For example, an `add` instruction uses two input operands and an output operand, but on most CISC machines an `add` instruction really has only two operands, one of them an input-output operand:
- ```
addl #35,r12
```
- Matching constraints are used in these circumstances. More precisely, the two operands that match must include one input-only operand and one output-only operand. Moreover, the digit must be a smaller number than the number of the operand that uses it in the constraint.
- p** An operand that is a valid memory address is allowed. This is for load address and push address instructions. `p` in the constraint must be accompanied by `address_operand` as the predicate in the `match_operand`. This predicate interprets the mode specified in the `match_operand` as the mode of the memory reference for which the address would be valid.
- Q, R, S, ... U** Letters in the range `Q` through `U` may be defined in a machine-dependent fashion to stand for arbitrary operand types.

Multiple Alternative Constraints

Sometimes a single instruction has multiple alternative sets of possible operands. For example, on the 68000, a logical-or instruction can combine register or an immediate value into memory, or it can combine any kind of operand into a register; but it cannot combine one memory location into another.

These constraints are represented as multiple alternatives. An alternative can be described by a series of letters for each operand. The overall constraint for an operand is made from the letters for this operand from the first alternative, a comma, the letters for this operand from the second alternative, a comma, and so on until the last alternative.

If all the operands fit any one alternative, the instruction is valid. Otherwise, for each alternative, the compiler counts how many instructions must be added to copy the operands so that that alternative applies. The alternative requiring the least copying is chosen. If two alternatives need the same amount of copying, the one that comes first is chosen. These choices can be altered with the `?` and `!` characters:

`?`

Disparage slightly the alternative that the `?` appears in, as a choice when no alternative applies exactly. The compiler regards this alternative as one unit more costly for each `?` that appears in it.

`!`

Disparage severely the alternative that the `!` appears in. This alternative can still be used if it fits without reloading, but if reloading is needed, some other alternative will be used.

Constraint Modifier Characters

The following are constraint modifier characters.

`=`

Means that this operand is write-only for this instruction: the previous value is discarded and replaced by output data.

`+`

Means that this operand is both read and written by the instruction.

When the compiler fixes up the operands to satisfy the constraints, it needs to know which operands are inputs to the instruction and which are outputs from it. `=` identifies an output; `+` identifies an operand that is both input and output; all other operands are assumed to be input only.

`&`

Means (in a particular alternative) that this operand is an *earlyclobber* operand which is modified before the instruction is finished using the input operands. Therefore, this operand may not lie in a register that is used as an input operand or as part of any memory address.

& applies only to the alternative in which it is written. In constraints with multiple alternatives, sometimes one alternative requires & while others do not. See, for example, the `movdf` insn of the 68000.

& does not obviate the need to write =.

%

Declares the instruction to be commutative for this operand and the following operand. This means that the compiler may interchange the two operands if that is the cheapest way to make all operands fit the constraints.

#

Says that all following characters, up to the next comma, are to be ignored as a constraint. They are significant only for choosing register preferences.

Constraints for Particular Machines

Whenever possible, you should use the general-purpose constraint letters in `asm` arguments, since they will convey meaning more readily to people reading your code. Failing that, use the constraint letters that usually have very similar meanings across architectures. The most commonly used constraints are `m` and `r` (for memory and general-purpose registers respectively; see “Simple Constraints” on page 181), and `i`, usually the letter indicating the most common immediate-constant format. For each machine architecture, the `config/machine.h` file defines additional constraints. These constraints are used by the compiler itself for instruction generation, as well as for `asm` statements; therefore, some of the constraints are not particularly interesting for `asm`.

The constraints are defined through the following macros.

`REG_CLASS_FROM_LETTER`

Register class constraints (usually lower case).

`CONST_OK_FOR_LETTER_P`

Immediate constant constraints, for non-floating point constants of word size or smaller precision (usually upper case).

`CONST_DOUBLE_OK_FOR_LETTER_P`

Immediate constant constraints, for all floating point constants and for constants of greater than word size precision (usually upper case).

`EXTRA_CONSTRAINT`

Special cases of registers or memory. This macro is not required, and is only defined for some machines.

Inspecting these macro definitions in the compiler source for your machine is the best way to be certain you have the right constraints. However, the following is a summary of the machine-dependent constraints available on some particular machines.

ARM Family—`arm.h` Constraints

The following is a summary of the machine-dependent constraints available on ARM machines.

- `f` Floating-point register
- `F` One of the floating-point constants 0.0, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0 or 10.0
- `G` Floating-point constant that would satisfy the constraint `F` if it were negated
- `I` Integer that is valid as an immediate operand in a data processing instruction. That is, an integer in the range 0 to 255 rotated by a multiple of 2
- `J` Integer in the range -4095 to 4095
- `K` Integer that satisfies constraint `I` when inverted (ones complement)
- `L` Integer that satisfies constraint `I` when negated (twos complement)
- `M` Integer in the range 0 to 32
- `Q` A memory reference where the exact address is in a single register (`m` is preferable for `asm` statements)
- `R` An item in the constant pool
- `S` A symbol in the text segment of the current file

AMD 29000 Family—`a29k.h` Constraints

The following is a summary of the machine-dependent constraints available on AMD 29K machines.

- `l` Local register 0
- `b` Byte Pointer (`BP`) register
- `q` `Q` register
- `h` Special purpose register

A	First accumulator register
a	Other accumulator register
f	Floating point register
I	Constant greater than 0, less than 0x100
J	Constant greater than 0, less than 0x10000
K	Constant whose high 24 bits are on (1)
L	16 bit constant whose high 8 bits are on (1)
M	32 bit constant whose high 16 bits are on (1)
N	32 bit negative constant that fits in 8 bits
O	The constant 0x80000000 or, on the 29050, any 32 bit constant whose low 16 bits are 0.
P	16 bit negative constant that fits in 8 bits
G, H	A floating point constant (in <code>asm</code> statements, use the machine independent <code>E</code> or <code>F</code> instead)

IBM RS6000—rs6000.h Constraints

The following is a summary of the machine-dependent constraints available on IBM RS6000 machines.

b	Address base register
f	Floating point register
h	MQ, CTR, or LINK register
q	MQ register
c	CTR register

<code>l</code>	LINK register
<code>x</code>	CR register (condition register) number 0
<code>Y</code>	CR register (condition register)
<code>I</code>	Signed 16 bit constant
<code>J</code>	Constant whose low 16 bits are 0
<code>K</code>	Constant whose high 16 bits are 0
<code>L</code>	Constant suitable as a mask operand
<code>M</code>	Constant larger than 31
<code>N</code>	Exact power of 2
<code>O</code>	Zero
<code>P</code>	Constant whose negation is a signed 16 bit constant
<code>G</code>	Floating point constant that can be loaded into a register with one instruction per word
<code>Q</code>	Memory operand that is an offset from a register (<code>m</code> is preferable for <code>asm</code> statements)
<code>R</code>	AIX TOC entry
<code>S</code>	Windows NT SYMBOL REF
<code>T</code>	Windows NT LABEL REF
<code>U</code>	System V Release 4 small data area reference

Intel 386—i386.h Constraints

The following is a summary of the machine-dependent constraints available on Intel 386 machines.

<code>q</code>	a, b, c, or d register
<code>A</code>	a, or d register (for 64-bit ints)
<code>f</code>	Floating point register
<code>t</code>	First (top of stack) floating point register
<code>u</code>	Second floating point register
<code>a</code>	a register
<code>b</code>	b register
<code>c</code>	c register
<code>d</code>	d register
<code>D</code>	di register
<code>S</code>	si register
<code>I</code>	Constant in range 0 to 31 (for 32 bit shifts)
<code>J</code>	Constant in range 0 to 63 (for 64 bit shifts)
<code>K</code>	<code>0xff'</code>
<code>L</code>	<code>0xffff'</code>
<code>M</code>	0, 1, 2, or 3 (shifts for <code>leaq</code> instruction)
<code>N</code>	Constant in range 0 to 255 (for <code>out</code> instruction)
<code>G</code>	Standard 80387 floating point constant

Intel 960—`i960.h` Constraints

The following is a summary of the machine-dependent constraints available on Intel 960 machines.

<code>f</code>	Floating point register (<code>fp0</code> to <code>fp3</code>)
<code>l</code>	Local register (<code>r0</code> to <code>r15</code>)
<code>b</code>	Global register (<code>g0</code> to <code>g15</code>)
<code>d</code>	Any local or global register
<code>I</code>	Integers from 0 to 31
<code>J</code>	0
<code>K</code>	Integers from -31 to 0
<code>G</code>	Floating point 0
<code>H</code>	Floating point 1

MIPS—`mips.h` Constraints

The following is a summary of the machine-dependent constraints available on MIPS machines.

<code>d</code>	General-purpose integer register
<code>f</code>	Floating-point register (if available)
<code>h</code>	<code>Hi</code> register
<code>l</code>	<code>Lo</code> register
<code>x</code>	<code>Hi</code> or <code>Lo</code> register
<code>y</code>	General-purpose integer register
<code>z</code>	Floating-point status register

I	Signed 16 bit constant (for arithmetic instructions)
J	Zero
K	Zero-extended 16-bit constant (for logic instructions)
L	Constant with low 16 bits zero (can be loaded with <code>lui</code>)
M	32 bit constant which requires two instructions to load (a constant which is not I, K, or L)
N	Negative 16 bit constant
O	Exact power of two
P	Positive 16 bit constant
G	Floating point zero
Q	Memory reference that can be loaded with more than one instruction (<code>'m</code> is preferable for <code>asm</code> statements)
R	Memory reference that can be loaded with one instruction (<code>'m</code> is preferable for <code>asm</code> statements)
S	Memory reference in external OSF/rose PIC for-format (<code>'m</code> is preferable for <code>asm</code> statements)

Motorola 68000—m68k.h Constraints

The following is a summary of the machine-dependent constraints available on Motorola 68K machines.

a	Address register
d	Data register
f	68881 floating-point register, if available
x	Sun FPA (floating-point) register, if available

Y	First 16 Sun FPA registers, if available
I	Integer in the range 1 to 8
J	16 bit signed number
K	Signed number whose magnitude is greater than 0x80
L	Integer in the range -8 to -1
M	Signed number whose magnitude is greater than 0x100.
G	Floating point constant that is not a 68881 constant
H	Floating point constant that can be used by Sun FPA

SPARC—`sparc.h` Constraints

The following is a summary of the machine-dependent constraints available on SPARC machines.

f	Floating-point register
e	Floating point register that can hold 64 or 128 bit values.
I	Signed 13 bit constant
J	Zero
K	32 bit constant with the low 12 bits clear (a constant that can be loaded with the <code>sethi</code> instruction)
G	Floating-point zero
H	Signed 13 bit constant, sign-extended to 32 or 64 bits
Q	Memory reference that can be loaded with one instruction (<code>'m</code> is more appropriate for <code>asm</code> statements)
S	Constant, or memory address

T	Memory address aligned to an 8-byte boundary
U	Even register

Controlling Names Used in Assembler Code

You can specify the name to be used in the assembler code for a C function or variable by writing the `asm` (or `__asm__`) keyword after the declarator, as the following example shows; this specifies that the name to be used for the variable `foo` in the assembler code should be `myfoo` rather than the usual `_foo`.

```
int foo asm ("myfoo") = 2;
```

On systems where an underscore is normally prepended to the name of a C function or variable, this feature allows you to define names for the linker that do not start with an underscore. You cannot use `asm` in this way in a function definition; but you can get the same effect by writing a declaration for the function before its definition and putting `asm` there, like the following example shows.

```
extern func () asm ("FUNC");
```

```
func (x, y)
int x, y;
```

It is up to you to make sure that the assembler names you choose do not conflict with any other assembler symbols. Also, you must not use a register name; that would produce completely invalid assembler code. GCC does not as yet have the ability to store static variables in registers.

Variables in Specified Registers

GCC allows you to put a few global variables into specified hardware registers. You can also specify the register in which an ordinary register variable should be allocated.

- Global register variables reserve registers throughout the program. This may be useful in programs such as programming language interpreters which have a couple of global variables that are accessed very often.
- Local register variables in specific registers do not reserve the registers. The compiler's data flow analysis is capable of determining where the specified registers contain live values, and where they are available for other uses.

These local variables are sometimes convenient for use with the extended `asm` feature (see "Assembler Instructions with C Expression Operands" on page 177), if you want to write one output of the assembler instruction directly into a

particular register. (This will work provided the register you specify fits the constraints specified for that operand in the `asm`.)

Defining Global Register Variables

You can define a global register variable in GCC, using the following example's input.

```
register int *foo asm ("a5");
```

In the previous example, `a5` is the name of the register which should be used. Choose a register which is normally saved and restored by function calls on your machine, so that library routines will not clobber it.

Naturally the register name is CPU-dependent, so you would need to conditionalize your program according to CPU type. The register `a5` would be a good choice on a 68000 for a variable of pointer type. On machines with register windows, be sure to choose a *global* register that is not affected magically by the function call mechanism.

In addition, operating systems on one type of CPU may differ in how they name the registers; then you would need additional conditionals. For example, some 68000 operating systems call the register, `%a5`.

Eventually there may be a way of asking the compiler to choose a register automatically, but no solution is evident.

Defining a global register variable in a certain register reserves that register entirely for this use, at least within the current compilation. The register will not be allocated for any other purpose in the functions in the current compilation. The register will not be saved and restored by these functions. Stores into this register are never deleted even if they would appear to be dead, but references may be deleted or moved or simplified. It is not safe to access the global register variables from signal handlers, or from more than one thread of control, because the system library routines may temporarily use the register for other things (unless you recompile them specially for the task at hand).

It is not safe for one function that uses a global register variable to call another such function, `foo`, by way of a third function, `lose`, that was compiled without knowledge of this variable (i.e., in a different source file in which the variable wasn't declared). This is because `lose` might save the register and put some other value there. For example, you can't expect a global register variable to be available in the comparison-function that you pass to `qsort`, since `qsort` might have put something else in that register. (If you are prepared to recompile `qsort` with the same global register variable, you can solve this problem.)

If you want to recompile `qsort` or other source files which do not actually use your global register variable, so that they will not use that register for any other purpose, then it suffices to specify the compiler option, `-ffixed-reg`. You need not actually add a global register declaration to their source code.

A function which can alter the value of a global register variable cannot safely be called from a function compiled without this variable, because it could clobber the value the caller expects to find there on return. Therefore, the function which is the entry point into the part of the program that uses the global register variable must explicitly save and restore the value which belongs to its caller.

On most machines, `longjmp` will restore to each global register variable the value it had at the time of the `setjmp`. On some machines, however, `longjmp` will not change the value of global register variables.

To be portable, the function that called `setjmp` should make other arrangements to save the values of the global register variables, and to restore them in a `longjmp`. This way, the same thing will happen regardless of what `longjmp` does. All global register variable declarations must precede all function definitions. If such a declaration could appear after function definitions, the declaration would be too late to prevent the register from being used for other purposes in the preceding functions.

Global register variables may not have initial values, because an executable file has no means to supply initial contents for a register. On the Sparc, there are reports that `g3 ... g7` are suitable registers, but certain library functions, such as `getwd`, as well as the subroutines for division and remainder, modify `g3` and `g4`. `g1` and `g2` are local temporaries. On the 68000, `a2 ... a5` should be suitable, as should `d2 ... d7`. Of course, it will not do to use more than a few of those.

Specifying Registers for Local Variables

You can define a local register variable with a specified register like the following.

```
register int *foo asm ("a5");
```

Here `a5` is the name of the register which should be used. This is the same syntax used for defining global register variables, but for a local variable it would appear within a function.

Naturally the register name is CPU-dependent, but this is not a problem, since specific registers are most often useful with explicit assembler instructions (see “Assembler Instructions with C Expression Operands” on page 177). Both of these things generally require that you conditionalize your program according to CPU type.

In addition, operating systems on one type of CPU may differ in how they name the registers; then you would need additional conditionals. For example, some 68000 operating systems call this register, `%a5`.

Eventually there may be a way of asking the compiler to choose a register automatically, but no solution is evident.

Defining such a register variable does not reserve the register; it remains available for other uses in places where flow control determines the variable’s value is not live. However, these registers are made unavailable for use in the reload pass. I would not

be surprised if excessive use of this feature leaves the compiler too few available registers to compile certain functions.

Alternate Keywords

The option, `-traditional`, disables certain keywords; `-ansi` disables certain others. This causes trouble when you want to use GCC extensions, or ANSI C features, in a general-purpose header file that should be usable by all programs, including ANSI C programs and traditional ones.

The keywords, `asm`, `typeof` and `inline`, cannot be used since they won't work in a program compiled with `-ansi`, while the keywords, `const`, `volatile`, `signed`, `typeof` and `inline`, won't work in a program compiled with `-traditional`.

The way to solve these problems is to put `__` at the beginning and end of each problematical keyword. For example, use `__asm__` instead of `asm`, `__const__` instead of `const`, and `__inline__` instead of `inline`.

Other C compilers won't accept these alternative keywords; if you want to compile with another compiler, you can define the alternate keywords as macros to replace them with the customary keywords.

It looks like the following declaration..

```
#ifndef __GNUC__
#define __asm__ asm
#endif
```

`-pedantic` causes warnings for many GCC extensions. You can prevent such warnings within one expression by writing `__extension__` before the expression. `__extension__` has no effect aside from this problem.

Incomplete enum Types

You can define an `enum` tag without specifying its possible values. This results in an incomplete type, much like what you get if you write `struct foo` without describing the elements. A later declaration which does specify the possible values completes the type.

You can't allocate variables or storage using the type while it is incomplete. However, you can work with pointers to that type.

This extension may not be very useful, but it makes the handling of `enum` more consistent with the way `struct` and `union` are handled.

This extension is not supported by GNU C++.

Function Names as Strings

GCC predefines two string variables to be the name of the current function. The variable, `__FUNCTION__`, is the name of the function as it appears in the source. The variable, `__PRETTY_FUNCTION__`, is the name of the function pretty printed in a language specific fashion. These names are always the same in a C function; in a C++ function, they may be different, like the following program.

```
extern "C" {
extern int printf (char *, ...);
}
class a {
public:
    sub (int i)
    {
        printf ("__FUNCTION__ = %s\n", __FUNCTION__);
        printf ("__PRETTY_FUNCTION__ = %s\n", __PRETTY_FUNCTION__);
    }
};
int
main (void)
{
    a ax;
    ax.sub (0);
    return 0;
}
```

The program, then, gives the following output.

```
__FUNCTION__ = sub
__PRETTY_FUNCTION__ = int a::sub (int)
```

These names are not macros: they are predefined string variables. For example, `#ifdef __FUNCTION__` does not have any special meaning inside a function, since the preprocessor does not do anything special with the identifier, `__FUNCTION__`.

Getting the Return or Frame Address of a Function

The following calls may be used to get information about the callers of a function.

```
__builtin_return_address (level)
```

This function returns the return address of the current function, or of one of its callers. The *level* argument is number of frames to scan up the call stack. A value of 0 yields the return address of the current function, a value of 1 yields the return address of the caller of the current function, and so forth. The *level* argument must be a constant integer. On some machines it may be impossible to determine the return address of any function other than the current one; in such cases, or when the top of the stack has been reached, this function will return 0. This function should only be used with a non-zero argument for debugging purposes.

`__builtin_frame_address (level)`

This function is similar to `__builtin_return_address`, but it returns the address of the function frame rather than the return address of the function. Calling `__builtin_frame_address` with a value of 0 yields the frame address of the current function, a value of 1 yields the frame address of the caller of the current function, and so forth. The frame is the area on the stack which holds local variables and saved registers. The frame address is normally the address of the first word pushed on to the stack by the function. However, the exact definition depends upon the processor and the calling convention. If the processor has a dedicated frame pointer register, and the function has a frame, then `__builtin_frame_address` will return the value of the frame pointer register. The caveats that apply to `__builtin_return_address` apply to this function as well.

Extensions to the C++ Language Family

The GNU compiler provides extensions to the C++ language (and you can also use most of the C language extensions in your C++ programs). If you want to write code that checks whether these features are available, you can test for the GNU compiler the same way as for C programs: check for a predefined macro, `__GNUC__`. You can also use `__GNUG__` to test specifically for the GNU C++ compiler tools, G++ (see “Standard Predefined Macros” on page 233).

- ““Named Return Values in C++”” (below)
- “Minimum and Maximum Operators in C++” on page 201
- “The `goto` and Destructors in GNU C++” on page 202
- “Declarations and Definitions in One Header” on page 202
- “Where’s the Template?” on page 204
- “Type Abstraction Using Signatures” on page 207

Named Return Values in C++

GNU C++ extends the function-definition syntax, which allows you to specify a name for the result of a function outside the body of the definition, in C++ programs, as the following example shows.

```
type
functionname (args) return resultname;
{
    ...
    body
    ...
}
```

You can use this feature to avoid an extra constructor call when a function result has a class type. For example, consider a function `m`, declared as `x v =m ();`, whose result is of class `x`:

```
x
m ()
{
    x b;
    b.a = 23;
    return b;
}
```

Although `m` appears to have no arguments, in fact it has one implicit argument: the address of the return value. At invocation, the address of enough space to hold `v` is sent in as the implicit argument. Then `b` is constructed and its `a` field is set to the value 23. Finally, a copy constructor (a constructor of the form `x(x&)`) is applied to `b`, with the (implicit) return value location as the target, so that `v` is now bound to the return value.

But this is wasteful. The local `b` is declared just to hold something that will be copied right out. While a compiler that combined an *elision* algorithm with interprocedural data flow analysis could conceivably eliminate all of this, it is much more practical to allow you to assist the compiler in generating efficient code by manipulating the return value explicitly, thus avoiding the local variable and copy constructor altogether.

Using the extended GNU C++ function-definition syntax, you can avoid the temporary allocation and copying by naming `r` as your return value at the outset, and assigning to its `a` field directly the following declaration.

```
x
m () return r;
{
    r.a = 23;
}
```

The declaration of `r` is a standard, proper declaration, whose effects are executed *before* any of the body of `m`.

Functions of this type impose no additional restrictions; in particular, you can execute return statements, or return implicitly by reaching the end of the function body (*falling off the edge*). Cases like the following declaration (or even the statement, `xm ();`)

`return r (23) ; { }` are unambiguous, since the return value `r` has been initialized in either case.

```
X
m () return r (23);
{
    return;
}
```

The following code may be hard to read, but also works predictably.

```
X
m () return r;
{
    X b;
    return b;
}
```

The return value slot denoted by `r` is initialized at the outset, but the statement `return b;` overrides this value. The compiler deals with this by destroying `r` (calling the destructor if there is one, or doing nothing if there is not), and then reinitializing `r` with `b`.

This extension is provided primarily to help people who use overloaded operators, where there is a great need to control not just the arguments, but the return values of functions. For classes where the copy constructor incurs a heavy performance penalty (especially in the common case where there is a quick default constructor), this is a major savings. The disadvantage of this extension is that you do not control when the default constructor for the return value is called: it is always called at the beginning.

Minimum and Maximum Operators in C++

It is very convenient to have operators which return the *minimum* or the *maximum* of two arguments. For instance, in GNU C++ (but not in GNU C), operations perform the following returns.

`a <? b` is the *minimum*, returning the smaller of the numeric values `a` and `b`;

`a >? b` is the *maximum*, returning the larger of the numeric values `a` and `b`.

These operations are not primitive in ordinary C++, since you can use a macro to return the minimum of two things in C++, as in the following example.

```
#define MIN(X,Y) ((X) < (Y) ? (X) : (Y))
```

You might then use `int min = MIN (i, j);` to set `min` to the minimum value of variables `i` and `j`. However, side effects in `x` or `y` may cause unintended behavior. For example, `MIN (i++, j++)` will fail, incrementing the smaller counter twice. A GNU C extension allows you to write safe macros that avoid this kind of problem (see “Naming an Expression’s Type” on page 152). However, writing `MIN` and `MAX` as

macros also forces you to use function-call notation for a fundamental arithmetic operation.

Using GNU C++ extensions, you can write `int min = i <? j;` instead.

Since `<?` and `>?` are built into the compiler, they properly handle expressions with side-effects; `int min = i++ <? j++;` works correctly.

The `goto` and Destructors in GNU C++

In C++ programs, you can safely use the `goto` statement. When you use it to exit a block which contains aggregates requiring destructors, the destructors will run before the `goto` transfers control. (In ANSI C++, `goto` is restricted to targets within the current function.) The compiler still forbids using `goto` to *enter* a scope that requires constructors.

Declarations and Definitions in One Header

C++ object definitions can be quite complex. In principle, your source code will need two kinds of things for each object that you use across more than one source file.

First, you need an *interface* specification, describing its structure with type declarations and function prototypes. Second, you need the *implementation* itself. It can be tedious to maintain a separate interface description in a header file, in parallel to the actual implementation. It is also dangerous, since separate interface and implementation definitions may not remain parallel. With GNU C++, you can use a single header file for both purposes.

WARNING! The mechanism to specify this is in transition. For the nonce, you must use one of two `#pragma` commands; in a future release of GNU C++, an alternative mechanism will make these `#pragma` commands unnecessary.

The header file contains the full definitions, but is marked with `#pragma interface` in the source code. This allows the compiler to use the header file only as an interface specification when ordinary source files incorporate it with `#include`. In the single source file where the full implementation belongs, you can use either a naming convention or `#pragma implementation` to indicate this alternate use of the header file.

```
#pragma interface
#pragma interface "subdir/objects.h"
```

Use this directive in *header files* that define object classes, to save space in most of the object files that use those classes. Normally, local copies of certain

information (backup copies of inline member functions, debugging information, and the internal tables that implement virtual functions) must be kept in each object file that includes class definitions. You can use this pragma to avoid such duplication.

When a header file containing `#pragma interface` is included in a compilation, this auxiliary information will not be generated (unless the main input source file itself uses `#pragma implementation`). Instead, the object files will contain references to be resolved at link time.

The second form of this directive is useful for the case where you have multiple headers with the same name in different directories. If you use this form, you must specify the same string to `#pragma implementation`.

```
#pragma implementation
#pragma implementation "objects.h"
```

Use this pragma in a *main input file*, when you want full output from included header files to be generated (and made globally visible). The included header file, in turn, should use `#pragma interface`. Backup copies of inline member functions, debugging information, and the internal tables used to implement virtual functions are all generated in implementation files.

If you use `#pragma implementation` with no argument, it applies to an include file with the same *basename*[†] as your source file. For example, in `allclass.cc`, `#pragma implementation`, by itself, is equivalent to `#pragma implementation "allclass.h"`.

In versions of GNU C++ prior to 2.6.0, `allclass.h` was treated as an implementation file whenever you would include it from `allclass.cc` even if you never specified `#pragma implementation`. This was deemed to be more trouble than it was worth, however, and disabled. If you use an explicit `#pragma implementation`, it must appear in your source file before you include the affected header files.

Use the string argument if you want a single implementation file to include code from multiple header files. (You must also use `#include` to include the header file; `#pragma implementation` only specifies how to use the file—it doesn't actually include it.) There is no way to split up the contents of a single header file into multiple implementation files.

`#pragma implementation` and `#pragma interface` also have an effect on function inlining.

If you define a class in a header file marked with `#pragma interface`, the effect on a function defined in that class is similar to an explicit `extern` declaration—the compiler emits no code at all to define an independent version of the function. Its

[†] A file's *basename* was the name stripped of all leading path information and of trailing suffixes (such as: `.h` or `.C` or `.cc`).

definition is used only for inlining with its callers.

Conversely, when you include the same header file in a main source file that declares it as `#pragma implementation`, the compiler emits code for the function itself; this defines a version of the function that can be found via pointers (or by callers compiled without inlining). If all calls to the function can be inlined, you can avoid emitting the function by compiling with `-fno-implementation-inlines`. If any calls were not inlined, you will get linker errors.

Where's the Template?

C++ templates are the first language feature to require more intelligence from the environment than one usually finds on a UNIX system. Somehow the compiler and linker have to make sure that each template instance occurs exactly once in the executable if it is needed, and not at all otherwise.

There are two basic approaches to this problem, which we will refer to as the *Borland model* and the *Cfront model*.

Borland model

Borland C++ solved the template instantiation problem by adding the code equivalent of common blocks to their linker; the compiler emits template instances in each translation unit that uses them, and the linker collapses them together at run time. The advantage of this model is that the linker only has to consider the object files themselves; there is no external complexity to worry about. This disadvantage is that compilation time is increased because the template code is being compiled repeatedly. Code written for this model tends to include definitions of all member templates in the header file, since they must be seen to be instantiated.

Cfront model

The AT&T C++ translator, “Cfront”, solved the template instantiation problem by creating the notion of a template repository, an automatically maintained place where template instances are stored. A more modern version of the repository works as follows.

As individual object files are built, the compiler places any template definitions and instantiations encountered in this repository. At link time, the link wrapper adds in the objects in the repository and compiles any needed instances that were not previously emitted. The advantages of this model are more optimal compilation speed and the ability to use the system linker; to implement the Borland model, a compiler vendor also needs to replace the linker. The disadvantages are vastly increased complexity, and thus potential for error; for some code, this can be just as transparent, but in practice it can be very difficult to build multiple programs in one directory and one program in multiple directories using Cfront. Code written for this model tends to separate definitions of non-

inline member templates into a separate file, which should be separately compiled.

When used with GNU `ld` version 2.8 or later on an ELF system such as Linux/GNU or Solaris 2, or on Microsoft Windows, `g++` supports the Borland model.

A future version of `g++` will support a hybrid model whereby the compiler will emit any instantiations for which the template definition is included in the compile, and store template definitions and instantiation context information into the object file for the rest. The link wrapper will extract that information as necessary and invoke the compiler to produce the remaining instantiations. The linker will then combine duplicate instantiations.

In the mean time, you have the following options for dealing with template instantiations.

- Compile your template-using code with `-frepo`. The compiler will generate files with the extension `.rpo` listing all of the template instantiations used in the corresponding object files which could be instantiated there; the link wrapper, `collect2`, will then update the `.rpo` files to tell the compiler where to place those instantiations and rebuild any affected object files. The link-time overhead is negligible after the first pass, as the compiler will continue to place the instantiations in the same files.

This is your best option for application code written for the Borland model, as it will just work. Code written for the Cfront model will need to be modified so that the template definitions are available at one or more points of instantiation; usually this is as simple as adding `#include tmethods.cc` to the end of each template header.

For library code, if you want the library to provide all of the template instantiations it needs, just try to link all of its object files together; the link will fail, but cause the instantiations to be generated as a side effect. Be warned, however, that this may cause conflicts if multiple libraries try to provide the same instantiations. For greater control, use explicit instantiation as described in the next option.

- Compile your code with `-fno-implicit-templates` to disable the implicit generation of template instances, and explicitly instantiate all the ones you use. This approach requires more knowledge of exactly which instances you need than do the others, but it's less mysterious and allows greater control. You can scatter the explicit instantiations throughout your program, perhaps putting them in the translation units where the instances are used or the translation units that define the templates themselves; you can put all of the explicit instantiations you need into one big file; or you can create small files for each of the instances you need, like the following examples define, and create a template instantiation library from those files.

```
#include "Foo.h"
#include "Foo.cc"

template class Foo<int>;
template ostream& operator <<
    (ostream&, const Foo<int>&);
```

If you are using Cfront-model code, you can probably get away with not using `-fno-implicit-templates` when compiling files that don't #include the member template definitions.

If you use one big file to do the instantiations, you may want to compile it without `-fno-implicit-templates` so you get all of the instances required by your explicit instantiations (but not by any other files) without having to specify them as well.

g++ has extended the template instantiation syntax outlined in the Working Paper to allow forward declaration of explicit instantiations, explicit instantiation of members of template classes and instantiation of the compiler support data for a template class (such as the vtable) without instantiating any of its members as the following example shows.

```
extern template int max (int, int);
template void Foo<int>::f ();
inline template class Foo<int>;
```

- Do nothing. Pretend G++ does implement automatic instantiation management. Code written for the Borland model will work fine, but each translation unit will contain instances of each of the templates it uses. In a large program, this can lead to an unacceptable amount of code duplication.
- Add `#pragma interface` to all files containing template definitions. For each of these files, add `#pragma implementation "filename"` to the top of some .C file which #include's it. Then compile everything with `-fexternal-templates`. The templates will then only be expanded in the translation unit which implements them (i.e., the translation unit has a `#pragma implementation` line for the file where they live); all other files will use external references. If you are lucky, everything should work properly. If you get undefined symbol errors, you need to make sure that each template instance which is used in the program is used in the file which implements that template. If you do not have any use for a particular instance in that file, you can just instantiate it explicitly, using the syntax from the latest C++ working paper:

```
template class A<int>;
template ostream& operator << (ostream&, const A<int>&);
```

This strategy will work with code written for either model. If you are using code written for the Cfront model, the file containing a class template and the file containing its member templates should be implemented in the same translation unit. A slight variation on this approach is to use `-falt-external-templates` instead; this causes template instances to be emitted in the translation unit that

implements the header where they are first instantiated, rather than the one which implements the file where the templates are defined. This header must be the same in all translation units, or things are likely to break.

See “Declarations and Definitions in One Header” on page 202 for more discussion of these pragmas.

Type Abstraction Using Signatures

In GNU C++, you can use the keyword `signature` to define a completely abstract class interface as a datatype. You can connect this abstraction with actual classes using signature pointers. If you want to use signatures, run the GNU compiler with the `-fhandle-signatures` command-line option. (With this option, the compiler reserves a second keyword, `sigof`, as well, for a future extension.)

Roughly, signatures are type abstractions or interfaces of classes. Some other languages have similar facilities. C++ signatures are related to ML’s signatures, Haskell’s type classes, definition modules in Modula-2, interface modules in Modula-3, abstract types in Emerald, type modules in Trellis/Owl, categories in Scratchpad II, and types in POOL-I. For a more detailed discussion of signatures, see *Signatures: A Language Extension for Improving Type Abstraction and Subtype Polymorphism in C++* by Gerald Baumgartner and Vincent F. Russo (Tech report CSD–TR–95–051, Dept. of Computer Sciences, Purdue University, August 1995, a slightly improved version appeared in *Software—Practice & Experience*, **25**(8), pp. 863–889, August 1995). You can get the tech report by anonymous FTP from `ftp.cs.purdue.edu` in `pub/gb/Signature-design.ps.gz`.

Syntactically, a signature declaration is a collection of member function declarations and nested type declarations. For example, the following signature declaration defines a new abstract type `s` with member functions, `int foo()` and `int bar(int)`.

```
signature S
{
    int foo ();
    int bar (int);
};
```

Since signature types do not include implementation definitions, you cannot write an instance of a signature directly. Instead, you can define a pointer to any class that contains the required interfaces as a *signature pointer*. Such a class *implements* the signature type.

To use a class as an implementation of `s`, you must ensure that the class has public member functions `int foo()` and `int bar(int)`. The class can have other member functions as well, public or not; as long as it offers what is declared in the signature, it is suitable as an implementation of that signature type.

For example, suppose that `C` is a class that meets the requirements of signature `S` (`C` conforms to `S`). Then the following statement defines a signature pointer `p` and initializes it to point to an object of type `C`.

```
C obj;  
S * p = &obj;
```

The member function call, `int i = p->foo ();`, executes `obj.foo ()`.

Abstract virtual classes provide somewhat similar facilities in standard C++. There are two main advantages to using signatures instead:

- Subtyping becomes independent from inheritance.

A class or signature type `T` is a subtype of a signature type `S` independent of any inheritance hierarchy as long as all the member functions declared in `S` are also found in `T`. So you can define a subtype hierarchy that is completely independent from any inheritance (implementation) hierarchy, instead of being forced to use types that mirror the class inheritance hierarchy.

- Signatures allow you to work with existing class hierarchies as implementations of a signature type. If those class hierarchies are only available in compiled form, you are out of luck with abstract virtual classes, since an abstract virtual class cannot be retrofitted on top of existing class hierarchies.

So you would be required to write interface classes as subtypes of the abstract virtual class.

There is one more detail about signatures. A signature declaration can contain member function *definitions* as well as member function declarations. A signature member function with a full definition is called a *default implementation*; classes need not contain that particular interface in order to conform.

For example, a class `C` can conform to the following signature.

```
signature T  
{  
    int f (int);  
    int f0 () { return f (0); };  
};
```

This happens whether `C` implements the member function, `int f0()`, or not. If you define `C::f0`, that definition takes precedence; otherwise, the default implementation, `S::f0`, applies.

gcov, a Test Coverage Program

`gcov` is a tool for working with GCC. The following documentation describes `gcov`.

- “Introduction to `gcov` Test Coverage” (below)
- “Invoking the `gcov` Program” on page 210
- “Using `gcov` with GCC Optimization” on page 213
- “Brief Description of `gcov` Data Files” on page 213

Introduction to `gcov` Test Coverage

`gcov` is a test coverage program (Jim Wilson wrote `gcov`, and the original form of this documentation; Pat McGregor edited the documentation). Use it with GCC to analyze your programs for creating more efficient, faster running code, to use as a profiling tool, to discover where your optimization efforts will best affect your code, and to use with the other profiling tool, `gprof`, for assessing which parts of your code use the greatest amount of computing time. Profiling tools help you analyze your code’s performance to find out some basic performance statistics, such as:

- how often each line of code executes
- what lines of code are actually executed
- how much computing time each section of code uses

Once you know how your code works when compiled, you can look at each module to

see which modules should be optimized. `gcov` helps you determine where to work on optimization.

Software developers also use coverage testing in concert with testsuites, to make sure software is actually good enough for a release. Testsuites can verify that a program works as expected; a coverage program tests to see how much of the program is exercised by the testsuite. You can then determine what kinds of test cases need to be added to the testsuites to create both better testing and a better final product.

You should compile your code without optimization if you plan to use `gcov`, because the optimization, by combining some lines of code into one function, may not give you as much information as you need to look for *hot spots*, where the code is using a great deal of computer time. Likewise, because `gcov` accumulates statistics by line (at the lowest resolution), it works best with a programming style that places only one statement on each line.

If you use complicated macros that expand to loops or to other control structures, the statistics are less helpful, since they only report on the line where the macro call appears. If your complex macros behave like functions, you can replace them with inline functions to solve this problem.

`gcov` creates a logfile called `sourcefile.gcov` which indicates how many times each line of a source file, `sourcefile.c`, has executed. You can use these logfiles in conjunction with `gprof` to aid in fine-tuning the performance of your programs. `gprof` gives timing information you can use along with the information you get from `gcov`.

`gcov` works only on code compiled with GCC; it is not compatible with any other profiling or test coverage mechanism.

Invoking the `gcov` Program

The following declaration is an example of invoking the `gcov` program.

```
gcov [-b] [-v] [-n] [-l] [-f] [-o directory] sourcefile
```

`-b`

Write branch frequencies to the output file. Write branch summary info to standard output. This option allows you to see how often each branch was taken.

`-v`

Display the `gcov` version number (on the standard error stream).

`-n`

Do not create the `gcov` output file.

`-l`

Create long filenames for included source files. For example, if the `x.h` header file contains code, and was included in the `a.c` file, then running `gcov` on the `a.c` file will produce an `a.c.x.h.gcov` output file instead of `x.h.gcov`. This can be useful if `x.h` is included in multiple source files.

- f Output summaries for each function in addition to the file level summary.
- o The directory where the object files belong. `gcov` will search for `.bb`, `.bbg`, and `.da` files in this directory.

To use `gcov`, first compile your program with `-fprofile-arcs` and `-ftest-coverage`, two special GCC options, so that the compiler generates additional information needed by `gcov` (basically a flow graph of the program) and also includes additional code in the object files for generating the extra profiling information needed by `gcov`. These additional files are placed in the directory where the source code is located.

Running `gcov` will cause profile output to be generated. For each source file compiled with `-fprofile-arcs`, an accompanying `.da` file will be placed in the source directory.

Running `gcov` with your program's source file names as arguments will now produce a listing of the code along with frequency of execution for each line. For example, if your program is called `tmp.c`, you will see the following when you use the `gcov` facility.

```
gcc -fprofile-arcs -ftest-coverage tmp.c
a.out
gcov tmp.c
    87.50% of 8 source lines executed in file tmp.c
    Creating tmp.c.gcov.
```

The `tmp.c.gcov` file contains output from `gcov`. The following is a sample of `gcov` output.

```

        main()
        {
1         int i, total;

        1         total = 0;
11        for (i = 0; i < 10; i++)
10         total += i;
1         if (total != 45)
#####         printf ("Failure\n");
        else
1         printf ("Success\n");
1     }

```

When you use the `b` option, your output looks like the following statement.

```
% gov -b tmp.c
% a.out
% gcov tmp.c
    87.50% of 8 source lines executed in file tmp.c
    80.00% of 5 branches executed in file tmp.c
```

```

80.00% of 5 branches taken at least once in file tmp.c
50.00% of 2 calls executed in file tmp.c
Creating tmp.c.gcov.

```

The following is an example of a resulting `tmp.c.gcov` file.

```

        main()
        {
1         int i, total;

1         total = 0;

11        for (i = 0; i < 10; i++)
branch 0 taken = 91%
branch 1 taken = 100%
branch 2 taken = 100%
10        total += i;

1         if (total != 45)
branch 0 taken = 100%
#####          printf ("Failure\n");
call 0 never executed
branch 1 never executed
            else
1            printf ("Success\n");

call 0 returns = 100%
1        }

```

For each basic block, a line is printed after the last line of the basic block describing the branch or call that ends the basic block. There can be multiple branches and calls listed for a single source line if there are multiple basic blocks that end on that line. In this case, the branches and calls are each given a number. There is no simple way to map these branches and calls back to source constructs. In general, though, the lowest numbered branch or call will correspond to the leftmost construct on the source line.

For a branch, if it was executed at least once, then a percentage indicating the number of times the branch was taken divided by the number of times the branch was executed will be printed. Otherwise, a “**never executed**” message is printed.

For a call, if it was executed at least once, then a percentage indicating the number of times the call returned divided by the number of times the call was executed will be printed.

This will usually be 100%, but may be less for `exit` or `longjmp` functions, and thus may not return every time they are called.

The execution counts are cumulative. If the example program were executed again without removing the `.da` file, the count for the number of times each line in the source was executed would be added to the results of the previous run(s). This is potentially useful in several ways. For example, it could be used to accumulate data over a number of program runs as part of a test verification suite, or to provide more

accurate long-term information over a large number of program runs.

The data in the `.da` files is saved immediately before the program exits. For each source file compiled with `-fprofile-arcs`, the profiling code first attempts to read in an existing `.da` file; if the file does not match the executable (because of differing number of basic block counts), it will ignore the contents of the file. It then adds in the new execution counts and finally writes the data to the file.

Using `gcov` with GCC Optimization

If you plan to use `gcov` to help optimize your code, you must first compile your program with the `-fprofile-arcs` and `-ftest-coverage` options. Aside from that, you can use any other GCC options; but if you want to prove that every single line in your program was executed, you should not compile with optimization at the same time. On some machines the optimizer can eliminate some simple code lines by combining them with other lines. Like the following example shows, code can be compiled into one instruction on some machines.

```
if (a != b)
    c =1;
else
    c =0;
```

In such a case, there is no way for `gcov` to calculate separate execution counts for each line because there is not separate code for each line. The `gcov` output then looks like the following declaration if you compiled the program with optimization.

```
100 if (a != b)
100 c = 1;
100 else
100 c = 0;
```

Such output shows that this block of code, combined by optimization, executed 100 times. In one sense this result is correct, because there was only one instruction representing all four of these lines. However, the output does not indicate how many times the result was 0 and how many times the result was 1.

Brief Description of `gcov` Data Files

`gcov` uses three files for doing profiling. The names of these files are derived from the original *source* file by substituting the file suffix with either `.bb`, `.bbg`, or `.da`. All of these files are placed in the same directory as the source file, and contain data stored in a platform-independent method.

The `.bb` and `.bbg` files are generated when the source file is compiled with the GCC `-ftest-coverage` option. The `.bb` file contains a list of source files (including headers), functions within those files, and line numbers corresponding to each basic

block in the source file.

The `.bb` file format consists of several lists of 4-byte integers which correspond to the line numbers of each basic block in the file. Each list is terminated by a line number of 0. A line number of -1 is used to designate that the source file name (padded to a 4-byte boundary and followed by another -1) follows. In addition, a line number of -2 is used to designate that the name of a function follows (also padded to a 4-byte boundary and followed by a -2).

The `.bbg` file is used to reconstruct the program flow graph for the source file. It contains a list of the program flow arcs (possible branches taken from one basic block to another) for each function which, in combination with the `.bb` file, enables `gcov` to reconstruct the program flow.

In the `.bbg` file, the format is the following declaration.

```
number of basic blocks for function #0 (4-byte number)
total number of arcs for function #0 (4-byte number)
count of arcs in basic block #0 (4-byte number)
destination basic block of arc #0 (4-byte number)
flag bits (4-byte number)
destination basic block of arc #1 (4-byte number)
flag bits (4-byte number)
...
destination basic block of arc #N (4-byte number)
flag bits (4-byte number)
count of arcs in basic block #1 (4-byte number)
destination basic block of arc #0 (4-byte number)
flag bits (4-byte number)
...
```

A -1 (stored as a 4-byte number) is used to separate each function's list of basic blocks, and to verify that the file has been read correctly.

The `.da` file is generated when a program containing object files built with the GCC `-fprofile-arcs` option is executed. A separate `.da` file is created for each source file compiled with this option, and the name of the `.da` file is stored as an absolute pathname in the resulting object file. This path name is derived from the source file name by substituting a `.da` suffix. The format of the `.da` file is fairly simple. The first 8-byte number is the number of counts in the file, followed by the counts (stored as 8-byte numbers). Each count corresponds to the number of times each arc in the program is executed. The counts are cumulative; each time the program is executed, it attempts to combine the existing `.da` files with the new counts for this invocation of the program. It ignores the contents of any `.da` files whose number of arcs does not correspond to the current program, and merely overwrites them instead.

All three of these files use the functions in `gcov-io.h` to store integers; the functions in this header provide a machine-independent mechanism for storing and retrieving data from a stream.

The C Preprocessor

1

Overview of the C Preprocessor

The C preprocessor is a *macro processor* that is used automatically for C compiling to transform your program before actual compilation. It is called a macro processor because it allows you to define macros, which are abbreviations for longer constructs. The following documentation discusses the GNU C preprocessor, the C-compatible compiler preprocessor.

- “What the C Preprocessor Provides” on page 218
- “Transformations Made Globally” on page 219
- “Preprocessing Directives” on page 221
- “Header Files” on page 223
- “Macros” on page 229
- “Conditionals” on page 249
- “Combining Source Files” on page 257
- “Other Preprocessing Directives” on page 259
- “C Preprocessor Output” on page 261
- “Invoking the C Preprocessor” on page 263

What the C Preprocessor Provides

The C preprocessor provides the following four separate facilities that you can use as you see fit.

- ***Inclusion of header files***

These are files of declarations that can be substituted into your program.

- ***Macro expansion***

You can define *macros*, which are abbreviations for arbitrary fragments of C code, and then the C preprocessor will replace the macros with their definitions throughout the program.

- ***Conditional compilation***

Using special preprocessing directives, you can include or exclude parts of the program according to various conditions.

- ***Line control***

If you use a program to combine or rearrange source files into an intermediate file, which is then compiled, you can use line control to inform the compiler of each source line's origin.

C preprocessors vary in some details. The GNU C preprocessor provides a superset of the features of ANSI Standard C.

ANSI Standard C requires the rejection of many harmless constructs commonly used by today's C programs. Such incompatibility would be inconvenient for users, so the GNU C preprocessor is configured to accept these constructs by default. To get ANSI Standard C, you must use the options, `-trigraphs`, `-undef` and `-pedantic`, but, in practice, the consequences of having strict ANSI Standard C make it undesirable to follow this practice. For more details, see “Invoking the C Preprocessor” on page 263.

2

Transformations Made Globally

Most C preprocessor features are inactive unless you give specific directives to request their use. (Preprocessing directives are lines starting with #; see “Preprocessing Directives” on page 221). But there are three transformations that the preprocessor always makes on all the input it receives, even in the absence of directives.

- All C comments are replaced with single spaces.
- Backslash-Newline sequences are deleted, no matter where. This feature allows you to break long lines for cosmetic purposes without changing their meaning.
- Predefined macro names are replaced with their expansions (see “Predefined Macros” on page 233).

The first two transformations are done *before* nearly all other parsing and before preprocessing directives are recognized. Thus, for example, you can split a line cosmetically with Backslash-Newline anywhere (except when trigraphs are in use; see the following example and its description).

```
/*  
*/ # /*  
*/ defi\  
ne FO\  
  0 10\  
  20
```

This input has the equivalent of `#define FOO 1020`. You can split an escape sequence

with Backslash-Newline. For example, you can split "foo\bar" between the \ and the b to get the following sequence.

```
"foo\  
bar"
```

This behavior is unclean: in all other contexts, a Backslash can be inserted in a string constant as an ordinary character by writing a double Backslash, and this creates an exception. But the ANSI C standard requires it. (Strict ANSI C does not allow Newlines in string constants, so they do not consider this a problem.)

There are a few exceptions to all three transformations.

- C comments and predefined macro names are not recognized inside a `#include` directive in which the file name is delimited with `<` and `>`.
- C comments and predefined macro names are never recognized within a character or string constant. (Strictly speaking, this is the rule, not an exception, but it is worth noting here anyway.)
- Backslash-Newline may not safely be used within an ANSI trigraph. Trigraphs are converted before Backslash-Newline is deleted. If you write what looks like a trigraph with a Backslash-Newline inside, the Backslash-Newline is deleted as usual, but it is then too late to recognize the trigraph.

This exception is relevant only if you use the `-trigraphs` option to enable trigraph processing. See “Invoking the C Preprocessor” on page 263.

3

Preprocessing Directives

Most preprocessor features are active only if you use preprocessing directives to request their use. Preprocessing directives are lines in your program that start with `#`. The `#` is followed by an *identifier*, which is the directive name. For example, `#define` is the directive that defines a macro. Whitespace is also allowed before and after the `#`. The set of valid directive names is fixed. Programs cannot define new preprocessing directives. Some directive names require arguments; these make up the rest of the directive line and must be separated from the directive name by whitespace. For example, `#define` must be followed by a macro name and the intended expansion of the macro. See “Simple Macros” on page 229.

A preprocessing directive cannot be more than one line in normal circumstances. It may be split cosmetically with Backslash-Newline, but that has no effect on its meaning. Comments containing Newlines can also divide the directive into multiple lines, but the comments are changed to Spaces before the directive is interpreted. The only way a significant Newline can occur in a preprocessing directive is within a string constant or character constant.

IMPORTANT! C compilers that are applied to the output from the preprocessor do not accept string or character constants containing Newlines.

The `#` and the directive name cannot come from a macro expansion; if `foo` is defined as a macro expanding to `define`, `#foo` does not become a valid preprocessing directive.

4

Header Files

A header file is a file containing C declarations and macro definitions (see “Macros” on page 229) to be shared between several source files. You request the use of a header file in your program with the C preprocessing `#include` directive.

The following documentation describes more about header files.

- ““Uses of Header Files”” (below)
- “The `#include` Directive” on page 224
- “How `#include` Works” on page 225
- “Once-only Include Files” on page 226
- “Inheritance and Header Files” on page 227

Uses of Header Files

Header files serve two kinds of purposes.

- System header files declare the interfaces to parts of the operating system. You include them in your program to supply the definitions and declarations you need to invoke system calls and libraries.

- Your own header files contain declarations for interfaces between the source files of your program. Each time you have a group of related declarations and macro definitions all or most of which are needed in several different source files, it is a good idea to create a header file for them.

Including a header file produces the same results in C compilation as copying the header file into each source file that needs it. But such copying would be time-consuming and error-prone. With a header file, the related declarations appear in only one place. If they need to be changed, they can be changed in one place, and programs that include the header file will automatically use the new version when next recompiled. The header file eliminates the labor of finding and changing all the copies as well as the risk that a failure to find one copy will result in inconsistencies within a program.

The usual convention is to give header files names that end with `.h`. Avoid unusual characters in header file names, as they reduce portability.

The #include Directive

Both user and system header files are included using the preprocessing directive `#include`. It has three variants:

- `#include <file>`
This variant is used for system header files. It searches for a file named *file* in a list of directories specified by you, then in a standard list of system directories. You specify directories to search for header files with the command option `-I` (see “Invoking the C Preprocessor” on page 263). The option `-nostdinc` inhibits searching the standard system directories; in this case only the directories you specify are searched.

The parsing of this form of `#include` is slightly special because comments are not recognized within the `<...>`.

Thus, in `#include <x/*y>`, the `/*` does not start a comment and the directive specifies inclusion of a system header file named `x/*y`. Of course, a header file with such a name is unlikely to exist on UNIX, where shell wildcard features would make it hard to manipulate.

The argument *file* may not contain a `>` character. It may, however, contain a `<` character.

- `#include "file"`
This variant is used for header files of your own program. It searches for a file named *file*, first in the current directory, then in the same directories used for system header files. The current directory is the directory of the current input file. It is tried first because it is presumed to be the location of the files that the current input file refers to. (If the `-I-` option is used, the special treatment of the current directory is inhibited.)

The argument, *file*, may not contain ‘\’ characters. If backslashes occur within *file*, they are considered ordinary text characters, not escape characters. None of the character escape sequences appropriate to string constants in C are processed. Thus, `#include "x\n\\y"` specifies a filename containing three backslashes. It is not clear why this behavior is ever useful, although the ANSI standard specifies it.

- `#include anything else`
This variant is called a computed `#include`. Any `#include` directive whose argument does not fit the above two forms is a computed include. The text *anything else*, is checked for macro calls, which are expanded (see “Macros” on page 229). When this is done, the result must fit one of the previous two variants—in particular, the expanded text must in the end be surrounded by either quotes or angle braces.

This feature allows you to define a macro which controls the file name to be used at a later point in the program. One application of this is to allow a site-specific configuration file for your program to specify the names of the system `include` files to be used. This can help in porting the program to various operating systems in which the necessary system header files are found in different places.

How `#include` Works

The `#include` directive works by directing the C preprocessor to scan the specified file as input before continuing with the rest of the current file. The output from the preprocessor contains the output already generated, followed by the output resulting from the included file, followed by the output that comes from the text after the `#include` directive. The example, given a header file `header.h`, follows.

```
char *test ();
```

Then, there is a main program called `program.c` that uses the header file, like the following.

```
int x;
#include "header.h"

main ()
{
    printf (test ());
}
```

The output generated by the C preprocessor for `program.c` as input would be as follows.

```
int x;
char *test ();
main ()
{
    printf (test ());
}
```

Included files are not limited to declarations and macro definitions; those are merely the typical uses. Any fragment of a C program can be included from another file. The include file could even contain the beginning of a statement that is concluded in the containing file, or the end of a statement that was started in the including file.

However, a comment or a string or character constant may not start in the included file and finish in the including file. An unterminated comment, string constant or character constant in an included file is considered to end (with an error message) at the end of the file.

It is possible for a header file to begin or end a syntactic unit such as a function definition, but that would be very confusing, so don't do it.

The line following the `#include` directive is always treated as a separate line by the C preprocessor even if the included file lacks a final newline.

Once-only Include Files

Very often, one header file includes another header file. It can easily result that a certain header file is included more than once in a file. This may lead to errors, if the header file defines structure types or typedefs, and is certainly wasteful. Therefore, we often try to prevent multiple inclusion of a header file. The standard way to do this is to enclose the entire real contents of the file in a conditional, like the following example's input demonstrates.

```
#ifndef FILE_FOO_SEEN
#define FILE_FOO_SEEN

    the entire file

#endif /* FILE_FOO_SEEN */
```

The macro, `FILE_FOO_SEEN`, indicates that the file has been included *once already*. In a user header file, the macro name should not begin with an underscore, `'_'`. In a system header file, this name should begin with `'__'` (two underscores) to avoid conflicts with user programs. In any kind of header file, the macro name should contain the name of the file and some additional text, to avoid conflicts with other header files.

The GNU C preprocessor is programmed to notice when a header file uses this particular construct and handle it efficiently. If a header file is contained entirely in a `#ifndef` conditional, then it records that fact. If a subsequent `#include` specifies the same file, and the macro in the `#ifndef` is already defined, then the file is entirely skipped, without even reading it.

There is also an explicit directive to tell the preprocessor that it need not include a file more than once. This is called `#pragma once`, and was used in addition to the `#ifndef` conditional around the contents of the header file. `#pragma once` is now obsolete and

should not be used at all.

In the Objective C language, there is a variant of `#include` called `#import` which includes a file, but does so at most once. If you use `#import` instead of `#include`, then you don't need the conditionals inside the header file to prevent multiple execution of the contents. `#import` is obsolete because it is not a well designed feature. Using `#ifndef` meets the requirement in a more straightforward manner.

Inheritance and Header Files

Inheritance is what happens when one object or file derives some of its contents by virtual copying from another object or file. In the case of C header files, inheritance means that one header file includes another header file and then replaces or adds something.

If the inheriting header file and the base header file have different names, then inheritance is straightforward: simply write `#include "base"` in the inheriting file (where *base* stands for the base file in use).

Sometimes it is necessary to give the inheriting file the same name as the base file. This is less straightforward.

For example, suppose an application program uses the system header file `sys/signal.h`, but the version of `/usr/include/sys/signal.h` on a particular system doesn't do what the application program expects. It might be convenient to define a *local* version, perhaps under the name `/usr/local/include/sys/signal.h`, to override or add to the one supplied by the system. Use the option, `-I.`, for compilation, and writing a file `sys/signal.h` that does what the application program expects. But making this file include the standard `sys/signal.h` is not so easy—writing `#include <sys/signal.h>` in that file doesn't work, because it includes your own version of the file, not the standard system version.

Used in that file itself, this leads to an infinite recursion and a fatal error in compilation.

`#include </usr/include/sys/signal.h>` would find the proper file, but that is not clean, since it makes an assumption about where the system header file is found. This is bad for maintenance, since it means that any change in where the system's header files are kept requires a change somewhere else.

The clean way to solve this problem is to use `#include_next`, which means, "Include the *next* file with this name." This directive works like `#include` except in searching for the specified file: it starts searching the list of header file directories after the directory in which the current file was found.

Suppose you specify `-I /usr/local/include`, and the list of directories to search also includes `/usr/include`; and suppose that both directories contain a file named

`sys/signal.h`. Ordinary `#include <sys/signal.h>` finds the file under `/usr/local/include`. If that file contains `#include_next <sys/signal.h>`, it starts searching after that directory, and finds the file in `/usr/include`.

5

Macros

A *macro* is a sort of abbreviation that you can define once and later use. There are many complicated features associated with macros in the C preprocessor. The following documentation describes more about macros.

- ““Simple Macros”” (below)
- “Macros with Arguments” on page 231
- “Predefined Macros” on page 233
- “Stringification” on page 237
- “Concatenation” on page 238
- “Undefining Macros” on page 239
- “Redefining Macros” on page 240
- “Pitfalls and Subtleties of Macros” on page 240

Simple Macros

A *simple macro* is a kind of abbreviation, a name which stands for a fragment of code. Some people refer to these as *manifest constants*. Before you can use a macro, you must define it explicitly with the `#define` directive. `#define` is followed by the name of the macro and then the code for which it should abbreviate; for example:

```
#define BUFFER_SIZE 1020
```

This input defines a macro named `BUFFER_SIZE` as an abbreviation for the text `1020`. If somewhere after this `#define` directive there comes a C statement of the form of the following example, then the C preprocessor will recognize and expand the macro `BUFFER_SIZE`.

```
foo = (char *) xmalloc (BUFFER_SIZE);
```

This then gives the following result.

```
foo = (char *) xmalloc (1020);
```

The use of all uppercase for macro names is a standard convention so that it is possible to tell at a glance which names are macros.

Normally, a macro definition must be a single line, like all C preprocessing directives. (You can split a long macro definition cosmetically with Backslash-Newline.) There is one exception: *newlines* can be included in the macro definition if within a string or character constant. This is because it is not possible for a macro definition to contain an unbalanced quote character; the definition automatically extends to include the matching quote character that ends the string or character constant. Comments within a macro definition may contain newlines, which make no difference since the comments are entirely replaced with spaces regardless of their contents.

Aside from the previous explanation, there is no restriction on what can go in a macro body. Parentheses need not balance. The body need not resemble valid C code. (But if it does not, you may get error messages from the C compiler when you use the macro.)

The C preprocessor scans your program sequentially, so macro definitions take effect at the place you write them. First, input the following.

```
foo = X;
#define X 4
bar = X;
```

This produces the following output.

```
foo = X;

bar = 4;
```

After the preprocessor expands a macro name, the macro's definition body is appended to the front of the remaining input, and the check for macro calls continues. Therefore, the macro body can contain calls to other macros.

Use the following, for example, as input.

```
#define BUFSIZE 1020
#define TABLESIZE BUFSIZE
```

Afterwards, name `TABLESIZE` when used in the program would go through two stages of expansion, resulting ultimately in `1020`.

This is not at all the same as defining `TABLESIZE` to be `1020`. The `#define` for `TABLESIZE` uses exactly the body you specify—in this case, `BUFSIZE`—and does not check to see whether it too is the name of a macro. It's only when you use `TABLESIZE` that the result of its expansion is checked for more macro names. See .

Macros with Arguments

A simple macro always stands for exactly the same text, each time it is used. Macros can be more flexible when they accept *arguments*. Arguments are fragments of code that you supply each time the macro is used. These fragments are included in the expansion of the macro according to the directions in the macro definition. A macro that accepts arguments is called a function-like macro because the syntax for using it looks like a function call.

To define a macro that uses arguments, you write a `#define` directive with a list of *argument names* in parentheses after the name of the macro. The argument names may be any valid C identifiers, separated by commas and optionally whitespace. The open parenthesis must follow the macro name immediately, with no space in between. For example, here is a macro that computes the minimum of two numeric values, as it is defined in many C programs:

```
#define min(X, Y) ((X) < (Y) ? (X) : (Y))
```

This is not the best way to define a *minimum* macro in GNU C; see “Duplication of Side Effects” on page 243 for more information. To use a macro that expects arguments, you write the name of the macro followed by a list of *actual arguments* in parentheses, separated by commas. The number of actual arguments you give must match the number of arguments the macro expects.

Examples of using the macro `min` include `min (1, 2)` and `min (x + 28, *p)`.

The expansion text of the macro depends on the arguments you use. Each of the argument names of the macro is replaced, throughout the macro definition, with the corresponding actual argument. Using the same macro, `min` (as defined in the previous input example), `min (1, 2)` expands into the following output.

```
((1) < (2) ? (1) : (2))
```

`1` has been substituted for `x` and `2` for `y`. Likewise, `min (x + 28, *p)` expands into the following output.

```
((x + 28) < (*p) ? (x + 28) : (*p))
```

Parentheses in the actual arguments must balance; a comma within parentheses does not end an argument. However, there is no requirement for brackets or braces to balance, and they do not prevent a comma from separating arguments. Use the following input as an example.

```
macro (array[x = y, x + 1])
```

This passes two arguments to `macro`: `array[x = y` and `x + 1]`. If you want to supply `array[x = y, x + 1]` as an argument, you must write it as equivalent C code with `array[(x = y, x + 1)]`.

After the actual arguments are substituted into the macro body, the entire result is appended to the front of the remaining input, and the check for macro calls continues. Therefore, the actual arguments can contain calls to other macros, either with or

without arguments, or even to the same macro. The macro body can also contain calls to other macros.

For example, `min (min (a, b), c)` expands into the following output.

```
((((a) < (b) ? (a) : (b))) < (c)
 ? (((a) < (b) ? (a) : (b)))
 : (c))
```

If a macro `foo` takes one argument, and you want to supply an empty argument, you must write at least some whitespace between the parentheses, like this: `foo ()`. Just `foo ()` is providing no arguments, which is an error if `foo` expects an argument.

But `foo0 ()` is the correct way to call a macro defined to take zero arguments, like the following example for input.

```
#define foo0() : : :
```

If you use the macro name followed by something other than an open-parenthesis (after ignoring any spaces, tabs and comments that follow), it is not a call to the macro, and the preprocessor does not change what you have written. Therefore, it is possible for the same name to be a variable or function in your program as well as a macro, and you can choose in each instance whether to refer to the macro (if an actual argument list follows) or the variable or function (if an argument list does not follow).

Such dual use of one name could be confusing and should be avoided except when the two meanings are effectively synonymous: that is, when the name is both a macro and a function and the two have similar effects. You can think of the name simply as a function; use of the name for purposes other than calling it (such as, to take the address) will refer to the function, while calls will expand the macro and generate better but equivalent code. For example, you can use a function named `min` in the same source file that defines the macro. If you write `&min` with no argument list, you refer to the function. If you write `min (x, bb)`, with an argument list, the macro is expanded. If you write `(min) (a, bb)`, where the name `min` is not followed by an open-parenthesis, the macro is not expanded, so you wind up with a call to the function `min`.

You may not define the same name as both a simple macro and a macro with arguments. In the definition of a macro with arguments, the list of argument names must follow the macro name immediately with no space in between your input. If there is a space after the macro name, the macro is defined as taking no arguments, and all the rest of the line is taken to be the *expansion*. The reason for this is that it is often useful to define a macro that takes no arguments and whose definition begins with an identifier in parentheses. This rule about spaces would use input like either one of the two following examples.

```
#define FOO(x) - 1 / (x)
```

The previous example defines `FOO` to take an argument and expand into minus the reciprocal of that argument.

The following example defines `BAR` to take no argument and always expand into `(x)`

```
-1 / (x)).
#define BAR (x) - 1 / (x)
```

IMPORTANT! The uses of a macro with arguments can have spaces before the left parenthesis; it's the *definition* where it matters whether there is a space.

Predefined Macros

Several simple macros are predefined. You can use them without giving definitions for them. They fall into two classes: standard macros and system-specific macros.

Standard Predefined Macros

The standard predefined macros are available with the same meanings regardless of the machine or operating system on which you are using GNU C. Their names all start and end with double underscores. Those preceding `__GNUC__` in this table are standardized by ANSI C; the rest are GNU C extensions.

`__FILE__`

This macro expands to the name of the current input file, in the form of a C string constant. The precise name returned is the one that was specified in `#include` or as the input file name argument.

`__LINE__`

This macro expands to the current input line number, in the form of a decimal integer constant. While we call it a predefined macro, it's a pretty strange macro, since its "definition" changes with each new line of source code.

This and `__FILE__` are useful in generating an error message to report an inconsistency detected by the program; the message can state the source line at which the inconsistency was detected.

Use the following output as an example.

```
fprintf (stderr, "Internal error: "
          "negative string length "
          "%d at %s, line %d.",
          length, __FILE__, __LINE__);
```

A `#include` directive changes the expansions of `__FILE__` and `__LINE__` to correspond to the included file. At the end of that file, when processing resumes on the input file that contained the `#include` directive, the expansions of `__FILE__` and `__LINE__` revert to the values they had before the `#include` (but `__LINE__` is then incremented by one as processing moves to the line after the `#include`).

The expansions of both `__FILE__` and `__LINE__` are altered if a `#line` directive is used. See "Combining Source Files" on page 257.

__DATE__

This macro expands to a string constant that describes the date on which the preprocessor is being run. The string constant contains eleven characters and looks like "Jan 29 1987" or "Apr 1 1905".

__TIME__

This macro expands to a string constant that describes the time at which the preprocessor is being run. The string constant contains eight characters and looks like "23:59:01".

__STDC__

This macro expands to the constant 1, to signify that this is ANSI Standard C. (Whether that is actually true depends on what C compiler will operate on the output from the preprocessor.)

__STDC_VERSION__

This macro expands to the C Standard's version number, a long integer constant of the form `yyyymmL` where `yyyy` and `mm` are the year and month of the Standard version. This signifies which version of the C Standard the preprocessor conforms to. Like `__STDC__`, whether this version number is accurate for the entire implementation depends on what C compiler will operate on the output from the preprocessor.

__GNUC__

This macro is defined if and only if this is GNU C. This macro is defined only when the entire GNU C compiler is in use; if you invoke the preprocessor directly, `__GNUC__` is undefined. The value identifies the major version number of GNU CC (1 for GNU CC version 1, which is now obsolete, and 2 for version 2).

__GNUC_MINOR__

The macro contains the minor version number of the compiler. This can be used to work around differences between different releases of the compiler (for example, if GCC 2.6.3 supports a feature, you can test for `__GNUC__ > 2 || (__GNUC__ == 2 && __GNUC_MINOR__ >= 6)`). The last number, 3, denotes the bugfix level of the compiler; no macro contains this value.

__GNUG__

The GNU C compiler defines this when the compilation language is C++; use `__GNUG__` to distinguish between GNU C and GNU C++.

__cplusplus

The draft ANSI standard for C++ *used to* require predefining this variable. Though it is no longer required, GNU C++ continues to define it, as do other popular C++ compilers. You can use `__cplusplus` to test whether a header is compiled by a C compiler or a C++ compiler.

__STRICT_ANSI__

This macro is defined if and only if the `-ansi` switch was specified when GNU C was invoked. Its definition is the null string. This macro exists primarily to direct certain GNU header files not to define certain traditional UNIX constructs which are incompatible with ANSI C.

__BASE_FILE__

This macro expands to the name of the main input file, in the form of a C string constant. This is the source file that was specified as an argument when the C compiler was invoked.

__INCLUDE_LEVEL__

This macro expands to a decimal integer constant that represents the depth of nesting in `include` files. The value of this macro increments on every `#include` directive and decremented at every end of file. For input files specified by command line arguments, the nesting level is zero.

__VERSION__

This macro expands to a string which describes the version number of GNU C. The string is normally a sequence of decimal numbers separated by periods, such as "2.6.0".

__OPTIMIZE__

This macro is defined in optimizing compilations. It causes certain GNU header files to define alternative macro definitions for some system library functions. It is unwise to refer to or test the definition of this macro unless you make very sure that programs will execute with the same effect regardless.

__CHAR_UNSIGNED__

This macro is defined if and only if the data type `char` is unsigned on the target machine. It exists to cause the standard header file `limits.h` to work correctly. It is bad practice to refer to this macro yourself; instead, refer to the standard macros defined in `limits.h`. The preprocessor uses this macro to determine whether or not to sign-extend large character constants written in octal; see “The `#if` Directive” on page 250.

__REGISTER_PREFIX__

This macro expands to a string (not a string constant) describing the prefix applied to CPU registers in assembler code. It can be used to write assembler code that is usable in multiple environments. For example, in the `m68k-aout` environment it expands to the NULL string. In the `m68k-coff` environment, it expands to the string, "%".

__USER_LABEL_PREFIX__

This macro expands to a string describing the prefix applied to user generated labels in assembler code. It can be used to write assembler code that is usable in multiple environments. For example, in the `m68k-aout` environment it expands to the string "_", but in the `m68k-coff` environment, it expands to the string "".

Non-standard Predefined Macros

The C preprocessor normally has several predefined macros that vary between machines because their purpose is to indicate what type of system and machine is in use.

This documentation, being for all systems and machines, cannot tell you exactly what

their names are. Instead, we offer a list of some typical ones.

You can use `cpp -dM` to see the values of predefined macros; for more information, see “Invoking the C Preprocessor” on page 263. Some nonstandard predefined macros describe the operating system in use, with more or less specificity, as in the following two examples.

`unix`

`unix` is normally predefined on all UNIX systems.

`BSD`

`BSD` is predefined on recent versions of Berkeley UNIX (perhaps only in version 4.3).

Other nonstandard predefined macros describe the kind of CPU, with more or less specificity, as in the following six examples.

`vax`

`vax` is predefined on Vax computers.

`mc68000`

`mc68000` is predefined on most computers whose CPU is a Motorola 68000, 68010 or 68020.

`m68k`

`m68k` is also predefined on most computers whose CPU is a 68000, 68010 or 68020; however, some makers use `mc68000` and some use `m68k` as the names for the macros. Some predefine both names. What happens in GNU C depends on the system you are using.

`M68020`

`M68020` has been observed to be predefined on some systems that use 68020 CPUs—in addition to `mc68000` and `m68k`, which are less specific.

`_AM29K`

`_AM29000`

Both `_AM29K` and `_AM29000` are predefined for the AMD 29000 CPU family.

`ns32000`

`ns32000` is predefined on computers which use the National Semiconductor 32000 series CPU.

Yet other nonstandard predefined macros describe the manufacturer of the system, as in the following three examples.

`sun`

`sun` is predefined on all models of Sun computers.

`pyr`

`pyr` is predefined on all models of Pyramid computers.

`sequent`

`sequent` is predefined on all models of Sequent computers.

These predefined symbols are not only nonstandard, they are contrary to the ANSI standard because their names do not start with underscores. Therefore, the option `-ansi` inhibits the definition of these symbols.

This tends to make `-ansi` useless, since many programs depend on the customary nonstandard predefined symbols. Even system header files check them and will generate incorrect declarations if they do not find the names that are expected. You might think that the header files supplied for the Uglix computer would not need to test what machine they are running on, because they can simply assume it is the Uglix; but often they do, and they do so using the customary names. As a result, very few C programs will compile with `-ansi`. We intend to avoid such problems on the GNU system.

What, then, should you do in an ANSI C program to test the type of machine it will run on? GNU C offers a parallel series of symbols for this purpose, whose names are made from the customary ones by adding ‘`__`’ at the beginning and end. Thus, the symbol `__vax__` would be available on a Vax, and so on. The set of nonstandard predefined names in the GNU C preprocessor is controlled (when `cpp` is itself compiled) by the macro `CPP_PREDEFINES`, which should be a string containing `-D` options, separated by spaces. For example, on the Sun 3, we use the following definition.

```
#define CPP_PREDEFINES "-Dmc68000 -Dsun -Dunix -Dm68k"
```

This macro is usually specified in `tm.h`.

Stringification

Stringification means turning a code fragment into a string constant whose contents are the text for the code fragment. For example, stringifying `foo (z)` results in `"foo (z)"`.

In the C preprocessor, stringification is an option available when macro arguments are substituted into the macro definition. In the body of the definition, when an argument name appears, the `#` character before the name specifies stringification of the corresponding actual argument when it is substituted at that point in the definition. The same argument may be substituted in other places in the definition without stringification if the argument name appears in those places with no `#`.

What follows is an example of a macro definition that uses stringification.

```
#define WARN_IF(EXP) \
do { if (EXP) \
    fprintf (stderr, "Warning: " #EXP "\n"); } \
while (0)
```

The actual argument for `EXP` is substituted once as given, into the `if` statement, and once as stringified, into the argument to `fprintf`. The `do` and `while (0)` are a work-around to make it possible to write `WARN_IF(arg);`, which the resemblance of `WARN_IF` to a function would make C programmers want to do; see “Swallowing the Semicolon” on page 242.

The stringification feature is limited to transforming one macro argument into one

string constant; there is no way to combine the argument with other text and then stringify it all together. The previous example shows how an equivalent result can be obtained in ANSI Standard C using the feature that adjacent string constants are concatenated as one string constant. The preprocessor stringifies the actual value of `EXP` into a separate string constant, resulting in text like the following output.

```
do {if (x==0)\
    fprintf (stderr, "Warning: " "x == 0" "\n"); } \
while (0)
```

The C compiler then sees three consecutive string constants and concatenates them into one, producing, effectively, the following output.

```
do {if (x==0)\
    fprintf (stderr, "Warning: x == 0\n"); } \
while (0)
```

Stringification in C involves more than putting doublequote characters around the fragment; it is necessary to put backslashes in front of all doublequote characters, and all backslashes in string and character constants, in order to get a valid C string constant with the proper contents. Thus, stringifying `p="foo\n"` results in `"p=\"foo\\n\""`. However, backslashes that are not inside of string or character constants are not duplicated: `\n` by itself stringifies to `"\n"`.

Whitespace (including comments) in the text being stringified is handled according to precise rules. All leading and trailing whitespace is ignored. Any sequence of whitespace in the middle of the text is converted to a single space in the stringified result.

Concatenation

In the context of macro expansion, *concatenation* refers to joining two lexical units or two strings into one longer one. Specifically, an actual argument to the macro can be concatenated with another actual argument or with fixed text to produce a longer name. The longer name might be the name of a function, variable or type, or a C keyword; it might even be the name of another macro, in which case it will expand.

When you define a macro, you request concatenation with the special operator `##` in the macro body. When the macro is called, after actual arguments are substituted, all `##` operators are deleted, and any whitespace next to them (including whitespace that was part of an actual argument). The result is to concatenate the syntactic tokens on either side of the `##`. Consider a C program interpreting named commands. There likely needs to be a table of commands, or an array of structures, declared as follows.

```
struct command
{
    char *name;
    void (*function) ();
};
```

```

struct command commands[] =
{
    { "quit", quit_command},
    { "help", help_command},
    ...
};

```

It would be cleaner not to have to give each command name twice, once in the string constant and once in the function name. A macro which takes the name of a command as an argument can make this unnecessary. The string constant can be created with stringification, and the function name by concatenating the argument with `_command`. What follows is an example of how it is done.

```

#define COMMAND(NAME)    { #NAME, NAME ## _command }

struct command commands[] =
{
    COMMAND (quit),
    COMMAND (help),
    ...
};

```

The usual case of concatenation is concatenating two names (or a name and a number) into a longer name. But this isn't the only valid case. It is also possible to concatenate two numbers (or a number and a name, such as `1.5` and `e3`) into a number. Also, multi-character operators such as `+=` can be formed by concatenation. In some cases it is even possible to piece together a string constant. However, two pieces of text that don't together form a valid lexical unit cannot be concatenated. For example, concatenation with `x` on one side and `+` on the other is not meaningful because those two characters can't fit together in any lexical unit of C. The ANSI standard says that such attempts at concatenation are undefined, but in the GNU C preprocessor it is well defined: it puts `x` and `+` side by side with no particular special results.

Keep in mind that the C preprocessor converts comments to whitespace before macros are even considered. Therefore, you cannot create a comment by concatenating `/` and `*`: the `/*` sequence that starts a comment is not a lexical unit, but rather the beginning of a "long" space character. Also, you can freely use comments next to `##` in a macro definition, or in actual arguments that will concatenate, because the comments will be convert to spaces at first sight, and concatenation will later discard the spaces.

Undefining Macros

To *undefine* a macro means to cancel its definition. This is done with the `#undef` directive. `#undef` is followed by the macro name to be undefined.

Like definition, *undefinition* occurs at a specific point in the source file, and it applies starting from that point. The name ceases to be a macro name, and from that point on

it is treated by the preprocessor as if it had never been a macro name. For clarification, use the following example.

```
#define FOO 4 x = FOO; #undef FOO x = FOO;
```

This input expands into the following output.

```
x = 4;  
x = FOO;
```

In the previous example, `FOO` had better be a variable or function as well as (temporarily) a macro, in order for the result of the expansion to be valid C code. The same form of `#undef` directive will cancel definitions with arguments or definitions that don't expect arguments. The `#undef` directive has no effect when used on a name not currently defined as a macro.

Redefining Macros

Redefining a macro means defining (with `#define`) a name that is already defined as a macro. A redefinition is trivial if the new definition is transparently identical to the old one. You probably wouldn't deliberately write a trivial redefinition, but they can happen automatically when a header file is included more than once (see "Header Files" on page 223), so they are accepted silently and without effect.

Nontrivial redefinition is considered likely to be an error, so it provokes a warning message from the preprocessor. However, sometimes it is useful to change the definition of a macro in mid-compilation. You can inhibit the warning by undefining the macro with `#undef` before the second definition. In order for a redefinition to be trivial, the new definition must exactly match the one already in effect, with two possible exceptions, as in the following.

- Whitespace may be added or deleted at the beginning or the end.
- Whitespace may be changed in the middle (but not inside strings). However, it may not be eliminated entirely, and it may not be added where there was no whitespace at all.

Recall that a comment counts as whitespace.

Pitfalls and Subtleties of Macros

In the following documentation, there is discussion of some special rules that apply to macros and macro expansion, and point out certain cases in which the rules have counterintuitive consequences that require exercising caution.

- "“Improperly Nested Constructs”" (below)
- "Unintended Grouping of Arithmetic" on page 241
- "Swallowing the Semicolon" on page 242

- “Duplication of Side Effects” on page 243
- “Self-referential Macros” on page 243
- “Separate Expansion of Macro Arguments” on page 244
- “Cascaded Use of Macros” on page 246

Improperly Nested Constructs

Recall that when a macro is called with arguments, the arguments are substituted into the macro body and the result is checked, together with the rest of the input file, for more macro calls. It is possible to piece together a macro call coming partially from the macro body and partially from the actual arguments. Use the following input as an example.

```
#define double(x) (2*(x))
#define call_with_1(x) x(1)
```

This input would expand `call_with_1 (double)` into `(2*(1))`.

Macro definitions do not have to have balanced parentheses. By writing an unbalanced open parenthesis in a macro body, it is possible to create a macro call that begins inside the macro body but ends outside of it. Use the following input as an example.

```
#define strange(file) fprintf (file, "%s %d",
. . .
strange(stderr) p, 35)
```

The previous bizarre example expands to the following output.

```
fprintf (stderr, "%s %d", p, 35)
```

Unintended Grouping of Arithmetic

You may have noticed that in most of the macro definition examples, each occurrence of a macro argument name had parentheses around it. In addition, another pair of parentheses usually surround the entire macro definition. The following is why it is best to write macros that way.

Suppose you define a macro as follows.

```
#define ceil_div(x, y) (x + y - 1) / y
```

This produces macro output whose purpose is to divide, rounding up. (One use for this operation is to compute how many `int` objects are needed to hold a certain number of `char` objects.) Then suppose it is used as follows.

```
a = ceil_div (b & c, sizeof (int));
```

This expands into output like the following.

```
a = (b & c + sizeof (int) - 1) / sizeof (int);
```

This output does not do what is intended. The operator-precedence rules of C make it equivalent to an operation like the following.

```
a = (b & (c + sizeof (int) - 1)) / sizeof (int);
```

But what we want is the following result.

```
a = ((b & c) + sizeof (int) - 1) / sizeof (int);
```

This would mean defining the macro as the following.

```
#define ceil_div(x, y) ((x) + (y) - 1) / (y)
```

This provides the desired result. However, unintended grouping can result in another way. Consider `sizeof ceil_div(1, 2)`. That has the appearance of a C expression that would compute the size of the type of `ceil_div(1, 2)`, but in fact it means something very different. Use the following output as an example of how it expands.

```
sizeof ((1) + (2) - 1) / (2)
```

This would take the size of an integer and divide it by two. The precedence rules have put the division outside the `sizeof` when it was intended to be inside.

Parentheses around the entire macro definition can prevent such problems. What follows, then, is the recommended way to define `ceil_div`.

```
#define ceil_div(x, y) (((x) + (y) - 1) / (y))
```

Swallowing the Semicolon

Often it is desirable to define a macro that expands into a compound statement.

Consider, for example, the following macro, that advances a pointer (the argument `p` says where to find it) across whitespace characters:

```
#define SKIP_SPACES (p, limit) \  
{ register char *lim = (limit); \  
  while (p != lim) { \  
    if (*p++ != ' ') { \  
      p--; break; }}}
```

*Note:*In the previous example, Backslash-Newline is used to split the macro definition, which must be a single line, so that it resembles the way such C code would be laid out if not part of a macro definition.

A call to this macro might be `SKIP_SPACES (p, lim)`. Strictly speaking, the call expands to a compound statement, which is a complete statement with no need for a semicolon to end it. But it looks like a function call. So it minimizes confusion if to use it like a function call, writing a semicolon afterward, as in the following example.

```
SKIP_SPACES (p, lim);
```

But this can cause trouble before `else` statements, because the semicolon is actually a null statement. Suppose you have the following input.

```
if (*p != 0)  
    SKIP_SPACES (p, lim);  
else. . .
```

The presence of two statements—the compound statement and a null statement—in between the `if` condition and the `else` makes invalid C code. The definition of the macro `SKIP_SPACES` can be altered to solve this problem, using a `do...while` statement. Use the following input as an example.

```
#define SKIP_SPACES (p, limit) \  
do
```

```
do { register char *lim = (limit); \
    while (p != lim) { \
        if (*p++ != ` `) { \
            p--; break; } } \
while (0)
```

Now `SKIP_SPACES (p, lim)`; expands into one output statement as the following example shows.

```
do { : : :
} while (0);
```

Duplication of Side Effects

Many C programs define a macro `min` (standing for *minimum*) like the following.

```
#define min(X, Y) ((X) < (Y) ? (X) : (Y))
```

When you use this macro with an argument containing a *side effect* like the statement, `next = min (x + y, foo (z));`, it expands into the following output.

```
next = ((x + y) < (foo (z)) ? (x + y) : (foo (z)));
```

`x + y` has been substituted for `x` and `foo (z)` for `y`. The function `foo` is used only once in the statement as it appears in the program, but the expression `foo (z)` has been substituted twice into the macro expansion. As a result, `foo` might be called two times when the statement is executed. If it has side effects or if it takes a long time to compute, the results might not be what you intended. We say that `min` is an unsafe macro. The best solution to this problem is to define `min` in a way that computes the value of `foo (z)` only once. The C language offers no standard way to do this, but it can be done with GNU C extensions as the following example shows.

```
#define min(X, Y) \
({ typeof (X) __x = (X), __y = (Y); (__x < __y) ? __x : __y; })
```

If you do not wish to use GNU C extensions, the only solution is to be careful when using the macro, `min`. For instance, you can calculate the value of `foo (z)`, save it in a variable, and use that variable in `min`, as in the following example.

```
#define min(X, Y) ((X) < (Y) ? (X) : (Y))
. . .
{ int tem = foo (z); next = min (x + y, tem); }
```

This operation assumes that `foo` returns type `int`.

Self-referential Macros

A self-referential macro is one whose name appears in its definition. A special feature of ANSI Standard C is that the self-reference is not considered a macro call. It is passed into the preprocessor output unchanged.

Let's consider an example.

```
#define foo (4 + foo)
```

`foo`, then, is also a variable in your program.

Following the ordinary rules, each reference to `foo` will expand into `(4+foo)`; then

this will be re-scanned and will expand into $(4+(4+f_{oo}))$; and so on until it causes a fatal error (memory full) in the preprocessor.

However, the special rule about self-reference cuts this process short after one step at $(4 + f_{oo})$. So, this macro definition has the possibly useful effect of causing the program to add 4 to the value of f_{oo} wherever f_{oo} is referred to. In most cases, it is a bad idea to take advantage of this feature. A person reading the program who sees that f_{oo} is a variable will not expect that it is a macro as well. The reader will come across the identifier f_{oo} in the program and think its value should be that of the variable f_{oo} , whereas in fact the value is four greater. The special rule for self-reference applies also to *indirect* self-reference. This is the case where a macro x expands to use a macro y , and the expansion of y refers to the macro x . The resulting reference to x comes indirectly from the expansion of x , so it is a self-reference and is not further expanded. Suppose you used the following input.

```
#define x (4 + y)
#define y (2 * x)
```

x , then, would expand into $(4+(2*x))$.

But suppose y is used elsewhere, not from the definition of x . Then the use of x in the expansion of y is not a self-reference because x is not *in progress*. So it does expand. However, the expansion of x contains a reference to y , and that is an indirect self-reference now because y is in progress. The result is that y expands to $(2*(4+y))$. It is not clear that this behavior would ever be useful, but it is specified by the ANSI C standard, so you may need to understand it.

Separate Expansion of Macro Arguments

We have explained that the expansion of a macro, including the substituted actual arguments, is scanned over again for macro calls to be expanded.

What really happens is more subtle: first each actual argument text is scanned separately for macro calls. Then the results of this are substituted into the macro body to produce the macro expansion, and the macro expansion is scanned again for macros to expand. The result is that the actual arguments are scanned twice to expand macro calls in them. Most of the time, this has no effect. If the actual argument contained any macro calls, they are expanded during the first scan. The result therefore contains no macro calls, so the second scan does not change it.

If the actual argument were substituted as given, with no pre-scan, the single remaining scan would find the same macro calls and produce the same results. You might expect the double scan to change the results when a self-referential macro is used in an actual argument of another macro (see “Self-referential Macros” on page 243). The self-referential macro would be expanded once in the first scan, and a second time in the second scan. But this is not what happens. The self-references that do not expand in the first scan are marked so that they will not expand in the second scan either.

The pre-scan is not done when an argument is stringified or concatenated. Use the following input as an example.

```
#define str(s) #s
#define foo 4
str (foo)
```

This, then, expands to "foo". Once more, prescan has been prevented from having any noticeable effect. More precisely, stringification and concatenation use the argument as written, in un-prescanned form. The same actual argument would be used in pre-scanned form if it is substituted elsewhere without stringification or concatenation.

```
#define str(s) #s lose(s)
#define foo 4
str (foo)
```

This, then, expands to "foo" lose(4).

You might now ask, "Why mention the pre-scan, if it makes no difference? And why not skip it and make the preprocessor faster?" The answer is that the pre-scan does make a difference in three special cases:

- Nested calls to a macro.
- Macros that call other macros that stringify or concatenate.
- Macros whose expansions contain unshielded commas.

Nested calls to a macro occur when a macro's actual argument contains a call to that very macro. For example, if `f` is a macro expecting one argument, `f (f (1))` is a nested pair of calls to `f`. The desired expansion is made by expanding `f (1)` and substituting that into the definition of `f`. The pre-scan causes the expected result to happen.

Without the prescan, `f (1)` itself would be substituted as an actual argument, and the inner use of `f` would appear during the main scan as an indirect self-reference and would not be expanded. Here, the pre-scan cancels an undesirable side effect (in the medical, not computational, sense of the term) of the special rule for self-referential macros. But pre-scan causes trouble in certain other cases of nested macro calls, as in the following example.

```
#define foo a,b
#define bar(x) lose(x)
#define lose(x) (1 + (x))

bar(foo)
```

We would like `bar(foo)` to turn into `(1 + (foo))`, which would then turn into `(1 + (a,b))`. But instead, `bar(foo)` expands into `lose(a,b)`, and you get an error because `lose` requires a single argument. In this case, the problem is easily solved by the same parentheses that ought to be used to prevent mis-nesting of arithmetic operations.

```
#define foo (a,b)
#define bar(x) lose((x))
```

The problem is more serious when the operands of the macro are not expressions: for example, this is true when they are statements. Then parentheses are unacceptable because they would make for invalid C code.

```
#define foo { int a, b; . . .
}
```

In GNU C, you can shield the commas using the `({ . . . })` construct, which turns a compound statement into an expression like the following example shows.

```
#define foo ({ int a, b; . . .
})
```

Or you can rewrite the macro definition to avoid such commas, using the following input.

```
#define foo { int a; int b; . . .
}
```

There is also one case where pre-scan is useful. It is possible to use pre-scan to expand an argument and then stringify it—if you use two levels of macros. Add a new macro, `xstr`, to the previous definition.

```
#define xstr(s) str(s)
#define str(s) #s
#define foo 4
xstr (foo)
```

This expands into `4`, not `foo`. The reason for the difference is that the argument of `xstr` is expanded at pre-scan (because `xstr` does not specify stringification or concatenation of the argument). The result of pre-scan then forms the actual argument for `str`. `str` uses its argument without pre-scan because it performs stringification; but it cannot prevent or undo the pre-scanning already done by `xstr`.

Cascaded Use of Macros

A cascade of macros is when one macro's body contains a reference to another macro. This is very common practice, as in the following example.

```
#define BUFSIZE 1020
#define TABLESIZE BUFSIZE
```

This is not at all the same as defining `TABLESIZE` to be `1020`. The `#define` for `TABLESIZE` uses exactly the body you specify—in this case, `BUFSIZE`—and does not check to see whether it too is the name of a macro.

It's only when you use `TABLESIZE` that the result of its expansion is checked for more macro names. This makes a difference if you change the definition of `BUFSIZE` at some point in the source file. `TABLESIZE`, defined as in the following example, will always expand using the definition of `BUFSIZE` that is currently in effect.

```
#define BUFSIZE 1020
#define TABLESIZE BUFSIZE
#undef BUFSIZE #define BUFSIZE 37
```

Now `TABLESIZE` expands (in two stages) to `37`. (The `#undef` is to prevent any

warning about the nontrivial redefinition of `BUFSIZE`.)

Newlines in Macro Arguments

Traditional macro processing carries forward all newlines in macro arguments into the expansion of the macro. This means that, if some of the arguments are substituted more than once, or not at all, or out of order, newlines can be duplicated, lost, or moved around within the expansion. If the expansion consists of multiple statements, then the effect is to distort the line numbers of some of these statements. The result can be incorrect line numbers, in error messages or displayed in a debugger.

The GNU C preprocessor operating in ANSI C mode adjusts appropriately for multiple use of an argument—the first use expands all the newlines, and subsequent uses of the same argument produce no new-lines. But even in this mode, it can produce incorrect line numbering if arguments are used out of order, or not used at all. What follows is an example illustrating this problem.

```
#define ignore_second_arg(a,b,c) a; c

ignore_second_arg (foo (),
                  ignored (),
                  syntax error);
```

The syntax error triggered by the tokens “**syntax error**” results in an error message citing *the line containing* “`ignored ()`,” even though the statement of “`syntax error`;`;`” is *the line containing the error*.

6

Conditionals

In a macro processor, a *conditional* is a directive that allows a part of the program to be ignored during compilation, on some conditions. In the C preprocessor, a conditional can test either an arithmetic expression or whether a name is defined as a macro. A conditional in the C preprocessor resembles in some ways an `if` statement in C, but it is important to understand the difference between them.

The condition in an `if` statement is tested during the execution of your program. Its purpose is to allow your program to behave differently from run to run, depending on the data it is operating on. The condition in a preprocessing conditional directive is tested when your program is compiled. Its purpose is to allow different code to be included in the program depending on the situation at the time of compilation.

See the following documentation for more details.

- “Why Conditionals are Useful” on page 250
- “Syntax of Conditionals” on page 250
- “Keeping Deleted Code for Future Reference” on page 252
- “Conditionals and Macros” on page 252
- “Assertions” on page 254
- “The `#error` and `#warning` Directives” on page 255

Why Conditionals are Useful

Generally there are three kinds of reason to use a conditional.

- A program may need to use different code depending on the machine or operating system it is to run on. In some cases the code for one operating system may be erroneous on another operating system; for example, it might refer to library routines that do not exist on the other system. When this happens, it is not enough to avoid executing the invalid code: merely having it in the program makes it impossible to link the program and run it. With a preprocessing conditional, the offending code can be effectively excised from the program when it is not valid.
- You may want to be able to compile the same source file into two different programs. Sometimes the difference between the programs is that one makes frequent time-consuming consistency checks on its intermediate data, or prints the values of those data for debugging, while the other does not.
- A conditional whose condition is always false is a good way to exclude code from the program but keep it as a sort of comment for future reference.

Most simple programs that are intended to run on only one machine will not need to use preprocessing conditionals.

Syntax of Conditionals

A conditional in the C preprocessor begins with one of three conditional directives: `#if`, `#ifdef` or `#ifndef`. See ““The `#if` Directive”” (below), “The `#else` Directive” on page 251 and “The `#elif` Directive” on page 251. See “Conditionals and Macros” on page 252 for information on `#ifdef` and `#ifndef`.

The `#if` Directive

The `#if` directive in its simplest form consists of the following statement.

```
#if expression
  controlled text
#endif /* expression */
```

The comment following `#endif` is not required, but it is a good practice because it helps people match the `#endif` to the corresponding `#if`. Such comments should always be used, except in short conditionals that are not nested. In fact, you can put anything at all after the `#endif` and it will be ignored by the GNU C preprocessor, but only comments are acceptable in ANSI Standard C. *expression* is a C expression of integer type, subject to stringent restrictions. It may contain:

- *Integer constants*, which are all regarded as `long` or `unsigned long`.

- *Character constants*, which are interpreted according to the character set and conventions of the machine and operating system on which the preprocessor is running. The GNU C preprocessor uses the C data type, `char`, for these character constants; therefore, whether some character codes are negative is determined by the C compiler used to compile the preprocessor. If it treats `char` as signed, then character codes large enough to set the sign bit will be considered negative; otherwise, no character code is considered negative.
- *Arithmetic operators* for addition, subtraction, multiplication, division, bitwise operations, shifts, comparisons, and logical operations (`&&` and `||`).
- *Identifiers that are not macros*, which are all treated as zero(!).
- *Macro calls*; all macro calls in the expression are expanded before actual computation of the expression's value begins.

IMPORTANT! `sizeof` operators and `enum`-type values are not allowed. `enum`-type values, like all other identifiers that are not taken as macro calls and expanded, are treated as zero.

The *controlled text* inside of a conditional can include preprocessing directives. Then the directives inside the conditional are obeyed only if that branch of the conditional succeeds. The text can also contain other conditional groups. However, the `#if` and `#endif` directives must balance.

The `#else` Directive

The `#else` directive can be added to a conditional to provide alternative text to be used if the condition is false. The following code example is what it resembles.

```
#if expression
text-if-true
#else /* Not expression */
text-if-false
#endif /* Not expression */
```

If *expression* is nonzero, and thus the *text-if-true* is active, then `#else` acts like a failing conditional and the *text-if-false* is ignored. Contrarily, if the `#if` conditional fails, the *text-if-false* is considered included.

The `#elif` Directive

One common case of nested conditionals is used to check for more than two possible alternatives. For example, you might have the following statement.

```
#if X == 1
. . .
#else /* X != 1 */ #if X == 2
. . .
#else /* X != 2 */
. . .
```

```
#endif /* X != 2 */
#endif /* X != 1 */
```

Another conditional directive, `#elif`, allows this to be abbreviated as in the following example.

```
#if X == 1
. . .
#elif X == 2
. . .
#else /* X != 2 and X != 1*/
. . .
#endif /* X != 2 and X != 1*/
```

`#elif` stands for *else if*. Like `#else`, it goes in the middle of “`#if`” and “`#endif`” pairs, subdividing the pair; `#elif` does not require a matching `#endif`. Like `#if`, the `#elif` directive includes an expression to be tested.

The text following the `#elif` is processed only if the original `#if` condition failed and the `#elif` condition succeeds. More than one `#elif` can go in the same “`#if`”-“`#endif`” group. Then the text after each `#elif` is processed only if the `#elif` condition succeeds after the original `#if` and any previous `#elif` directives within it have failed. `#else` is equivalent to `#elif 1`, and `#else` is allowed after any number of `#elif` directives, but `#elif` may not follow `#else`.

Keeping Deleted Code for Future Reference

If you replace or delete a part of the program but want to keep the old code around as a comment for future reference, the easy way to do this is to put `#if 0` before it and `#endif` after it. This is better than using comment delimiters `/*` and `*/` since those won’t work if the code already contains comments (C comments do not nest).

This works even if the code being turned off contains conditionals, but they must be entire conditionals (balanced `#if` and `#endif`).

Conversely, do not use `#if 0` for comments which are not C code. Use the comment delimiters `/*` and `*/` instead. The interior of `#if 0` must consist of complete tokens; in particular, single quote characters must balance. But comments often contain unbalanced single quote characters (known in English as apostrophes). These confuse `#if 0`. They do not confuse `/*`.

Conditionals and Macros

Conditionals are useful in connection with macros or assertions, because those are the only ways that an expression’s value can vary from one compilation to another. A `#if`

directive whose expression uses no macros or assertions is equivalent to `#if1` or `#if0`; you might as well determine which one, by computing the value of the expression yourself, and then simplify the program. For example, what follows is a conditional statement that tests the expression `BUFSIZE==1020`, where `BUFSIZE` must be a macro.

```
#if BUFSIZE == 1020
    printf ("Large buffers!\n");
#endif /* BUFSIZE is large */
```

IMPORTANT! Programmers often wish they could test the size of a variable or data type in `#if`, but this does not work. The preprocessor does not understand `sizeof`, or `typedef` names, or even the type keywords such as `int`.)

The special operator `defined` is used in `#if` expressions to test whether a certain name is defined as a macro. Either `defined name` or `defined(name)` is an expression whose value is 1 if `name` is defined as macro at the current point in the program, and 0 otherwise. For the `defined` operator it makes no difference what the definition of the macro is; all that matters is whether there is a definition, as in the following example.

```
#if defined (vax) || defined (ns16000)
```

This statement would succeed if either of the names `vax` and `ns16000` is defined as a macro. You can test the same condition using assertions (see “Assertions” on page 254), like the following example shows.

```
#if #cpu (vax) || #cpu (ns16000)
```

If a macro is defined and later undefined with `#undef`, subsequent use of the `defined` operator returns 0, because the name is no longer defined. If the macro is defined again with another `#define`, `defined` will recommence returning 1.

Conditionals that test whether just one name is defined are very common, so there are two special short conditional directives for this case.

- `#ifdef name` is equivalent to `#if defined (name)`.
- `#ifndef name` is equivalent to `#if ! defined (name)`.

Macro definitions can vary between compilations for several reasons.

- Some macros are predefined on each kind of machine. For example, on a Vax, the name `vax` is a predefined macro. On other machines, it would not be defined.
- Many more macros are defined by system header files. Different systems and machines define different macros, or give them different values. It is useful to test these macros with conditionals to avoid using a system feature on a machine where it is not implemented.
- Macros are a common way of allowing users to customize a program for different machines or applications. For example, the macro `BUFSIZE` might be defined in a configuration file for your program that is included as a header file in each source file. You would use `BUFSIZE` in a preprocessing conditional in order to generate different code depending on the chosen configuration.

- Macros can be defined or undefined with `-D` and `-U` command options when you compile the program. You can arrange to compile the same source file into two different programs by choosing a macro name to specify which program you want, writing conditionals to test whether or how this macro is defined, and then controlling the state of the macro with compiler command options. See “Invoking the C Preprocessor” on page 263.

Assertions

Assertions are a more systematic alternative to macros in writing conditionals to test what sort of computer or system the compiled program will run on. Assertions are usually predefined, but you can define them with preprocessing directives or command-line options. The macros traditionally used to describe the type of target are not classified in any way according to which question they answer; they may indicate a hardware architecture, a particular hardware model, an operating system, a particular version of an operating system, or specific configuration options. These are jumbled together in a single namespace. In contrast, each assertion consists of a named question and an answer. The question is usually called the predicate. An assertion looks like the following statement.

```
#predicate (answer)
```

You must use a properly formed identifier for *predicate*. The value of *answer* can be any sequence of words; all characters are significant except for leading and trailing whitespace, and differences in internal whitespace sequences are ignored. Thus, `x +y` is different from `x+y` but equivalent to `x +y`. `)` is not allowed in an answer. What follows is a conditional to test whether the *answer* is asserted for *predicate*.

```
#if #predicate (answer)
```

There may be more than one answer asserted for a given predicate. If you omit the answer, you can test whether *any* answer is asserted for *predicate*:

```
#if #predicate
```

Most of the time, the assertions you test will be predefined assertions. GNU C provides three predefined predicates: `system`, `cpu`, and `machine`. `system` is for assertions about the type of software, `cpu` describes the type of computer architecture, and `machine` gives more information about the computer. For example, on a GNU system, the following assertions would be true.

```
#system (gnu)
#system (mach)
#system (mach 3)
#system (mach 3. subversion)
#system (hurd)
#system (hurd version)
```

Perhaps there are others. The alternatives with more or less version information let you ask more or less detailed questions about the type of system software. On a UNIX

system, you would find `#system (unix)` and perhaps one of: `#system (aix)`, `#system (bsd)`, `#system (hpux)`, `#system (lynx)`, `#system (mach)`, `#system (posix)`, `#system (svr3)`, `#system (svr4)`, or `#system (xpg4)` with possible version numbers following.

Other values for `system` are `#system (mvs)` and `#system (vms)`.

NOTE Many UNIX C compilers provide only one answer for the `system` assertion, `#system (unix)`, if they support assertions at all. This is less than useful.

An assertion with a multi-word answer is completely different from several assertions with individual single-word answers. For example, the presence of `system (mach 3.0)` does not mean that `system (3.0)` is true. It also does not directly imply `system (mach)`, but in GNU C, that last will normally be asserted as well. The current list of possible assertion values for `cpu` is: `#cpu (a29k)`, `#cpu (alpha)`, `#cpu (arm)`, `#cpu (clipper)`, `#cpu (convex)`, `#cpu (elxsi)`, `#cpu (tron)`, `#cpu (h8300)`, `#cpu (i370)`, `#cpu (i386)`, `#cpu (i860)`, `#cpu (i960)`, `#cpu (m68k)`, `#cpu (m88k)`, `#cpu (mips)`, `#cpu (ns32k)`, `#cpu (hppa)`, `#cpu (pyr)`, `#cpu (ibm032)`, `#cpu (rs6000)`, `#cpu (sh)`, `#cpu (sparc)`, `#cpu (spur)`, `#cpu (tahoe)`, `#cpu (vax)`, `#cpu (we32000)`.

You can create assertions within a C program using `#assert`, with the following input.

```
#assert predicate (answer)
```

NOTE `#` does not appear before *predicate*.

Each time you do this, you assert a new true answer for *predicate*. Asserting one answer does not invalidate previously asserted answers; they all remain true. The only way to remove an assertion is with `#unassert`. `#unassert` has the same syntax as `#assert`. You can also remove all assertions about *predicate* using the following example's statement.

```
#unassert predicate
```

You can also add or cancel assertions using command options when you run `gcc` or `cpp`. See “Invoking the C Preprocessor” on page 263.

The `#error` and `#warning` Directives

The `#error` directive causes the preprocessor to report a fatal error.

The rest of the line that follows `#error` is used as the error message.

You would use `#error` inside of a conditional that detects a combination of parameters which you know the program does not properly support. For example, if you know that the program will not run properly on a Vax, you might use the following input.

```
#ifdef __vax__
#error Won't work on Vaxen. See comments at get_last_object.
#endif
```

See “Non-standard Predefined Macros” on page 235 for a description of why this works.

If you have several configuration parameters that must be set up by the installation in a consistent way, you can use conditionals to detect an inconsistency and report it with `#error`. For clarification, see the following example.

```
#if HASH_TABLE_SIZE % 2 == 0 || HASH_TABLE_SIZE % 3 == 0 \
    || HASH_TABLE_SIZE % 5 == 0
#error HASH_TABLE_SIZE should not be divisible by a small prime
#endif
```

The directive, `#warning`, is like the directive, `#error`, but causes the preprocessor to issue a warning and continue preprocessing. The rest of the line that follows `#warning` is used as the warning message.

You might use `#warning` in obsolete header files, with a message directing the user to the header file which should instead be used.

7

Combining Source Files

One of the jobs of the C preprocessor is to inform the C compiler of where each line of C code has as its origin, which source file and which line number.

C code can come from multiple source files if you use `#include`; both `#include` and the use of conditionals and macros can cause the line number of a line in the preprocessor output to be different from the line's number in the original source file. You will appreciate the value of making both the C compiler (in error messages) and symbolic debuggers such as GDB use the line numbers in your source file.

The C preprocessor builds on this feature by offering a directive by which you can control the feature explicitly. This is useful when a file for input to the C preprocessor is the output from another program such as the `bison` parser generator, which operates on another file that is the true source file. Parts of the output from `bison` are generated from scratch, other parts come from a standard parser file. The rest are copied nearly verbatim from the source file, but their line numbers in the `bison` output are not the same as their original line numbers. Naturally you would like compiler error messages and symbolic debuggers to know the original source file and line number of each line in the `bison` input.

`bison` arranges this by writing `#line` directives into the output file. `#line` is a directive that specifies the original line number and source file name for subsequent input in the current preprocessor input file. `#line` has three variants:

`#line linenum`

linenum is a decimal integer constant, for specifying that the line number of the following line of input, in its original source file, was *linenum*, that same designation.

`#line linenum filename`

Here *linenum* is a decimal integer constant and *filename* is a string constant. This specifies that the following line of input came originally from source file *filename* and its line number there was *linenum*. Keep in mind that *filename* is not just a file, *filename*.

`#line anything else`

anything else is checked for macro calls, which are expanded. The result should be a decimal integer constant followed optionally by a string constant.

`#line` directives alter the results of the `__FILE__` and `__LINE__` predefined macros from that point on. See “Standard Predefined Macros” on page 233.

The output of the preprocessor (which is the input for the rest of the compiler) contains directives that look much like `#line` directives. They start with just `#` instead of `#line`, but this is followed by a line number and file name as in `#line`. See “C Preprocessor Output” on page 261.

8

Other Preprocessing Directives

The following documentation describes three additional preprocessing directives.

- The *null directive* consists of a # followed by a Newline, with only whitespace (including comments) in between. A null directive is understood as a preprocessing directive but has no effect on the preprocessor output. The primary significance of the existence of the null directive is that an input line consisting of just a # will produce no output, rather than a line of output containing just a #. Supposedly some old C programs contain such lines.
- The ANSI standard specifies that the #pragma directive has an arbitrary, implementation-defined effect. In the GNU C preprocessor, #pragma directives are not used, except for #pragma once (see “Once-only Include Files” on page 226). However, they are left in the preprocessor output, so they are available to the compilation pass.
- The #ident directive is supported for compatibility with certain other systems. It is followed by a line of text. On some systems, the text is copied into a special place in the object file; on most systems, the text is ignored and this directive has no effect. Typically, #ident is only used in header files supplied with those systems where it is meaningful.

9

C Preprocessor Output

The output from the C preprocessor looks much like the input, except that all preprocessing directive lines have been replaced with blank lines and all comments with spaces. Whitespace within a line is not altered; however, a space is inserted after the expansions of most macro calls. Source file name and line number information is conveyed by lines of the following form.

```
# linenum filename flags
```

Spaces are inserted as needed into the middle of the input (but never within a string or character constant). Such a line means that the line following it originated in file, *filename*, at line, *linenum*. After the file name comes zero or more flags, which are 1, 2, 3, or 4. If there are multiple flags, spaces separate them. What the flags mean:

- 1 indicates the start of a new file.
- 2 indicates returning to a file (after having included another file).
- 3 indicates that the text following comes from a system header file, so certain warnings should be suppressed.
- 4 indicates that its subsequent text should be treated as C.

10

Invoking the C Preprocessor

The following documentation discusses the commands as options accepted by the C preprocessor. Most often when using the C preprocessor, you will not have to invoke it explicitly: the C compiler will do so automatically. However, the preprocessor is sometimes useful on its own. The C preprocessor expects two file names as arguments, referred to in this documentation as *infile* and *outfile*. The preprocessor reads *infile* together with any other files that it specifies with `#include`. All the output generated by the combined input files is written in a file that the documentation refers to as *outfile*. Either *infile* or *outfile* may use a preceding hyphen or dash, which *infile* means to read from standard input and, as *outfile*, means to write to standard output. Also, if *outfile* or both file names are omitted, the standard output and standard input are used for the omitted file names.

What follows is a list of command options accepted by the C preprocessor. These options can also be given when compiling a C program; they are passed along automatically to the preprocessor when it is invoked by the compiler.

-P

Inhibit generation of `#`-lines with line-number information in the output from the preprocessor (see “C Preprocessor Output” on page 261). This might be useful when running the preprocessor on something that is not C code and is sent to a program that might be confused with the `#`-lines.

-C

Do not discard comments: pass them through to the output file. Comments

appearing in arguments of a macro call will be copied to the output before the expansion of the macro call.

`-traditional`

Try to imitate the behavior of old-fashioned C, as opposed to ANSI C.

- Traditional macro expansion pays no attention to single-quote or double-quote characters; macro argument symbols are replaced by the argument values even when they appear within apparent string or character constants.
- Traditionally, it is permissible for a macro expansion to end in the middle of a string or character constant. The constant continues into the text surrounding the macro call.
- However, traditionally the end of the line terminates a string or character constant, with no error.
- In traditional C, a comment is equivalent to no text at all. (In ANSI C, a comment counts as whitespace.)
- Traditional C does not have the concept of a preprocessing number. It considers `1.0e+4` to be three tokens: `1.0e`, `+`, and `4`.
- A macro is not suppressed within its own definition, in traditional C. Thus, any macro that is used recursively inevitably causes an error.
- The `#` character has no special meaning within a macro definition in traditional C.
- In traditional C, the text at the end of a macro expansion can run together with the text after the macro call, to produce a single token. This is impossible in ANSI C.
- Traditionally, `\` inside a macro argument suppresses the syntactic significance of the following character.

Use the `-traditional` option when preprocessing Fortran code, so that singlequotes and doublequotes within Fortran comment lines (which are generally not recognized as such by the preprocessor) do not cause diagnostics about unterminated character or string constants.

However, this option does not prevent diagnostics about unterminated comments when a C-style comment appears to start, but not end, within Fortran-style commentary. So, the following Fortran comment lines are accepted with

`-traditional`:

```
C   This isn't an unterminated character constant
C   Neither is "20000000000, an octal constant
C   in some dialects of Fortran
```

However, this type of comment line will likely produce a diagnostic, or at least unexpected output from the preprocessor, due to the unterminated comment:

```
C   Some Fortran compilers accept /* as starting
```

- C an inline comment.
- `g77` automatically supplies the `-traditional` option when it invokes the preprocessor. However, a future version of `g77` might use a different, more Fortran-aware preprocessor in place of `cpp`.
- `-trigraphs`
Process ANSI standard trigraph sequences. These are three-character sequences, all starting with `??`, that are defined by ANSI C to stand for single characters. For example, `??/` stands for `\`, so `??/n` is a character constant for a new line. Strictly speaking, the GNU C preprocessor does not support all programs in ANSI C Standard unless `-trigraphs` is used.
- `-pedantic`
Issue warnings required by the ANSI C standard in certain cases such as when text other than a comment follows `#else` or `#endif`.
- `-pedantic-errors`
Like `-pedantic`, except that errors are produced rather than warnings.
- `-Wtrigraphs`
Warn if any trigraphs are encountered (assuming they are enabled).
- `-Wcomment`
Warn whenever a comment-start sequence `/*` appears in a comment.
- `-Wall`
Requests both `-Wtrigraphs` and `-Wcomment` (but not `-Wtraditional`).
- `-Wtraditional`
Warn about certain constructs that behave differently in traditional and ANSI C.
- `-I directory`
Add the directory *directory* to the head of the list of directories to be searched for header files (see “The `#include` Directive” on page 224). This can be used to over-ride a system header file, substituting your own version, since these directories are searched before the system header file directories. If you use more than one `-I` option, the directories are scanned in left-to-right order; the standard system directories come after.
- `-I-`
Any directories specified with `-I` options before the `-I-` option are searched only for the case of `#include "file"`; they are not searched for `#include <file>`. If additional directories are specified with `-I` options after the `-I-`, these directories are searched for all `#include` directives. In addition, the `-I-` option inhibits the use of the current directory as the first search directory for `#include "file"`. Therefore, the current directory is searched only if `-I.` is requested explicitly with it. Specifying both `-I-` and `-I.` allows you to control precisely which directories are searched before and which after the current one is searched.
- `-nostdinc`
Do not search the standard system directories for header files. Only the directories you have specified with `-I` options (and the current directory, if appropriate) are searched.

`-nostdinc++`

Do not search for header files in the C++-specific standard directories, but do still search the other standard directories. (This option is used when building the C++ library.)

`-D name`

Predefine *name* as a macro, with definition 1.

`-D name=definition`

Predefine *name* as a macro, with definition *definition*. There are no restrictions on the contents of *definition*, but if you are invoking the preprocessor from a shell or shell-like program you may need to use the shell's quoting syntax to protect characters such as spaces that have a meaning in the shell syntax. If you use more than one `-D` for the same *name*, the rightmost definition takes effect.

`-U name`

Do not predefine *name*. If both `-U` and `-D` are specified for one name, the `-U` beats the `-D` and the name is not predefined.

`-undef`

Do not predefine any nonstandard macros.

`-A predicate(answer)`

Make an assertion with the predicate *predicate* and answer *answer*. See “Assertions” on page 254. You can use `-A-` to disable all predefined assertions; it also undefines all predefined macros that identify the type of target system.

`-dM`

Instead of outputting the result of preprocessing, output a list of `#define` directives for all the macros defined during the execution of the preprocessor, including predefined macros. This gives you a way of finding out what is predefined in your version of the preprocessor; assuming you have no file `foo.h`, the following command will show the values of any predefined macros.

```
touch foo.h; cpp -dM foo.h
```

`-dD`

Like `-dM` except in two respects: it does *not* include the pre-defined macros, and it outputs *both* the `#define` directives and the result of preprocessing. Both kinds of output go to the standard output file.

`-M [-MG]`

Instead of outputting the result of preprocessing, output a rule suitable for `make` describing the dependencies of the main source file. The preprocessor outputs one `make` rule containing the object file name for that source file, a colon, and the names of all the included files. If there are many included files then the rule is split into several lines using `\-newline`. `-MG` says to treat missing header files as generated files and assume they live in the same directory as the source file. It must be specified in addition to `-M`. This feature is used in automatic updating of makefiles.

- `-MM [-MG]`
Like `-M`, mentions only the files included with `#include "file"`. System header files included with `#include<file>` are omitted.
- `-MD file`
Like `-M` but the dependency information is written to *file*. This is in addition to compiling the file as specified; `-MD` does not inhibit ordinary compilation the way `-M` does. When invoking `gcc`, do not specify the *file* argument. `gcc` will create file names made by replacing “.c” with “.d” at the end of the input file names. In Mach, you can use the utility `md` to merge multiple dependency files into a single dependency file suitable for using with the `make` command.
- `-MMD file`
Like `-MD` except mention only user header files, not system header files.
- `-H`
Print the name of each header file used, in addition to other normal activities.
- `-imacros file`
Process *file* as input, discarding the resulting output, before processing the regular input file. Because the output generated from *file* is discarded, the only effect of `-imacros file` is to make the macros defined in *file* available for use in the main input.
- `-include file`
Process *file* as input, and include all the resulting output, before processing the regular input file.
- `-idirafter dir`
Add the directory *dir* to the second include path. The directories on the second include path are searched when a header file is not found in any of the directories in the main include path (the one that `-I` adds to).
- `-iprefix prefix`
Specify *prefix* as a prefix for using subsequent `-iwithprefix` options.
- `-iwithprefix dir`
Add a directory to the second include path. The directory’s name is made by concatenating *prefix* and *dir*, where *prefix* was previously specified with `-iprefix`.
- `-isystem dir`
Add a directory to the beginning of the second include path, marking it as a system directory, so that it gets the same special treatment as is applied to the standard system directories.
- `-lang-c`
`-lang-c89`
`-lang-c++`
`-lang-objc`
`-lang-objc++`
Specify the source language. `-lang-c` is the default; it allows recognition of C++ comments (comments that begin with `//` and end at end of line), since this is a

common feature and it will most likely be in the next C standard. `-lang-c89` disables recognition of C++ comments. `-lang-c++` handles C++ comment syntax and includes extra default include directories for C++. `-lang-objc` enables the Objective C `#import` directive. `-lang-objc++` enables both C++ and Objective C extensions. These options are generated by the compiler driver `gcc`, but not passed from the `gcc` command line unless you use the driver's `-wp` option.

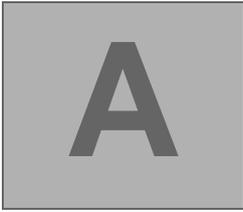
`-lint`

Look for commands to the program checker `lint` embedded in comments, and emit them preceded by `#pragma lint`. For example, `/* NOTREACHED */` becomes `#pragma lint NOTREACHED`. This option is available only when you call `cpp` directly; `gcc` will not pass it from its command line.

`-$`

Forbid the use of `$` in identifiers. This is required for ANSI conformance. `gcc` automatically supplies this option to the preprocessor when you specify `-ansi`, but `gcc` does not recognize the `-$` option itself; to use it without the other effects of `-ansi`, you must call the preprocessor directly.

Appendices



GNU General Public License

Version 2, June 1991

Copyright © 1989, 1991 Free Software Foundation, Inc.
59 Temple Place / Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software— to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute

copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all. The precise terms and conditions for copying, distribution and modification follow.

Terms and Conditions for Copying, Distribution and Modification

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2)

in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or

works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a

version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
one line: the program's name and a brief idea of what it does.  
Copyright © 19yy name of author
```

```
This program is free software; you can redistribute it and/or  
modify it under the terms of the GNU General Public License as  
published by the Free Software Foundation; either version 2 of  
the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be  
useful, but WITHOUT ANY WARRANTY; without even the implied  
warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR  
PURPOSE. See the GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public  
License along with this program; if not, write to the Free  
Software Foundation, Inc., 59 Temple Place - Suite 330,  
Boston, MA 02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like the following example when it starts in an interactive mode:

```
Gnomovision version 69, Copyright © 19yy name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details  
type 'show w'. This is free software, and you are welcome to  
redistribute it under certain conditions; type 'show c' for  
details.
```

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they can be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. The following is a sample (when copying, *alter the names*).

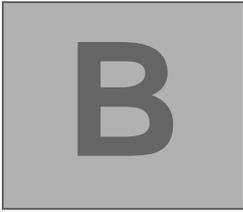
```
Yoyodyne, Inc., hereby disclaims all copyright interest in the  
program 'Gnomovision' (which makes passes at compilers)
```

written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.



Contributors to GNU CC

In addition to Richard Stallman, several people have written parts of the GNU compiler, GNU CC. The GCC project would like to thank its many contributors. Without them the project would not have been nearly as successful as it has been. Any omissions in this list are accidental. Feel free to contact law@cygnus.com if you have been left out or some of your contributions are not listed. Please keep this list in alphabetical order.

- Analog Devices helped implement the support for complex data types and iterators.
- James van Artsdalen wrote the code that makes efficient use of the Intel 80387 register stack.
- Alasdair Baird for various bugfixes.
- Gerald Baumgartner added the signature extension to the C++ front-end.
- Per Bothner for various improvements to our infrastructure for supporting new languages such as Chill and Java.
- Devon Bowen helped port GCC to the Tahoe.
- Don Bowman for `mips-vxworks` contributions.
- Dave Brolley for work on `cpp1ib` and Chill.

-
- Robert Brown implemented the support for Encore 32000 systems.
 - Christian Bruel for improvements to local store elimination.
 - Herman A.J. Brugge for various fixes.
 - Joe Buck for his guidance and leadership via the steering committee.
 - Craig Burley for leadership of the Fortran effort.
 - John Carr for his alias work, SPARC hacking, infrastructure improvements, direction via the steering committee, loop optimizations, etc.
 - Steve Chamberlain wrote the support for the Hitachi SH and H8 processors and the PicoJava processor.
 - Scott Christley for his Objective C contributions.
 - Branko Cibej for more warning contributions.
 - Nick Clifton for ARM, MCore, FR30, V850, M32R work, `--help`, and other random hacking.
 - Ralf Corsepilus for SH testing and minor bugfixing.
 - Stan Cox for care and feeding of the x86 port and lots of behind the scenes hacking.
 - Alex Crain provided changes for the 3b1.
 - Ian Dall for major improvements to the NS32k port.
 - Dario Dariol contributed the four varieties of sample programs that print a copy of their source.
 - Ulrich Drepper for his work on the C++ runtime libraries, `glibc`, testing of GCC using `glibc`, ISO C9X support, CFG dumping support, etc.
 - Richard Earnshaw for his ongoing work with the ARM.
 - David Edelsohn for his direction via the steering committee, ongoing work with the RS6000/PowerPC port, and help cleaning up Haifa loop changes.
 - Mark Elbrecht for various DJGPP improvements.
 - Ben Elliston for his work to move the Objective-C runtime into its own subdirectory and for his work on `autoconf`.
 - Paul Eggert for random hacking all over GCC.
 - Marc Espie for OpenBSD support.
 - Doug Evans for much of the global optimization framework, ARC, M32R, and SPARC work.
 - Fred Fish for BeOS support and Ada fixes.
 - Peter Gerwinski for various bugfixes and the Pascal front end.
 - Anthony Green for his `-Os` contributions and Java front end work.
 - Kaveh Ghazi for overall direction via the steering committee and amazing work to make `-w -Wall` useful.
 - Judy Goldberg for C++ contributions.

- Torbjorn Granlund for various fixes and the c-torture testsuite, multiply- and divide-by-constant optimization, improved `long long` support, and improved leaf function register allocation.
- Michael K. Gschwind contributed the port to the PDP-11.
- Ron Guilmette implemented the `protoize` and `unprotoize` tools, the support for Dwarf symbolic debugging information, and much of the support for System V Release 4. He has also worked heavily on the Intel 386 and 860 support.
- Bruno Haible for improvements in the runtime overhead for EH, new warnings and assorted bugfixes.
- Andrew Haley for his Java work.
- Chris Hanson assisted in making GCC work on HP-UX for the 9000 series 300.
- Michael Hayes for various thankless work he's done trying to get the c30/c40 ports functional. Lots of loop and unroll improvements and fixes.
- Kate Hedstrom for staking the g77 folks with an initial testsuite.
- Richard Henderson for his ongoing SPARC and alpha work, loop optimizations, and generally fixing lots of old problems that have been ignored for years, flow rewrite and others.
- Nobuyuki Hikichi of Software Research Associates, Tokyo, contributed the support for the Sony NEWS machine.
- Manfred Hollstein for his ongoing work to keep the M88K alive, lots of testing and bugfixing, particularly of our configury code.
- Steve Holmgren for MachTen patches.
- Jan Hubicka for his x86 port improvements.
- Christian Iseli for various bugfixes.
- Kamil Iskra for general M68K hacking.
- Lee Iverson for random fixes and MIPS testing.
- Andreas Jaeger for various fixes to the MIPS port
- Jakub Jelinek for his SPARC work and sibling call optimizations.
- J. Kean Johnston for OpenServer support.
- Klaus Kaempf for his ongoing work to make alpha-VMS a viable target.
- David Kashtan of SRI adapted GCC to VMS.
- Richard Kenner of the New York University Ultracomputer Research Laboratory wrote the machine descriptions for the AMD 29000, the DEC Alpha, the IBM RT PC, and the IBM RS/6000 as well as the support for instruction attributes. He also made changes to better support RISC processors including changes to common subexpression elimination, strength reduction, function calling sequence handling, and condition code support, in addition to generalizing the code for frame pointer elimination and delay slot scheduling. Richard Kenner was also the head maintainer of GCC for several years.

-
- Robin Kirkham for CPU32 support.
 - Mark Klein for PA improvements.
 - Geoffrey Keating for his ongoing work to make the PPC work for Linux.
 - Brendan Kehoe for his ongoing work with G++.
 - Oliver M. Kellogg of Deutsche Aerospace contributed the port to the MIL-STD-1750A.
 - Mumit Khan for various contributions to the Cygwin and maintaining binary releases for Windows hosts.
 - Thomas Koenig for various bugfixes.
 - Bruce Korb for the new and improved fixincludes code.
 - Benjamin Kosnik for his G++ work.
 - Ted Lemon wrote parts of the RTL reader and printer.
 - Charles LaBrec contributed the support for the Integrated Solutions 68020 system.
 - Jeff Law for coordinating the entire project, rolling out snapshots and releases, handling merges from GCC2, and random but extensive hacking.
 - Marc Lehmann for his guidance via the steering committee and helping with analysis and improvements of x86 performance.
 - Kriang Lerdsuwanakij for improvements to demangler and various C++ fixes.
 - Warren Levy major work on `libgjc` (Java Runtime Library) and random work on the Java front-end.
 - Alain Lichnewsky ported GCC to the MIPS CPU.
 - Robert Lipe for OpenServer support, new testsuites, testing, etc.
 - Weiwen Liu for testing and various bugfixes.
 - Martin von Loumlwis for internal consistency checking infrastructure, and various C++ improvements including namespace support.
 - Dave Love for his ongoing work with the Fortran front end and runtime libraries.
 - H.J. Lu for his contributions to the steering committee, many x86 bug reports, prototype patches, and keeping the Linux ports working.
 - Andrew MacLeod for his ongoing work in building a real EH system, various code generation improvements, work on the global optimizer, etc.
 - Bob Manson for his behind the scenes work on DejaGNU.
 - Vladimir Makarov for hacking some ugly i960 problems, PowerPC hacking improvements to compile-time performance and overall knowledge and direction in the area of instruction scheduling.
 - Greg McGary for random fixes and (someday) bounded pointers.
 - Michael Meissner for LRS framework, IA32, M32R, v850, M88K, MIPS PowerPC, Haifa, ECOFF debug support, and other assorted hacking.
 - Jason Merrill for leading the G++ effort.

- Jason Molenda for major help in the care and feeding of all the services on the `gcc.gnu.org` (formerly `egcs.cygnus.com`) machine: mail, web services, ftp services, etc., etc.
- David Mosberger-Tang for various Alpha improvements.
- Stephen Moshier contributed the floating point emulator that assists in cross-compilation and permits support for floating point numbers wider than 64 bits and for ISO C9X support.
- Bill Moyer for his behind the scenes work on various issues.
- David Miller for overall direction via the steering committee, lots of SPARC work, improvements in `jump.c` and interfacing with the Linux kernel developers.
- Gary Miller ported GCC to Charles River Data Systems machines.
- Mark Mitchell for mountains of C++ work, load/store hoisting out of loops and alias analysis improvements, ISO “restrict” support.
- Alan Modra for various Linux bits and testing.
- Toon Moene for overall leadership via the steering committee, and his ongoing work to make us make Fortran run fast.
- Catherine Moore for fixing various ugly problems, including the Haifa bug which was killing the Alpha & PowerPC Linux kernels.
- Philippe De Muyter for his work on the M68K port.
- Joseph Myers for his work on the PDP-11 port.
- Nathan Myers for his work on `libstdc++-v3` libraries.
- NeXT, Inc. donated the front end that supports the Objective C language.
- Hans-Peter Nilsson for improvements to the search engine setup, various documentation fixes and other small fixes.
- Geoff Noer for this work on getting Cygwin native builds working.
- Alexandre Oliva for various build infrastructure improvements, scripts and amazing testing work.
- Rainer Orth for random MIPS work, including improvements to our 032 ABI support, improvements to DejaGNU’s MIPS support, etc.
- Melissa O’Neill for various NeXT fixes.
- Paul Petersen wrote the machine description for the Alliant FX/8.
- Alexandre Petit-Bianco for his Java work.
- Matthias Pfaller for major improvements to the NS32k port.
- Gerald Pfeifer for maintenance of the web pages and pointing out lots of problems.
- Ovidiu Predescu for his work on the Objective C front end and runtime libraries.
- Ken Raeburn for various improvements to checker, MIPS ports and various cleanups in the compiler.
- David Reese of Sun Microsystems contributed to the Solaris on Power PC port.

-
- Gabriel Dos Reis for contributions and maintenance of `libstdc++-v3`, including `valarray` implementation and limits support.
 - Joern Rennecke for maintaining the SH port, `loop`, `regmove` and `reload` hacking.
 - Gavin Romig-Koch for lots of behind the scenes MIPS work.
 - Ken Rose for fixes to our delay slot filling code.
 - Paul Rubin wrote most of the preprocessor.
 - Andreas Schwab for his work on the M68K port.
 - Graham Stott for various infrastructure improvements.
 - Juha Sarlin for improvements to the H8 code generator.
 - Greg Satz assisted in making GCC work on HP-UX for the 9000 series 300.
 - Peter Schauer wrote the code to allow debugging to work on the Alpha.
 - William Schelter did most of the work on the Intel 80386 support.
 - Bernd Schmidt for various code generation improvements and major work in the `reload` pass.
 - Franz Sirl for his ongoing work with making the Power PC port stable for Linux.
 - Joel Sherrill for his direction via the steering committee, RTEMS contributions and RTEMS testing.
 - Nathan Sidwell for many C++ fixes/improvements.
 - Jeffrey Siegal for helping with the original design of GCC, some code which handles the parse tree and RTL data structures, constant folding and help with the original VAX and M68K ports.
 - Andrey Slepuhin for assorted AIX hacking.
 - Randy Smith finished the Sun FPA support.
 - Christopher Smith did the port for Convex machines.
 - Scott Snyder for various fixes.
 - Richard Stallman, for writing the original GCC and launching the GNU project.
 - Jan Stein of the Chalmers Computer Society provided support for Genix, as well as part of the 32000 machine description.
 - Nigel Stephens for various MIPS 16 related fixes and improvements.
 - Jonathan Stone wrote the machine description for the Pyramid computer.
 - Mike Stump for his Elxsi port, G++ contributions over the years and more recently his VxWorks contributions.
 - Shigeya Suzuki for this fixes for the BSDI platforms.
 - Ian Lance Taylor for his MIPS 16 work, general configury hacking, `fixincludes`, etc.
 - Holger Teutsch provided the support for the Clipper CPU.

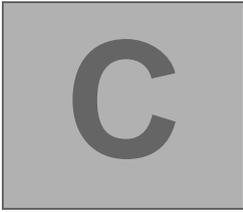
- Michael Tiemann for random bugfixes, the first instruction scheduler, initial C++ support, function integration, NS32K, SPARC and M88K machine description work, and delay slot scheduling.
- Gary Thomas for his ongoing work to make the Power PC work for Linux.
- Philipp Thomas for random bugfixes throughout the compiler
- Kresten Krab Thorup wrote the runtime support for the Objective C language.
- Teemu Torma for thread safe exception handling support.
- Leonard Tower wrote parts of the parser, RTL generator, and RTL definitions, and of the Vax machine description.
- Tom Tromeu for internationalization support and his Java work.
- Lassi Tuura for improvements to `config.guess` to determine HP processor types.
- Todd Vierling for contributions for NetBSD ports.
- Krister Walfridsson for random bugfixes.
- John Wehle for various improvements for the x86 code generator, related infrastructure improvements to help x86 code generation, value range propagation and other work, WE32k port.
- Dale Wiles helped port GCC to the Tahoe.
- Zack Weinberg for major work on `cpplib` and various other bugfixes.
- Jim Wilson for tackling hard problems in various places that nobody else wanted to work on, strength reduction and other loop optimizations.
- Carlo Wood for various fixes.
- Tom Wood for work on the M88K port.
- Masanobu Yuhara of Fujitsu Laboratories implemented the machine description for the Tron architecture (specifically, the Gmicro).
- Kevin Zachmann helped ported GCC to the Tahoe.

The following individuals have contributed time and energy in testing GCC:

- David Billingham
- Horst von Brand
- Rodney Brown
- Joe Buck
- Craig Burley
- Ulrich Drepper
- David Edelsohn
- Kaveh Ghazi
- Yung Shing Gene
- Richard Henderson
- Manfred Hollstein
- Kate Hedstrom

-
- Kamil Iskra
 - Christian Joensson
 - Jeff Law
 - Robert Lipe
 - Dave Love
 - Damon Love
 - H.J. Lu
 - Mumit Khan
 - Matthias Klose
 - Martin Knoblauch
 - Toon Moene
 - David Miller
 - Matthias Mueller
 - Alexandre Oliva
 - Richard Polton
 - David Rees
 - Peter Schmid
 - David Schuler
 - Vin Shelton
 - Franz Sirl
 - Mike Stump
 - Carlo Wood

And finally, we'd like to thank everyone who uses the compiler, submits bug reports and generally reminds us why we're doing this work in the first place.



Funding Free Software

If you want to have more free software a few years from now, it makes sense for you to help encourage people to contribute funds for its development. The most effective approach known is to encourage commercial redistributors to donate.

Users of free software systems can boost the pace of development by encouraging for-a-fee distributors to donate part of their selling price to free software developers—the Free Software Foundation, and others.

The way to convince distributors to do this is to demand it and expect it from them. So when you compare distributors, judge them partly by how much they give to free software development. Show distributors they must compete to be the one who gives the most.

To make this approach work, you must insist on numbers that you can compare, such as, “We will donate ten dollars to the Frobnitz project for each disk sold.” Don’t be satisfied with a vague promise, such as “A portion of the profits are donated,” since it doesn’t give a basis for comparison.

Even a precise fraction “of the profits from this disk” is not very meaningful, since creative accounting and unrelated business decisions can greatly alter what fraction of the sales price counts as profit. If the price you pay is \$50, ten percent of the profit is

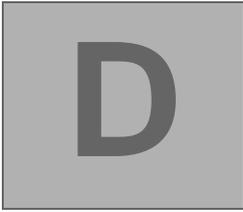
probably less than a dollar; it might be a few cents, or nothing at all.

Some redistributors do development work themselves. This is useful too; but to keep everyone honest, you need to inquire how much they do, and what kind. Some kinds of development make much more long-term difference than others. For example, maintaining a separate version of a program contributes very little; maintaining the standard version of a program for the whole community contributes much. Easy new ports contribute little, since someone else would surely do them; difficult ports such as adding a new CPU to the GNU C compiler contribute more; major new features or packages contribute the most.

By establishing the idea that supporting further development is “the proper thing to do” when distributing free software for a fee, we can assure a steady flow of resources into making more free software.

Copyright © 1994 Free Software Foundation, Inc.

Verbatim copying and redistribution of this section is permitted without royalty; alteration is not permitted.



Protect Your Freedom; Fight “Look and Feel”

This section is a political message from the League for Programming Freedom to the users of GCC. We have included it here because the issue of interface copyright is important to the GNU project.

Apple, Lotus, and now CDC have tried to create a new form of legal monopoly: a copyright on a user interface.

An interface is a kind of language—a set of conventions for communication between two entities, human or machine. Until a few years ago, the law seemed clear: interfaces were outside the domain of copyright, so programmers could program freely and implement whatever interface the users demanded. Imitating de-facto standard interfaces, sometimes with improvements, was standard practice in the computer field. These improvements, if accepted by the users, caught on and became the norm; in this way, much progress took place.

Computer users, and most software developers, were happy with this state of affairs. However, large companies such as Apple and Lotus would prefer a different system—one in which they can own interfaces and thereby rid themselves of all serious competitors. They hope that interface copyright will give them, in effect, monopolies on major classes of software.

Other large companies such as IBM and Digital also favor interface monopolies, for the same reason: if languages become property, they expect to own many de-facto standard languages. But Apple and Lotus are the ones who have actually sued. Apple's lawsuit was defeated, for reasons only partly related to the general issue of interface copyright.

Lotus won lawsuits against two small companies, which were thus put out of business. Then they sued Borland; they won in the trial court (no surprise, since it was the same court that had ruled for Lotus twice before), but the decision was reversed by the court of appeals, with help from the League for Programming Freedom in the form of a friend-of-the- court brief. We are now waiting to see if the Supreme Court will hear the case. If it does, the League for Programming Freedom will again submit a brief.

The battle is not over. Just this summer a company that produced a simulator for a CDC computer was shut down by a copyright lawsuit by CDC, which charged that the simulator infringed the copyright on the manuals for the computer.

If the monopolists get their way, they will hobble the software field:

- Gratuitous incompatibilities will burden users. Imagine if each car manufacturer had to design a different way to start, stop, and steer a car.
- Users will be “locked in” to whichever interface they learn; then they will be prisoners of one supplier, who will charge a monopolistic price.
- Large companies have an unfair advantage wherever lawsuits become commonplace. Since they can afford to sue, they can intimidate smaller developers with threats even when they don't really have a case.
- Interface improvements will come slower, since incremental evolution through creative partial imitation will no longer occur.

If interface monopolies are accepted, other large companies are waiting to grab theirs:

- Adobe is expected to claim a monopoly on the interfaces of various popular application programs, if Lotus ultimately wins the case against Borland.
- Open Computing magazine reported a Microsoft vice president as threatening to sue people who imitate the interface of Windows.

Users invest a great deal of time and money in learning to use computer interfaces. Far more, in fact, than software developers invest in developing and even implementing the interfaces. Whoever can own an interface, has made its users into captives, and misappropriated their investment.

To protect our freedom from monopolies like these, a group of programmers and users have formed a grass-roots political organization, the League for Programming Freedom.

The purpose of the League is to oppose monopolistic practices such as interface copyright and software patents. The League calls for a return to the legal policies of the recent past, in which programmers could program freely. The League is not concerned with free software as an issue, and is not affiliated with the Free Software

Foundation.

The League’s activities include publicizing the issues, as is being done here, and filing friend-of-the-court briefs on behalf of defendants sued by monopolists.

The League’s membership rolls include Donald Knuth, the foremost authority on algorithms, John McCarthy, inventor of Lisp, Marvin Minsky, founder of the MIT Artificial Intelligence lab, Guy L. Steele, Jr., author of well-known books on Lisp and C, as well as Richard Stallman, the developer of GCC. Please join and add your name to the list. Membership dues in the League are \$42 per year for programmers, managers and professionals; \$10.50 for students; \$21 for others.

Activist members are especially important, but members who have no time to give are also important. Surveys at major ACM conferences have indicated a vast majority of attendees agree with the League on both issues (interface copyrights and software patents). If just ten percent of the programmers who agree with the League join the League, we will probably triumph.

To join, or for more information, phone (617) 243-4091 or write to:

League for Programming Freedom
1 Kendall Square #143
P.O. Box 9171
Cambridge, MA 02139

You can also send electronic mail to lpf@uunet.uu.net.

In addition to joining the League, here are some suggestions from the League for other things you can do to protect your freedom to write programs:

- Tell your friends and colleagues about this issue and how it threatens to ruin the computer industry.
- Mention that you are a League member in your ‘.signature,’ and mention the League’s email address for inquiries.
- Ask the companies you consider working for or working with to make statements against software monopolies, and give preference to those that do.
- When employers ask you to sign contracts giving them copyright on your work, insist on a clause saying they will not claim the copyright covers imitating the interface.
- When employers ask you to sign contracts giving them patent rights, insist on clauses saying they can use these rights only defensively. Don’t rely on company policy, since policies can change at any time; don’t rely on an individual executive’s private word, since that person may be replaced. Get a commitment just as binding as the commitment they get from you.
- Write to Congress to explain the importance of these issues.

House Subcommittee on Intellectual Property
2137 Rayburn Building

Washington, DC 20515

Senate Subcommittee on Patents, Trademarks and Copyrights

United States Senate

Washington, DC 20510

(These committees have received lots of mail already; let's give them even more.)

Democracy means nothing if you don't use it. Stand up and be counted!



Installing GCC

The following documentation discusses installation and what you need to get the GNU compiler collection (GCC) to work for you, if you are using the compiler tools that are not part of the GNUPro Toolkit release. This material may be useful if you are installing a net release version of the GNU Compiler Collection.

- ““Installing GCC on Systems When It Exists”” (below)
- “Configurations That GCC Supports” on page 303
- “Compilation in a Separate Directory” on page 318
- “Building and Installing a Cross-compiler” on page 319
- “Steps of Cross-compilation” on page 319
- “Standard Header File Directories” on page 324
- “Actually Building the Cross-compiler” on page 324
- “collect2 and Cross-compiling” on page 325

See also “Installing GCC on Sun” on page 326 and “Installing GCC on VMS” on page 327.

Installing GCC on Systems When It Exists

The following documentation only applies if you are trying to do a rebuild of the GNUPro tools, if the tools do not work; it is only for when a compiler is already available on the system for which you intend to develop. In this documentation, we assume you compile in the same directory that contains the source files; see “Compilation in a Separate Directory” on page 318 to find out how to compile in a separate directory on UNIX systems. For Sun systems, see “Installing GCC on Sun” on page 326, and for VMS systems, see “Installing GCC on VMS” on page 327. For Windows, you need to get the complete compilation package, DJGPP, which includes binaries as well as sources along with all the necessary compilation tools and libraries.

Installing GCC on UNIX Systems

What follows is the procedure for installing GCC on a UNIX system.

1. If you have built GCC previously in the same directory for a different target machine, do `make distclean` to delete all files that might be invalid. One of the files `make distclean` deletes is `Makefile`; if `make distclean` complains that `Makefile` does not exist, it probably means that the directory is already suitably clean.
2. On a System V release 4 system, make sure `/usr/bin` precedes `/usr/ucb` in `PATH`. The `cc` command in `/usr/ucb` uses libraries which have bugs.
3. Specify the host, build and target machine configurations. You do this by running the file, `configure`.

The *build* machine is the system which you are using, the *host* machine is the system where you want to run the resulting compiler (normally the build machine), and the *target* machine is the system for which you want the compiler to generate code.

If you are building a compiler to produce code for the machine it runs on (a native compiler), you normally do not need to specify any operands to `configure`; it will try to guess the type of machine you are on and use that as the build, host and target machines. So you don't need to specify a configuration when building a native compiler unless `configure` cannot figure out what your configuration is or guesses wrong. In those cases, specify the build machine's *configuration name* with the option, `'--build'`; the host and target will default to be the same as the build machine. (If you are building a cross-compiler, see “Building and Installing a Cross-compiler” on page 319.) The following is an example.

```
./configure--build=sparc-sun-sunos4.1
```

A configuration name may be canonical or it may be more or less abbreviated. A canonical configuration name has three parts, separated by dashes, like the following example where *cpu* designates the processor, *company* designates the company who makes the processor, and *system* designates the actual system for which the processor is configured.

cpu-company-system

The three parts may themselves contain dashes; `configure` can figure out which dashes serve which purpose. For example, `m68k-sun-sunos4.1` specifies a Sun 3. You can also replace parts of the configuration by nicknames or aliases. For example, `sun3` stands for `m68k-sun`, so `sun3-sunos4.1` is another way to specify a Sun 3. You can also use simply `sun3-sunos`, since the version of SunOS is assumed by default to be version 4. You can specify a version number after any of the system types, and some of the CPU types. In most cases, the version is irrelevant, and will be ignored. So you might as well specify the version if you know it.

WARNING! See “Configurations That GCC Supports” for the supported configuration names and notes on many of the configurations. See also *Getting Started Guide*.

There are four additional options you can specify independently to describe variant hardware and software configurations.

These are: `--with-gnu-as`, `--with-gnu-ld`, `--with-stabs` and `--nfp`.

`--with-gnu-as`

If you will use GCC with the GNU assembler (GAS), you should declare this by using the `--with-gnu-as` option when you run `configure`.

Using this option does not install GAS. It only modifies the output of GCC to work with GAS. Building and installing GAS is up to you.

Conversely, if you do *not* wish to use GAS *and* do *not* specify `--with-gnu-as` when building GCC, it is up to you to make sure that GAS is *not* installed.

GCC searches for a program named `as` in various directories; if the program it finds is GAS, then it runs GAS. If you are not sure where GCC finds the assembler it is using, try specifying `-v` when you run it.

The systems where it makes a difference whether you use GAS are `hppa1.0-any-any`, `hppa1.1-any-any`, `i386-any-sysv`, `i386-any-isc`, `i860-any-bsd`, `m68k-bull-sysv`, `m68k-hp-hpux`, `m68k-sony-bsd`, `m68k-altos-sysv`, `m68000-hp-hpux`, `m68000-att-sysv`, `any-lynx-lynxos` and `mips-any`. (*any* in these designations refers to any version of the specified configuration.) On any other system, `--with-gnu-as` has no effect.

On the previously listed systems (except for the HPPA, for ISC on the 386, and for `mips-sgi-irix5.*`), if you use GAS, you should also use the GNU linker (specifying with the option, `--with-gnu-ld`).

`--with-gnu-ld`

Specify the option, `--with-gnu-ld`, if you plan to use the GNU linker with GCC.

This option does not cause the GNU linker to be installed; it just modifies the behavior of GCC to work with the GNU linker. Specifically, it inhibits the installation of `collect2`, a program which otherwise serves as a front-end for the system's linker on most configurations.

`--with-stabs`

On MIPS based systems and on Alphas, you must specify whether you want GCC to create the normal ECOFF debugging format, or to use BSD-style `stabs` passed through the ECOFF symbol table. The normal ECOFF debug format cannot fully handle languages other than C. BSD `stabs` format can handle other languages, but it only works with the GNU debugger GDB.

Normally, GCC uses the ECOFF debugging format by default; if you prefer BSD `stabs`, specify `--with-stabs` when you configure GCC.

No matter which default you choose when you configure GCC, you can use the `-gcoff` and `-gstabs+` options to specify explicitly the debug format for a particular compilation.

`--with-stabs` is meaningful on the ISC system on the 386 as well as if using `--with-gas`. It selects use of `stabs` debugging information embedded in COFF output. This kind of debugging information supports C++ well; ordinary COFF debugging information does not.

`--with-stabs` is also meaningful on 386 systems running SVR4. It selects use of `stabs` debugging information embedded in ELF output. The C++ compiler currently (2.6.0) does not support the DWARF debugging information normally used on 386 SVR4 platforms; `stabs` provide a workable alternative. This requires `gas` and `gdb`, as the normal SVR4 tools can not generate or interpret `stabs`.

`--nfp`

On certain systems, you must specify whether the machine has a floating point unit. These systems include `m68k-sun-sunosn` and `m68k-isi-bsd`. On any other system, `--nfp` currently has no effect, though perhaps there are other systems where it could usefully make a difference.

`--enable-objthreads=type`

Certain systems, notably Linux, can't be relied on to supply a threads facility for the Objective C runtime and so will default to single-threaded runtime.

They may, however, have a library threads implementation available, in which case threads can be enabled with this option by supplying a suitable *type*, probably 'posix'. The possibilities for *type* are 'single', 'posix', 'win32', 'solaris', 'irix' and 'mach'.

`--enable-leading-underscore`

```
--disable-leading-underscore
  --enable-leading-underscore forces generation of underscores for
  assembly output using a.out object file format to override the default user
  label prefix used by the compiler and recognized by the binary utilities and the
  debugger. --disable-leading-underscore will do the opposite.
```

The `configure` script searches subdirectories of the source directory for other compilers that are to be integrated into GCC.

The GNU compiler for C++, called `g++`, is in a subdirectory named `cp`. `configure` inserts rules into `Makefile` to build all of those compilers.

In the following, we clarify which files will be set up by `configure`. Normally you need not be concerned with these files.

- ❖ A file named `config.h` is created that contains a `#include` of the top-level configuration file for the machine you will run the compiler on (for a specific CPU and system, see “Configurations That GCC Supports” on page 303).

This file is responsible for defining information about the host machine. It includes `tm.h`.

The top-level configuration file is located in the subdirectory, `config`. Its name is always `xm- something.h`; usually, `xm- machine.h`, but there are some exceptions.

If your system does not support symbolic links, you might want to set up `config.h` to contain a `#include` command which refers to the appropriate file.

- ❖ A file named `tconfig.h` is created which includes the top-level configuration file for your target machine. This is used for compiling certain programs to run on that machine.
- ❖ A file named `tm.h` is created which includes the machine-description macro file for your target machine. It should be in the subdirectory `config` with the specific name of the system followed by `.h` (for a specific CPU and system, see “Configurations That GCC Supports” on page 303).
- ❖ The command file, `configure`, also constructs the file, `Makefile`, by adding some text to the template file, `Makefile.in`. The additional text comes from files in the `config` directory, named `t-target` and `x-host`. *If these files do not exist, it means nothing needs to be added for a given target or host.*

4. The standard directory for installing GCC is `/usr/local/lib`.

If you want to install its files somewhere else, specify `--prefix=dir` when you run `configure`. `dir` is a directory name to use instead of `/usr/local` for all purposes with one exception: the directory `/usr/local/include` is searched for header files no matter where you install the compiler. To override this name, use the `--local-prefix` option in the following documentation, Step 5.

5. Specify `--local-prefix=dir` if you want the compiler to search directory, `dir/include`, for locally installed header files instead of `/usr/local/include`. You should specify `--local-prefix` only if your site has a different convention (not `/usr/local`) for where to put site-specific files.
The default value for `--local-prefix` is `/usr/local` regardless of the value of `--prefix`. Specifying `--prefix` has no effect on which directory GCC searches for local header files. This may seem counterintuitive, but actually it is logical. The purpose of `--prefix` is to specify where to install GCC. The local header files in `/usr/local/include`—if you put any in that directory—are not part of GCC. They are part of other programs—perhaps many others. (GCC installs its own header files in another directory which is based on the `--prefix` value.)

WARNING! Do not specify `/usr` as the `--local-prefix`!

The directory you use for `--local-prefix` **must not** contain any of the system's standard header files. If it did contain them, certain programs would be miscompiled (including GNU Emacs, on certain targets), because this would override and nullify the header file corrections made by the `fixincludes` script.

Indications are that people use this option use it based on mistaken ideas of its purpose. They use it as if it specified where to install GCC, perhaps on the assumption that installing GCC creates this directory.

6. Make sure the Bison parser generator is installed. (This is unnecessary if the Bison output files `c-parse.c` and `cexp.c` are more recent than `c-parse.y` and `cexp.y` and you do not plan to change the `.y` files.)
Bison versions older than Sept 8, 1988 will produce incorrect output for `c-parse.c`.
7. If you have chosen a configuration for GCC which requires other GNU tools (such as GAS or the GNU linker) instead of the standard system tools, install the required tools in the build directory under the names `as`, `ld`, or whatever is appropriate. This will enable the compiler to find the proper tools for compilation of the program `enquire`.
Alternatively, you can do subsequent compilation using a value of the `PATH` environment variable such that the necessary GNU tools come before the standard system tools.
8. Build the compiler. Just type `make LANGUAGES=c` in the compiler directory. `LANGUAGES=c` specifies that only the C compiler should be compiled. The `Makefile` normally builds compilers for all the supported languages; currently, C, C++ and Objective C. However, C is the only language that is sure to work when you build with other non-GNU C compilers. In addition, building anything but C, at this stage, is a waste of time.

In general, you can specify the languages to build by typing the argument `'LANGUAGES="list"'` where *list* is one or more words from the list `c`, `c++`, and `objective-c`. If you have any additional GNU compilers as subdirectories of the GCC source directory, you may also specify their names in this list.

Ignore any warnings you may see about `"statement not reached"` in `'insn-emit.c'`; they are normal.

Also, warnings about `"unknown escape sequence"` are normal in `'genopinit.'` and perhaps some other files. Likewise, you should ignore warnings about `"constant is so large that it is unsigned"` in `'insn-emit.c'` and `'insn-recog.c'` and a warning about a comparison always being zero in `'enquire.o'` and warnings about shift counts exceeding type widths in `'cexp.y'`. Any other compilation errors may represent bugs in the port to your machine or operating system, and should be investigated and reported (see *Getting Started Guide*). Some commercial compilers fail to compile GCC because they have bugs or limitations. For example, the Microsoft compiler is said to run out of macro space. Some Ultrix compilers run out of expression space; then you need to break up the statement where the problem happens.

9. If you are building a cross-compiler, stop here. See "Building and Installing a Cross-compiler" on page 319.
10. Move the first-stage object files and executables into a subdirectory with the following command:

```
make stage1
```

The files are moved into a subdirectory named `stage1`. Once installation is complete, you may wish to delete these files with `rm -r stage1`.

11. If you have chosen a configuration for GCC which requires other GNU tools (such as GAS or the GNU linker) instead of the standard system tools, install the required tools in the `stage1` subdirectory under the names `as`, `ld` or whatever is appropriate. This will enable the stage 1 compiler to find the proper tools in the following stage.

Alternatively, you can do subsequent compilation using a value of the `PATH` environment variable such that the necessary GNU tools come before the standard system tools.

12. Recompile the compiler with itself, with this command:

```
make CC="stage1/xgcc -Bstage1/" CFLAGS="-g -O2"
```

This is called making the stage 2 compiler. The command shown in the previous example builds compilers for all the supported languages. If you don't want them all, you can specify the languages to build by typing the argument, `LANGUAGES="list"`, where *list* should contain one or more of the words `c`, `c++`, `objective-c` or `proto`. Separate the words with spaces. `proto` stands for the programs `protoize` and `unprotoize`; they are not a separate language, but you use `LANGUAGES` to enable or disable their installation. If you are going to build the

stage 3 compiler, then you might want to build only the C language in stage 2. Once you have built the stage 2 compiler, if you are short of disk space, you can delete the subdirectory `stage1`. On a 68000 or 68020 system lacking floating point hardware, unless you have selected a `‘tm.h’` file that expects by default to find no hardware; instead, use the following statement.

```
make CC="stage1/xgcc -Bstage1/" CFLAGS="-g -O2 -msoft-float"
```

- 13.** If you wish to test the compiler by compiling it with itself one more time, install any other necessary GNU tools (such as GAS or the GNU linker) in the `stage2` subdirectory as you did in the `stage1` subdirectory; then, use the following. This is called making the stage 3 compiler.

```
make stage2
make CC="stage2/xgcc -Bstage2/" CFLAGS="-g -O2"
```

Aside from the `‘-B’` option, the compiler options should be the same as when you made the stage 2 compiler. But the `LANGUAGES` option need not be the same. The command in the previous example builds compilers for all the supported languages; if you don’t want them all, you can specify the languages to build by typing the argument `‘LANGUAGES= list’` as described in Step 8. If you don’t have to install any additional GNU tools, you may use the following command instead of making `stage1`, `stage2`, and performing the two compiler builds.

```
make bootstrap LANGUAGES= language-list \
                BOOT_CFLAGS= option-list
```

- 14.** Then compare the latest object files with the stage 2 object files—they ought to be identical, aside from time stamps (if any). On some systems, meaningful comparison of object files is impossible; they always appear *different*. This is currently true on Solaris and some systems that use ELF object file format. On some versions of Irix on SGI machines and DEC UNIX (OSF/1) on Alpha systems, you will not be able to compare the files without specifying `-save-temps`; see the description of individual systems in the previous discussions to see if you get comparison failures. You may have similar problems on other systems. Use the following command to compare the files.

```
make compare
```

This will mention any object files that differ between stage 2 and stage 3. Any difference, no matter how innocuous, indicates that the stage 2 compiler has compiled GCC incorrectly, and is therefore a potentially serious bug which you should investigate and report (see *Getting Started with GNUPro Toolkit*). If your system does not put time stamps in the object files, then use the following as a faster way to compare them (using the Bourne shell).

```
for file in *.o; do
  cmp $file stage2/$file
done
```

IMPORTANT! If you built the compiler on MIPS machines with the option,

-mno-mips-tfile, you can't compare files.

15. Install the compiler driver, the compiler's passes and run-time support with 'make install'. Use the same value for CC, CFLAGS and LANGUAGES that you used when compiling the files that are being installed. One reason this is necessary is that some versions of Make have bugs and recompile files gratuitously when you use this step. If you use the same variable values, those files will be properly recompiled.

For example, if you have built the stage 2 compiler, you can use the following command.

```
make install CC="stage2/xgcc -Bstage2/" CFLAGS="-g -O" \
             LANGUAGES=" list"
```

This copies the files 'cc1', 'cpp' and 'libgcc.a' to files 'cc1', 'cpp' and 'libgcc.a' in the directory /usr/local/lib/gcc-lib/target/version, which is where the compiler driver program looks for them. Here, target is the target machine type specified when you ran configure, and version is the version number of GCC.

This naming scheme permits various versions and/or cross-compilers to coexist. It also copies the executables for compilers for other languages (e.g., 'cc1plus' for C++) to the same directory.

This also copies the driver program 'xgcc' into /usr/local/bin/gcc, so that it appears in typical execution search paths. It also copies 'gcc.1' into /usr/local/man/man1 and info pages into /usr/local/info.

On some systems, this command causes recompilation of some files. This is usually due to bugs in make. You should either ignore this problem, or use GNU Make.

WARNING! There is a bug in `alloca` in the Sun library. To avoid this bug, be sure to *install the executables of GCC that were compiled by GCC*. (The executables from stage 2 or 3, *not* stage 1.) They use `alloca` as a built-in function, *never* the one in the library.

(It is usually better to install GCC executables from stage 2 or 3, as they usually run faster than ones compiled with some other compiler.)

16. If you're going to use C++, it's likely that you also need to install a C++ runtime library. Just as GNU C does not distribute a C runtime library, it also does not include a C++ run-time library. All I/O functionality, special class libraries, etc., are provided by the C++ run-time library.

The standard C++ runtime library for GNU CC is called 'libstdc++'. An obsolescent library 'libg++' may also be available, but it's necessary only for older software that hasn't been converted yet; if you don't know whether you need 'libg++' then you probably don't need it.

The following procedure shows one way to build and install 'libstdc++' for

GNU CC.

- a. Build and install GNU CC, so that invoking `'gcc'` obtains the GNU CC that was just built.
- b. Obtain a copy of a compatible `'libstdc++'` distribution. For example, the `'libstdc++-2.8.0.tar.gz'` distribution should be compatible with GCC 2.8.0. GCC distributors normally distribute `'libstdc++'` as well.
- c. Set the `'CXX'` environment variable to `'gcc'` while running the `'libstdc++'` distribution's `'configure'` command. Use the same `'configure'` options that you used when you invoked GCC's `'configure'` command.
- d. Invoke `'make'` to build the C++ runtime.
- e. Invoke `'make install'` to install the C++ runtime.

To summarize, after building and installing GNU CC, invoke the following shell commands in the topmost directory of the C++ library distribution. For *configure-options*, use the same options that you used to configure GNU CC.

```
$ CXX=gcc ./configure configure-options
$ make
$ make install
```

6. GCC includes a runtime library for Objective-C because it is an integral part of the language. You can find the files associated with the library in the subdirectory `'objc'`. The GNU Objective-C Runtime Library requires header files for the target's C library in order to be compiled, and also requires the header files for the target's thread library if you want thread support. See "Cross-compilers and Header Files" on page 323 for discussion about header files issues for cross-compilation.

When you run `'configure'`, it picks the appropriate Objective-C thread implementation file for the target platform. In some cases, you may wish to choose a different back-end as some platforms support multiple thread implementations or you may wish to disable thread support completely. To do this, specify a value for the `OBJC_THREAD_FILE` makefile variable on the command line when you run `make` with something like the following input.

```
make CC="stage2/xgcc -Bstage2/" CFLAGS="-g -O2"
      OBJC_THREAD_FILE=thr-single
```

The following list shows the currently available back-ends.

- ✦ `thr-single`
Disable thread support, should work for all platforms.
- ✦ `thr-decosf1`
DEC OSF/1 thread support.
- ✦ `thr-irix`
SGI IRIX thread support.

- ❖ `thr-mach`
Generic MACH thread support, known to work on NEXTSTEP.
- ❖ `thr-os2`
IBM OS/2 thread support.
- ❖ `thr-posix`
Generic POSIX thread support.
- ❖ `thr-pthreads`
PThreads on Linux-based GNU systems.
- ❖ `thr-solaris`
SUN Solaris thread support.
- ❖ `thr-win32`
Microsoft Win32 API thread support.

Configurations That GCC Supports

The following are the possible CPU types:

`1750a`, `a29k`, `alpha`, `arm`, `cn`, `clipper`, `dsp16xx`, `elxsi`, `h8300`, `hppa1.0`,
`hppa1.1`, `i370`, `i386`, `i486`, `i586`, `i860`, `i960`, `m68000`, `m68k`, `m88k`, `mips`, `mipsel`,
`mips64`, `mips64el`, `ns32k`, `powerpc`, `powerpcle`, `pyramid`, `romp`, `rs6000`, `sh`,
`sparc`, `sparclite`, `sparc64`, `vax`, `we32k`.

The following are the recognized company names. As you can see, customary abbreviations are used rather than the longer official names.

`acorn`, `alliant`, `altos`, `apollo`, `apple`, `att`, `bull`, `cbm`, `convergent`, `convex`,
`crds`, `dec`, `dg`, `dolphin`, `elxsi`, `encore`, `harris`, `hitachi`, `hp`, `ibm`, `intergraph`,
`isi`, `mips`, `motorola`, `ncr`, `next`, `ns`, `omron`, `plexus`, `sequent`, `sgi`, `sony`, `sun`,
`tti`, `unicom`, `wrs`.

The company name is meant only to clarify when the rest of the information supplied is insufficient. You can omit it, substituting '`cpu-system`' (where `cpu` stands for your processor and `system` for your operating system), if the company name is not needed. For example, `vax-ultrix4.2` is equivalent to `vax-dec-ultrix4.2`.

The following is a list of system types:

`386bsd`, `aix`, `acis`, `amigados`, `aos`, `aout`, `aux`, `bosx`, `bsd`, `clix`, `coff`, `ctix`, `cxux`,
`dgux`, `dynix`, `ebmon`, `ecoff`, `elf`, `esix`, `freebsd`, `hms`, `genix`, `gnu`, `gnu/linux`,
`hiux`, `hpux`, `iris`, `irix`, `isc`, `luna`, `lynxos`, `mach`, `minix`, `msdos`, `mvs`, `netbsd`,
`newsos`, `nindy`, `ns`, `osf`, `osfrose`, `ptx`, `riscix`, `riscos`, `rtu`, `sco`, `sim`, `solaris`,
`sunos`, `sym`, `sysv`, `udi`, `ultrix`, `unicos`, `uniplus`, `unos`, `vms`, `vsta`, `vxworks`,
`winnt`, `xenix`.

You can omit the system type; then `configure` guesses the operating system from the CPU and company.

You can add a version number to the system type; this may or may not make a difference. For example, you can write 'bsd4.3' or 'bsd4.4' to distinguish versions of BSD. In practice, the version number is most needed for `sysv3` and `sysv4`, which are often treated differently.

If you specify an impossible combination such as `i860-dg-vms`, then you may get an error message from `configure`, or it may ignore part of the information and do the best it can with the rest. `configure` always prints the canonical name for the alternative that it used. GCC does not support all possible alternatives.

Often a particular model of machine has a name. Many machine names are recognized as aliases for CPU/company combinations. Thus, the machine name `sun3`, mentioned previously, is an alias for `m68k-sun`.

Sometimes we accept a company name as a machine name, when the name is popularly used for a particular machine. The following are the known machine names:

```
3300, 3b1, 3bn, 7300, altos3068, altos, apollo68, att-7300, balance,
convex-cn, crds, decstation-3100, decstation, delta, encore, fx2800,
gmicro, hp7nn, hp8nn, hp9k2nn, hp9k3nn, hp9k7nn, hp9k8nn, iris4d, iris,
isi68, m3230, magnum, merlin, miniframe, mmax, news-3600, news800, news,
next, pbd, pc532, pmax, powerpc, powerpcle, ps2, risc-news, rtpc, sun2,
sun386i, sun386, sun3, sun4, symmetry, tower-32, tower.
```

Remember that a machine name specifies both the CPU type *and* the company name.

If you want to install your own homemade configuration files, you can use 'local' as the company name to access them.

If you use configuration `cpu-local`, the configuration name without the `cpu` prefix is used to form the configuration filenames. Thus, if you specify 'm68k-local', configuration uses files `m68k.md`, `local.h`, `m68k.c`, `xm-local.h`, `t-local`, and `x-local`, all in the directory, `config/m68k`.

What follows is a list of configurations that have special treatment or special things you must know.

1750a-**-*

MIL-STD-1750A processors.

The MIL-STD-1750A cross configuration produces output for `as1750`, an assembler/linker available under the GNU Public License for the 1750A.

Download from `ftp://ftp.fta-berlin.de/pub/crossgcc/1750gals/` to get `as1750`. A similarly licensed simulator for the 1750A is available from the same address.

You should ignore a fatal error during the building of `libgcc` (`libgcc` is not yet implemented for the 1750A.)

The `as1750` assembler requires the file, `ms1750.inc`, which is found in the

directory `config/1750a`.

GCC produced the same sections as the Fairchild F9450 C Compiler; namely, they are the following sections.

- `Normal`
The program code section.
- `Static`
The read/write (RAM) data section.
- `Konst`
The read-only (ROM) constants section.
- `Init`
Initialization section (code to copy KREL to SREL).

The smallest addressable unit is 16 bits (BITS PER UNIT is 16). This means that type `char` is represented with a 16-bit word per character. The 1750A's "Load/Store Upper/Lower Byte" instructions are not used by GCC.

`alpha-*-osf1`

Systems using processors that implement the DEC Alpha architecture and are running the DEC UNIX (OSF/1) operating system, for example the DEC Alpha AXP systems. (VMS on the Alpha is not currently supported by GCC.)

GCC writes a `.verstamp` directive to the assembler output file unless it is built as a cross-compiler. It gets the version to use from the system header file, `/usr/include/stamp.h`. If you install a new version of DEC UNIX, you should rebuild GCC to pick up the new version stamp.

IMPORTANT! Since the Alpha is a 64-bit architecture, cross-compilers from 32-bit machines will not generate code as efficient as that generated when the compiler is running on a 64-bit machine. That is because many optimizations that depend on being able to represent a word on the target in an integral value on the host cannot be performed.

Building cross-compilers on the Alpha for 32-bit machines has only been tested in a few cases and may not work properly.

`make compare` may fail on old versions of DEC UNIX unless you add `savetemps` to `CFLAGS`. On these systems, the name of the assembler input file is stored in the object file, and that makes comparison fail if it differs between the `stage1` and `stage2` compilations. The option, `-save-temps`, forces a fixed name to be used for the assembler input file, instead of a randomly chosen name in `/tmp`.

Do not add `-save-temps` unless the comparisons fail without that option. If you add `-save-temps`, you will have to manually delete the `.i` and `.s` files after each series of compilations.

GCC now supports both the native (ECCOFF) debugging format used by DBX and GDB and an encapsulated STABS format for use only with GDB. See Step 3 and its discussion of the `--with-stabs` option for `configure` on page 296 for more

information on these formats and how to select them.

There is a bug in DEC's assembler that produces incorrect line numbers for ECOFF format when the `.align` directive is used. To work around this problem, GCC will not emit such alignment directives while writing ECOFF format debugging information even if optimization is being performed. Unfortunately, this has the very undesirable side-effect that code addresses when `-o` is specified are different depending on whether or not `-g` is also specified.

To avoid this behavior, specify `-gstabs+` and use GDB instead of DBX. DEC is now aware of this problem with the assembler and hopes to provide a fix shortly. See "Options Controlling Debugging" on page 45.

`arm-*-aout`

Advanced RISC Machines ARM-family processors. These are often used in embedded applications. There are no standard UNIX configurations. This configuration corresponds to the basic instruction sequences and will produce `a.out` format object modules.

You may need to make a variant of the file, `arm.h`, for your particular configuration.

`arm-*-linuxaout`

Any of the ARM family processors running the Linux-based GNU system with the `a.out` binary format (ELF is not yet supported). You must use version 2.8.1.0.7 or later of the Linux binutils; download it from `'sunsite.unc.edu:/pub/Linux/GCC'` and other mirror sites for Linux-based GNU systems.

`arm-*-riscix`

The ARM2 or ARM3 processor running RISC iX, Acorn's port of BSD UNIX. If you are running a version of RISC iX prior to 1.2, then you must specify the version number during configuration.

IMPORTANT! The assembler shipped with RISC iX does not support `stabs` debugging information; a new version of the assembler, with `stabs` support included, is now available from Acorn.

`a29k`

AMD Am29k-family processors. These are normally used in embedded applications. There are no standard UNIX configurations. This configuration corresponds to AMD's standard calling sequence and binary interface and is compatible with other 29k tools.

You may need to make a variant of the file `a29k.h` for your particular configuration.

`a29k-*-bsd`

AMD Am29050 used in a system running a variant of BSD UNIX.

`decstation-*`

DECstations can support three different personalities: Ultrix, DEC OSF/1, and OSF/rose. To configure GCC for these platforms use the following configurations:

- ❖ `decstation-ultrix`
Ultrix configuration.
- ❖ `decstation-osf1`
Dec's version of OSF/1.
- ❖ `decstation-osfrose`
Open Software Foundation reference port of OSF/1 which uses the OSF/rose object file format instead of ECOFF. Normally, you would not select this configuration.

The MIPS C compiler needs to be told to increase its table size for switch statements with the `-wf, -XNg1500` option in order to compile `cp/parse.c`. If you use the `'-O2'` optimization option, you also need to use `-Olimit 3000`. Both of these options are automatically generated in the `'Makefile'` that the shell script `'configure'` builds. If you override the `CC` `make` variable and use the MIPS compilers, you may need to add `-wf, -XNg1500 -Olimit 3000`.

`elxsi-elxsi-bsd`

The Elxsi's C compiler has known limitations that prevent it from compiling GNU C. Please contact `mrs@cygnus.com` for more details.

`dsp16xx`

A port to the AT&T DSP1610 family of processors.

`h8300-*-*`

The calling convention and structure layout has changed in release 2.6. All code must be recompiled. The calling convention now passes the first three arguments in function calls in registers. Structures are no longer a multiple of 2 bytes.

`hppa*-*-*`

There are several variants of the HP-PA processor which run a variety of operating systems. GCC must be configured to use the correct processor type and operating system, or GCC will not function correctly. The easiest way to handle this problem is to avoid specifying a target when configuring GCC. The `'configure'` script will try to automatically determine the right processor type and operating system.

`-g` does not work on HP/UX, since that system uses a peculiar debugging format about which GCC does not know. However, `-g` will work if you also use GAS and GDB in conjunction with GCC. We highly recommend using GAS for all HPPA configurations.

You should be using GAS-2.6 (or later) along with GDB-4.16 (or later). These can be retrieved from all the traditional GNU ftp archive sites. Install GAS into a directory before `/bin`, `/usr/bin`, and `/usr/ccs/bin` in your search path.

To enable debugging, configure GCC with the `--with-gnu-as` option before

building.

i370-*-*

This port is very preliminary and has many known bugs. We hope to have a higher-quality port for this machine soon.

i386-*-linuxoldld

Use this configuration to generate a.out binaries on Linux if you do not have gas/binutils version 2.5.2 or later installed. This is an obsolete configuration.

i386-*-linuxaout

Use this configuration to generate 'a.out' binaries on Linux. This configuration is being superseded. You must use gas/binutils version 2.5.2 or later.

i386-*-linux-gnu

Use this configuration to generate ELF binaries on Linux. You must use gas/binutils version 2.5.2 or later.

i386-*-sco

Compilation with RCC is recommended. Also, it may be a good idea to link with GNU malloc instead of the malloc that comes with the system.

i386-*-sco3.2v4

Use this configuration for SCO release 3.2 version 4.0.

i386-*-sco3.2v5

Use this for the SCO OpenServer Release family including 5.0.0, 5.0.2, 5.0.4, 5.0.5, Internet FastStart 1.0, and Internet FastStart 1.1.

GNU CC can generate COFF binaries if you specify '-mcoff' or ELF binaries, the default. A full 'make bootstrap' is recommended so that an ELF compiler that builds ELF is generated.

You must have TLS597 from ftp://ftp.sco.com/TLS installed for ELF C++ binaries to work correctly on releases before 5.0.4.

The native SCO assembler that is provided with the OS at no charge is normally required. If, however, you must need to use the GNU assembler (perhaps you have complex asms) you must configure this package '--with-gnu-as'. To do this, install (using the commands, cp or symlink) gcc/as to your copy of the GNU assembler. You must use a recent version of GNU binutils; version 2.9.1 seems to work well.

If you select this option, you will be unable to build COFF images. Trying to do so will result in non-obvious failures. In general, the '-with-gnu-as' option isn't as well tested as the native assembler.

IMPORTANT! If you are building C++, you must follow the instructions about invoking 'make bootstrap' because the native OpenServer compiler may build a 'cc1plus' that will not correctly parse many valid C++ programs. You must do a 'make bootstrap' if you are building with the native compiler.

`i386-*-isc`

It may be a good idea to link with GNU `malloc` instead of the `malloc` that comes with the system. In ISC version 4.1, `sed` core dumps when building `'deduced.h'`. Use the version of `sed` from version 4.0.

`i386-*-esix`

It may be good idea to link with GNU `malloc` instead of the `malloc` that comes with the system.

`i386-ibm-aix`

You need to use GAS version 2.1 or later, and LD from GNU binutils version 2.2 or later.

`i386-sequent-bsd`

Go to the Berkeley universe before compiling.

`i386-sequent-ptr1*``i386-sequent-ptr2*`

You must install GNU `'sed'` before running `'configure'`.

`i386-sun-sunos4`

You may find that you need another version of GCC to begin bootstrapping with, since the current version when built with the system's own compiler seems to get an infinite loop compiling part of `'libgcc2.c'`. GCC version 2 compiled with GCC (any version) seems not to have this problem. See "Installing GCC on Sun" on page 326 for information on installing GCC on Sun systems.

`i[345]86-*-winnt3.5`

This version requires a GAS that has not yet been released. Until it is, you can get a pre-built binary version via anonymous ftp from `'cs.washington.edu:pub/gnat'` or `'cs.nyu.edu:pub/gnat'`. You must also use the Microsoft header files from the Windows NT 3.5 SDK. Find these on the CDROM in the `/mstools/h` directory dated September 4, 1994. You must use a fixed version of Microsoft linker made especially for NT 3.5, which is also available on the NT 3.5 SDK CDROM. If you do not have this linker, can you also use the linker from Visual C/C++ 1.0 or 2.0.

Installing GCC for NT builds a wrapper linker, called `'ld.exe'`, which mimics the behavior of UNIX `'ld'` in the specification of libraries (`'-L'` and `'-l'`). `'ld.exe'` looks for both UNIX and Microsoft named libraries. For example, if you specify `'-lfoo'`, `'ld.exe'` will look first for `'libfoo.a'` and then for `'foo.lib'`. You may install GCC for Windows NT in one of two ways, depending on whether or not you have a UNIX-like shell and various UNIX-like utilities.

- ✦ If you do not have a UNIX-like shell and few UNIX-like utilities, you will use a DOS style batch script called `'configure.bat'`.

Invoke it as `'configure winnt'` from an MSDOS console window or from the program manager dialog box. `'configure.bat'` assumes you have already installed and have in your path a UNIX-like `sed` program which is used to create a working Makefile from `'Makefile.in'`. Makefile uses the

Microsoft `Nmake` program maintenance utility and the Visual C/C++ V8.00 compiler to build GCC.

You only need the utilities, `sed` and `touch`, to use this installation method, which only automatically builds the compiler itself. You must then examine what `'fixinc.winnt'` does, edit the header files by hand and build `'libgcc.a'` manually.

- ✦ The second type of installation assumes you are running a UNIX-like shell, have a complete suite of UNIX-like utilities in your path, and have a previous version of GCC already installed, either through building it via the previous installation method or acquiring a pre-built binary. In this case, use the `configure` script in the normal fashion.

`i860-intel-osfl`

This is the Paragon. If you have version 1.0 of the operating system, see “Installation Problems” on page 338 for special things you need to do to compensate for peculiarities in the system.

`*-lynx-lynxos`

LynxOS 2.2 and earlier comes with GCC 1.x already installed as `/bin/gcc`. You should compile with `/bin/gcc` instead of `/bin/cc`. You can tell GCC to use the GNU assembler and linker, by specifying the following declaration when configuring.

```
--with-gnu-as --with-gnu-ld
```

These will produce COFF format object files and executables; otherwise GCC will use the installed tools, which produce `'a.out'` format executables.

`m32r-*-elf`

Mitsubishi M32R processor, a configuration intended for embedded systems.

`m68000-hp-bsd`

HP 9000 series 200 running BSD.

IMPORTANT! The C compiler that comes with this system cannot compile GCC; contact `'law@cs.utah.edu'` to get binaries of GCC for bootstrapping.

`m68k-altos`

Altos 3068. You must use the GNU assembler, linker and debugger. Also, you must fix a kernel bug. Details in the file, `'README.ALTOS'`.

`m68k-apple-aux`

Apple Macintosh running A/UX. You may configure GCC to use either the system assembler and linker or the GNU assembler and linker.

You should use the GNU configuration if you can, especially if you also want to use G++. You enabled that configuration with the options, `--with-gnu-as` and `--with-gnu-ld`, to configure.

IMPORTANT! The C compiler that comes with this system cannot compile GCC. You can find binaries of GCC for bootstrapping on `jagubox.gsfc.nasa.gov`. You

will also a patched version of `‘/bin/ld’` there that raises some of the arbitrary limits found in the original.

m68k-att-sysv

AT&T 3b1, a.k.a. 7300 PC. Special procedures are needed to compile GCC with this machine’s standard C compiler, due to bugs in that compiler. You can bootstrap it more easily with previous versions of GCC if you have them.

Installing GCC on the 3b1 is difficult if you do not already have GCC running, due to bugs in the installed C compiler. However, the following procedure *might* work. (We were unable to test it.)

1. Comment out `‘#include “config.h”` line near the start of `‘cccp.c’` and do `‘make cpp’`. This makes a preliminary version of GNU `cpp`.
2. Save the old `‘/lib/cpp’` and copy the preliminary GNU `cpp` to that filename.
3. Undo your change in `‘cccp.c’`, or reinstall the original version, and do `‘make cpp’` again.
4. Copy this final version of GNU `cpp` into `‘/lib/cpp’`.
5. Replace every occurrence of `obstack_free` in the file, `‘tree.c’`, with `_obstack_free`.
6. Run `make` to get the first-stage GCC.
7. Reinstall the original version of `/lib/cpp`.
8. Now compile GCC with itself and install it in the normal fashion.

m68k-bull-sysv

Bull DPX/2 series 200 and 300 with BOS2.00.45 up to BOS-2.01. GCC works either with native assembler or GNU assembler.

You can use GNU assembler with native COFF generation by providing the `‘--with-gnu-as’` option to the configure script or use GNU assembler with the `‘dbx-in-coff’` encapsulation by providing the `‘--with-gnu-as-stabs’` option. For any problem with native assembler or for availability of the DPX/2 port of GAS, contact: F.Pierresteguy@frcl.bull.fr.

m68k-crds-unox

Use `‘configure unos’` for building on Unos. The Unos assembler is named `casm` instead of `as`.

For some strange reason, linking `/bin/as` to `/bin/casm` changes the behavior, and does not work.

So, when installing GCC, you should install the following script as `as` in the subdirectory where the passes of GCC are installed:

```
#!/bin/sh
casm $*
```

The default Unos library is named `‘libunos.a’` instead of `‘libc.a’`. To allow GCC to function, either change all references to `‘-lc` in `‘gcc.c’` to `‘-lunos’` or

link `‘/lib/libc.a’` to `‘/lib/libunos.a’`.

When compiling GCC with the standard compiler, to overcome bugs in the support of `alloca`, do not use `‘-o’` when making stage 2. Then use the stage 2 compiler with `‘-o’` to make the stage 3 compiler. This compiler will have the same characteristics as the usual stage 2 compiler on other systems. Use it to make a stage 4 compiler and compare that with stage 3 to verify proper compilation.

Unos uses memory segmentation instead of demand paging, so you will need a lot of memory. 5 Mb is barely enough if no other tasks are running. If linking `cc1` fails, try putting the object files into a library and linking from that library.

`m68k-hp-hpux`

HP 9000 series 300 or 400 running HP/UX. HP/UX version 8.0 has a bug in the assembler that prevents compilation of GCC. To fix it, get patch PHCO 4484 from HP.

In addition, if you wish to use the `‘--with-gnu-as’` function, you must use GAS (version 2.1 or later), and you must use the GNU linker (version 2.1 or later). Earlier versions of `gas` relied upon a program which converted the `gas` output into the native HP/UX format, but that program has not been kept up to date. GDB does not understand that native HP/UX format, so you must use GAS if you wish to use GDB.

`m68k-sun`

Sun 3. We do not provide a configuration file to use the Sun FPA by default, because programs that establish signal handlers for floating point traps inherently cannot work with the FPA. See “Installing GCC on Sun” on page 326 for information on installing GCC on Sun systems.

`m88k-*-svr3`

Motorola m88k running the AT&T/Unisoft/Motorola V.3 reference port. These systems tend to use the Green Hills C, revision 1.8.5, as the standard C compiler. There are apparently bugs in this compiler that result in object files differences between stage 2 and stage 3. If this happens, make the stage 4 compiler and compare it to the stage 3 compiler. If the stage 3 and stage 4 object files are identical, this suggests you encountered a problem with the standard C compiler; the stage 3 and 4 compilers may be usable. It is best, however, to use an older version of GCC for bootstrapping if you have one.

`m88k-*-dgux`

Motorola m88k running DG/UX. To build 88open BCS native or cross compilers on DG/UX, specify the configuration name as `'m88k-*-dguxbcs'` and build in the 88open BCS software development environment. To build ELF native or cross compilers on DG/UX, specify `'m88k-*-dgux'` and build in the DG/UX ELF development environment. You set the software development environment by issuing `'sde-target'` command and specifying either `'m88kbcs'` or `'m88kdguxelf'` as the operand. If you do not specify a configuration name, `configure` guesses the configuration based on the current software development environment.

`m88k-tektronix-sysv3`

Tektronix XD88 running UTekV 3.2e. Do not turn on optimization while building stage1 if you bootstrap with the buggy Green Hills compiler. Also, The bundled LAI System V NFS is buggy so if you build in an NFS mounted directory, start from a fresh reboot, or avoid NFS all together. Otherwise you may have trouble getting clean comparisons between stages.

`mips-mips-bsd`

MIPS machines running the MIPS operating system in BSD mode. It's possible that some old versions of the system lack the functions `memcpy`, `memcmp`, and `memset`. If your system lacks these, you must remove or undo the definition of `TARGET_MEM_FUNCTIONS` in `'mips-bsd.h'`.

The MIPS C compiler needs to be told to increase its table size for switch statements with the `'-wf,-XNg1500'` option in order to compile `'cp/parse.c'`. If you use the `'-O2'` optimization option, you also need to use `'-Olimit 3000'`. Both of these options are automatically generated in the `Makefile` that the shell script `configure` builds. If you override the `CC` make variable and use the MIPS compilers, you may need to add `'-wf,-XNg1500 -Olimit 3000'`.

`mips-mips-riscos*`

The MIPS C compiler needs to be told to increase its table size for switch statements with the `'-wf,-XNg1500'` option in order to compile `'cp/parse.c'`.

If you use the `'-O2'` optimization option, you also need to use `'-Olimit 3000'`. Both of these options are automatically generated in the `'Makefile'` that the shell script `configure` builds.

If you override the `CC` make variable and use the MIPS compilers, you may need to add `'-wf,-XNg1500 -Olimit 3000'`. MIPS computers running RISC-OS can support four different personalities: default, BSD 4.3, System V.3, and System V.4 (older versions of RISC-OS don't support V.4). To configure GCC for these platforms use the following configurations:

- ✦ `mips-mips-riscosrev`

Default configuration for RISC-OS, revision, *rev*.

- ❖ `mips-mips-riscosrevbsd`
BSD 4.3 configuration for RISC-OS, revision, *rev*.
- ❖ `mips-mips-riscosrevsysv4`
System V.4 configuration for RISC-OS, revision, *rev*.
- ❖ `mips-mips-riscosrevsysv`
System V.3 configuration for RISC-OS, revision, *rev*.

The revision, *rev* (mentioned in the previous paragraphs for the option, `mips-mips-riscos*`), is the version of RISC-OS to use. You must reconfigure GCC when going from a RISC-OS revision 4 to RISC-OS revision 5. This has the effect of avoiding a linker bug (see “Installation Problems” on page 338).

`mips-sgi-*`

In order to compile GCC on an SGI running IRIX 4, the `c.hdr.lib` option must be installed from the CD-ROM supplied from Silicon Graphics. This is found on the second CD in release 4.0.1.

In order to compile GCC on an SGI running IRIX 5, the `compiler dev.hdr` subsystem must be installed by the IDO CD-ROM, supplied by Silicon Graphics. `make compare` may fail on version 5 IRIX unless you add `-save-temps` to `CFLAGS`. On these systems, the name of the assembler input file is stored in the object file, and that makes `compare` fail if it differs between the `stage1` and `stage2` compilations. The `-save-temps` option forces a fixed name to be used for the assembler input file, instead of a randomly chosen name in `/tmp`.

Do not add `'-save-temps'` unless the comparisons fail without that option. If you do you `'-save-temps'`, you will have to manually delete the `'i'` and `'s'` files after each series of compilations. The MIPS C compiler needs to be told to increase its table size for switch statements with the `-wf, -XNg1500` option in order to compile `cp/parse.c`. If you use the `-O2` optimization option, you also need to use `-Olimit 3000`. Both of these options are automatically generated in the `'Makefile'` that the shell script `configure` builds.

If you override the `CC` `make` variable and use the MIPS compilers, you may need to add `-wf, -XNg1500 -Olimit 3000`. On Irix version 4.0.5F, and perhaps on some other versions as well, there is an assembler bug that reorders instructions incorrectly. To work around it, specify `mips-sgi-irix4loser` as the target configuration. This configuration inhibits assembler optimization. In a compiler configured with target, `mips-sgi-irix4`, you can turn off assembler optimization by using the `'-noasmopt'` option.

This compiler option passes the option, `-O0`, to the assembler, to inhibit reordering. The `-noasmopt` option can be useful for testing whether a problem is due to erroneous assembler reordering.

Even if a problem does not go away with `-noasmopt`, it may still be due to assembler reordering—perhaps GCC itself was miscompiled as a result. To enable debugging under Irix 5, you must use GNU `as` 2.5 or later, and use the

`--with-gnu-as` configure option when configuring gcc. GNU `as` is distributed as part of the `binutils` package.

mips-sony-sysv

Sony MIPS NEWS. This works in NEWSOS 5.0.1, but not in 5.0.2 (which uses ELF instead of COFF). Support for 5.0.2 will probably be provided soon by volunteers. In particular, the linker does not like the code generated by GCC when shared libraries are linked in.

ns32k-encore

Encore NS32000 system. Encore systems are supported only under BSD.

ns32k-*-genix

National Semiconductor NS32000 system. Genix has bugs in `alloca` and `malloc`; you must get the compiled versions of these from GNU Emacs.

ns32k-sequent

Go to the Berkeley universe before compiling. In addition, you probably need to create a file named `'string.h'` containing just one line:

```
#include <strings.h>
```

ns32k-utek

UTEK NS32000 system (“merlin”). The C compiler that comes with this system cannot compile GCC; contact `'tektronix!reed!mason'` to get binaries of GCC for bootstrapping.

romp-*-aos

romp-*-mach

The only operating systems supported for the IBM RT PC are AOS and MACH. GCC does not support AIX running on the RT.

We recommend you compile GCC with an earlier version of itself; if you compile GCC with `hc`, the Metaware compiler, it will work, but you will get mismatches between the stage 2 and stage 3 compilers in various files. These errors are minor differences in some floating-point constants and can be safely ignored; the stage 3 compiler is correct.

rs6000-*-aix

powerpc-*-aix

Various early versions of each release of the IBM XLC compiler will not bootstrap GCC. Symptoms include differences between the stage2 and stage3 object files, and errors when compiling `'libgcc.a'` or `'enquire'`. Known problematic releases include: `xlc-1.2.1.8`, `xlc-1.3.0.0` (distributed with AIX 3.2.5), and `xlc-1.3.0.19`. Both `xlc-1.2.1.28` and `xlc-1.3.0.24` (PTF 432238) are known to produce working versions of GCC, but most other recent releases correctly bootstrap GCC. Also, releases of AIX prior to AIX 3.2.4 include a version of the IBM assembler which does not accept debugging directives: assembler updates are available as PTFs. Also, if you are using AIX 3.2.5 or greater and the GNU assembler, you must have a version modified after October 16, 1995 in order for

the GNU C compiler to build. See the file ‘README.RS6000’ for more details on of these problems.

GCC does not yet support the 64-bit PowerPC instructions.

Objective C does not work on this architecture because it makes assumptions that are incompatible with the calling conventions.

AIX on the RS/6000 provides support (NLS) for environments outside of the United States. Compilers and assemblers use NLS to support locale-specific representations of various objects including floating-point numbers (“.” vs “,” for separating decimal fractions). There have been problems reported where the library linked with GCC does not produce the same floating-point formats that the assembler accepts. If you have this problem, set the LANG environment variable to “C” or “En_US”.

Due to changes in the way that GCC invokes the binder (`linker`) for AIX 4.1, you may now receive warnings of duplicate symbols from the link step that were not reported before. The assembly files generated by GCC for AIX have always included multiple symbol definitions for certain global variable and function declarations in the original program. The warnings should not prevent the linker from producing a correct library or runnable executable. By default, AIX 4.1 produces code that can be used on either Power or PowerPC processors. You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option, `--with-cpu-cpu_type`.

`powerpc-*-elf`

`powerpc-*-sysv4`

PowerPC system in big endian mode, running System V.4. You can specify a default version for the `-mcpu=cpu_type` switch using the option,

`--with-cpu-cpu_type`.

`powerpc-*-linux-gnu`

PowerPC system in big endian mode, running Linux. You can specify a default version for the ‘`-mcpu=`’`cpu_type` switch using the option,

`--with-cpu-cpu_type`.

`powerpc-*-eabiaix`

Embedded PowerPC system in big endian mode with ‘`-mcall-aix`’ selected as the default. You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option, `--with-cpu-cpu_type`.

`powerpc-*-eabisim`

Embedded PowerPC system in big endian mode for use in running under the PSIM simulator. You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option, `--with-cpu-cpu_type`.

`powerpc-*-eabi`

Embedded PowerPC system in big endian mode. You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option,

`--with-cpu-cpu_type`.

powerpcle-*-elf

powerpcle-*-sysv4

PowerPC system in little endian mode, running System V.4. You can specify a default version for the `-mcpu=cpu_type` switch using the option,

`--with-cpu-cpu_type`.

powerpcle-*-solaris2*

PowerPC system in little endian mode for use in running Solaris 2.5.1 or higher.

You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option, `--with-cpu-cpu_type`. Beta versions of the Sun 4.0 compiler do not seem to be able to build GCC correctly. There are also problems with the host assembler and linker that are fixed by using the GNU versions of these tools.

powerpcle-*-eabisim

Embedded PowerPC system in little endian mode for use in running under the PSIM simulator.

powerpcle-*-eabi

Embedded PowerPC system in little endian mode. You can specify a default version for the `-mcpu=cpu_type` switch by using the option,

`--with-cpu-cpu_type`.

powerpcle-*-winnt

powerpcle-*-pe

PowerPC system in little endian mode running Windows NT. You can specify a default version for the `-mcpu=cpu_type` switch by using the configure option,

`--with-cpu-cpu_type`.

vax-dec-ultrix

Don't try compiling with Vax C (`vcc`). It produces incorrect code in some cases (for example, when `alloca` is used).

Meanwhile, compiling `cp/parse.c` with `pcc` does not work because of an internal table size limitation in that compiler. To avoid this problem, compile just the GNU C compiler first, and use it to recompile building all the languages that you want to run.

sparc-sun-*

sparc64-sun-*

See "Installing GCC on Sun" on page 326 for information on installing GCC on Sun systems.

vax-dec-vms

See "Installing GCC on VMS" on page 327 for details on how to install GCC on VMS.

we32k-*-*

These computers are also known as the 3b2, 3b5, 3b20 and other similar names. (However, the 3b1 is actually a 68000; see “Configurations Supported by GCC”) Don’t use ‘-g’ when compiling with the system’s compiler. The system’s linker seems to be unable to handle such a large program with debugging information. The system’s compiler runs out of capacity when compiling ‘stmt.c’ in GCC. You can work around this by building ‘cpp’ in GCC first, then use that instead of the system’s preprocessor with the system’s C compiler to compile ‘stmt.c’. Use the following.

```
mv /lib/cpp /lib/cpp.att
cp cpp /lib/cpp.gnu
echo ` /lib/cpp.gnu -traditional ${1+"$@"} ` > /lib/cpp
chmod +x /lib/cpp
```

The system’s compiler produces bad code for some of the GCC optimization files. So you must build the stage 2 compiler without optimization. Then build a stage 3 compiler with optimization. That executable should work. Use the following.

```
make LANGUAGES=c CC=stage1/xgcc CFLAGS="-Bstage1/ -g"
make stage2
make CC=stage2/xgcc CFLAGS="-Bstage2/ -g -O"
```

You may need to raise the `ULIMIT` setting to build a C++ compiler, as the file ‘cc1plus’ is larger than one megabyte.

Compilation in a Separate Directory

If you wish to build the object files and executables in a directory other than the one containing the source files, use the following.

1. Make sure you have a version of `make` that supports the `VPATH` feature. (GNU `make` supports it, as do `make` versions on most BSD systems.)
2. If you have ever run ‘`configure`’ in the source directory, you must undo the configuration. Do this by running `make distclean`.
3. Go to the directory in which you want to build the compiler before running `configure`.

```
mkdir gcc-sun3
cd gcc-sun3
```

On systems that do not support symbolic links, this directory must be on the same file system as the source code directory.

4. Specify where to find `configure` when you run it.

```
../gcc/configure...
```

This also tells `configure` where to find the compiler sources; `configure` takes the directory from the filename that was used to invoke it.

And, if you want to be sure, you can specify the source directory with the

'`--srcdir`' option like the following example demonstrates.

```
../gcc/configure--srcdir= ../gcc other options
```

The directory you specify with '`--srcdir`' need not be the same as the one in which `configure` is found.

Now, you can run `make` in that directory. You need not repeat the configuration steps shown previously when ordinary source files change. You must, however, run `configure` again when the configuration files change, if your system does not support symbolic links.

Building and Installing a Cross-compiler

GCC can function as a cross-compiler for many machines, but not all.

- Cross-compilers for the MIPS as target using the MIPS assembler currently do not work, because the auxiliary programs '`mips-tdump.c`' and '`mips-tfile.c`' can't be compiled on anything but a MIPS. It does work to cross compile for a MIPS if you use the GNU assembler and linker.
- Cross-compilers between machines with different floating point formats have not all been made to work. GCC now has a floating point emulator with which these can work, but each target machine description needs to be updated to take advantage of it.
- Cross-compilation between machines of different word sizes is some what problematic and sometimes does not work.

Since GCC generates assembler code, you probably need a cross-assembler that GCC can run, in order to produce object files. If you want to link on other than the target machine, you need a cross-linker as well. You also need header files and libraries suitable for the target machine that you can install on the host machine.

Steps of Cross-compilation

To compile and run a program using a cross-compiler involves several steps:

1. Run the cross-compiler on the host machine to produce assembler files for the target machine. This requires header files for the target machine.
2. Assemble the files produced by the cross-compiler. You can do this either with an assembler on the target machine, or with a cross-assembler on the host machine.

3. Link those files to make an executable. You can do this either with a linker on the target machine, or with a cross-linker on the host machine. Whichever machine you use, you need libraries and certain startup files (typically, `crt...o`) for the target machine.

It is most convenient to do all of these steps on the same host machine, since then you can do it all with a single invocation of GCC. This requires a suitable cross-assembler and cross-linker. For some targets, the GNU assembler and linker are available.

Configuring a Cross-compiler

To build GCC as a cross-compiler, you start out by running `configure`. Use the `--target=target` option to specify the target type. If `configure` was unable to correctly identify the system you are running on, specify the option, `'--build=build'`. For instance, the following example shows how to configure for a cross-compiler that produces code for an HP 68030 system running BSD on a system that `configure` can correctly identify.

```
./configure--target=m68k-hp-bsd4.3
```

Tools and Libraries for a Cross-compiler

If you have a cross-assembler and cross-linker available, you should install them now. Put them in the directory, `/usr/local/target/bin`. What follows are the tools you should put in this directory.

- `as`
This should be the cross-assembler.
- `ld`
This should be the cross-linker.
- `ar`
This should be the cross-archiver: a program which can manipulate archive files (linker libraries) in the target machine's format.
- `ranlib`
This should be a program to construct a symbol table in an archive file.

The installation of GCC will find these programs in that directory, and copy or link them to the proper place to for the cross-compiler to find them when run later.

The easiest way to provide these files is to build the `binutils` package and `gas`.

Configure them with the same `'--host'` and `'--target'` options that you use for configuring GCC, then build and install them. They install their executables automatically into the proper directory. Alas, they do not support all the targets that

GCC supports.

If you want to install libraries to use with the cross-compiler, such as a standard C library, put them in the directory `‘/usr/local/target/lib’`; installation of GCC copies all the files in that subdirectory into the proper place for GCC to find them and link with them.

What follows is an example of copying some libraries from a target machine.

```
ftp target-machine
lcd /usr/local/target/lib
cd /lib
get libc.a
cd /usr/lib
get libg.a
get libm.a
quit
```

The precise set of libraries you’ll need, and their locations on the target machine, vary depending on its operating system. Many targets require “*start files*” such as `‘crt0.o’` and `‘crtn.o’` which are linked into each executable. These too should be placed in `‘/usr/local/target/lib’`.

There may be several alternatives for `‘crt0.o’`, for use with profiling or other compilation options.

Check your target’s definition of `STARTFILE_SPEC` to find out what start files it uses. The following is an example of copying these files from a target machine.

```
ftp target-machine
lcd /usr/local/target/lib
prompt
cd /lib
mget *crt*.o
cd /usr/lib
mget *crt*.o
quit
```

libgcc.a and Cross-compilers

Code compiled by GCC uses certain runtime support functions implicitly. Some of these functions can be compiled successfully with GCC itself, but a few cannot be. These problem functions are in the source file, `‘libgcc1.c’`; the library made from them is called `‘libgcc1.a’`.

When you build a native compiler, these functions are compiled with some other compiler—the one that you use for bootstrapping GCC. Presumably it knows how to open code these operations, or else knows how to call the run-time emulation facilities that the machine comes with. But this approach doesn’t work for building a cross-compiler. The compiler that you use for building knows about the host system,

not the target system.

So, when you build a cross-compiler you have to supply a suitable library 'libgcc1.a' that does the job it is expected to do.

To compile 'libgcc1.c' with the cross-compiler itself does not work. The functions in this file are supposed to implement arithmetic operations that GCC does not know how to open code for your target machine. If these functions are compiled with GCC itself, they will compile into infinite recursion.

On any given target, most of these functions are not needed. If GCC can open code an arithmetic operation, it will not call these functions to perform the operation. It is possible that on your target machine, none of these functions is needed. If so, you can supply an empty library as 'libgcc1.a'.

Many targets need library support only for multiplication and division. If you are linking with a library that contains functions for multiplication and division, you can tell GCC to call them directly by defining the macros `MULSI3_LIBCALL`, and the like. These macros need to be defined in the target description macro file. For some targets, they are defined already. This may be sufficient to avoid the need for 'libgcc1.a'; if so, you can supply an empty library.

Some targets do not have floating point instructions; they need other functions in 'libgcc1.a', which do floating arithmetic. Recent versions of GCC have a file which emulates floating point. With a certain amount of work, you should be able to construct a floating point emulator that can be used as 'libgcc1.a'. Perhaps future versions will contain code to do this automatically and conveniently. That depends on whether someone wants to implement it.

Some embedded targets come with all the necessary 'libgcc1.a' routines written in C or assembler. These targets build 'libgcc1.a' automatically and you do not need to do anything special for them. Other embedded targets do not need any 'libgcc1.a' routines since all the necessary operations are supported by the hardware.

If your target system has another C compiler, you can configure GCC as a native compiler on that machine.

Build just 'libgcc1.a' with 'make libgcc1.a' on that machine, and use the resulting file with the cross-compiler. To do this, execute the following on the target machine.

```
cd target-build-dir
./configure--host=sparc--target=sun3
make libgcc1.a
```

And then, execute the following on the host machine:

```
ftp target-machine
binary
cd target-build-dir
get libgcc1.a
quit
```

Another way to provide the functions you need in `'libgcc1.a'` is to define the appropriate `perform_...` macros for those functions. If these definitions do not use the C arithmetic operators that they are meant to implement, you should be able to compile them with the cross-compiler you are building. (If these definitions already exist for your target file, then you are all set.) To build `'libgcc1.a'` using the `perform` macros, use `LIBGCC1=libgcc1.a OLDCC=./xgcc` when building the compiler. Otherwise, you should place your replacement library under the name `'libgcc1.a'` in the directory in which you will build the cross-compiler, before you run `make`.

Cross-compilers and Header Files

If you are cross-compiling a stand-alone program or a program for an embedded system, then you may not need any header files except the few that are part of GCC (and those of your program). However, if you intend to link your program with a standard C library such as `'libc.a'`, then you probably need to compile with the header files that go with the library you use.

The GNU C compiler does not come with these files, because (1) they are system-specific, and (2) they belong in a C library, not in a compiler.

If the GNU C library supports your target machine, then you can get the header files from there (assuming you actually use the GNU library when you link your program).

If your target machine comes with a C compiler, it probably comes with suitable header files also. If you make these files accessible from the host machine, the cross-compiler can use them also. Otherwise, you're on your own in finding header files to use when cross-compiling.

When you have found suitable header files, put them in `'/usr/local/target/include'`, before building the cross compiler. Then installation will run `fixincludes` properly and install the corrected versions of the header files where the compiler will use them. Provide the header files before you build the cross-compiler, because the build stage actually runs the cross-compiler to produce parts of `'libgcc.a'`. (These are the parts that can be compiled with GCC.) Some of them need suitable header files.

To copy header files from a target machine, use the following example's input.

```
(cd /usr/include; tar cf - .) > tarfile
```

Then, on the host machine, use the following example's input where *target-machine* represents your intended target machine.

```
ftp target-machine
lcd /usr/local/target/include
get tarfile
quit
tar xf tarfile
```

Standard Header File Directories

`GCC_INCLUDE_DIR` means the same thing for native and cross. It is where GCC stores its private include files, and also where GCC stores the fixed `include` files. A cross compiled GCC runs `fixincludes` on the header files in `'$(tooldir)/include'`. (If the cross compilation header files need to be fixed, they must be installed before GCC is built. If the cross compilation header files are already suitable for ANSI C and GCC, nothing special need be done).

`GPLUS_INCLUDE_DIR` means the same thing for native and cross. It is where `g++` looks first for header files. `libg++` installs only target independent header files in that directory.

`LOCAL_INCLUDE_DIR` is used only for a native compiler. It is normally `'/usr/local/include'`. GCC searches this directory so that users can install header files in `'/usr/local/include'`.

`CROSS_INCLUDE_DIR` is used only for a cross compiler. GCC doesn't install anything there.

`TOOL_INCLUDE_DIR` is used for both native and cross compilers. It is the place for other packages to install header files that GCC will use. For a cross-compiler, this is the equivalent of `'/usr/include'`. When you build a cross-compiler, `fixincludes` processes any header files in this directory.

```
(cd /usr/include; tar cf - .) > tarfile
```

Then, on the host machine, use the following example's input where *target-machine* represents your intended target machine.

```
ftp target-machine
lcd /usr/local/target/include
get tarfile
quit
tar xf tarfile
```

Actually Building the Cross-compiler

Now you can proceed just as for compiling a single-machine compiler through the step of building stage 1. If you have not provided some sort of `'libgcc1.a'`, then compilation will give up at the point where it needs that file, printing a suitable error message. If you do provide `'libgcc1.a'`, then building the compiler will automatically compile and link a test program called `'libgcc1-test'`; if you get errors in the linking, it means that not all of the necessary routines in `'libgcc1.a'` are available.

You must provide the header file `float.h`. One way to do this is to compile `enquire` and run it on your target machine. The job of `enquire` is to run on the target machine and figure out by experiment the nature of its floating point representation. `enquire` records its findings in the header file `float.h`. If you can't produce this file by running `enquire` on the target machine, then you will need to come up with a suitable `float.h` in some other way (or else, avoid using it in your programs).

Do not try to build stage 2 for a cross-compiler. It doesn't work to rebuild GCC as a cross-compiler using the cross-compiler, because that would produce a program that runs on the target machine, not on the host. For example, if you compile a 386 to 68030 cross-compiler with itself, the result will not be right either for the 386 (because it was compiled into 68030 code) or for the 68030 (because it was configured for a 386 as the host). If you want to compile GCC into 68030 code, whether you compile it on a 68030 or with a cross-compiler on a 386, you must specify a 68030 as the host when you configure it.

To install the cross-compiler, use `'make install'`, as usual.

collect2 and Cross-compiling

Many target systems do not have support in the assembler and linker for “constructors”—initialization functions to be called before the official “start” of `main`. On such systems, GCC uses a utility called `collect2` to arrange to call these functions at start time.

The program `collect2` works by linking the program once and looking through the linker output file for symbols with particular names indicating they are constructor functions. If it finds any, it creates a new temporary `.c` file containing a table of them, compiles it, and links the program a second time including that file.

The actual calls to the constructors are carried out by a subroutine called `__main`, which is called (automatically) at the beginning of the body of `main` (provided `main` was compiled with GCC).

Calling `__main` is necessary, even when compiling C code, to allow linking C and C++ object code together. (If you use `'-nostdlib'`, you get an unresolved reference to `__main`, since it's defined in the standard GCC library. Include `-lgcc` at the end of your compiler command line to resolve this reference.)

The program, `collect2`, is installed as `'ld'` in the directory where the passes of the compiler are installed. When `collect2` needs to find the *real* `'ld'`, it tries the following filenames.

- `real-ld` in the directories listed in the compiler's search directories.
- `real-ld` in the directories listed in the environment variable, `PATH`.
- The file specified in the `REAL_LD_FILE_NAME` configuration macro, if specified.

- `ld` in the compiler's search directories, except that `collect2` will not execute itself recursively.
- `ld` in `PATH`.

The “compiler's search directories” means all the directories where `gcc` searches for compiler passes, including directories that you specify with `'-B'`.

Cross-compilers search a little differently than normal configurations, using the following filenames.

- `real-ld` in the compiler's search directories.
- `target-real-ld` in `PATH`.
- The file specified in the `REAL_LD_FILE_NAME` configuration macro, if specified.
- `ld` in the compiler's search directories.
- `target-ld` in `PATH`.

`collect2` explicitly avoids running `ld` using the filename under which `collect2` itself was invoked. In fact, it remembers up a list of such names—in case one copy of `collect2` finds another copy (or version) of `collect2` installed as `'ld'` in a second place in the search path.

`collect2` searches for the utilities `nm` and `strip` using the same algorithm as previous installation for `'ld'`.

Installing GCC on Sun

On Solaris (version 2.1), do not use the linker or other tools in `'/usr/ucb'` to build GCC. Use `'/usr/ccs/bin'`.

Make sure the environment variable `FLOAT_OPTION` is not set when you compile `'libgcc.a'`. If this option were set to `f68881` when `'libgcc.a'` is compiled, the resulting code would demand to be linked with a special startup file and would not link properly without special pains.

There is a bug in `alloca` in certain versions of the Sun library. To avoid this bug, install the binaries of GCC that were compiled by GCC. They use `alloca` as a built-in function and never the one in the library.

Some versions of the Sun compiler crash when compiling GCC. The problem is a segmentation fault in `c++`. This problem seems to be due to the bulk of data in the environment variables. You may be able to avoid it by using the following command to compile GCC with Sun CC:

```
make CC="TERMCAP=x OBJS=x LIBFUNCS=x STAGESTUFF=x cc"
```

SunOS 4.1.3 and 4.1.3 U1 have bugs that can cause intermittent core dumps when compiling GCC. A common symptom is an internal compiler error which does not recur if you run it again. To fix the problem, install Sun recommended patch 100726 (for SunOS 4.1.3) or 101508 (for SunOS 4.1.3 U1), or upgrade to a later SunOS

release.

Installing GCC on VMS

The VMS version of GCC is distributed in a backup “save set” containing both source code and precompiled binaries. To install the `gcc` command so you can use the compiler easily, in the same manner as you use the VMS C compiler, you must install the VMS CLD file for GCC as follows.

1. Define the VMS logical names `GNU_CC` and `GNU_CC_INCLUDE` to point to the directories where the GCC executables (`'gcc-cpp.exe'`, `'gcc-cc1.exe'`, etc.) and the C include files are kept respectively. This should be done with the following commands.

```
assign /system /translation=concealed -disk:[gcc.] gnu_cc
assign /system /translation=concealed -
disk:[gcc.include.]
gnu_cc_include
```

Include the appropriate disk and directory names. These commands can be placed in your system startup file so they will be executed whenever the machine is rebooted. You may, if you choose, do this using the `GCC_INSTALL.COM` script in the `[GCC]` directory.

2. Install the `gcc` command with the following declaration.

```
set command /table=sys$common:[syslib]dcltables \
-/ output=sys$common:[syslib]dcltables \
gnu_cc:[000000]gcc
install replace sys$common:[syslib]dcltables
```

3. To install the help file, use the following declaration.

```
library/help sys$library:helplib.hlb gcc.hlp
```

Now, invoke the compiler with a command like `'gcc /verbose file.c'`, which is equivalent to the command `'gcc -v -c file.c'` in UNIX.

If you wish to use G++, you must first install GCC, and then perform the following steps.

1. Define the VMS logical name `GNU_GXX_INCLUDE` to point to the directory where the preprocessor will search for the C++ header files. This can be done with the following command.

```
assign /system /translation=concealed -disk:[ \
gcc.gxx_include.]
gnu_gxx_include
```

Include the appropriate disk and directory name. If you are going to be using `libg++`, this is where the `libg++` install procedure will install the `libg++` header files.

2. Obtain the file ‘gcc-cc1plus.exe’, and place this in the same directory that ‘gcc-cc1.exe’ is kept. The G++ compiler can be invoked with a command like ‘gcc /plus /verbose file.cc’, which is equivalent to the command ‘g++ -v -c file.cc’ in UNIX.

We try to put corresponding binaries and sources on the VMS distribution tape. But sometimes the binaries will be from an older version than the sources, because we don’t always have time to update them. (Use the ‘/version’ option to determine the version number of the binaries and compare it with the source file ‘version.c’ to tell whether this is so.) In this case, you should use the binaries you get to recompile the sources. If you must recompile, use the following steps.

1. Execute the command procedure, ‘vmsconfig.com’ to set up the files, ‘tm.h’, ‘config.h’, ‘aux-output.c’, and ‘md. ’, and to create files, ‘tconfig.h’ and ‘hconfig.h’. This procedure also creates several linker option files used by ‘make-cc1.com’ and a data file used by ‘make-l2.com’.

```
@vmsconfig.com
```

2. Setup the logical names and command tables as defined in Step 1. In addition, define the VMS logical name GNU_BISON to point to the directories where the Bison executable is kept. This should be done with the following command.

```
assign /system /translation=concealed -disk:[ bison.] \
                                                    gnu_bi
son
```

If you want, use the INSTALL_BISON.COM script in the [BISON] directory.

3. Install the BISON command with the following command.

```
set command /table=sys$common:[syslib]dcltables \
            -output=sys$common:[syslib]dcltables \
            -gnu_ bison:[000000]bison
install replace sys$common:[syslib]dcltables
```

4. Type ‘@make-gcc’ to recompile everything (alternatively, submit the file ‘make-gcc.com’ to a batch queue). If you wish to build the G++ compiler as well as the GCC compiler, you must first edit make-gcc.com and follow the instructions that appear in the comments.
5. In order to use GCC, you need a library of functions which GCC compiled code will call to perform certain tasks, and these functions are defined in the file ‘libgcc2.c’.

To compile this you should use the command procedure, ‘make-l2.com’, which will generate the library ‘libgcc2.olb’. Build ‘libgcc2.olb’ using the compiler built from the same distribution that ‘libgcc2.c’ came from, and make-gcc.com will automatically do all of this for you. To install the library, use the following commands.

```

library gnu_cc:[000000]gcclib/delete=(new,eprintf)
library gnu_cc:[000000]gcclib/delete=L_*
library libgcc2/extract=*/output=libgcc2.obj
library gnu_cc:[000000]gcclib libgcc2.obj

```

The first command simply removes old modules that will be replaced with modules from `libgcc2` under different module names. The modules `new` and `eprintf` may not actually be present in your `gcclib.olb`—if the VMS librarian complains about those modules not being present, simply ignore the message and continue on with the next command. The second command removes the modules that came from the previous version of the library `libgcc2.c`. Whenever you update the compiler on your system, you should also update the library with the previous procedure.

6. You may wish to build GCC in such a way that no files are written to the directory where the source files reside. An example would be the when the source files are on a read-only disk. In these cases, execute the following DCL commands (substituting your actual path names) where the `'dua1:[gcc.source_dir]'` directory contains the source code, and the `'dua0:[gcc.build_dir]'` directory is meant to contain all of the generated object files and executables.

```

assign dua0:[gcc.build_dir.]/translation=concealed, \
-dual:[ gcc.source_dir.]/translation=concealed
gcc_build
set default gcc_build:[000000]

```

Once you have done this, you can proceed building GCC as previously described. (Keep in mind that `'gcc_build'` is a rooted logical name, and thus the device names in each element of the search list must be an actual physical device name rather than another rooted logical name).

7. If you are building GCC with a previous version of GCC, you also should check to see that you have the newest version of the assembler. In particular, GCC version 2 treats global constant variables slightly differently from GCC version 1, and GAS version 1.38.1 does not have the patches required to work with GCC version 2. If you use GAS 1.38.1, then extern const variables will not have the read-only bit set, and the linker will generate warning messages about mismatched psect attributes for these variables. These warning messages are merely a nuisance, and can safely be ignored.

If you are compiling with a version of GCC older than 1.33, specify `'/DEFINE=("inline=")'` as an option in all the compilations. This requires editing all the `gcc` commands in `'make-cc1.com'`. (The older versions had problems supporting `inline`.) Once you have a working 1.33 or newer GCC, you can change this file back.

8. If you want to build GCC with the VAX C compiler, you will need to make minor changes in `'make-cccp.com'` and `'make-cc1.com'` to choose alternate definitions of `CC`, `CFLAGS`, and `LIBS`. See comments in those files. However, you must also

have a working version of the GNU assembler (GNU `as`, also known as GAS) as it is used as the back-end for GCC to produce binary object modules and is not included in the GCC sources. GAS is also needed to compile `'libgcc2'` in order to build `'gcclib'` (see Step 5); `'make-l2.com'` expects to be able to find it operational in `'gnu_cc:[000000]gnu-as.exe'`.

To use GCC on VMS, you need the VMS driver programs `'gcc.exe'`, `'gcc.com'`, and `'gcc.cld'`. They are distributed with the VMS binaries (`gcc-vms`) rather than the GCC sources. GAS is also included in `'gcc-vms'`, as is `bison`.

Once you have successfully built GCC with VAX C, you should use the resulting compiler to rebuild itself. Before doing this, be sure to restore the `CC`, `CFLAGS`, and `LIBS` definitions in `'make-ccc.com'` and `'make-cc1.com'`. The second generation compiler will be able to take advantage of many optimizations that must be suppressed when building with other compilers.

Under previous versions of GCC, the generated code would occasionally give strange results when linked with the sharable `'VAXCRTL'` library. Now this should work.

Even with this version, however, GCC itself should not be linked with the sharable `'VAXCRTL'`. The version of `qsort` in `'VAXCRTL'` has a bug (known to be present in VMS versions V4.6 through V5.5) which causes the compiler to fail.

The executables are generated by `'make-cc1.com'` and `'make-ccc.com'` uses the object library version of `'VAXCRTL'` in order to make use of the `qsort` routine in `'gcclib.olb'`.

If you wish to link the compiler executables with the shareable image version of `'VAXCRTL'`, you should edit the file `'tm.h'` (created by `'vmsconfig.com'`) to define the macro `QSORT_WORKAROUND`.

`QSORT_WORKAROUND` is always defined when GCC is compiled with VAX C, to avoid a problem in case `'gcclib.olb'` is not yet available.

Using GCC on VMS

See the following documentation for how to use GCC on VMS.

Include Files and VMS

Due to the differences between the file systems of UNIX and VMS, GCC attempts to translate filenames in `'#include'` into names that VMS will understand. The basic strategy is to prepend a prefix to the specification of the include file, convert the whole filename to a VMS filename, and then try to open the file. GCC tries various prefixes one by one until one of them succeeds:

- The first prefix is the `'GNU_CC_INCLUDE:'` logical name: this is where GNU C header files are traditionally stored. If you wish to store header files in non-standard locations, then you can assign the logical `'GNU_CC_INCLUDE'` to be a search list, where each element of the list is suitable for use with a rooted logical.

- The next prefix tried is ‘SYS\$SYSROOT:[SYSLIB.]’. This is where VAX-C header files are traditionally stored.
- If the include file specification by itself is a valid VMS filename, the preprocessor then uses this name with no prefix in an attempt to open the include file.
- If the file specification is not a valid VMS filename (i.e., the specification does not contain a device or a directory specifier, and contains a ‘/’ character), the preprocessor tries to convert it from UNIX syntax to VMS syntax. Conversion works like this: the first directory name becomes a device, and the rest of the directories are converted into VMS-format directory names. For example, the name ‘x11/foobar.h’ is translated to ‘x11:[00000]foobar.h’ or ‘x11:foobar.h’, whichever one can be opened. This strategy allows you to assign a logical name to point to the actual location of the header files.
- If none of these strategies succeeds, the ‘#include’ fails. Include directives of the following form.

```
#include foobar
```

Such directives are a common source of incompatibility between VAX-C and GCC.

VAX-C treats them much like a standard ‘#include <foobar.h>’ directive. That is incompatible with the ANSI C behavior implemented by GCC, to expand the name ‘foobar’ as a macro.

Macro expansion should eventually yield one of the two standard formats for ‘#include’:

```
#include "file"
#include <file>
```

If you have this problem, the best solution is to modify the source to convert the ‘#include’ directives to one of the two standard forms. That will work with either compiler. If you want a quick and dirty fix, define the filenames as macros with the proper expansion, like the following example.

```
#define stdio <stdio.h>
```

This will work, as long as the name doesn’t conflict with anything else in the program. Another source of incompatibility is that VAX-C assumes the following.

```
#include "foobar"
```

The program is actually asking for the file ‘foobar.h’. GCC does not make this assumption, and instead takes what you ask for literally; it tries to read the file ‘foobar’. The best way to avoid this problem is to always specify the desired file extension in your include directives. GCC for VMS is distributed with a set of include files that is sufficient to compile most general purpose programs. Even though the GCC distribution does not contain header files to define constants and structures for some VMS system-specific functions, there is no reason why you cannot use GCC with any of these functions. You first may have to generate or create header files, either by using the public domain utility UNSDL (which can be

found on a DECUS tape), or by extracting the relevant modules from one of the system macro libraries, and using an editor to construct a C header file. A ‘#include’ filename cannot contain a DECNET node name. The preprocessor reports an I/O error if you attempt to use a node name, whether explicitly, or implicitly via a logical name.

Global Declarations and VMS

GCC does not provide the `globalref`, `globaldef` and `globalvalue` keywords of VAX-C. You can get the same effect with an obscure feature of GAS, the GNU assembler. (This requires GAS version 1.39 or later.) The following macros allow you to use this feature in a fairly natural way.

```

#ifdef __GNUC__
#define
GLOBALREF(TYPE,NAME) \
    TYPE \
    NAME \
    asm ("_$$PsectAttributes_GLOBALSYMBOL$$" #NAME)
#define
GLOBALDEF(TYPE,NAME,VALUE) \
    TYPE \
    NAME \
    asm ("_$$PsectAttributes_GLOBALSYMBOL$$" \
        #NAME \
        = VALUE)
#define
GLOBALVALUEREf(TYPE,NAME) \
    const TYPE \
    NAME[1] \
    asm ("_$$PsectAttributes_GLOBALVALUE$$" #NAME)
#define
GLOBALVALUEDEF(TYPE,NAME,VALUE) \
    const TYPE \
    NAME[1] \
    asm ("_$$PsectAttributes_GLOBALVALUE$$" \
        #NAME) \
        = {VALUE}
#else
#define
GLOBALREF(TYPE,NAME) \
    globalref TYPE NAME
#define
GLOBALDEF(TYPE,NAME,VALUE) \
    globaldef TYPE NAME = VALUE
#define
GLOBALVALUEDEF(TYPE,NAME,VALUE) \
    globalvalue TYPE NAME = VALUE
#define
GLOBALVALUEREf(TYPE,NAME) \

```

```

    globalvalue TYPE NAME
#endif

```

(The ‘`_${PsectAttributes_GLOBALSYMBOL}`’ prefix at the start of the name is removed by the assembler, after it has modified the attributes of the symbol). These macros are provided in the VMS binaries distribution in a header file ‘`GNU_HACKS.H`’. An example of the usage is the following.

```

GLOBALREF (int, ijk);
GLOBALDEF (int, jkl, 0);

```

The macros, `GLOBALREF` and `GLOBALDEF`, cannot be used straightforwardly for arrays, since there is no way to insert the array dimension into the declaration at the right place. However, declare an array with these macros if you first define a typedef for the array type, like the following declaration.

```

typedef int intvector[10];
GLOBALREF (intvector, foo);

```

Array and structure initializers will also break the macros; you can define the initializer to be a macro of its own, or you can expand the `GLOBALDEF` macro by hand. You may find a case where you wish to use the `GLOBALDEF` macro with a large array, but you are not interested in explicitly initializing each element of the array. In such cases you can use an initializer like: ‘`{0,}`’, which will initialize the entire array to ‘0’.

A shortcoming of this implementation is that a variable declared with `GLOBALVALUEREDEF` or `GLOBALVALUEDEF` is always an array. For example, the following declaration gives the variable ‘`ijk`’ as an array of type ‘`int [1]`’.

```

GLOBALVALUEREDEF(int, ijk);

```

This is done because a `globalvalue` is actually a constant; its *value* is what the linker would normally consider an address. That is not how an integer value works in C, but it is how an array works. So treating the symbol as an array name gives consistent results—with the exception that the value seems to have the wrong type. ***Don’t try to access an element of the array.*** It doesn’t have any elements. The array *address* may not be the address of actual storage.

The fact that the symbol is an array may lead to warnings where the variable is used. Insert type casts to avoid the warnings. Here is an example; it takes advantage of the ANSI C feature allowing macros that expand to use the same name as the macro itself.

```

GLOBALVALUEREDEF (int, ss$_normal);
GLOBALVALUEDEF (int, xyzzy,123);
#ifdef __GNUC__
#define ss$_normal ((int) ss$_normal)
#define xyzzy ((int) xyzzy)
#endif

```

Don’t use `GLOBALDEF` or `GLOBALREF` with a variable whose type is an enumeration type; this is not implemented. Instead, make the variable an integer, and use a `GLOBALVALUEDEF` for each of the enumeration values. An example of this would be the

following declaration.

```
#ifdef __GNUC__
GLOBALDEF (int, color, 0);
GLOBALVALUEDEF (int, RED, 0);
GLOBALVALUEDEF (int, BLUE, 1);
GLOBALVALUEDEF (int, GREEN, 3);
#else
enum globaldef color {RED, BLUE, GREEN = 3};
#endif
```

Other VMS Issues

GCC automatically arranges for `main` to return 1 by default if you fail to specify an explicit return value. This will be interpreted by VMS as a status code indicating a normal successful completion. Version 1 of GCC did not provide this default. GCC on VMS works only with the GNU assembler, GAS. You need version 1.37 or later of GAS in order to produce value debugging information for the VMS debugger. Use the ordinary VMS linker with the object files produced by GAS. Under previous versions of GCC, the generated code would occasionally give strange results when linked to the sharable ‘`VAXCRTL`’ library. Now this should work.

A caveat for use of `const` global variables: the `const` modifier must be specified in every external declaration of the variable in all of the source files that use that variable. Otherwise the linker will issue warnings about conflicting attributes for the variable. Your program will still work despite the warnings, but the variable will be placed in writable storage.

Although the VMS linker does distinguish between upper and lower case letters in global symbols, most VMS compilers convert all such symbols into upper case and most run-time library routines also have upper case names. To be able to reliably call such routines, GCC (by means of the assembler, GAS) converts global symbols into upper case like other VMS compilers. However, since the usual practice in C is to distinguish case, GCC (using GAS) tries to preserve usual C behavior by augmenting each name that is not all lower case. This means truncating the name to at most 23 characters and then adding more characters at the end which encode the case pattern of those 23. Names which contain at least one dollar sign are an exception; they are converted directly into upper case without augmentation.

Name augmentation yields bad results for programs that use precompiled libraries (such as Xlib) which were generated by another compiler. You can use the compiler option ‘`/NOCASE_HACK`’ to inhibit augmentation; it makes external C functions and variables case-independent as is usual on VMS. Alternatively, you could write all references to the functions and variables in such libraries using lower case; this will work on VMS, but is not portable to other systems. The compiler option ‘`/NAMES`’ also provides control over global name handling.

Function and variable names are handled somewhat differently with G++. The G++

compiler performs *name mangling* on function names, which means that it adds information to the function name to describe the data types of the arguments that the function takes. One result of this is that the name of a function can become very long. Since the VMS linker only recognizes the first 31 characters in a name, special action is taken to ensure that each function and variable has a unique name that can be represented in 31 characters. If the name (plus a name augmentation, if required) is less than 32 characters in length, then no special action is performed. If the name is longer than 31 characters, the assembler (GAS) will generate a hash string based upon the function name, truncate the function name to 23 characters, and append the hash string to the truncated name. If the `'/VERBOSE'` compiler option is used, the assembler will print both the full and truncated names of each symbol that is truncated.

The `'/NOCASE_HACK'` compiler option should not be used when you are compiling programs that use `libg++`. `libg++` has several instances of objects (i.e., `Filebuf` and `filebuf`) which become indistinguishable in a case-insensitive environment. This leads to cases where you need to inhibit augmentation selectively (if you were using `libg++` and `Xlib` in the same program, for example). There is no special feature for doing this, but you can get the result by defining a macro for each mixed case symbol for which you wish to inhibit augmentation. The macro should expand into the lower case equivalent of itself, as in the following example (such macro definitions can be placed in a header file to minimize the number of changes to your source code).

```
#define StudlyCapS studlycaps
```




Known Problems with GCC

The following documentation describes known problems that affect users of GCC. Most of these are not GCC bugs; if they were, there would be fixes for them. Their results may be like those of a bug.

- “Actual Bugs Not Fixed Yet” on page 338
- “Installation Problems” on page 338
- “Cross-compiler Problems” on page 343
- “Interoperation” on page 344
- “Problems Compiling Certain Programs” on page 349
- “Incompatibilities of GCC” on page 350
- “Fixed Header Files” on page 353
- “Standard Libraries” on page 353
- “Disappointments and Misunderstandings” on page 354
- “Common Misunderstandings with GNU C++” on page 356
- “protoize and unprotoize Warnings” on page 357
- “Certain Changes GCC Will Not Use” on page 358
- “Warning Messages and Error Messages” on page 361

Some of these problems were due to bugs in other software, some are missing features

too problematic to add, and some are due to conflicts of opinion.

Actual Bugs Not Fixed Yet

The following documentation describes known problems that affect users of GCC. Most of these are not GCC bugs; if they were, there would be fixes for them. But the results may be like those of a bug.

Some of these problems are due to bugs in other software, some are missing features that are too much work to add, and some are places where people's opinions differ as to what is best.

- The `fixincludes` script interacts badly with automounters; if the directory of system header files is automounted, it tends to be unmounted while `fixincludes` is running. This would seem to be a bug in the automounter. We do not know any good way to work around it.
- The `fixproto` script will sometimes add prototypes for the `sigsetjmp` and `siglongjmp` functions that reference the `jmp_buf` type before that type is defined. To work around this, edit the offending file and place the `typedef` in front of the prototypes.
- There are several obscure cases of misusing `struct`, `union`, and `enum` tags that are not detected as errors by the compiler.
- When `-pedantic-errors` is specified, GNU C will incorrectly give an error message when a function name is specified in an expression involving the comma operator.
- Loop unrolling does not work properly for certain C++ programs. This is a bug in the C++ front end. It sometimes emits incorrect debug information, and the loop unrolling code is unable to recover from this error.

```
#!/bin/ksh
```

Installation Problems

The following documentation describes problems (and some apparent problems that do not really mean anything is wrong) showing up during installation of GCC.

- On certain systems, defining certain environment variables such as `CC` can interfere with the functioning of `make`.
- If you encounter seemingly strange errors when trying to build the compiler in a directory other than the source directory, it could be because you have previously configured the compiler in the source directory. Make sure you have done all the necessary preparations. See “Compilation in a Separate Directory” on page 318.

- If you build GCC on a BSD system using a directory stored in a System V file system, problems may occur in running `fixincludes` if the System V file system does not support symbolic links. These problems result in a failure to fix the declaration of `size_t` in `sys/types.h`. If you find that `size_t` is a signed type and that type mismatches occur, this could be the cause. The solution is to use a different directory for building GCC.
- In previous versions of GCC, the `gcc` driver program looked for `as` and `ld` in various places; for example, in files beginning with `/usr/local/lib/gcc-`. GCC version 2 looks for them in the following path, signifying machine for the *machine* and *version* that you need to specify.

```
/usr/local/lib/gcc-lib/target/version
```

To use a version of the assembler or linker that is not the default, for example, `gas` or `ld`, you must put them in that path (or make links to them from that path).
- Some commands executed when making the compiler may fail (return a non-zero status) and be ignored by `make`. These failures, which are often due to files that were not found, are expected, and can safely be ignored.
- It is normal to have warnings in compiling certain files about unreachable code and about enumeration type clashes. The names of these files have `insn-` beginnings. Also, `real.c` may get some warnings that you can ignore.
- Sometimes `make` recompiles parts of the compiler when installing the compiler. In one case, this was traced down to a bug in `make`. Either ignore the problem or switch to GNU `make`.
- If you have installed a program known as `purify`, you may find that it causes errors while linking `enquire`, which is part of building GCC. The fix is to get rid of the file `real-ld` which `purify` installs—so that GCC will not try to use it.
- On Linux SLS 1.01, there is a problem with `libc.a`; it does not contain the `obstack` functions. However, GCC assumes that the `obstack` functions are in `libc.a` when it is the GNU C library. To work around this problem, change the `__GNU_LIBRARY__` conditional (around line 31) to `#if 1`.
- On some 386 systems, building the compiler never finishes because `enquire` hangs due to a hardware problem in the motherboard—it reports floating point exceptions to the kernel incorrectly. You can install GCC except for `float.h` by patching out the command to run `enquire`. You may also be able to fix the problem for `real` by getting a replacement motherboard. This problem was observed in Revision E of the Micronics motherboard, and is fixed in Revision F. It has also been observed in the MYLEX MXA-33 motherboard. If you encounter this problem, you may also want to consider removing the FPU from the socket during the compilation. Alternatively, if you are running SCO UNIX, you can reboot and force the FPU to be ignored. To do this, type `hd(40)unix auto ignorefpu`.

- On some 386 systems, GCC crashes trying to compile the `enquire.c` class file. This happens on machines that do not have a 387 FPU chip. On 386 machines, the system kernel is supposed to emulate the 387 when you do not have one. The crash is due to a bug in the emulator.

One of these systems is the UNIX from Interactive Systems: 386/ix. On this system, an alternate emulator is provided, and it does work. To use it, execute the following command as super-user and then reboot the system. (The default emulator file remains present under the name, `emulator.dflt`.)

```
In /etc/emulator.rell /etc/emulator
```

Try using `/etc/emulator.att`, if you have such a problem on the SCO system.

Another system which has this problem is Esix, although it may not have a working alternate emulator.

On NetBSD 0.8, a similar problem manifests itself as the following error messages.

```
enquire.c: In function fprop:
enquire.c:2328: floating overflow
```

- On SCO systems, when compiling GCC with the system's compiler, do not use `-o`. Some versions of the system's compiler miscompile GCC with `-o`.
- Sometimes on a Sun 4 you may observe a crash in the program `genflags` or `genoutput` while building GCC. This is said to be due to a bug in `sh`. You can probably get around it by running `genflags` or `genoutput` manually and then retrying the `make`.
- On Solaris 2, executables of GCC version 2.0.2 are commonly available, but they have a bug that shows up when compiling current versions of GCC: undefined symbol errors occur during assembly if you use the `-g` option. The solution is to compile the current version of GCC without the `-g` option. That makes a working compiler which you can use to recompile with the `-g` option.
- Solaris 2 comes with a number of optional OS packages. Some of these packages are needed to use GCC fully. If you did not install all optional packages when installing Solaris, you will need to verify that the packages that GCC needs are installed.

To check whether an optional package is installed, use the `pkginfo` command. To add an optional package, use the `pkgadd` command. For further details, see the Solaris documentation. For Solaris 2.0 and 2.1, GCC needs six packages:

`SUNWarc`, `SUNWbtool`, `SUNWesu`, `SUNWhea`, `SUNWlibm`, and `SUNWtoo`. For Solaris 2.2, GCC needs an additional seventh package: `SUNWsprot`.

- On Solaris 2, trying to use the linker and other tools in `/usr/ucb` to install GCC has been observed to cause trouble. For example, the linker may hang indefinitely. The fix is to remove `/usr/ucb` from your `PATH`.

- If you use the 1.31 version of the MIPS assembler (such as was shipped with Ultrix 3.1), you need to use the `-fno-delayed-branch` switch when optimizing floating point code. Otherwise, the assembler will complain when the GCC compiler fills a branch delay slot with a floating point instruction, such as `add.d`.
- If, on a MIPS system, you get a “does not have gp sections for all its sections” error message, disregard it. This happens whenever you use the GNU assembler with the MIPS linker; there is not really anything wrong, and it is appropriate to use the output file. You can stop such warnings by installing the GNU linker. It would be nice to extend the assembler to produce the `gp` tables, but they are optional, and there should not be a warning about their absence.
- In Ultrix 4.0 on the MIPS machine, `stdio.h` does not work with GCC at all unless it has been fixed with `fixincludes`. This causes problems in building GCC. Once GCC is installed, the problems go away. To work around this problem, when making the stage 1 compiler, specify the following option to `make`.

```
GCC_FOR_TARGET="./xgcc -B./ -I./include"
```

When making stage 2 and stage 3, specify the following option.

```
CFLAGS="-g -I./include"
```

- Users have reported some problems with version 2.0 of the MIPS compiler tools that were shipped with Ultrix 4.1. Version 2.10 which came with Ultrix 4.2 seems to work fine.

Users have also reported some problems with version 2.20 of the MIPS compiler tools that were shipped with RISC/OS 4.x. The earlier version 2.11 seems to work fine.

- Some versions of the MIPS linker will issue an assertion failure when linking code that uses `alloca` against shared libraries on RISC-OS 5.0, and DEC’s OSF/1 systems.

This is a bug in the linker that is supposed to be fixed in future revisions. To protect against this, GCC passes `-non_shared` to the linker unless you pass an explicit `-shared` or `-call_shared` switch.

- On System V release 3, you may get the following error message while linking:

```
ld fatal: failed to write symbol name something
in strings table for file whatever
```

This probably indicates that the disk is full or your `ULIMIT` will not allow the file to be as large as it needs to be. This problem can also result because the kernel parameter `MAXUMEM` is too small. If so, regenerate the kernel and make the value much larger. The default value is reported to be 1024; a value of 32768 is said to work. Smaller values may also work.

- On System V, you may get an error like the following message:

```
/usr/local/lib/bison.simple: In function yyparse:
/usr/local/lib/bison.simple:625: virtual memory exhausted
```

This indicates a problem with disk space, `ULIMIT`, or `MAXUMEM`.

- Current GCC versions probably do not work on version 2 of the NeXT operating system.
- On NeXTStep 3.0, the Objective C compiler does not work, due, apparently, to a kernel bug that it happens to trigger. This problem does not happen on 3.1.
- On the Tower models 4n0 and 6n0, by default a process is not allowed to have more than one megabyte of memory. GCC cannot compile itself (or many other programs) with `-o` in that much memory.

To solve this problem, reconfigure the kernel adding the following line to the configuration file:

```
MAXUMEM = 4096
```

- On HP 9000 series 300 or 400 running HP/UX release 8.0, there is a bug in the assembler that must be fixed before GCC can be built. This bug manifests during the first stage of compilation, while building `libgcc2.a`:

```
_floatdisf
ccl: warning: -g option not supported on this of GCC
ccl: warning: -gl option not supported on this version of GCC
./xgcc: Internal compiler error: program as got fatal signal 11
```

`archive/cph/hpux-8.0-assembler`, a patched version of the assembler, is available by anonymous ftp from `altdorf.ai.mit.edu`. If you have HP software support, the patch can also be obtained directly from HP, as described in the following note:

```
This is the patched assembler, to patch SR#1653-010439,
where the assembler aborts on floating point constants.
```

```
The bug is not really in the assembler, but in the shared
library version of the function "cvtnum(3c)". The bug
on "cvtnum(3c)" is SR#4701-078451. Anyway, the attached
assembler uses the archive library version of "cvtnum(3c)" and
thus does not exhibit the bug.
```

This patch is also known as PHCO 4484.

- On HP-UX version 8.05, but not on 8.07 or more recent versions, the `fixproto` shell script triggers a bug in the system shell. If you encounter this problem, upgrade your operating system or use BASH (the GNU shell) to run `fixproto`.
- Some versions of the Pyramid C compiler are reported to be unable to compile GCC. You must use an older version of GCC for bootstrapping. One indication of this problem is if you get a crash when GCC compiles the function, `muldi3`, in file, `libgcc2.c`.

You may be able to succeed by getting GCC version 1, installing it, and using it to compile GCC version 2. The bug in the Pyramid C compiler does not seem to affect GCC version 1.

- On the Intel Paragon (an i860 machine), if you are using operating system version 1.0, you will get warnings or errors about redefinition of `va_arg` when you build GCC. If this happens, then you need to link most programs with the library, `iclib.a`. You must also modify `stdio.h` as follows.

```
#if defined(__i860__) && !defined(_VA_LIST)
#include <va_list.h>
```

Between the previous lines, insert the line, `#if __PGC__`. Then, after the following lines, insert the line, `#endif /* __PGC__ */`.

```
extern int vprintf(const char *, va_list );
extern int vsprintf(char *, const char *, va_list );
#endif
```

These problems do not exist in operating system version 1.1.

- On the Altos 3068, programs compiled with GCC will not work unless you fix a kernel bug. This happens using system versions V.2.2 1.0gT1 and V.2.2 1.0e and perhaps later versions as well. See the file, `README.ALTOS`.
- You will get several sorts of compilation and linking errors on the we32k if you do not follow the special instructions. See “Configurations That GCC Supports” on page 303.
- A bug in the HP/UX 8.05 (and earlier) shell will cause the `fixproto` program to report an error of the following form.

```
./fixproto: sh internal 1K buffer overflow
```

To fix this, change the first line of the `fixproto` script to look like the following declaration.

```
#!/bin/ksh
```

Cross-compiler Problems

You may run into problems with cross compilation on certain machines, for many reasons.

- Cross compilation can run into trouble for certain machines because some assemblers on some target machines require floating point numbers to be written as integer constants in certain contexts.

The compiler writes these integer constants by examining the floating point value as an integer and printing that integer, because this is simple to write and independent of the details of the floating point representation. But this does not work if the compiler is running on a different machine with an incompatible floating point format, or even a different byte-ordering. In addition, correct constant folding of floating point values requires representing them in the target machine’s format. (The C standard does not quite require this, but in practice it is the only way to win.) It is now possible to overcome these problems by defining macros such as `REAL_VALUE_TYPE`. But doing so is a substantial amount of work

- for each target machine. See “`libgcc.a` and Cross-compilers” on page 321 and “Actually Building the Cross-compiler” on page 324 for more on macros.
- At present, the program, `mips-tfile`, which adds debug support to object files on MIPS systems, does not work in a cross compile environment.

Interoperation

The following documentation discusses various difficulties in using GNU C or GNU C++ together with other compilers or with the assemblers, linkers, libraries and debuggers on certain systems.

- Objective C does not work on the RS/6000.
- GNU C++ does not do name mangling in the same way as other C++ compilers. This means that object files compiled with one compiler cannot be used with another compiler.

This effect is intentional, to protect you from more subtle problems. Compilers differ as to many internal details of C++ implementation, including: how class instances are laid out, how multiple inheritance is implemented, and how virtual function calls are handled.

If the name encoding were made the same, your programs would link against libraries provided from other compilers—but the programs would then crash when run. Incompatible libraries are then detected at link time, rather than at run time.

- Older GDB versions sometimes fail to read the output of GCC version 2. If you have trouble, get GDB version 4.4 or later.
- DBX rejects some files produced by GCC, though it accepts similar constructs in output from PCC. Until someone can supply a coherent description of what is valid DBX input and what is not, there is nothing we can do about these problems. You are on your own.
- The GNU assembler (GAS) does not support PIC. To generate PIC code, you must use some other assembler, such as `/bin/as`.
- On some BSD systems, including some versions of Ultrix, use of profiling causes static variable destructors (currently used only in C++) not to be run.
- Use of `-I/usr/include` may cause trouble.

Many systems come with header files that will not work with GCC unless corrected by `fixincludes`. The corrected header files go in a new directory; GCC searches this directory before `/usr/include`. If you use `-I/usr/include`, this tells GCC to search `/usr/include` earlier on, before the corrected headers. The result is that you get the uncorrected header files.

Instead, you should use the following options when compiling C programs (where *target* and *version* signify what you must specify for the particular machine

and its version).

```
-I/usr/local/lib/gcc-lib/target/version/include
-I/usr/include
```

For C++ programs, GCC also uses a special directory that defines C++ interfaces to standard C subroutines. This directory is meant to be searched *before* other standard include directories, so that it takes precedence. If you are compiling C++ programs and specifying include directories explicitly, use this option first, then the two previous options.

```
-I/usr/local/lib/g++-include
```

- On some SGI systems, when you use `-lg1_s` as an option, it gets translated magically to `-lg1_s -lx11_s -lc_s`. Naturally, this does not happen when you use GCC. You must specify all three options explicitly.
- On a SPARC, GCC aligns all values of type `double` on an 8-byte boundary, and it expects every `double` to be so aligned. The Sun compiler usually gives `double` values 8-byte alignment, with one exception: function arguments of type `double` may not be aligned.

As a result, if a function compiled with Sun CC takes the address of an argument of type `double` and passes this pointer of type `double*` to a function compiled with GCC, de-referencing the pointer may cause a fatal signal.

One way to solve this problem is to compile your entire program with GCC.

Another solution is to modify the function that is compiled with Sun CC to copy the argument into a local variable; local variables are always properly aligned.

A third solution is to modify the function that uses the pointer to de-reference it using the function, `access_double`, instead of directly with `*` as in the following declaration.

```
inline double
access_double (double *unaligned_ptr)
{
    union d2i { double d; int i[2]; };

    union d2i *p = (union d2i *) unaligned_ptr;
    union d2i u;

    u.i[0] = p->i[0];
    u.i[1] = p->i[1];

    return u.d;
}
```

Storing into the pointer can be done likewise with the same union.

- On Solaris, the `malloc` function in the `libmalloc.a` library may allocate memory that is only 4 byte aligned. Since GCC on the SPARC assumes that doubles are 8 byte aligned, this may result in a fatal signal if doubles are stored in memory allocated by the `libmalloc.a` library.

The solution is not to use the `libmalloc.a` library. Instead, use `malloc` and related functions from `libc.a`; they do not have this problem.

- Sun forgot to include a static version of `libdl.a` with some versions of SunOS (mainly 4.1). This results in undefined symbols when linking static binaries (that is, if you use `-static`). If you see undefined symbols `_dlclose`, `_dlsym` or `_dlopen` when linking, compile and link against the file, `mit/util/misc/dlsym.c`, from the MIT version of X windows.
- The 128-bit `long double` format that the Sparc port supports currently works by using the architecturally defined quad-word floating point instructions. Since there is no hardware that supports these instructions they must be emulated by the operating system.

Long doubles do not work in Sun OS versions 4.0.3 and earlier, because the kernel emulator uses an obsolete and incompatible format. Long doubles do not work in Sun OS version 4.1.1 due to a problem in a Sun library. Long doubles do work on Sun OS versions 4.1.2 and higher, but GCC does not enable them by default. Long doubles appear to work in Sun OS 5.x (Solaris 2.x).

- On HP-UX version 9.01 on the HPPA, the HP compiler, `cc`, does not compile GCC correctly. We do not yet know why. However, GCC compiled on earlier HP-UX versions works properly on HP-UX 9.01 and can compile itself properly on 9.01.
- On the HPPA machine, ADB sometimes fails to work on functions compiled with GCC. Specifically, it fails to work on functions that use `alloca` or variable-size arrays. This is because GCC does not generate HP-UX unwind descriptors for such functions. It may even be impossible to generate them.
- Debugging `-g` is not supported on the HPPA machine, unless you use the preliminary GNU tools (see “Installing GCC” on page 293 for descriptions of the `--with-gnu-as` and `--with-gnu-ld`).
- Taking the address of a label may generate errors from the HP-UX PA assembler. GAS for the PA does not have this problem.
- Using floating point parameters for indirect calls to static functions will not work when using the HP assembler. There simply is no way for GCC to specify what registers hold arguments for static functions when using the HP assembler. GAS for the PA does not have this problem.
- In extremely rare cases involving some very large functions you may receive errors from the HP linker complaining about an out of bounds unconditional branch offset. This used to occur more often in previous versions of GCC, but is now exceptionally rare. If you should run into it, you can work around by making your function smaller.
- GCC compiled code sometimes emits warnings from the HP/UX assembler of the following form; such warnings can be safely ignored.

(warning) Use of GR3 when frame >= 8192 may cause conflict.

- The current version of the assembler (`/bin/as`) for the RS/6000 has certain problems that prevent the `-g` option in GCC from working. Note that `Makefile.in` uses `-g` by default when compiling `libgcc2.c`.

IBM has produced a fixed version of the assembler. The upgraded assembler unfortunately was not included in any of the AIX 3.2 update PTF releases (3.2.2, 3.2.3, or 3.2.3e). Users of AIX 3.1 should request PTF U403044 from IBM and users of AIX 3.2 should request PTF U416277. See the file `README.RS6000` for more details on these updates.

Test for the presence of a fixed assembler using the following command.

```
as -u < /dev/null
```

If the command exits normally, the assembler fix already is installed. If the assembler complains that `-u` is an unknown flag, you need to order the fix.

- On the IBM RS/6000, compiling code of the following form will cause the linker to report an undefined symbol, `foo`.

```
extern int foo;

... foo ...

static int foo;
```

Although this behavior differs from most other systems, it is not a bug because redefining an `extern` variable as `static` is undefined in ANSI C.

- AIX on the RS/6000 provides support (NLS) for environments outside of the United States. Compilers and assemblers use NLS to support locale-specific representations of various objects including floating-point numbers (`.` vs `,` for separating decimal fractions). There have been problems reported where the library linked with GCC does not produce the same floating-point formats that the assembler accepts. If you have this problem, set the `LANG` environment variable to `C` or `En_US`.
- Even if you specify `-fdollars-in-identifiers`, you cannot successfully use `$` in identifiers on the RS/6000 due to a restriction in the IBM assembler. GAS supports these identifiers.
- On the RS/6000, XLC version 1.3.0.0 will miscompile `jump.c`. XLC version 1.3.0.1 or later fixes this problem. You can obtain XLC-1.3.0.2 by requesting PTF 421749 from IBM.
- There is an assembler bug in versions of DG/UX prior to 5.4.2.01 that occurs when the `fldcr` instruction is used. GCC uses `fldcr` on the 88100 to serialize volatile memory references. Use the option, `-mno-serialize-volatile`, if your version of the assembler has this bug.
- On VMS, GAS versions 1.38.1 and earlier may cause spurious warning messages from the linker. These warning messages complain of mismatched `psect` attributes. You can ignore them. See “Installing GCC on VMS” on page 327.

- On NewsOS version 3, if you include both of the files, `stddef.h`, and `sys/types.h`, you get an error because there are two typedefs of `size_t`. You should change `sys/types.h` by adding the following lines around the definition of `size_t`.

```
#ifndef _SIZE_T
#define _SIZE_T
    /* actual typedef here */
#endif
```

- On the Alliant, the system's own convention for returning structures and unions is unusual, and is not compatible with GCC no matter what options are used.
- On the IBM RT PC, the MetaWare HighC compiler (`hc`) uses a different convention for structure and union returning.

Use the option, `-mhc-struct-return`, to tell GCC to use a convention compatible with it.

- On Ultrix, the Fortran compiler expects registers 2 through 5 to be saved by function calls. However, the C compiler uses conventions compatible with BSD UNIX: registers 2 through 5 may be clobbered by function calls. GCC uses the same convention as the Ultrix C compiler. Use these options to produce code compatible with the Fortran compiler.

```
-fcall-saved-r2      -fcall-saved-r3
-fcall-saved-r4      -fcall-saved-r5
```

- On the WE32k, you may find that programs compiled with GCC do not work with the standard shared C library. You may need to link with the ordinary C compiler. If you do so, you must specify the following options:

```
-L/usr/local/lib/gcc-lib/we32k-att-sysv/2.7.1 -lgcc -lc_s
```

The first specifies where to find the library, `libgcc.a`, specified with the `-lgcc` option. GCC does linking by invoking `ld`, just as `cc` does, and there is no reason why it should matter which compilation program invokes `ld`. If someone tracks this problem down, it can probably be fixed easily.

- On the Alpha, you may get assembler errors about invalid syntax as a result of floating point constants. This is due to a bug in the C library functions, `ecvt`, `fcvt` and `gcvt`. Given valid floating point numbers, they sometimes print `NaN`.
- On Irix 4.0.5F (and perhaps in some other versions), an assembler bug sometimes reorders instructions incorrectly when optimization is turned on. If you think this may be happening to you, try using the GNU assembler; GAS version 2.1 supports ECOFF on Irix.

Or use the `-noasmopt` option when you compile GCC with itself, and then again when you compile your program. (This is a temporary kludge to turn off assembler optimization on Irix.) If this proves to be what you need, edit the assembler spec in the file, `specs` so that it unconditionally passes `-O0` to the assembler, and never passes `-O2` or `-O3`.

Problems Compiling Certain Programs

Certain programs have problems compiling.

- Parse errors may occur compiling X11 on a Decstation running Ultrix 4.2 because of problems in DEC's versions of the X11 header files `x11/xlib.h` and `x11/xutil.h`. People recommend adding `-I/usr/include/mit` to use the MIT versions of the header files, using the `-traditional` switch to turn off ANSI C, or fixing the header files by adding the following:

```
#ifdef __STDC__
#define NeedFunctionPrototypes 0
#endif
```

- If you have trouble compiling Perl on a SunOS 4 system, it may be because Perl specifies `-I/usr/ucbinclude`. This accesses the unfixed header files. Perl specifies the following options most of which are unnecessary with GCC 2.4.5 and newer versions.

```
-traditional -Dvolatile=__volatile__
-I/usr/include/sun -I/usr/ucbinclude
-fpcc-struct-return
```

You can make a properly working Perl by setting `ccflags` to `-fwritable-strings` (implied by the option, `-traditional`, in the original options) and `cppflags` to empty in `config.sh`, then using the following declaration.

```
./doSH; make depend; make
```

- On various 386 UNIX systems derived from System V, including SCO, ISC, and ESIX, you may get error messages about running out of virtual memory while compiling certain programs.

You can prevent this problem by linking GCC with the GNU `malloc` (which thus replaces the `malloc` that comes with the system). GNU `malloc` is available as a separate package, and also in the file, `src/gmalloc.c` in the GNU Emacs distribution.

If you have installed GNU `malloc` as a separate library package, use the following option when you relink GCC.

```
MALLOCC=/usr/local/lib/libgmalloc.a
```

Alternatively, if you have compiled `gmalloc.c` from Emacs, copy the object file to `gmalloc.o` and use the following option when you relink GCC.

```
MALLOCC=gmalloc.o
```

Incompatibilities of GCC

There are several noteworthy incompatibilities between GNU C and most existing (non-ANSI) versions of C. The `-traditional` option eliminates many of these incompatibilities, *but not all*, by telling GNU C to behave like the other C compilers.

- GCC normally makes string constants read-only. If several identical-looking string constants are used, GCC stores only one copy of the string.

One consequence is that you cannot call `mktemp` with a string constant argument. The function, `mktemp`, always alters the string its argument points to. Another consequence is that `sscanf` does not work on some systems when passed a string constant as its format control string or input. This is because `sscanf` incorrectly tries to write into the string constant. Likewise `fscanf` and `scanf`.

The best solution to these problems is to change the program to use `char`-array variables with initialization strings for these purposes instead of string constants. But if this is not possible, you can use the `-fwritable-strings` flag, which directs GCC to handle string constants the same way most C compilers do. `-traditional` also has this effect, among others.

- `-2147483648` is positive.

This is because `2147483648` cannot fit in the type, `int`; so (following the ANSI C rules), its data type is `unsigned long int`. Negating this value yields `2147483648` again. GCC does not substitute macro arguments when they appear within string constants. For example, the following macro in GCC will produce output `"a"` regardless of what the `a` argument is.

```
#define foo(a) "a"
```

The `-traditional` option directs GCC to handle such cases (among others) in the old-fashioned (non-ANSI) fashion.

- When you use `setjmp` and `longjmp`, the only automatic variables guaranteed to remain valid are those declared `volatile`. This is a consequence of automatic register allocation. Consider the following function statement. In the statement, `a` may or may not be restored to its first value when the `longjmp` occurs.

```
jmp_buf j;

foo ()
{
    int a, b;

    a = fun1 ();

    if (setjmp (j))
        return a;

    a = fun2 ();
```

```

    /* longjmp (j) may occur in fun3. */
    return a + fun3 ();
}

```

If `a` is allocated in a register, then its first value is restored; otherwise, it keeps the last value stored in it.

If you use the `-w` option with the `-O` option, you will get a warning when GCC thinks such a problem might be possible.

The `-traditional` option directs GNU C to put variables in the stack by default, rather than in registers, in functions that call `setjmp`. This results in the behavior found in traditional C compilers.

- Programs that use preprocessing directives in the middle of macro arguments do not work with GCC. For example, a program like the following will not work:

```

foobar (
#define luser
        hack)

```

ANSI C does not permit such a construct. It would make sense to support it when `-traditional` is used, but it is too much work to implement.

- Declarations of external variables and functions within a block apply only to the block containing the declaration. In other words, they have the same scope as any other declaration in the same place.

In some other C compilers, an `extern` declaration affects all the rest of the file even if it happens within a block. The `-traditional` option directs GNU C to treat all `extern` declarations as global, like traditional compilers.

- In traditional C, you can combine `long`, etc., with a typedef name, as shown by the following declaration.

```

typedef int foo;
typedef long foo bar;

```

In ANSI C, this is not allowed: `long` and other type modifiers require an explicit `int`. Because this criterion is expressed by Bison grammar rules rather than C code, the `-traditional` flag cannot alter it.

- PCC allows typedef names to be used as function parameters. The difficulty described previously applies here too.
- PCC allows whitespace in the middle of compound assignment operators such as `+=`. GCC, following the ANSI standard, does not allow this. The difficulty described previously applies here too.
- GCC complains about unterminated character constants inside of preprocessing conditionals that fail. Some programs have English comments enclosed in conditionals that are guaranteed to fail; if these comments contain apostrophes, GCC will probably report an error. For example, the following such code would produce an error.

```

#if 0
...

```

```
#endif
```

The best solution to such a problem is to put the text into an actual C comment delimited by `/*...*/`. However, `-traditional` suppresses these error messages.

- Many user programs contain the declaration `long time()`; In the past, the system header files on many systems did not actually declare `time`, so it did not matter what type your program declared it to return. But in systems with ANSI C headers, `time` is declared to return `time_t`, and if that is not the same as `long`, then `long time()` is erroneous. The solution is to change your program to use `time_t` as the return type of `time`.
- When compiling functions that return `float`, PCC converts it to a `double`. GCC actually returns a `float`. If you are concerned with PCC compatibility, you should declare your functions to return `double`; you might as well say what you mean.
- When compiling functions that return structures or unions, GCC output code normally uses a method different from that used on most versions of UNIX. As a result, code compiled with GCC cannot call a structure-returning function compiled with PCC, and vice versa.

The method used by GCC is as follows: a structure or union which is 1, 2, 4 or 8 bytes long is returned like a scalar. A structure or union with any other size is stored into an address supplied by the caller (usually in a special, fixed register, but on some machines it is passed on the stack). The machine-description macros, `STRUCT_VALUE` and `STRUCT_INCOMING_VALUE`, tell GCC where to pass this address.

By contrast, PCC on most target machines returns structures and unions of any size by copying the data into an area of static storage, and then returning the address of that storage as if it were a pointer value. The caller must copy the data from that memory area to the place where the value is wanted. GCC does not use this method because it is slower and non re-entrant.

On some newer machines, PCC uses a reentrant convention for all structure and union returning. GCC on most of these machines uses a compatible convention when returning structures and unions in memory, but still returns small structures and unions in registers.

You can tell GCC to use a compatible convention for all structure and union returning with the option `-fpcc-struct-return`.

- GNU C complains about program fragments such as `0x74ae-0x4000` which appear to be two hexadecimal constants separated by the minus operator. Actually, this string is a single preprocessing token. Each such token must correspond to one token in C. Since this does not, GNU C prints an error message. Although it may appear obvious that what is meant, is an operator and two values, the ANSI C standard specifically requires that this be treated as erroneous.

A *preprocessing token* is a *preprocessing number* if it begins with a digit and is followed by letters, underscores, digits, periods and `e+`, `e-`, `E+`, or `E-` character

sequences. To make the previous program fragment valid, place whitespace in front of the minus sign. This whitespace will end the preprocessing number.

Fixed Header Files

GCC needs to install corrected versions of some system header files. This is because most target systems have some header files that will not work with GCC unless they are changed. Some have bugs, some are incompatible with ANSI C, and some depend on special features of other compilers.

Installing GCC automatically creates and installs the fixed header files, by running a program called `fixincludes` (or for certain targets an alternative such as `fixinc.svr4`). Normally, you do not need to pay attention to this. But there are cases where it does not do the right thing automatically.

- If you update the system's header files, such as by installing a new system version, the fixed header files of GCC are not automatically updated. The easiest way to update them is to reinstall GCC. (If you want to be clever, look in the makefile and you can find a shortcut.)
- On some systems, in particular SunOS 4, header file directories contain machine-specific symbolic links in certain places. This makes it possible to share most of the header files among hosts running the same version of SunOS 4 on different machine models.

The programs that fix the header files do not understand this special way of using symbolic links; therefore, the directory of fixed header files is good only for the machine model used to build it.

In SunOS 4, only programs that look inside the kernel will notice the difference between machine models. Therefore, for most purposes, you need not be concerned about this.

It is possible to make separate sets of fixed header files for the different machine models, and arrange a structure of symbolic links so as to use the proper set, but you will have to do this by hand.

- On Lynxos, GCC by default does not fix the header files. This is because bugs in the shell cause the `fixincludes` script to fail.

This means you will encounter problems due to bugs in the system header files. It may be no comfort that they are not GCC's fault, but it does mean that there is nothing to do about them.

Standard Libraries

GCC by itself attempts to be what the ISO/ANSI C standard calls a *conforming*

freestanding implementation. This means all ANSI C language features are available, as well as the contents of `float.h`, `limits.h`, `stdarg.h`, and `stddef.h`. The rest of the C library is supplied by the vendor of the operating system. If that C library does not conform to the C standards, then your programs might get warnings (especially when using `-Wall`) that you do not expect.

For example, the `sprintf` function on SunOS 4.1.3 returns `char *` while the C standard says that `sprintf` returns an `int`. The `fixincludes` program could make the prototype for this function match the Standard, but that would be wrong, since the function will still return `char *`.

If you need a Standard compliant library, then you need to find one, as GCC does not provide one. The GNU C library (called `glibc`) has been ported to a number of operating systems, and provides ANSI/ISO, POSIX, BSD and SystemV compatibility. You could also ask your operating system vendor if newer libraries are available.

Disappointments and Misunderstandings

These problems are regrettable, without any practical way to avoid them.

- Certain local variables are not recognized by debuggers when you compile with optimization.

This occurs because sometimes GCC optimizes the variable out of existence. There is no way to tell the debugger how to compute the value such a variable would have had and it is not clear that would be desirable anyway. So GCC simply does not mention the eliminated variable when it writes debugging information.

You have to expect a certain amount of disagreement between the executable and your source code, when you use optimization.

- Users often think it is a bug when GCC reports an error for code like the following example.

```
int foo (struct mumble *);

struct mumble { ... };

int foo (struct mumble *x)
{ ... }
```

This code really is erroneous, because the scope of `struct mumble` in the prototype is limited to the argument list containing it. (It does not refer to the `struct mumble` defined with file scope in the following descriptions—they are two unrelated types with similar names in different scopes.)

In the definition of `foo`, the file-scope type is used because that is available to be

inherited. Thus, the definition and the prototype do not match, and you get an error.

This behavior may seem silly, but it is what the ANSI standard specifies. It is easy enough for you to make your code work by moving the definition of `struct mumble` above the prototype. It is not worth being incompatible with ANSI C just to avoid an error like the previous example .

- Accesses to bitfields even in volatile objects works by accessing larger objects, such as a byte or a word. You cannot rely on what size of object is accessed in order to read or write the bitfield; it may even vary for a given bitfield according to the precise usage.

If you care about controlling the amount of memory that is accessed, use volatile but do not use bitfields.

- GCC comes with shell scripts to fix certain known problems in system header files. They install corrected copies of various header files in a special directory where only GCC will normally look for them. The scripts adapt to various systems by searching all the system header files for the problem cases that we know about.

If new system header files are installed, nothing automatically arranges to update the corrected header files. You will have to reinstall GCC to fix the new header files. More specifically, go to the build directory and delete the files `stamp-fixinc` and `stamp-headers`, and the subdirectory `include`; then do `make install` again.

- On 68000 systems and x86 systems, you can get paradoxical results if you test the precise values of floating point numbers. For example, you can find that a floating point value, which is not a NaN, is not equal to itself. This results from the fact that the floating point registers hold a few more bits of precision than fit in a `double` in memory. Compiled code moves values between memory and floating point registers at its convenience, and moving them into memory truncates them.

You can partially avoid this problem by using the `-ffloat-store` option (see “Options Controlling Optimization” on page 53).

- On the MIPS, variable argument functions using `varargs.h` cannot have a floating point value for the first argument. The reason for this is that in the absence of a prototype in scope, if the first argument is a floating point, it is passed in a floating point register, rather than an integer register.

If the code is rewritten to use the ANSI standard `stdarg.h` method of variable arguments, and the prototype is in scope at the time of the call, everything will work fine.

- On the H8/300 and H8/300H, variable argument functions must be implemented using the ANSI standard `stdarg.h` method of variable arguments. Furthermore, calls to functions using `stdarg.h` variable arguments must have a prototype for the called function in scope at the time of the call.

Common Misunderstandings with GNU C++

C++ is a complex language and an evolving one, and its standard definition (the ANSI C++ draft standard) is also evolving. As a result, your C++ compiler may occasionally surprise you, even when its behavior is correct. The following documentation discusses some areas that frequently give rise to questions of this nature.

Declare and Define Static Members

When a class has static data members, it is not enough to declare the static member; you must also define it. Use the following example, for instance, for such declarations.

```
class Foo
{
    : : :
    void method();
    static int bar;
};
```

This declaration only establishes that the class `Foo` has an `int` named `Foo::bar`, and a member function named `Foo::method`. But you still need to define both `method` and `bar` elsewhere. According to the draft ANSI standard, you must supply an initializer in one (and only one) source file, such as the following example shows.

```
int Foo::bar = 0;
```

Other C++ compilers may not correctly implement the standard behavior. As a result, when you switch to `g++` from one of these compilers, you may discover that a program that appeared to work correctly in fact does not conform to the standard: `g++` reports as undefined symbols any static data members that lack definitions.

Temporaries May Vanish

It is dangerous to use pointers or references to portions of a temporary object.

The compiler may very well delete the object before you expect it to, leaving a pointer to garbage. The most common place where this problem crops up is in classes like the `libg++ string` class, that define a conversion function to type, `char *`, or to type, `const char *`.

However, any class that returns a pointer to some internal structure is potentially subject to this problem. For instance, a program may use a function, `strfunc`, that returns `String` objects, and another function, `charfunc`, that operates on pointers to `char`, as in the following example's declaration.

```
String strfunc ();
void charfunc (const char *);
```

In such a situation, it may seem natural to write `charfunc (strfunc());`

based on the knowledge that class `String` has an explicit conversion to `char` pointers. However, what really happens is akin to `charfunc (strfunc().convert());` where the conversion method uses a function to do the same data conversion normally performed by a cast. Since the last use of the temporary `String` object is the call to the conversion function, the compiler may delete that object before actually calling `charfunc`. The compiler has no way of knowing that deleting the `String` object will invalidate the pointer. The pointer then points to garbage, so that by the time `charfunc` is called, it gets an invalid argument. Code like this may run successfully under some other compilers, especially those that delete temporaries relatively late. However, the GNU C++ behavior is also standard-conforming, so if your program depends on late destruction of temporaries it is not portable.

If you think this is surprising, you should be aware that the ANSI C++ committee continues to debate the lifetime-of-temporaries problem.

For now, at least, for instance, the following declaration defines the safe way to write such code by giving the temporary a name, forcing it to remain until the end of the scope of the name.

```
String& tmp = strfunc ();
charfunc (tmp);
```

protoize and unprotoize Warnings

The conversion programs, `protoize` and `unprotoize`, can sometimes change a source file in a way that will not work unless you rearrange it.

- `protoize` can insert references to a type name or type tag before the definition, or in a file where they are not defined.

If this happens, compiler error messages should show you where the new references are, so fixing the file by hand is straightforward.

- There are some C constructs which `protoize` cannot figure out. For example, it can not determine argument types for declaring a pointer-to- function variable; this you must do by hand. `protoize` inserts a comment containing ??? each time it finds such a variable; so you can find all such variables by searching for this string. ANSI C does not require declaring the argument types of pointer-to-function types.
- Using `unprotoize` can easily introduce bugs. If the program relied on prototypes to bring about conversion of arguments, these conversions will not take place in the program without prototypes. One case in which you can be sure `unprotoize` is safe is when you are removing prototypes that were made with `protoize`; if the program worked before without any prototypes, it will work again without them. You can find all the places where this problem might occur by compiling the

program with the `-wconversion` option. It prints a warning whenever an argument is converted.

- Both conversion programs can be confused if there are macro calls in and around the text to be converted. In other words, the standard syntax for a declaration or definition must not result from expanding a macro. This problem is inherent in the design of C and cannot be fixed. If only a few functions have confusing macro calls, you can easily convert them manually.
- `protoize` cannot get the argument types for a function whose definition was not actually compiled due to preprocessing conditionals. When this happens, `protoize` changes nothing in regard to such a function. `protoize` tries to detect such instances and warn about them.

You can generally work around this problem by using `protoize` step by step, each time specifying a different set of `-D` options for compilation, until all of the functions have been converted. There is no automatic way to verify that you have got them all, however.

- Confusion may result if there is an occasion to convert a function declaration or definition in a region of source code where there is more than one formal parameter list present. Thus, attempts to convert code containing multiple (conditionally compiled) versions of a single function header (in the same vicinity) may not produce the desired (or expected) results.

If you plan on converting source files which contain such code, it is recommended that you first make sure that each conditionally compiled region of source code which contains an alternative function header also contains at least one additional follower token (past the final right parenthesis of the function header). This should circumvent the problem.

- `unprotoize` can become confused when trying to convert a function definition or declaration which contains a declaration for a pointer-to-function formal argument which has the same name as the function being defined or declared. We recommend you avoid such choices of formal parameter names.
- You might also want to correct some of the indentation by hand and break long lines. (The conversion programs do not write lines longer than eighty characters in any case.)

Certain Changes GCC Will Not Use

The following documentation lists changes that people frequently request.

- Checking the number and type of arguments to a function which has an old-fashioned definition and no prototype.

Such a feature would work only occasionally—only for calls that appear in the same file as the called function, following the definition. The only way to check

all calls reliably is to add a prototype for the function. But adding a prototype eliminates the motivation for this feature. So the feature is not worthwhile.

- Warning about using an expression whose type is signed as a shift count.
Shift count operands are probably signed more often than unsigned. Warning about this would cause far more annoyance than good.
- Warning about assigning a signed value to an unsigned variable.
Such assignments must be very common; warning about them would cause more annoyance than good.
- Warning about unreachable code.
It is very common to have unreachable code in machine-generated programs. For example, this happens normally in some files of GNU C itself.
- Warning when a non-void function value is ignored.
Coming from a Lisp background, the idea seems silly that there is something dangerous about discarding a value. There are functions that return values which some callers may find useful; it makes no sense to clutter the program with a cast to `void` whenever the value is not useful.
- Assuming (for optimization) that the address of an external symbol is never zero.
This assumption is false on certain systems when `#pragma weak` is used.
- Making `-fshort-enums` the default.
This would cause storage layout to be incompatible with most other C compilers. And it does not seem very important, given that you can get the same result in other ways. The case where it matters most is when the enumeration-valued object is inside a structure, and in that case you can specify a field width explicitly.
- Making bitfields unsigned by default on particular machines where “the ABI standard” says to do so.

The ANSI C standard leaves it up to the implementation whether a bitfield declared plain `int` is signed or not. This in effect creates two alternative dialects of C.

The GNU C compiler supports both dialects; you can specify the signed dialect with `-fsigned-bitfields` and the unsigned dialect with `funsigned-bitfields`. However, this leaves open the question of which dialect to use by default.

Currently, the preferred dialect makes plain bitfields signed, because this is simplest. Since `int` is the same as `signed int` in every other context, it is cleanest for them to be the same in bitfields as well.

Some computer manufacturers have published Application Binary Interface standards which specify that plain bitfields should be unsigned. It is a mistake, however, to say anything about this issue in an ABI. This is because the handling of plain bitfields distinguishes two dialects of C. Both dialects are meaningful on every type of machine. Whether a particular object file was compiled using signed

bitfields or unsigned is of no concern to other object files, even if they access the same bitfields in the same data structures.

A given program is written in one or the other of these two dialects. The program stands a chance to work on most any machine if it is compiled with the proper dialect. It is unlikely to work at all if compiled with the wrong dialect.

Many users appreciate the GNU C compiler because it provides an environment that is uniform across machines. These users would be inconvenienced if the compiler treated plain bitfields differently on certain machines.

Occasionally users write programs intended only for a particular machine type. On these occasions, the users would benefit if the GNU C compiler were to support by default the same dialect as the other compilers on that machine. But such applications are rare. And users writing a program to run on more than one type of machine cannot possibly benefit from this kind of compatibility.

This is why GCC does and will treat plain bitfields in the same fashion on all types of machines (by default). There are some arguments for making bitfields unsigned by default on all machines. If, for example, this becomes a universal de facto standard, it would make sense for GCC to go along with it. This is a future consideration.

(Of course, users strongly concerned about portability should indicate explicitly in each bitfield whether it is signed or not. In this way, they write programs which have the same meaning in both C dialects.)

- Undefined `__STDC__` when `-ansi` is not used.

Currently, GCC defines `__STDC__` as long as you do not use `-traditional`. This provides good results in practice.

Programmers normally use conditionals on `__STDC__` to ask whether it is safe to use certain features of ANSI C, such as function prototypes or ANSI token concatenation. Since plain GCC supports all the features of ANSI C, the correct answer to these questions is “yes.” Some users try to use `__STDC__` to check for the availability of certain library facilities. This is actually incorrect usage in an ANSI C program, because the ANSI C standard says that a conforming freestanding implementation should define `__STDC__` even though it does not have the library facilities. `gcc -ansi -pedantic` is a conforming freestanding implementation, and it is therefore required to define `__STDC__`, even though it does not come with an ANSI C library.

Sometimes people say that defining `__STDC__` in a compiler that does not completely conform to the ANSI C standard somehow violates the standard. This is illogical. The standard is a standard for compilers that claim to support ANSI C, such as `gcc -ansi`, not for other compilers such as plain `gcc`. Whatever the ANSI C standard says is relevant to the design of plain `gcc` without `-ansi` only for pragmatic reasons, not as a requirement.

- **Undefining `__STDC__` in C++:** programs written to compile with C++-to-C translators get the value of `__STDC__` that goes with the C compiler that is subsequently used. These programs must test `__STDC__` to determine what kind of C preprocessor that compiler uses: whether they should concatenate tokens in the ANSI C fashion or in the traditional fashion. These programs work properly with GNU C++ if `__STDC__` is defined. They would not work otherwise.

In addition, many header files are written to provide prototypes in ANSI C but not in traditional C. Many of these header files can work without change in C++ provided `__STDC__` is defined. If `__STDC__` is not defined, they will all fail, and will all need to be changed to test explicitly for C++ as well.

- **Deleting `empty` loops:** GCC does not delete `empty` loops because the most likely reason you would put one in a program is to have a delay. Deleting them will not make real programs run any faster, so it would be pointless.

It would be different if optimization of a non-empty loop could produce an empty one. But this generally can not happen.

- **Making side effects happen in the same order as in some other compiler.** It is never safe to depend on the order of evaluation of side effects. For example, a function call like this may very well behave differently from one compiler to another:

```
void func (int, int);

int i = 2;
func (i++, i++);
```

There is no guarantee (in either the C or the C++ standard language definitions) that the increments will be evaluated in any particular order. Either increment might happen first. `func` might get the arguments 2, 3, or it might get 3, 2, or even 2, 2.

- **Not allowing structures with volatile fields in registers.** Strictly speaking, there is no prohibition in the ANSI C standard against allowing structures with volatile fields in registers, but it does not seem to make any sense and is probably not what you wanted to do. So the compiler will give an error message in this case.

Warning Messages and Error Messages

The GNU compiler can produce two kinds of diagnostics: errors and warnings. Each kind has a different purpose:

- **Errors** report problems that make it impossible to compile your program. GCC reports errors with the source filename and line number where the problem is apparent.

- Warnings report other unusual conditions in your code that may indicate a problem, although compilation can (and does) proceed. Warning messages also report the source filename and line number, but include the text, `warning:`, to distinguish them from error messages.

Warnings may indicate danger points where you should check to make sure that your program really does what you intend; or the use of obsolete features; or the use of nonstandard features of GNU C or C++. Many warnings are issued only if you ask for them, with one of the `-w` options (for instance, `-Wall` requests a variety of useful warnings). GCC always tries to compile your program if possible; it never gratuitously rejects a program whose meaning is clear merely because (for instance) it fails to conform to a standard. In some cases, however, the C and C++ standards specify that certain extensions are forbidden, and a diagnostic must be issued by a conforming compiler. The `-pedantic` option tells GCC to issue warnings in such cases; `-pedantic-errors` says to make them errors instead. This does not mean that all non-ANSI constructs get warnings or errors. See “Options Requesting or Suppressing Warnings” on page 35 for more detail on these and related command line options.



Reporting Bugs

Your reporting of problems (*bugs*) plays an essential role in making GCC reliable. When you encounter a problem, the first thing to do is to see if it is already known. See “Known Problems with GCC” on page 337. If it isn’t known, report the problem. For those with valid support agreements with GNUPro Toolkit, see “Reporting Problems” in *Getting Started Guide*.

Reporting a bug may have a solution, or someone may have to fix it. If it does not have a solution, look in the service directory; see “How to Get Help with GCC” on page 373. In any case, the principal function of a bug report is to help the entire community by making the next version of GCC work better. Bug reports are your contribution to the maintenance of GCC.

Since the maintainers are very overloaded, we cannot respond to every bug report. However, if the bug has not been fixed, we are likely to send you a patch and ask you to tell us whether it works.

In order for a bug report to serve its purpose, you must include the information that makes for fixing the bug.

Have You Found a Bug?

If you are not sure whether you have found a bug, the following guidelines will help:

- If the compiler gets a fatal signal, for any input whatever, that is a compiler bug. Reliable compilers never crash.
- If the compiler produces invalid assembly code, for any input whatever (except an `asm` statement), that is a compiler bug, unless the compiler reports errors (not just warnings) which would ordinarily prevent the assembler from being run.
- If the compiler produces valid assembly code that does not correctly execute the input source code, that is a compiler bug.

However, you must double-check to make sure, because you may have run into an incompatibility between GCC and traditional C (see “Incompatibilities of GCC” on page 350). These incompatibilities might be considered bugs, but they are inescapable consequences of valuable features. Or you may have a program whose behavior is undefined, which happened by chance to give the desired results with another C or C++ compiler.

For example, in many nonoptimizing compilers, you can write `x;` at the end of a function instead of `return x;`, with the same results. But the value of the function is undefined if `return` is omitted; it is not a bug when GCC produces different results.

Problems often result from expressions with two increment operators, as in `f(*p++, *p++)`. Your previous compiler might have interpreted that expression the way you intended; GCC might interpret it another way. Neither compiler is wrong. The bug is in your code.

After you have localized the error to a single source line, it should be easy to check for these things. If your program is correct and well defined, you have found a compiler bug.

- If the compiler produces an error message for valid input, that is a compiler bug.
- If the compiler does not produce an error message for invalid input, that is a compiler bug. However, you should note that your idea of “invalid input” might be my idea of “an extension” or “support for traditional practice.”
- If you are an experienced user of C or C++ compilers, your suggestions for improvement of GCC or G++ are welcome in any case.

Where to Report Bugs

For those with valid support agreements, see “Reporting Problems” in *Getting Started Guide*.

For others, send bug reports for GCC to:

`bug-gcc@prep.ai.mit.edu`

Also, send bug reports for G++ to:

`bug-g++@prep.ai.mit.edu`

If your bug involves the C++ class library `libg++`, send mail to:

`bug-lib-g++@prep.ai.mit.edu`

If you're not sure, you can send the bug report to both lists.

Do not send bug reports to `help-gcc@prep.ai.mit.edu` ***or to the*** `gnu.gcc.help` ***newsgroup***, . Most users of GCC do not want to receive bug reports. Those that do, have asked to be on `bug-gcc` and/or `bug-g++`.

The mailing lists, `bug-gcc` and `bug-g++`, both have newsgroups which serve as repeaters: `gnu.gcc.bug` and `gnu.g++.bug`. Each mailing list and its newsgroup carry exactly the same messages.

Often people think of posting bug reports to the newsgroup instead of mailing them. This appears to work, but it has one problem which can be crucial: *a newsgroup posting does not contain a mail path back to the sender*. Thus, if maintainers need more information, they may be unable to reach you. For this reason, you should always send bug reports by mail to the proper mailing list.

As a last resort, send bug reports on paper to:

GNU Compiler Bugs
Free Software Foundation
59 Temple Place Suite 330
Boston, MA 02111-1307
USA

How to Report Bugs

The fundamental principle of reporting bugs usefully is this: ***report all the facts***. If you are not sure whether to state a fact or leave it out, state it!

Often people omit facts because they think they know what causes the problem and they conclude that some details don't matter. Thus, you might assume that the name of the variable you use in an example does not matter. Well, probably it doesn't, but one cannot be sure. Perhaps the bug is a stray memory reference which happens to fetch from the location where that name is stored in memory; perhaps, if the name were different, the contents of that location would fool the compiler into doing the right thing despite the bug. Play it safe and give a specific, complete example. That is the easiest thing for you to do, and the most helpful.

Keep in mind that the purpose of a bug report is to enable someone to fix the bug if it is not known. It isn't very important what happens if the bug is already known. Therefore, always write your bug reports on the assumption that the bug is not known.

Sometimes people give a few sketchy facts and ask, "Does this ring a bell?" This cannot help us fix a bug, so it is basically useless. We respond by asking for enough

details to enable us to investigate. You might as well expedite matters by sending them to begin with.

Try to make your bug report self-contained. If we have to ask you for more information, it is best if you include all the previous information in your response, as well as the information that was missing.

Please report each bug in a separate message. This makes it easier for us to track which bugs have been fixed and to forward your bugs reports to the appropriate maintainer.

Do not compress and encode any part of your bug report using programs such as `uuencode`. If you do so it will slow down the processing of your bug. If you must submit multiple large files, use `shar`, which allows us to read your message without having to run any decompression programs.

To enable someone to investigate the bug, you should include all the following things.

- ***The version of GCC***

You can get this by running the program with the `-v` option.

Without this, we won't know whether there is any point in looking for the bug in the current version of GCC.

- ***A complete input file that will reproduce the bug***

If the bug is in the C preprocessor, send a source file and any header files that it requires. If the bug is in the compiler proper (`cc1`), run your source file through the C preprocessor by doing `gcc -E sourcefile > outfile`, then include the contents of `outfile` in the bug report. (When you do this, use the same `-I`, `-D` or `-U` options that you used in actual compilation.)

A single statement is not enough of an example. In order to compile it, it must be embedded in a complete file of compiler input; and the bug might depend on the details of how this is done.

Without a real example one can compile, all anyone can do about your bug report is wish you luck. It would be futile to try to guess how to provoke the bug. For example, bugs in register allocation and reloading frequently depend on every little detail of the function they happen in.

Even if the input file that fails comes from a GNU program, you should still send the complete test case. Don't ask the **GCC** maintainers to do the extra work of obtaining the program in question—they are all overworked as it is. Also, the problem may depend on what is in the header files on your system; it is unreliable for the **GCC** maintainers to try the problem with the header files available to them. By sending CPP output, you can eliminate this source of uncertainty and save us a certain percentage of wild goose chases.

- ***The command arguments you gave GCC or G++ to compile that example and observe the bug***

For example, did you use `-o`? To guarantee you won't omit something important,

list all the options.

If we were to try to guess the arguments, we would probably guess wrong and then we would not encounter the bug.

- ***The type of machine you are using, and the operating system name and version number***
- ***The operands you gave to the configure command when you installed the compiler***
- ***A complete list of any modifications you have made to the compiler source***
There is no promise to investigate the bug unless it happens in an unmodified compiler. But if you've made modifications and don't tell us, then you are sending us on a wild goose chase.

Be precise about these changes. A description in English is not enough—send a context `diff` for them.

Adding files of your own (such as a machine description for a machine we don't support) is a modification of the compiler source.

- ***Details of any other deviations from the standard procedure for installing GCC***
- ***A description of what behavior you observe that you believe is incorrect***
For example, “The compiler gets a fatal signal” or “The assembler instruction at line 208 in the output is incorrect” are messages of a fatal signal.

Of course, if the bug is of the compiler getting a fatal signal, then one can't miss it. But if the bug is incorrect output, the maintainer might not notice unless it is glaringly wrong. None of us has time to study all the assembler code from a 50-line C program just on the chance that one instruction might be wrong. *We need you to do this part!*

Even if the problem you experience is a fatal signal, you should still say so explicitly. Suppose something strange is going on, such as, your copy of the compiler is out of synch, or you have encountered a bug in the C library on your system. (This has happened!) Your copy might crash and the copy here would not. If you *said* to expect a crash, then when the compiler here fails to crash, we would know that the bug was not happening. If you don't say to expect a crash, then we would not know whether the bug was happening. We would not be able to draw any conclusion from our observations.

If the problem is a diagnostic when compiling GCC with some other compiler, say whether it is a warning or an error.

Often the observed symptom is incorrect output when your program is run. Sad to say, this is not enough information unless the program is short and simple. None of us has time to study a large program to figure out how it would work if compiled correctly, much less which line of it was compiled wrong. So you will have to do that. Tell us which source line it is, and what incorrect result happens

when that line is executed. A person who understands the program can find this as easily as finding a bug in the program itself.

- ***If you send examples of assembler code output from GCC or G++, please use `-g` when you make them.***

The debugging information includes source line numbers which are essential for correlating the output with the input.

- ***If you wish to mention something in the GCC source, refer to it by context, not by line number.***

The line numbers in the development sources don't match those in your sources. Your line numbers would convey no useful information to the maintainers.

- ***Additional information from a debugger.***

This might enable someone to find a problem on a machine which he does not have available. However, you need to think when you collect this information if you want it to have any chance of being useful.

For example, many people send just a backtrace, but that is never useful by itself. A simple backtrace with arguments conveys little about GCC because the compiler is largely data-driven; the same functions are called over and over for different RTL insns, doing different things depending on the details of the insn.

Most of the arguments listed in the backtrace are useless because they are pointers to RTL list structure. The numeric values of the pointers, which the debugger prints in the backtrace, have no significance whatever; all that matters is the contents of the objects they point to (and most of the contents are other such pointers).

In addition, most compiler passes consist of one or more loops that scan the RTL insn sequence. The most vital piece of information about such a loop—which insn it has reached—is usually in a local variable, not in an argument.

What you need to provide in addition to a backtrace are the values of the local variables for several stack frames up. When a local variable or an argument is an RTX, first print its value and then use the GDB command `pr` to print the RTL expression that it points to. (If GDB doesn't run on your machine, use your debugger to call the function `debug_rtx` with the RTX as an argument.) In general, whenever a variable is a pointer, its value is no use without the data it points to.

What follows are some things that are not necessary.

- ***A description of the envelope of the bug***

Often people who encounter a bug spend a lot of time investigating which changes to the input file will make the bug go away and which changes will not affect it. This is often time consuming and not very useful, because the way we will find the bug is by running a single example under the debugger with breakpoints, not by pure deduction from a series of examples. You might as well save your time

for something else.

Of course, if you can find a simpler example to report instead of the original one, that is a convenience.

Errors in the output will be easier to spot, running under the debugger will take less time, etc.

Most GCC bugs involve just one function, so the most straightforward way to simplify an example is to delete all the function definitions except the one where the bug occurs. Those earlier in the file may be replaced by external declarations if the crucial function depends on them. (Exception: inline functions may affect compilation of functions defined later in the file.)

However, simplification is not vital; if you don't want to do this, report the bug anyway and send the entire test case you used.

- ***Conditionals***

In particular, some people insert conditionals (`#ifdef BUG`) around a statement which, if removed, makes the bug not happen. These are just clutter; we won't pay any attention to them anyway. Besides, you should send us cpp output, and that can't have conditionals.

- ***A patch for the bug***

A patch for the bug is useful if it is a good one. But don't omit the necessary information, such as the test case, on the assumption that a patch is all we need.

We might see problems with your patch and decide to fix the problem another way, or we might not understand it at all.

Sometimes with a program as complicated as GCC it is very hard to construct an example that will make the program follow a certain path through the code.

If you don't send the example, we won't be able to construct one, so we won't be able to verify that the bug is fixed.

And if we can't understand what bug you are trying to fix, or why your patch should be an improvement, we won't install it. A test case will help us to understand.

See "Sending Patches for GCC" on page 370 for guidelines on how to make it easy for us to understand and install your patches.

- ***A guess about what the bug is or what it depends on***

Such guesses can be wrong.

- ***A core dump file***

There is no way of examining a core dump for your type of machine without an identical system.

Sending Patches for GCC

If you would like to write bug fixes or improvements for the GCC compiler, that is very helpful.

Send suggested fixes to the bug report mailing list:

`bug-gcc@prep.ai.mit.edu`

Please follow these guidelines so we can study your patches efficiently. If you don't follow these guidelines, your information might still be useful, but using it will take extra work. Maintaining GCC is a lot of work in the best of circumstances, and we can't keep up unless you do your best to help.

- Send an explanation with your changes of what problem they fix or what improvement they bring about. For a bug fix, just include a copy of the bug report, and explain why the change fixes the bug.

(Referring to a bug report is not as good as including it, because then we will have to look it up, and we have probably already deleted it if we've already fixed the bug.)

- Always include a proper bug report for the problem you think you have fixed. We need to convince ourselves that the change is right before installing it. Even if it is right, we might have trouble judging it if we don't have a way to reproduce the problem.
- Include all the comments that are appropriate to help people reading the source in the future understand why this change was needed.
- Don't mix together changes made for different reasons. Send them *individually*.

If you make two changes for separate reasons, then we might not want to install them both. We might want to install just one. If you send them all jumbled together in a single set of diffs, we have to do extra work to disentangle them—to figure out which parts of the change serve which purpose. If we don't have time for this, we might have to ignore your changes entirely.

If you send each change as soon as you have written it, with its own explanation, then the two changes never get tangled up, and we can consider each one properly without any extra work to disentangle them. Ideally, each change you send should be impossible to subdivide into parts that we might want to consider separately, because each of its parts gets its motivation from the other parts.

- Send each change as soon as that change is finished. Sometimes people think they are helping us by accumulating many changes to send them all together. As explained above, this is absolutely the worst thing you could do.

Since you should send each change separately, you might as well send it right away. That gives us the option of installing it immediately if it is important.

- Use `diff -c` to make your diffs. Diffs without context are hard for us to install

reliably. More than that, they make it hard for us to study the diffs to decide whether we want to install them. Unidiff format is better than diffs without context, but not as easy to read as `-c` format.

If you have GNU `diff`, use `diff -cp`, in order to show the name of the function in which each change occurs.

- Write the change log entries for your changes. We get lots of changes, and we don't have time to do all the change log writing ourselves.

Read the `ChangeLog` file to see what sorts of information to put in, and to learn the style that we use. The purpose of the change log is to show people where to find what was changed. So you need to be specific about what functions you changed; in large functions, it's often helpful to indicate where within the function the change was.

On the other hand, once you have shown people where to find the change, you need not explain its purpose. Thus, if you add a new function, all you need to say about it is that it is new. If you feel that the purpose needs explaining, it probably does—but the explanation will be much more useful if you put it in comments in the code.

If you would like your name to appear in the header line for who made the change, send us the header line.

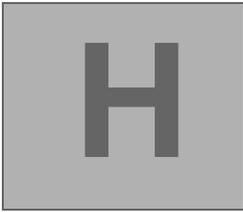
- When you write the fix, keep in mind that we can't install a change that would break other systems.

People often suggest fixing a problem by changing machine-independent files such as `toplev.c` to do something special that a particular system needs. Sometimes it is totally obvious that such changes would break GCC for almost all users. We can't possibly make a change like that. At best it might tell us how to write another patch that would solve the problem acceptably.

Sometimes people send fixes that *might* be an improvement in general—but it is hard to be sure of this. It's hard to install such changes because we have to study them very carefully. Of course, a good explanation of the reasoning by which you concluded the change was correct can help convince us.

The safest changes are changes to the configuration files for a particular machine. These are safe because they can't create new bugs on other machines.

Please help us keep up with the workload by designing the patch in a form that is good to install.



How to Get Help with GCC

If you still need help installing, using or changing GCC, there are several ways to find it, if they are not in this documentation.

- Send a message to a suitable network mailing list. First try:

`bug-gcc@prep.ai.mit.edu`

If that brings no response, try:

`help-gcc@prep.ai.mit.edu`.

- Look in the service directory for someone who might help you for a fee. The service directory is found in the file named `SERVICE` in the GCC distribution.

Index

Symbols

" (quotes), in #include statements 225
23, 202, 221, 223, 229, 250, 255, 264
158, 238–239
#assert 255
#define 221, 229–230, 240, 266
#elif directive 252
#else 251, 265
#endif 250, 252, 265
#error 255
#foo 221
#ident 259
#if 250, 252
#ifndef 226
#import 227, 268
#include 223–225, 227, 257, 265
#include_next 227
#line 257–258
#pragma 42, 167, 202–203, 206, 259
#pragma interface 206
#system 255
#unassert 255
#undef 240
\$ (dollar sign) 168, 268
&& (unary operators) 148
() (parentheses), as expressions in C 146
{ } (parentheses with curly braces) 147
+= (plus equals usage) 154
, (commas) 169, 179
.bb 49
.bbg 49
.ii 32
.save files 143
.verstamp 305
/* and */ 252, 265
// (forward slashes) 168, 267
/usr/local/lib 297
< (angle brackets), in #include statements 225
??, trigraph sequences 265
\ (backslashes) 227, 265
_ (underscores), in macro names 226
__aligned__ 173
__alignof__ 169
__arg 180
__attribute__ 161
__attribute__((aligned (alignment))) 169
__BASE_FILE__ 235
__bb_hidecall__ 46
__bb_jumps__ 46
__bb_showret__ 46
__bb_trace__ 46
__builtin_apply () 151
__builtin_apply_args () 151
__builtin_return () 151
__byte__ 171
__CHAR_UNSIGNED__ 235
__cplusplus 234
__DATE__ 234
__extension__ 196
__FILE__ 233
__GNUC__ 23, 145, 234
__GNUC_MINOR__ 234
__GNUG__ 234
__imp 165
__INCLUDE_LEVEL__ 235
__LINE__ 233
__OPTIMIZE__ 235
__REGISTER_PREFIX__ 235
__STDC__ 23, 234, 360
__STDC_VERSION__ 23, 234
__STRICT_ANSI__ 234

__TIME__ 234
 __typeof__ 152
 __USER_LABEL_PREFIX__ 235
 __VERSION__ 235
 __word__ 171
 _command, used with naming 239
 _exit 20
 _IEEE_FP 86
 {}, as expressions in C 146
 ~ (tilde), with complex type 155

Numerics

0, in constraint 178
 386 295
 64 94
 68000 options
 -m68000 114
 -m68020-40 115
 -m68020-60 115
 -m68030 114
 -m68040 114
 -m68060 114
 -m68881 114
 -malign-int 116
 -mbitfield 115
 -mc68000 114
 -mfpa 115
 -mno-align-int 116
 -mno-bitfield 115
 -mno-strict-align 116
 -mpcrel 116
 -mrtd 115
 -mshort 115
 -msoft-float 115
 -mstrict-align 116
 see Motorola 68000 options
 68020 options
 -m68020 114
 -mc68020 114
 88000 options
 see Motorola 88000 options

A

-A 61
 -a 45
 abbreviations 303
 abort 20, 162
 abs 21
 abstract datatypes 207
 ADB 346
 addition operations 158
 addressable unit 305
 Advanced RISC Machines (ARM) 78
 aggregate initializer 159
 AIX 315

AIX Threads 96
 aligned 170, 173
 alignment
 maximum 174
 of types, of variables 169
 allclass.cc 203
 allclass.h 203
 alloca 20, 156, 326, 341, 346
 alloca, for Unos 312
 Alpha 296
 alternative keywords, compiling 196
 Altos 343
 AM33 options
 -mam33 107
 -mmult-bug 107
 -mno-am33 107
 -mno-mult-bug 107
 -mrelax 107
 AMD
 K6 102
 AMD options
 K6 102
 AMD29K options
 -m29000 76
 -m29050 76
 -mbw 76
 -mdw 76
 -mkernel-registers 77
 -mlarge 77
 -mnbw 76
 -mndw 76
 -mno-reuse-arg-regs 77
 -mnormal 77
 -mno-stack-check 77
 -mno-storem-bug 77
 -mreuse-arg-regs 77
 -msmall 77
 -mstack-check 77
 -mstorem-bug 77
 -muser-registers 77
 angle braces 225
 -ansi 19, 237, 360
 ANSI C standard
 comments = whitespace 264
 exceptions with using backslash 225
 naming definitions 236
 requirements 220
 string constants 238
 using -trigraphs 265
 ANSI conformance 268
 ANSI or GNU C conventions 20
 ANSI Standard C 218
 ANSI support 19
 -ansi with keywords 196
 answer 254
 AOS 315
 APCS 78–79
 ar 320
 ARC options

- EB 78
- EL 78
- mcpu= 78
- mdata= 78
- mmangle-cpu 78
- mrodata= 78
- mtext=-mtext= 78
- archive file 68
- arcs during compilation 49
- arcs, gcov 49
- args 157
- arguments 357
 - types 357
 - with braces 231
 - with brackets 231
 - with commas 231
- arithmetic operators 251
- ARM 78
 - PIC addressing 83
- ARM options
 - mabort-on-noreturn 82
 - malignment-traps 80
 - mapcs 78
 - mapcs-26 78
 - mapcs-32 78
 - mapcs-float 79
 - mapcs-frame 78
 - mapcs-reentrant 79
 - mapcs-stack-check 79
 - march= 81
 - matpcs 79, 84
 - mbig-endian 80
 - mbsd 82
 - mbuggy-return-in-memory 79
 - mcpu= 81
 - mfpe= 81
 - mhard-float 81
 - mlittle-endian 80
 - mlong-calls 82
 - mnoabort-on-noreturn 82
 - mno-alignment-traps 80
 - mno-apcs-frame 78
 - mno-apcs-reentrant 79
 - mno-apcs-stack-check 79
 - mno-atpcs 79
 - mno-buggy-return-in-memory 79
 - mno-long-calls 82
 - mnop-fun-dllimport 83
 - mnop-nop-fun-dllimport 83
 - mno-sched-prolog 80
 - mno-short-load-bytes 82
 - mno-short-load-words 82
 - mno-single-pic-base 83
 - mno-symrename 83
 - mpic-register= 83
 - msched-prolog 80
 - mshort-load-bytes 82
 - mshort-load-words 82
 - msingle-pic-base 83
 - msoft-float 81
 - mstructure-size-boundary= 81
 - mthumb-interwork 81
 - mwords-little-endian 80
 - mxopen 82
- ARM Procedure Call Standard (APCS) 78
- ARM targets
 - libraries 81
- ARM THUMB Procedure Call Standard (ATPCS) 79, 84
- array 148–149, 158–160, 238
- as 320, 339
- as C++ constructors and destructors 129
- asm 196
- asm operands 181
- assembler
 - labels 235
 - naming process 193
- assembler options
 - Wa 65
- assertions 254
- assessment of computing time 9, 209
- AT&T C++ translator, Cfront 204
- ATPCS (ARM THUMB Procedure Call Standard) 79, 84
- attribute 162, 167
 - alias 165
 - aligned() 170
 - cdecl 165
 - const 162
 - constructor 162
 - destructor 162
 - dllexport 165
 - dllimport 165
 - format 162
 - longcall 165
 - mode 171
 - nocommon 171
 - noreturn 162
 - packed 171
 - regparm 165
 - section 162, 171
 - stdcall 165
 - transparent_union 172
 - types 173
 - underscore, specifying with 162
 - unused 162, 172
 - variables 169
 - weak 162, 172
- auto 150
- automatic arrays 156
- automatic variables 23, 350
- automounters 338
- awrning options
 - Wunknown-pragmas 30
- ax 46
- AXP systems (DEC Alpha) 305

B

- B 71, 138, 142
- b 73, 76
- back-ends 302
- backslashes 219, 227, 238
- backtrace 368
- basename 203
- basic block 49, 212
- Berkeley 315
- Binutils 320
- Bison parser generator 257, 298
- bit field 147
- bitfield 359
- blocks to linkers 204
- Borland model 204
- braces 146, 174, 225
- BSD 296, 304, 339, 344
- BUFSIZE 246
- bugs 363–364
- build 294
- build machine 294
- building a cross-compiler 324
- building in separate directories 318
- built-in functions 151
- byte 76, 171

C

- C 61, 142
- c 67, 142
- C code 257
- C dialect options
 - ansi 19
 - fallow-single-precision 20
 - fcond-mismatch 20
 - ffreestanding 21
 - fhosted 21
 - fno-asm 21
 - fno-builtin 21
 - fno-signed-bitfields 21
 - fno-unsigned-bitfields 21
 - fsigned-bitfields 21
 - fsigned-char 22
 - fstd 20
 - funsigned-bitfields 21
 - funsigned-char 22
 - fwritable-strings 22
 - pedantic 19
 - traditional 22
 - traditional-cpp 23
 - trigraphs 24
- C expression
 - arithmetic operators 251
 - character constants 251
 - identifiers 251
 - integer constants 250
- C language extensions 199

- C preprocessor
 - #include 223
 - changes to input
 - deleted backslash-newline sequences 219
 - exceptions 220
 - replaced predefined macro names 219
 - spaces 219
 - conditional compilation 218
 - CPP_PREDEFINES 237
 - definition 217
 - directives to activate 219
 - header files 218, 223
 - infile, outfile arguments 263
 - invoking 263
 - line control 218
 - macro 217–218, 230
 - necessity of -trigraphs 265
 - output 261
 - parsing 219
 - recognizing directives 219
 - source file requirements 224
 - source files, combining 257
 - stringification 237
 - variants of #include 225
- C thread implementation 302
- C++ 10, 234
 - templates (Borland, Cfront models) 204
 - undefining __STDC__ 361
- C++ comments 267
- C++ dialect options
 - falt-external-templates 26
 - fcheck-new 25
 - fconserve-space 26
 - fdollars-in-identifiers 26
 - fembedded-cxx 26
 - fexternal-templates 26
 - ffor-scope 26
 - fguiding-decls 29
 - fhandle-signatures 27
 - fhonor-std 27
 - fhuge-objects 32
 - finit-priority 27
 - fname-mangling-version-n 28
 - fno-access-control 25
 - fno-default-inline 28
 - fno-elide-constructors 26
 - fno-for-scope 26
 - fno-gnu-keywords 28
 - fno-implement-inlines 27
 - fno-implicit-templates 27
 - fno-nonnull-objects 29
 - fno-optional-diags 27
 - fno-squangle 27
 - foperator-names 29
 - fpermissive 27
 - frepo 27
 - fsquangle 27
 - fstrict-prototype 29
 - ftemplate-depth-n 28

- fthis-is-variable 29
- fuse-cxa-atexit 28
- fvtbl-thunks 28
- nostdinc 29
- W 30
- Wall 30
- Wctor-dtor-privacy 30
- Weffc++ 30
- Wno-deprecated 30
- Wno-non-template-friend 30
- Wnon-virtual-dtor 30
- Wno-pmf-conversions 30
- Wold-style-cast 30
- Woverloaded-virtual 30
- Wreorder 30
- Wsign-promo 31
- Wsynth 31
- Wunknown-pragmas 30
- C++ linkage 29
- C++ run-time library 301
- C++ suffixes 32
- C_INCLUDE_PATH 139
- C9X support for C standards 20
- canonical configuration name 295
- case ranges 161
- cast 153, 159, 359
- cdecl 103
- Cdialect options
 - flang-isoc9x 20
- ceil_div 242
- cexp.c 298
- cexp.y 298
- cfront 44
- Cfront model 204
- changes, unnecessary 358
- character constants 251, 261
- charfunc 356
- class function declaration 44
- class scope 28
- classes 356, 364
- CLipper optionms
 - mc300 85
- CLipper options
 - mc400 85
- Clipper options 85
- clobbering 23
- code generation options 127
 - +e0 133
 - +e1 133
 - fcall-saved 127
 - fcall-used 127
 - ffixed 128
 - finhibit-size-directive 128
 - fno-common 129
 - fno-gnu-linker 129
 - fno-ident 129
 - fpack-struct 132
 - fpcc-struct-return 129
 - fPIC 130
 - fpic 130
 - freg-struct-return 130
 - fshared-data 130
 - fshort-double 131
 - fshort-enums 130
 - funaligned-pointers 131
 - funaligned-struct-hack 131
 - fverbose-asm 132
 - fvolatile 131
 - fvolatile-global 131
 - fvolatile-static 131
- code size reduction 60
- COFF 50
- collect2 129, 325
- colon after the input operands 178
- commas 158, 167, 169, 179, 231
- comments 168, 240, 252, 265, 268
- common sub-expression elimination (CSE) 54
- compare instructions 181
- compilation errors 299
- compiler
 - conditional preprocessing 218
 - driver 71
 - options 142
 - passes, statistics on 52
 - search directories 326
 - versions, specifying 73
- COMPILER_PATH 138
- complex automatic variables 155
- complex data types 155
- compound expressions 153
- compound statement 146
- concatenation 238–239
- conditional compilation 145, 218
- conditional expressions 153–154
- conditionals 195, 226, 251, 255, 257
- config.h 297
- configuration 137, 294, 304
- configure
 - custom naming 303
 - error messages 304
- conflicts with names in assembler code 193
- conforming freestanding implementation 354
- conjugation 155
- consecutive string constants 238
- const 196
- CONST_DOUBLE_OK_FOR_LETTER_P 185
- CONST_OK_FOR_LETTER_P 185
- constant addresses 130
- constant expressions 160
- constants 22
- constraints 178, 181, 184
- constructor call 200
- constructor expressions 159
- constructors 159, 325
- contacting Red Hat iii
- containing functions 149, 156
- contributors to GNU CC 279
- conversion pragmas 357

cos 21
 cp 297
 c-parse.c 298
 c-parse.y 298
 CPLUS_INCLUDE_PATH 139
 CPP_PREDEFINES 237
 CPU types, supported 303
 cross compilation problems 343
 CROSS_INCLUDE_DIR 324
 cross-assembler 320
 cross-compiling 73, 320
 cross-linker 320
 CSE 48, 54
 CSE (Common Sub-expression Elimination) 48, 58
 CSE (common sub-expression elimination) 54
 C-series 106
 curly braces 147
 customary abbreviations for configuring 303

D

-D 62, 237, 254
 -d 47
 D10V options
 -maccum 112
 -maddac3 112
 -masm-optimize 112
 -mbranch-cost= 113
 -mcond-exec= 113
 -mcond-move 112
 -mdouble32 112
 -mdouble64 112
 -mint16 112
 -mint32 112
 -mno-accum 112
 -mno-addac3 112
 -mno-asm-optimize 112
 -mno-cond-move 112
 -mno-small-insn 113
 -msmall-insn 113
 data flow analysis, interprocedural 200
 DBX 344
 -dD 61
 debugging 45, 155, 250, 296, 363
 RTL (see RTL)
 debugging options
 -a 45
 -ax 46
 -d 47
 -fdump-translation-unit-file 49
 -fdump-unnumbered 48
 -fpretend-float 49
 -fprofile-arcs 49
 -ftest-coverage 49
 -g 50
 -gcoff 50–51
 -gdwarf 50
 -gdwarf-1 51
 -gdwarf-1+ 51
 -gdwarf-2 51
 -gdwarf2 50
 -ggdb 50
 -gstabs 50
 -gstabs+ 51
 -gxcoff 50–51
 -gxcoff+ 51
 -p 51
 -pg 51
 -print-file-name= 51
 -print-libgcc-file-name 52
 -print-prog-name=program 51
 -print-search-dirs 52
 -Q 52
 -save-temps 52
 DEC Alph options
 -mmax 89
 DEC Alpha 305
 floating point operations 86
 options 85
 DEC ALpha options
 -mcpu= 88
 DEC Alpha options
 -fprm 87
 -malpha-as 86
 -mbuild-constants 88
 -mbwx 89
 -mcix 89
 -mfp-reg 85
 -mfp-rounding-mode= 87
 -mfp-trap-mode= 86
 -mgas 86
 -mieee 86
 -mieee-conformant 86
 -mieee-with-inexact 86
 -mmemory-latency= 89
 -mno-bwx 89
 -mno-cix 89
 -mno-fp-regs 85
 -mno-max 89
 -mno-soft-float 85
 -msoft-float 85
 -mtrap-precision= 87
 DEC Unix 305
 declarations as header files 218
 definition 239
 Delta 88, debugging 117
 dependencies, make 63
 DEPENDENCIES_OUTPUT 139
 derived class 44
 destructors 201–202
 DG/UX, debugging 117
 diagnostics 27
 diffs 370
 dir 267
 directive names, defined 221
 directives, miscellaneous 259
 directories, searching 71

-dM 62
 -dN 62
 do...while statement 242
 double precision 20
 double scan 244
 doublequote characters 238
 double-word integers 154
 DW bit (29K) 76
 DWARF 50–51

E

-E 61, 67
 EABI, PowerPC 99
 -EB 78, 111
 ECOFF symbol table 296
 -EL 78, 111
 elision algorithm 200
 else statements 242
 Emacs 315
 Embedded Applications Binary Interface 99
 Embedded C++ specification 26
 embedded comments 268
 Embedded PowerPC 316
 embedded statements, unallowed in constants 147
 Emerald type modules 207
 Encore ns32000 315
 enquire 325, 339
 enum 174, 196
 enumeration 147, 339
 environment variables 137, 338
 C_INCLUDE_PATH 139
 COMPILER_PATH 138
 CPLUS_INCLUDE_PATH 139
 crt0.o 138
 dependencies, specifying 139
 DEPENDENCIES_OUTPUT 139
 GCC_EXEC_PREFIX 138
 LANG 139
 languages, using 139
 LC_ALL 137
 LC_COLLATE 137
 LC_CTYPE 137
 LC_MESSAGES 137
 LC_MONETARY 137
 LC_NUMERIC 137
 LC_TIME 137
 libraries, using 138
 LIBRARY_PATH 138
 linking files 138
 locale information 139
 localization information 137
 OBJC_INCLUDE_PATH 139
 preprocessor output, using 138
 subprograms, using 138
 temporary files, using 138
 TMPDIR 138
 errors 361

ESC, in constants 169
 EV4, EV5 processors 88
 exception handling 26
 execution time, reducing 60
 exit 20, 162
 expanding a macro 358
 expressions 146, 152, 250, 299
 extensions 145, 199
 extern 22, 150, 351
 external variables 351
 EXTRA_CONSTRAINT 185

F

f68881 326
 fabs 21
 Fairchild F9450 C compiler 305
 -falign-functions 53
 -falign-functions= 53
 -falign-jumps 54
 -falign-jumps= 54
 -falign-labels 53
 -falign-labels= 53
 -falign-loops 54
 -falign-loops= 54
 -fallow-single-precision 20
 -falt-external-templates 26, 206
 -fbranch-probabilities 54
 -fcaller-saves 54
 -fcheck-new 25
 -fcond-mismatch 20
 -fcse-follow-jumps 54
 -fcse-skip-blocks 54
 -fdata-sections 56
 -fdelayed-branch 54
 -fdelete-null-pointer-checks 55
 -fdollars-in-identifiers 26, 347
 -fdump-translation-unit-file 49
 -fdump-unnumbered 48
 -fembedded-cxx 26
 -fexpensive-optimizations 55
 -fexternal-templates 26
 -ffast-math 55
 -ffloat-store 55
 -fforce-addr 55
 -fforce-mem 55
 -ffor-scope 26
 -ffreestanding 21
 ffs 21
 -ffunction-sections 56
 -fgcse 56
 -fguiding-decls 29
 -fhandle-signatures 27, 207
 -f Honor-std 27
 -fhonored 21
 -fhuge-objects 32
 filenames 67
 file-scope objects 27

- finit-priority 27
- finline-functions 56
- finline-limit= 56
- fixincludes 338, 344
- fixproto 338
- fkeep-inline-functions 56
- fkeep-static-consts 57
- flang-isoc9x 20
- flive-range 54
- float 352
- float.h 325
- FLOAT_OPTION 326
- floating point 20, 55, 343
- floating point values 81
- fmove-all-movables 56
- fname-mangling-version-n 28
- fno-asm 21
- fno-builtin 21
- fno-default-inline 28, 57
- fno-defer-pop 57
- fno-delayed-branch 341
- fno-elide-constructors 26
- fno-for-scope 26
- fno-function-cse 57
- fno-gnu-keywords 28
- fno-implement-inlines 27
- fno-implicit-templates 27, 205
- fno-inline 57
- fno-math-errno 55
- fno-nonnull-objects 29
- fno-optional-diags 27
- fno-peephole 57
- fno-signed-bitfields 21
- fno-squangle 27
- fno-unsigned-bitfields 21
- fomit-frame-pointer 57
- foperator-names 29
- foptimize-register-moves 58
- foptimize-sibling-calls 57
- FORTRAN options 58
- fpcc-struct-return 352
- fpermissive 27
- fPIC 130
- fpic 130
- fpretend-float 49
- fprm 87
- fprofile-arcs 49, 211
- FPU 102
- frame pointers 57
- FRAME_POINTER_REQUIRED 57
- freduce-all-givs 58
- fregmove 58
- frepo 27, 205
- frerun-cse-after-loop 58
- frerun-loop-opt 58
- friend functions 43
- fschedule-insns 58
- fschedule-insns2 58
- fshort-enums 359
- fsigned-bitfields 21
- fsigned-char 22
- fsquangle 27
- fssa 58
- fstd 20
- fstrength-reduce 58
- fstrict-aliasing 58
- fstrict-prototype 29
- fsyntax-only 35
- ftemplate-depth-n 28
- ftest-coverage 49, 211
- fthis-is-variable 29
- fthread-jumps 59
- Fujitsu SPARClike options 122–123
- function 156
 - attributes, declaring 161
 - block, within 351
 - calls 194
 - compiling for interrupt calls 168
 - constructing calls 151
 - declarations 207
 - definition extensions 199
 - expanded macros 147
 - inline 176
 - inlining 60, 203
 - jumping to a different function's code 149
 - noreturn 162
 - pointer 165
 - prototype 167
 - strings 197
 - value, non-void 359
- funroll-all-loops 59
- funroll-loops 59
- funsigned-bitfields 21
- funsigned-char 22
- fuse-cxa-atexit 28
- fvtbl-thunks 28, 32
- fwritable-strings 22

G

- G 100, 111, 113, 124
- g 50, 142
- G++, definition 10
- GAS 295, 332, 344
- gcc 267, 339
- GCC_EXEC_PREFIX 138
- GCC_INCLUDE_DIR 324
- gcoff 50–51
- gcov 49, 209–210
- gcov 49
- GCSE 48
- GDB 257, 344
- gdwarf 50
- gdwarf-1 51
- gdwarf-1+ 51
- gdwarf-2 50–51
- General Purpose group 94

genflags 340
 genoutput 340
 -ggdb 50
 glibc 354
 Global Common Sub-expression Elimination 48, 56
 global declarations 142, 351
 global initializations 129
 global offset table 130
 global pointer 165
 global symbols 69
 global variables 193
 globalref, globaldef, globalvalue 332
 GNU C
 builtin functions 22
 complex data types 155
 errors, warnings 362
 extensions 196
 function-definition syntax 200
 global variables 193
 incompatibilities 350
 library 339
 operators 201
 with other compilers 344
 GNU C++
 defining with `__STDC__` 361
 errors, warnings 362
 extensions 199
 function-definition syntax 200
 goto statements 202
 headers 202
 operators 201
 templates 204
 unsupported extensions 196
 unsupported nested functions 149
 with other compilers 344
 GNU CC
 BSD systems 339
 bugs 337–338, 369
 building 341
 building & installing a cross-compiler 319
 command options 11
 errors, warnings 362
 help services 373
 installation problems 338
 machine-description macros 352
 malloc 349
 MIPS 110
 on VMS 330
 optimization with `gcov` 213
 standard directory 297
 standard libraries 353
 static variables in registers 193
 string constants 350
 strings 197
 supported configurations 303
 system types 303
 GNU diff 371
 GNU extensions 35
 GNU linker 129, 295

GNU Make 339
 GOT 130
 goto 147–148, 202
 GPLUS_INCLUDE_DIR 324
 gprof 51
 grouping 242
 -gstabs 50
 -gstabs+ 51, 296
 -gxcoff 50–51
 -gxcoff+ 51
 gzip 46

H

-H 62
 H8/300 options 78, 89
 -mrelax 89
 hard registers 178
 hardware interrupts 168
 hardware, configuring 76
 Haskell type classes 207
 header 156
 header files 203, 218, 223, 298, 323–324
 base header files 227
 GCC_INCLUDE_DIR 324
 GPLUS_INCLUDE_DIR 324
 inheritance 227
 LOCAL_INCLUDE_DIR 324
 missing 266
 redefinition 240
 with `#ident` 259
 header files, searching 71
 header.h 225
 help 373
 heuristic 56
 Hewlett Packard HPPA options 90
 Hitachi H8/300 options 89
 -malign-300 90
 -mh 89
 -mint32 89
 -ms 89
 Hitachi SH options 90
 -m1 90
 -m2 90
 -m3 90
 -m3e 90
 -m4 90
 -mb 90
 -mdalign 90
 Hitachi SH options
 -mrelax 90
 host machine 294
 hot spots 210
 hot spots, code 210
 House Subcommittee on Intellectual Property 291
 HPPA 295, 346
 HPPA options 90
 -march= 90

- mbig-switch 91
- mdisable-fpregs 91
- mdisable-indexing 91
- mfast-indirect-calls 91
- mgas 91
- mjump-in-delay 91
- mlong-load-store 91
- mno-big-switch 91
- mno-disable-fpregs 91
- mno-disable-indexing 91
- mno-gas 91
- mno-jump-in-delay 91
- mno-long-load-store 91
- mno-portable-runtime 91
- mno-soft-float 92
- mno-space-regs 92
- mpa-risc-1-0 92
- mpa-risc-1-1 92
- mpa-risc-2-0 92
- mportable-runtime 92
- mschedule= 92
- msoft-float 92
- mspace-regs 92
- HPPA targets
 - libraries 92
- HPUX 91, 307, 346

I

- I 72, 227, 265
- I- 72
- i 143
- goto *array 148
- i960 106
- IBM 101
- IBM PowerPC options
 - G 100
 - mads 99
 - maix32 95
 - maix64 95
 - mbig 98
 - mbig-endian 98
 - mbit-align 97
 - mcall-aix 98
 - mcall-linux 98
 - mcall-solaris 98
 - mcall-sysv 98
 - mcall-sysv-eabi 98
 - mcall-sysv-noeabi 98
 - meabi 99
 - memb 99
 - mfull-toc 95
 - mfused-madd 97
 - mhard-float 96
 - mlittle 98
 - mlittle-endian 98
 - mminimal-toc 95
 - mmpc860c0= 99
 - mmultiple 96
 - mmvme 99
 - mnew-mnemonics 94
 - mno-bit-align 97
 - mno-eabi 99
 - mno-fp-in-toc 95
 - mno-fused-madd 97
 - mno-multiple 96
 - mno-power 94
 - mno-power2 94
 - mno-powerpc 94
 - mno-powerpc-gfxopt 94
 - mno-powerpc-gpopt 94
 - mno-prototype 98
 - mno-regnames 100
 - mno-relocatable 97
 - mno-relocatable-lib 97
 - mno-sdata 100
 - mno-strict-align 97
 - mno-string 96
 - mno-sum-in-toc 95
 - mno-toc 98
 - mno-traceback 100
 - mno-update 97
 - mno-xl-call 96
 - mold-mnemonics 94
 - mpower 94
 - mpower2 94
 - mpowerpc 94
 - mpowerpc-gfxopt 94
 - mpowerpc-gpopt 94
 - mprototype 98
 - mregnames 100
 - mrelocatable 97
 - mrelocatable-lib 97
 - msdata 99–100
 - msdata=default 100
 - msdata=eabi 100
 - msdata=none 100
 - msdata=sysv 100
 - msdata-data 100
 - msim 99
 - msoft-float 96
 - mstrict-align 97
 - mstring 96
 - mthreads 96
 - mtoc 98
 - mtraceback 100
 - mupdate 97
 - mxl-call 96
 - myellowknife 99
 - specs= 96
- IBM RS/600 and PowerPC
 - threads 96
- IBM RS/6000 and PowerPC 92
- IBM RS/6000 and PowerPC options 92
 - mcpu= 92
 - mcpu=403 93
 - mcpu=505 93

- mcpu=601 93
- mcpu=602 93
- mcpu=603 93
- mcpu=603e 93
- mcpu=604 93
- mcpu=620 93
- mcpu=821 93
- mcpu=860 93
- mcpu=common 93
- mcpu=power 93
- mcpu=power2 93
- mcpu=powerpc 93
- mcpu=rios1 93
- mcpu=rios2 93
- mcpu=rsc 93
- mpower 93
- mpowerpc 93
- msoft-float 93
- mtune= 93
- IBM RS/6000 options 94
 - mpe 96
 - mtune= 93
- IBM RS/6000 SP Parallel Environment 96
- IBM RT options
 - mcall-lib-mul 101
 - mfp-arg-in-fpregs 101
 - mfp-arg-in-gregs 101
 - mfull-fp-blocks 101
 - mhc-struct-return 101
 - min-line-mul 101
 - mminimum-fp-blocks 101
 - mnohc-struct-return 101
- IBM RT PC options 101
- IBM RT PC, supported operating systems 315
- IBM System V options
 - mbig 98
 - mbig-endian 98
 - mbit-align 97
 - mcall-aix 98
 - mcall-linux 98
 - mcall-solaris 98
 - mcall-sysv-eabi 98
 - mcall-sysv-noeab 98
 - meabi 99
 - mlittle 98
 - mlittle-endian 98
 - mno-bit-align 97
 - mno-eabi 99
 - mno-prototype 98
 - mno-regnames 100
 - mno-relocatable-lib 97
 - mno-strict-align 97
 - mno-toc 98
 - mno-traceback 100
 - mprototype 98
 - mregnames 100
 - mrelocatable 97
 - mrelocatable-lib 97
 - msdata 99–100
 - msdata=default 100
 - msdata=eabi 100
 - msdata=sysv 100
 - msdata-data 100
 - mstrict-align 97
 - mtoc 98
 - mtraceback 100
- IBM System, V options
 - mcall-sysv 98
- IBM XLC compiler 315
- identifier 168, 221, 251, 268
- idirafter 62
- idirafter dir 267
- IEEE compliant code 86
- IEEE floating point standard 86
- IEEE rounding mode 87
- imacros 62
- immediate integer operand 182
- implementation of object definition in headers 202
- include 62
- index values 160
- indexing 148, 160
- Inetl x86 options
 - mpush-args 105
- infinite recursion 227
- inheritance 227
- inherited variables 149
- initialization functions 325
- initialized global definition 172
- initializer 152
- inline 28, 38, 176, 196
- input operands 178
- inside a macro 264
- insn-emit.c 299
- Installation 294
- installing 297
 - binaries for Sun 326
 - compiler driver 301
 - Microsoft compilers 299
 - naming scheme 301
 - Ultrix compilers 299
- instantiation of templates 204
- instructions
 - compare 181
 - store 181
 - test 181
- int 154, 167, 243, 253, 350
- integer constants 154, 250, 343
- Intel
 - i386, i486, i586, i686 102
 - instructions, scheduling 102, 105
 - integer registers 103
 - leaf procedures 106
 - Pentium, or Pentium Pro 102
 - SSE (Streaming SIMD Extension) 104
- Intel 960 options 105
 - m 105
 - masm-compat 105
 - mclean-linkage 105

- mcode-align 106
 - mcomplex-addr 106
 - mic2.0-compat 106
 - mic3.0-compat 106
 - mic-compat 106
 - mintel-asm 106
 - mleaf-procedures 106
 - mno-code-align 106
 - mno-complex-addr 106
 - mno-leaf-procedures 106
 - mno-old-align 106
 - mno-strict-align 106
 - mno-tail-call 106
 - mnumerics 106
 - mold-align 106
 - msoft-float 106
 - mstrict-align 106
 - mtail-call 106
 - Intel compiler options 102
 - Intel x386 options
 - mno-ieee-fp 102
 - Intel x86 options
 - m386 102
 - m486 102
 - maccumulate-outgoing-args 105
 - malign-double 103
 - malign-functions= 104
 - malign-jumps= 104
 - malign-loops= 104
 - march= 102
 - mcpu= 102
 - mieee-fp 102
 - minline-all-stringops 105
 - mno-align-double 103
 - mno-align-stringops 105
 - mno-fancy-math-387 102
 - mno-fp-ret-in-387 102
 - mno-svr3-shlib 103
 - mno-wide-multiply 103
 - mpentium 102
 - mpentiumpro 102
 - mpreferred-stack-boundary= 104
 - mreg-alloc= 103
 - mregparm= 104
 - mrtd 103
 - msoft-float 102
 - msvr3-shlib 103
 - mthreads 105
 - mwide-multiply 103
 - intellectual property 291
 - interface specification 202
 - intermediate 149
 - internal linkage 150
 - interpreter function 148
 - invalid assembler code 193
 - invalid C code 242
 - iprefix 62
 - iprefix prefix 267
 - IRIX thread support 302
 - ISC 295
 - ISO/ANSI C standard 353
 - isystem 63
 - isystem dir 267
 - isystem-c++ 63
 - iwithprefix 62
 - iwithprefix dir 267
 - iwithprefixbefore 62
- ## J
- JALR instruction, for MIPS 109
 - jmp_buf 338
- ## K
- k 143
 - K&R C 20
 - K6 options 102
 - keyword with parentheses 162
 - K-series 106
- ## L
- L 72
 - l 67, 138, 143
 - label, with colon or period= syntax 160
 - labeled elements of an initializer 160
 - labels
 - addresses stored in automatic variables 149
 - indexing, with 148
 - values, as constants 148
 - labs 21
 - LANG 139
 - lang-c 267
 - lang-c++ 267
 - lang-c89 267
 - lang-objc 267
 - lang-objc++ 267
 - LANGUAGES=c 298
 - LC_COLLATE 137
 - LC_CTYPE 137
 - LC_MESSAGES 137
 - LC_MONETARY 137
 - LC_NUMERIC 137
 - LC_TIME 137
 - ld 320, 326, 339
 - leaf procedures 106
 - League for Programming Freedom 291
 - less than brackets 225
 - lexical scoping 149
 - lexical units, valid usage 239
 - lgcc 68
 - libg++ 32
 - libgcc.a 321
 - libgcc1.a 321
 - libgcc1.c 321

- libgcc1-test 324
- libraries, linking 67–68, 320
- library 67
- LIBRARY_PATH 138
- libstdc++ 301
- limit.h 235
- line control 218
- linkage 28
- linker 67
- linker errors 204
- linker options
 - c 67
 - E 67
 - l 67
 - lobjc 68
 - nodefaultlibs 68
 - nostartfiles 68
 - nostdlib 68
 - object-file-name 67
 - S 67
 - s 68
 - shared 69
 - static 68
 - symbolic 69
 - u 69
 - Wl 69
 - Xlinker 69
- linking options
 - libraries 67
 - object files 67
- lint 268
- Linux 303, 339
- LL, adding to an integer 154
- lobjc 68
- local labels as identifiers 147
- local register variables 193, 195
- local variables 152, 354
- local version of applications 227
- LOCAL_INCLUDE_DIR 324
- locale information 139
- long 250, 351
- long long int 154
- longjmp 23, 36, 350
- loops 210
 - deleting 361
 - empty 361
 - optimizer 58
 - unrolling 60
- lvalues 153, 158, 169, 178

M

- M 63
- m 105
- m1 90
- m2 90
- m29000 76
- m29050 76

- m3 90
- m32 124
- M32R/D/X options options
 - G 113
 - mcode-model=large 113
 - mcode-model=medium 113
 - mcode-model=small 113
 - mcond-exec= 114
 - msdata=none 113
 - msdata=sdata 113
 - msdata=use 113
- m386 102
- m3e 90
- m4 90
- m4650 110
- m486 102
- m64 124
- m68000 114
- m68020 114
- m68020-40 115
- m68020-60 115
- m68030 114
- m68040 114
- m68060 114
- m68881 114
- M68K options 114
 - m68000 114
 - m68020 114
 - m68020-40 115
 - m68020-60 115
 - m68030 114
 - m68040 114
 - m68060 114
 - m68881 114
 - malign-int 116
 - mbitfield 115
 - mc68000 114
 - mc68020 114
 - mfpa 115
 - mno-align-int 116
 - mno-bitfield 115
 - mno-strict-align 116
 - mpcrel 116
 - mrtld 115
 - mshort 115
 - msoft-float 115
 - mstrict-align 116
- m88000 116
- m88100 116
- m88110 116
- M88K options 116
 - m88000 116
 - m88100 116
 - m88110 116
 - mbig-pic 116
 - mcheck-zero-division 116
 - mhandle-large-shift 117
 - midentify-revision 117
 - mno-check-zero-division 116

- mno-ocs-debug-info 117
- mno-ocs-frame-position 117
- mocs-debug-info 117
- mtrap-large-shift 117
- M88k options
 - mno-optimize-arg-area 117
 - mno-serialize-volatile 117
 - mno-underscores 118
 - moptimize-arg-area 117
 - mserialize-volatile 117
 - mshort-data 118
 - msvr3 118
 - msvr4 118
 - muse-div-instruction 118
 - mversion 119
 - mwarn-passed-structs 119
- mabi=32 111
- mabi=64 111
- mabi=eabi 111
- mabi=n32 111
- mabi=n64 111
- mabicalls 107, 110
- mabort-on-noreturn 82
- maccum 112
- maccumulate-outgoing-args 105
- MACH 303, 315
- Mach 267
- machine dependent options 73
- machine names 304
- macros 229
 - AM29000K series 236
 - appending 231
 - arguments 157, 231, 244
 - assertions 252
 - calls 244, 251
 - combining source files 257
 - compound statements 242
 - concatenation 238
 - conditionals 249, 252
 - controlling file names 225
 - customizing a program 253
 - definitions 3, 147, 217, 233
 - double scan 244
 - errors 247, 264
 - expansion 218, 231, 244
 - functions 231–232
 - input string, character constants 261
 - long definitions 230
 - M68000 series 236
 - min 243
 - Motorola 236
 - names 230, 232
 - nested calls 241, 245
 - newlines 247
 - ns 32000 series 236
 - parentheses 241
 - parentheses, using 233
 - pitfalls, subtleties 240
 - predefined 233, 266
 - predefined, non-standard 236
 - Pyramid 236
 - Sequent 236
 - Sun 236
 - pre-scan 244
 - processor 217
 - quote characters 264
 - recursive 264
 - redefining 240
 - self-referential 243
 - simple 229
 - source files 266
 - spaces, using 233
 - standard 233
 - string or character constants 264
 - stringification 237
 - stringified, concatenated 244
 - system-specific 233
 - undefined 239
 - unintended grouping 242
 - using assertions 254
 - using parentheses 230
 - variables, functions 240
 - Vax 236, 253
 - with # 264
 - with side effect 243
 - with text after 264
- maddac3 112
- mads 99
- main 325
- maix32 95
- maix64 95
- make 63
 - environmental variables 338
- make install 301
- make LANGUAGES=c 298
- Makefile 297
- makefile variable 302
- malign-300 90
- malign-double 103
- malign-functions= 104, 112, 123
- malign-int 116
- malign-jumps= 104, 111, 123
- malign-loops= 104, 111, 123
- malignment-traps 80
- malloc 308, 349
- malph-as 86
- mam33 107
- mapcs 78
 - mapcs-26 78
 - mapcs-32 78
 - mapcs-float 79
 - mapcs-frame 78
 - mapcs-reentrant 79
 - mapcs-stack-check 79
- mapp-regs 120–121
- march= 81, 90, 102
- masm-compat 105
- masm-optimize 112

- matching quote characters 230
- matpcs 79, 84
- Matsushita AM33 options 107
 - mam33 107
 - mmult-bug 107
 - mno-am33 107
 - mno-mult-bug 107
 - mrelax 107
- Matsushita MN10200 options 107
 - mrelax 107
- Matsushita MN10300 options 107
 - mmult-bug 107
 - mno-mult-bug 107
 - mrelax 107
- maximum arguments in C operators 201
- mb 90
- mbig 98
- mbig-endian 80, 84, 98
- mbig-pic 116
- mbig-switch 91, 120
- mbit-align 97
- mbitfield 115
- mbranch-cost= 113
- mbroken-saverestore 123
- mbsd 82
- mbuggy-return-in-memory 79
- mbuild-constants 88
- mbw 76
- mbwx 89
- mc300 85
- mc400 85
- mc68000 114
- mc68020 114
- mcall-aix 98
- mcallee-super-interworking 84
- mcaller-super-interworking 84
- mcall-lib-mul 101
- mcall-linux 98
- mcall-solaris 98
- mcall-sysv 98
- mcall-sysv-eabi 98
- mcall-sysv-noeabi 98
- mcheck-zero-division 116
- mcix 89
- mclean-linkage 105
- mcmode= 124
- mcode-align 106
- mcode-model=large 113
- mcode-model=medium 113
- mcode-model=small 113
- mcomplex-addr 106
- mcond-exec= 113–114
- mcond-move 112
- mcpu= 78, 81, 88, 92, 102, 108, 120, 123
- mcpu=403 93
- mcpu=505 93
- mcpu=601 93
- mcpu=602 93
- mcpu=603 93
- mcpu=603e 93
- mcpu=604 93
- mcpu=620 93
- mcpu=821 93
- mcpu=860 93
- mcpu=common 93
- mcpu=power 93
- mcpu=power2 93
- mcpu=powerpc 93
- mcpu=rios1 93
- mcpu=rios2 93
- mcpu=rsc 93
- mcypress 123
- MD 63
- mdalign 90
- mdata= 78
- mdisable-callt 120
- mdisable-fpregs 91
- mdisable-indexing 91
- mdouble32 112
- mdouble64 112
- mdouble-float 110
- meabi 99
- memb 99
- membedded-data 108
- membedded-pic 108
- member functions 207
- memcmp 21
- memcpy 21
- memory model (29K) 77
- memory operand 181
- mep 119
- mepilogue 122
- MetaWare 315
- MetaWare HighC (hc) compiler 101
- mfast-indirect-calls 91
- mflat 122
- mfp32 108
- mfp64 108
- mfpa 115
- mfp-arg-in-fpregs 101
- mfp-arg-in-gregs 101
- mfpe= 81
- mfp-reg 85
- mfp-rounding-mode= 87
- mfp-trap-mode= 86
- mfpu 121
- mfull-fp-blocks 101
- mfull-toc 95
- mfused-madd 97
- MG 63
- mgas 86, 91, 108
- mgp32 108
- mgp64 108
- mgpopt 108
- mh 89
- mhalf-pic 109
- mhandle-large-shift 117
- mhard-float 81, 96, 109, 121

-
- mhard-quad-float 121
 - mhc-struct-return 101
 - mic2.0-compat 106
 - mic3.0-compat 106
 - mic-compat 106
 - midentify-revision 117
 - mieee 86
 - mieee-conformant 86
 - mieee-fp 102
 - mieee-with-inexact 86
 - min, unsafe macro 243
 - minimum arguments in C operators 201
 - minline-all-stringops 105
 - min-line-mul 101
 - mint16 112
 - mint32 89, 112
 - mint64 109
 - mintel-asm 106
 - MIPS 296, 341
 - BSD mode 313
 - cross-compiling 319
 - JALR instructions 109
 - requisite libraries for GNU CC 110
 - switch statements 313
 - MIPS options
 - EB 111
 - EL 111
 - G 111
 - m4650 110
 - mabi=32 111
 - mabi=64 111
 - mabi=eabi 111
 - mabi=n32 111
 - mabi=n64 111
 - mabiccalls 107, 110
 - malign-functions= 112
 - malign-jumps= 111
 - malign-loops= 111
 - mcpu= 108
 - mdouble-float 110
 - membedded-data 108
 - membedded-pic 108
 - mfix7000
 - mfix7000 111
 - mfp32 108
 - mfp64 108
 - mgas 108
 - mgp32 108
 - mgp64 108
 - mgpopt 108
 - mhalf-pic 109
 - mhard-float 109
 - mint64 109
 - mips1 109
 - mips2 109
 - mips3 109
 - mips4 109
 - mlong32 109
 - mlong64 109
 - mlong-calls 109
 - mmad 111
 - mmax-skip-functions= 112
 - mmax-skip-jumps= 112
 - mmax-skip-loops= 112
 - mmemcpy 109
 - mmips-as 109
 - mmips-tfile 110
 - mno-abiccalls 107, 110
 - mno-crt0 111
 - mno-embedded-data 108
 - mno-embedded-pic 108
 - mno-gpopt 108
 - mno-half-pic 109
 - mno-long-calls 109
 - mno-mad 111
 - mno-memcpy 109
 - mno-mips-tfile 110
 - mno-rnames 110
 - mno-split-addresses 110
 - mno-stats 110
 - mrnames 110
 - msingle-float 110
 - msoft-float 110
 - msplit-addresses 110
 - mstats 110
 - muninit-const-in-rodata 108
 - nocpp 111
 - mips1 109
 - mips2 109
 - mips3 109
 - mips4 109
 - mips-mips-riscosrev 313
 - mips-mips-riscosrevbsd 314
 - mips-mips-riscosrevsysv 314
 - mips-mips-riscosrevsysv4 314
 - mips-sgi-irix5.* 295
 - miscellaneous preprocessing directives 259
 - Mitsubishi 112
 - Mitsubishi D10V options 112
 - maccum 112
 - maddac3 112
 - masm-optimize 112
 - mbranch-cost= 113
 - mcond-exec= 113
 - mcond-move 112
 - mdouble32 112
 - mdouble64 112
 - mint16 112
 - mint32 112
 - mno-accum 112
 - mno-addac3 112
 - mno-asm-optimize 112
 - mno-cond-move 112
 - mno-small-insn 113
 - msmall-insn 113
 - Mitsubishi M32R/D/X options 113
 - Mitsubishi M32R/D/X options options
 - G 113

- mcode-model=large 113
- mcode-model=medium 113
- mcode-model=small 113
- mcond-exec= 114
- msdata=none 113
- msdata=sdata 113
- msdata=use 113
- mjump-in-delay 91
- mktemp 350
- ML signatures 207
- mleaf-procedures 106
- mlittle 98
- mlittle-endian 80, 84, 98, 123–124
- mlive-g0 123
- mlong32 109
- mlong64 109
- mlong-calls 82, 109, 119
- mlong-load-store 91
- MM 63
- mmad 111
- mmangle-cpu 78
- mmax 89
- mmax-skip-functions= 112
- mmax-skip-jumps= 112
- mmax-skip-loops= 112
- MMD 63
- mmemcpy 109
- mmemory-latency= 89
- mmminimal-toc 95
- mmminimum-fp-blocks 101
- mmips-as 109
- mmips-tfile 110
- mmpc860c0= 99
- mmult-bug 107
- mmultiple 96
- mmvme 99
- MN10300 options
 - mmult-bug 107
 - mno-mult-bug 107
 - mrelax 107
- mnbw 76
- mnew-mnemonics 93–94
- mno-abicalls 107, 110
- mnoabort-on-noreturn 82
- mno-accum 112
- mno-addac3 112
- mno-align-double 103
- mno-align-int 116
- mno-alignment-traps 80
- mno-align-stringops 105
- mno-am33 107
- mno-apcs-float 79
- mno-apcs-frame 78
- mno-apcs-leaf-frame 84
- mno-apcs-reentrant 79
- mno-apcs-stack-check 79
- mno-app-regs 120–121
- mno-asm-optimize 112
- mno-atpcs 79, 84
- mno-big-switch 91
- mno-bit-align 97
- mno-bitfield 115
- mno-buggy-return-in-memory 79
- mno-bwx 89
- mno-callee-super-interworking 84
- mno-caller-super-interworking 84
- mno-check-zero-division 116
- mno-cix 89
- mno-code-align 106
- mno-complex-addr 106
- mno-cond-move 112
- mno-crt0 111
- mno-disable-callt 120
- mno-disable-fpregs 91
- mno-disable-indexing 91
- mno-eabi 99
- mno-embedded-data 108
- mno-embedded-pic 108
- mno-ep 119
- mno-epilogue 122
- mno-fancy-math-387 102
- mno-flat 122
- mno-fp-in-toc 95
- mno-fp-regs 85
- mno-fp-ret-in-387 102
- mno-fpu 121
- mno-fused-madd 97
- mno-gas 91
- mno-gpopt 108
- mno-half-pic 109
- mnohc-struct-return 101
- mno-ieee-fp 102
- mno-jump-in-delay 91
- mno-leaf-procedures 106
- mno-long-calls 82, 109, 119
- mno-long-load-store 91
- mno-mad 111
- mno-max 89
- mno-memcpy 109
- mno-mips-tfile 110
- mno-mult-bug 107
- mno-multiple 96
- mno-ocs-debug-info 117
- mno-ocs-debug-info -mocs-frame-position 117
- mno-ocs-frame-position 117
- mno-old-align 106
- mno-optimize-arg-area 117
- mnop-fun-dllimport 83
 - THUMB options
 - mnop-fun-dllimport 84
- mno-pic-register= 85
- mnop-nop-fun-dllimport 83
 - THUMB options
 - mnop-nop-fun-dllimport 84
- mno-portable-runtime 91
- mno-power 94
- mno-power2 94
- mno-powerpc 94

- mno-powerpc-gfxopt 94
- mno-powerpc-gpopt 94
- mno-prolog-function 119
- mno-prototype 98
- mno-regnames 100
- mno-relocatable 97
- mno-relocatable-lib 97
- mno-rnames 110
- mno-sched-prolog 80
- mno-sdata 100
- mno-serialize-volatile 117, 347
- mno-short-load-bytes 82
- mno-short-load-words 82
- mno-single-pic-base 83
 - THUMB options
 - mno-single-pic-base 85
- mno-small-insn 113
- mno-soft-float 85, 92
- mno-space-regs 92
- mno-split-addresses 110
- mno-stack-bias 124
- mno-stats 110
- mno-strict-align 97, 106, 116
- mno-string 96
- mno-sum-in-toc 95
- mno-svr3-shlib 103
- mno-symrename 83
- mno-tail-call 106
- mno-thumb-interwork 84
- mno-toc 98
- mno-tpcs-frame 83
- mno-tpcs-leaf-frame 83
- mno-traceback 100
- mno-unaligned-doubles 122
- mno-underscores 118
- mno-update 97
- mno-wide-multiply 103
- mno-xl-call 96
- mnumerics 106
- mocs-debug-info 117
- mode 171
- Modula-2, Modula-3 207
- modules, optimizing 210
- mold-mnemonics 94
- moptimize-arg-area 117
- Motorola 236
- Motorola 68000 options 114
 - m68000 114
 - m68020 114
 - m68020-40 115
 - m68020-60 115
 - m68030 114
 - m68040 114
 - m68060 114
 - m68881 114
 - malign-int 116
 - mbitfield 115
 - mc68000 114
 - mc68020 114
 - mfpa 115
 - mno-align-int 116
 - mno-bitfield 115
 - mno-strict-align 116
 - mpcrel 116
 - mrtd 115
 - mshort 115
 - msoft-float 115
 - mstrict-align 116
- Motorola 88000 options 116
 - m88000 116
 - m88100 116
 - m88110 116
 - mbig-pic 116
 - mcheck-zero-division 116
 - mhandle-large-shift 117
 - midentify-revision 117
 - mno-check-zero-division 116
 - mno-ocs-debug-info 117
 - mno-ocs-frame-position 117
 - mno-optimize-arg-area 117
 - mno-serialize-volatile 117
 - mno-underscores 118
 - mocs-debug-info 117
 - moptimize-arg-area 117
 - mserialize-volatile 117
 - mshort-data 118
 - msvr3 118
 - msvr4 118
 - mtrap-large-shift 117
 - muse-div-instruction 118
 - mversion 119
 - mwarn-passed-structs 119
- Object Compatibility Standard (OCS) 117
 - OCS 117
 - mpa-risc-1-0 92
 - mpa-risc-1-1 92
 - mpa-risc-2-0 92
 - mpcrel 116
 - mpe 96
 - mpentium 102
 - mpentiumpro 102
 - mpic-register= 83, 85
 - mportable-runtime 92
 - mpowepc 93
 - mpower 93–94
 - mpower2 94
 - mpowerpc 94
 - mpowerpc-gfxopt 94
 - mpowerpc-gpopt 94
 - mpreferred-stack-boundary= 104
 - mprolog-function 119
 - mprototype 98
 - mpush-args 105
- MQ register, for PowerPC 94
- mreg-alloc= 103
- mregnames 100
- mregparm= 104
- mrelax 89–90, 107

- mrelocatable 97
- mrelocatable-lib 97
- mrnames 110
- mrodata= 78
- mrtd 103, 115
- ms 89
- msa1110-fix 83
- msched-prolog 80
- mschedule= 92
- msda= 120
- msdata 99–100
- msdata=default 100
- msdata=eabi 100
- msdata=none 100, 113
- msdata=sdata 113
- msdata=sysv 100
- msdata=use 113
- msdata-data 100
- mserialize-volatile 117
- mshort 115
- mshort-data 118
- mshort-load-bytes 82
- mshort-load-words 82
- msim 99
- msingle-float 110
- msingle-pic-base 83, 85
- msmall-insn 113
- msoft-float 81, 85, 92–93, 96, 102, 106, 110, 115, 121
- msoft-quad-float 122
- mspace 119
- mspace-regs 92
- msparlite 122
- msplit-addresses 110
- mstack-bias 124
- mstats 110
- mstrict-align 97, 106, 116
- mstring 96
- mstructure-size-boundary= 81, 84
- msupersparc 123
- msvr3 118
- msvr3-shlib 103
- msvr4 118
- mtail-call 106
- mtda= 120
- mtext= 78
- mthreads 96, 105
- mthumb-interwork 81, 84
- mtoc 98
- mtpcs-frame 83
- mtpcs-leaf-frame 83
- mtraceback 100
- mtrap-large-shift 117
- mtrap-precision= 87
- mtune= 93, 121
- multiple alignments 173
- multiple alternative operands 184
- multiple attributes 167, 172
- multiple basic blocks, branches 212

- multiple inheritance 26
- munaligned-doubles 122
- muninit-const-in-rodata 108
- mupdate 97
- muse-div-instruction 118
- mv8 122
- mv850 120
- mv850e 120
- mversion 119
- mwarn-passed-structs 119
- mwide-multiply 103
- mwords-little-endian 80
- mxl-call 96
- mxopen 82
- myellowknife 99
- mzda= 120

N

- N 143
- n 143
- name mangling 29, 335
- names, numbers, concatenation 239
- namespace 27
- National Semiconductor ns32000 315
- NEC V850 options 119
 - __v850e__ 120
 - mapp-regs 120
 - mbig-switch 120
 - mdisable-callt 120
 - mep 119
 - mlong-calls 119
 - mno-app-regs 120
 - mno-disable-callt 120
 - mno-ep 119
 - mno-long-calls 119
 - mno-prolog-function 119
 - mprolog-function 119
 - msda= 120
 - mspace 119
 - mtda= 120
 - mv850 120
 - mv850e 120
 - mzda= 120
 - preprocessor constants 120
- NEC V850 preprocessor constants
 - __v850__, __v850__, __v850e__ 120
- NEC V850e 120
- nested conditionals 251
- nested function 149
- nested loops 147
- nested type declarations 207
- NeXT operating system 342
- NEXTSTEP 303
 - nfp 295
 - nfp 296
- nocommon 171
- nocpp 111

- nodefaultlibs 68
- non-ANSI constructs, warnings and errors 362
- non-constant initializers 159
- noncontiguous complex variable 155
- non-GNU linker 129
- non-lvalue, subscripting 158
- nontemplate function 43
- non-void function value 359
- nostartfiles 68
- nostdinc 29, 224
- nostdlib 68
- ns16000 253
- null directive 242, 259
- number/pound symbol 23

O

- O 60, 340
- O0 60
- O1 60
- O2 60
- O3 60
- OBJC_INCLUDE_PATH 139
- OBJC_THREAD_FILE 302
- object definitions 202
- object files, linking 67
- Objective C 268, 302
- objects, temporary 356
- offset 149
- old-style (C-style) cast 44
- old-style function definitions 143, 168
- old-style non-prototype definition 167
- omitted operands, in conditionals 154
- once-only include files 226
- operands 147, 154
- operators, minimum, maximum 201
- optimiization options
 - falign-labels 53
 - falign-labels= 53
- optimization 9, 58, 209
- optimization options
 - falign-functions 53
 - falign-functions= 53
 - falign-jumps 54
 - falign-jumps= 54
 - falign-loops 54
 - falign-loops= 54
 - fbranch-probabilities 54
 - fcaller-saves 54
 - fcse-follow-jumps 54
 - fcse-skip-blocks 54
 - fdata-sections 56
 - fdelayed-branch 54
 - fdelete-null-pointer-checks 55
 - fexpensive-optimizations 55
 - ffast-math 55
 - ffloat-store 55
 - fforce-addr 55
 - fforce-mem 55
 - ffunction-sections 56
 - fgcse 56
 - finline-functions 56
 - finline-limit= 56
 - fkeep-inline-functions 56
 - fkeep-static-consts 57
 - flive-range 54
 - fno-default-inline 57
 - fno-defer-pop 57
 - fno-function-cse 57
 - fno-inline 57
 - fno-math-errno 55
 - fno-peephole 57
 - fomit-frame-pointer 57
 - foptimize-register-moves 58
 - foptimize-sibling-calls 57
 - freduce-all-givs 58
 - fregmove 58
 - frerun-cse-after-loop 58
 - frerun-loop-opt 58
 - fschedule-insns 58
 - fschedule-insns2 58
 - fssa 58
 - fstrength-reduce 58
 - fstrict-aliasing 58
 - fthread-jumps 59
 - funroll-all-loops 59
 - funroll-loops 59
 - O 60
 - O0 60
 - O1 60
 - O2 60
 - O3 60
 - Os 60
 - optimization options-fmove-all-movables 56
 - options 296
 - assembler 65, 295
 - COFF 296
 - compiling 142
 - gcc 142
 - Os 60
 - OS/2 thread support 303
 - OSF/1 108, 302
 - output file option 17
 - output operand expressions 178
 - output operands, write-only 178

P

- P 63
- p 51, 143
- PA systems and kernels 92
- packed 171
- parameter forward declaration 157
- parentheses 146–147, 167, 230, 241
- parser files 257
- parsing 219

-
- patches for bugs 369
 - Patents, Trademarks and Copyrights 292
 - PCC 351
 - PCThreads 303
 - PE (Parallel Environment) 96
 - pedantic 35, 38, 145, 196, 218, 265, 362
 - pedantic-errors 26, 36, 265, 338
 - Perl 349
 - pg 51
 - PIC 108, 130
 - PIC addressing 83, 85
 - pointer 171
 - pointer-to-function 357
 - Portable C Compiler 101
 - porting 225
 - position independent code 79
 - position-independent code 130
 - POSIX thread support 303
 - pound/number symbol 23
 - POWER 94
 - POWER instruction set 94
 - PowerPC 94
 - configuration 316
 - EABI options 99
 - ELF options 99
 - General Purpose group 94
 - instruction set 94
 - MQ register 94
 - running System Version4 316
 - running Windows NT 317
 - Table Of Contents, executable files 95
 - threads 96
 - PowerPC options 92
 - G 100
 - linker, using with MPC860 99
 - mads 99
 - maix32 95
 - maix64 95
 - mbig 98
 - mbig-endian 98
 - mbit-align 97
 - mcall-aix 98
 - mcall-linux 98
 - mcall-solaris 98
 - mcall-sysv 98
 - mcall-sysv-eabi 98
 - mcall-sysv-noeabi 98
 - mcpu= 92
 - mcpu=403 93
 - mcpu=505 93
 - mcpu=601 93
 - mcpu=602 93
 - mcpu=603 93
 - mcpu=603e 93
 - mcpu=604 93
 - mcpu=620 93
 - mcpu=821 93
 - mcpu=860 93
 - mcpu=common 93
 - mcpu=power 93
 - mcpu=power2 93
 - mcpu=powerpc 93
 - mcpu=rios1 93
 - mcpu=rios2 93
 - mcpu=rscc 93
 - meabi 99
 - memb 99
 - mfull-toc 95
 - mfused-madd 97
 - mhard-float 96
 - mlittle 98
 - mlittle-endian 98
 - mminimal-toc 95
 - mmpc860c0= 99
 - mmultiple 96
 - mmvme 99
 - mnew-mnemonics 94
 - mno-bit-align 97
 - mno-eabi 99
 - mno-fp-in-toc 95
 - mno-fused-madd 97
 - mno-multiple 96
 - mno-power 94
 - mno-power2 94
 - mno-powerpc 94
 - mno-powerpc-gfxopt 94
 - mno-powerpc-gpopt 94
 - mno-prototype 98
 - mno-regnames 100
 - mno-relocatable 97
 - mno-relocatable-lib 97
 - mno-sdata 100
 - mno-strict-align 97
 - mno-string 96
 - mno-sum-in-toc 95
 - mno-toc 98
 - mno-traceback 100
 - mno-update 97
 - mno-xl-call 96
 - mpower 93-94
 - mpower2 94
 - mpowerpc 93-94
 - mpowerpc-gfxopt 94
 - mpowerpc-gpopt 94
 - mprototype 98
 - mregnames 100
 - mrelocatable 97
 - mrelocatable-lib 97
 - msdata 99-100
 - msdata=default 100
 - msdata=eabi 100
 - msdata=none 100
 - msdata=sysv 100
 - msdata-data 100
 - msim 99
 - msoft-float 93, 96
 - mstrict-align 97
 - mstring 96

- mthreads 96
- mtoc 98
- mtraceback 100
- mtune= 93
- mupdate 97
- mxl-call 96
- myellowknife 99
- specs= 96
- pragmas, warning of unknown 30, 42
- predecrement addressing 181
- predefined macros 145
- prefix 267
- preincrement addressing 181
- preprocessing 351
- preprocessing conditionals 250
- preprocessing directives 251
- preprocessing number 264
- preprocessor 24
- preprocessor options 61
 - ## in rest arguments 158
 - A 61
 - C 61
 - D 62
 - dD 61
- directive name 221
- dM 62
- dN 62
- E 61
- H 62
- idirafter 62
- imacros 62
- include 62
- iprefix 62
- isystem 63
- isystem-c++ 63
- iwithprefix 62
- iwithprefixbefore 62
- M 63
- MD 63
- MG 63
- MM 63
- MMD 63
- nostdinc 63
- P 63
- trigraphs 63
- U 64
- undef 64
- Wp 64
- pre-scan 244
 - nested calls 245
 - self-referent macros 245
 - unshielded commas 245
- print-file-name= 51
- print-libgcc-file-name 52
- print-search-dirs 52
- private data 130
- problems iii, 363–364
- processor selection (29K) 76
- prof 51

- profiling 45–46, 344
- profiling tools 9, 209
- program checker 268
- program.c 225
- protoize 141–142, 357
- prototypes 141
- PSIM simulator 316
- pthreads 96
- purify 339
- Pyramid 236

Q

- Q 52
- q 143
- Qn 125
- qsort 194
- quotes 225, 230
- Qy 125

R

- r 200
- r4650 110
- ranlib 320
- read-write operand 178
- real.c 339
- REAL_LD_FILE_NAME 325
- real-ld 325
- Red Hat, contacting iii
- redefinition 240
- reentrant code 79
- references 68
- REG_CLASS_FROM_LETTER 185
- register transfer language (RTL) 48
- REGISTER_NAMES 128
- registers 179, 193, 195
 - MQ 94
- reporting a bug 363
- rest arguments 157
- RISC iX options 82
- RISC-OS, reconfiguring 314
- RS/600
 - threads 96
- RS/6000 344
- RS/6000 options 92
 - mcpu= 92
 - mcpu=403 93
 - mcpu=505 93
 - mcpu=601 93
 - mcpu=602 93
 - mcpu=603 93
 - mcpu=603e 93
 - mcpu=604 93
 - mcpu=620 93
 - mcpu=821 93
 - mcpu=860 93

- mcpu=common 93
- mcpu=power 93
- mcpu=power2 93
- mcpu=powerpc 93
- mcpu=rios1 93
- mcpu=rios2 93
- mcpu=rsc 93
- mpe 96
- mpower 93
- mpowerpc 93
- msoft-float 93
- mtune= 93
- rtd 115
- RTL (register transfer language) 48
- RTTI 26
- run-time varying elements 159

S

- S 67
- s 68
- saveset
 - for VMS 327
- save-temps 52
- scalar types 159
- scope
 - unallowed jumping 156
- scope, of variables 26
- Scratchpad II 207
- SDB 51
- search options
 - B 71
 - I 72
 - I- 72
 - L 72
 - specs= 72
- searching 71
- second include path 62
- self-reference, indirect 244
- self-referents 243
- semicolon, following an expression 147
- semicolon, swallowing 242
- semicolons in GNU assembler 179
- Senate Subcommittee on Patents, Trademarks and Copyrights 292
- Sequent 236
- setjmp 350
- sh 340
- SH options 90
- shared 69
- shared data 130
- shift count operands 359
- short int 32
- side effect 243
- signal.h 227
- signature 207
- signed 196
- signed integer 154
- Silicon Graphics, compiling GNU CC on IRIX 314
- simple constraints 181
- sin 21
- sizeof 152, 158, 169, 253
- SKIP_SPACES 242
- smallest addressable unit 305
- Solaris thread support 303
- Solaris, installing GNU CC 326
- SONY, compiling 315
- source file 239
- source line control 218
- SPARC options
 - malign-functions= 123
 - malign-jumps= 123
 - malign-loops= 123
 - mapp-regs 121
 - mcpu= 120
 - mcypress 123
 - mepilogue 122
 - mflat 122
 - mfpu 121
 - mhard-float 121
 - mhard-quad-float 121
 - mno-app-regs 121
 - mno-epilogue 122
 - mno-flat 122
 - mno-fpu 121
 - mno-stack-bias 124
 - mno-unaligned-doubles 122
 - msoft-float 121
 - msoft-quad-float 122
 - msparclite 122
 - mstack-bias 124
 - msupersparc 123
 - mtune= 121
 - munaligned-doubles 122
 - mv8 122
- SPARC V (options 124
- SPARC V9 options
 - m32 124
 - m64 124
 - mcmodel= 124
 - mlittle-endian 124
- SPARClet options 123
 - mbroken-saverstore 123
 - mlittle-endian 123
 - mlive-g0 123
- SPARClite code 123
- specify -ansi 268
- specifying targets 73
- specs= 72, 96
- sqrt 21
- SSA (Static Single Assignment) 58
- SSE (Streaming SIMD Extension) 104
- stabs 50, 296
- stabs debugging output 51
- standard C extensions 145
- start files 321
- STARTFILE_SPEC 321

- static 68
- static data members, declaring, defining 356
- Static Single Assesment (SSA) 58
- static variable 147, 193
- std and namespace 27
- stdcall 103
- stdio.h 341
- store 149
- store instructions 180
- storem bug (29) 77
- strcmp 21
- strcpy 21
- Streaming SIMD Extention 104
- strfunc 356
- string constants 24, 220, 238, 258, 261
- string variables 197
- stringification 237
- strlen 21
- StrongARM options
 - march= 81
 - mcpu= 81
 - mfpe= 81
 - msa1110-fix 83
 - mstructure-size-boundary= 81
- struct mumble 354
- STRUCT_VALUE 352
- structure and union returning 352
- structure initializer 160
- structure types, defining 226
- structures, arraying 238
- submodel options 76
- subscript 148
- subtraction operations 158
- suffixes 32
- Sun 236–237
- Sun SPARC options 120
- SVr4
 - default debugging 117
- switch statements 148
- symbolic 69
- symbolic links
 - unsupported 297
- symbols, assembler 193
- sys/signal.h 227
- System V 341, 349
- System V options 124
 - G 124
 - mbig 98
 - mbit-align 97
 - mcall-linux 98
 - mcall-solaris 98
 - mcall-sysv-eabi 98
 - mcall-sysv-noeab 98
 - meabi 99
 - mlittle 98
 - mlittle-endian 98
 - mno-bit-align 97
 - mno-eabi 99
 - mno-prototype 98

- mno-regnames 100
- mno-relocatable-lib 97
- mno-toc 98
- mno-traceback 100
- mprototype 98
- mregnames 100
- mrelocatable 97
- mrelocatable-lib 97
- msdata 99–100
- msdata=default 100
- msdata=eabi 100
- msdata=sysv 100
- msdata-data 100
- mstrict-align 97
- mtoc 98
- mtraceback 100
- Qn 125
- Qy 125
- Ym 125
- YP 125
- System, V options
 - mcall-sysv 98

T

- tail-recursive calls 106
- target machine, specifying 73, 294
- target options 73
 - V 73
- tconfig.h 297
- tcov 46
- Tektronix, install 313
- templates 26, 204
- test coverage 209
- test instructions 181
- testsuites 210
- thread implementation 148, 302
- thread support 303
- THUMB options
 - mno-callee-super-interworking> 84
- THUMB
 - leaf functions 84
 - mno-apcs-leaf-frame 84
 - PIC addressing 85
- THUMB options
 - march= 81
 - matpcs 84
 - mbig-endian 84
 - mcallee-super-interworking 84
 - mcaller-super-interworking 84
 - mcpu= 81
 - mfpe= 81
 - mlittle-endian 84
 - mno-atpcs 84
 - mno-caller-super-interworking 84
 - mno-pic-register= 85
 - mno-thumb-interwork 84
 - mno-tpcs-frame 83

- mno-tpcs-leaf-frame 83
- mpic-register= 85
- msingle-pic-base 85
- mstructure-size-boundary= 81, 84
- mthumb-interwork 84
- mtpcs-frame 83
- mtpcs-leaf-frame 83
- Thumb options 83
- THUMB Procedure Call Standard (TPCS) 83
- tilde operator, (~) 155
- time 352
- tm.h 237, 297
- tmp.c.gcov 211
- TMPDIR 138
- tokens, syntactic 238
- TOOL_INCLUDE_DIR 324
- tools cross-compiling 320
- TPCS (THUMB Procedure Call Standard) 83
- trace 46
- traditional 22, 196, 350
- traditional-cpp 23
- trampolines 150
- transparent_union 172–173
- Trellis/Owl type modules 207
- trigraphs 24, 63, 218, 265
- trigraphs 24, 40, 63, 219–220
- type attributes 173
- typedef 152, 174, 226, 338
- typeof 147, 152, 196
- types 357

U

- U 64, 254
- u 69
- U and -D 266
- uid_t 167
- ULL, adding to an integer 154
- unaligned addresses 131
- unary operator, && 148
- unary operators 148
- undef 64, 218
- undefined macros 239
- undefinition 239
- underscores 152, 193
- Unidiff format 371
- uninitialized global variables 129
- union 161, 352
- union type 160
- union types 159, 161
- unknown pragmas, warning 30, 42
- unprotoize 357
- unresolved references 68
- unsigned integer 154
- unsigned long 250
- unsigned long long int 154
- unused 172
- uppercase usage in macro names 230

- using '-g' when reporting bugs 368
- UTEK ns32000 315

V

- V 73
- v 143
- v with GAS 295
- variable attributes 169
- variable names 152
- variable-length arrays 156
- variables 26, 155, 243
- Vax 236–237, 253, 317
- Vax computers 236
- VAX-C 331–332
- VAXCRT library 330
- version of GNU CC 73
- virtual classes 208
- virtual function calls 32, 344
- virtual function definitions in classes 133
- virtual functions 44
- VMS 327, 330
- void and function pointers 158
- volatile 196, 350

W

- W 30, 36
- w 36
- W with -O 351
- Wa 65
- Waggregate-return 37
- Wall 30, 36, 265
- WARN_IF 237
- warning for synthesized methods 31
- warning for unknown pragmas 30, 42
- warning options
 - fsyntax-only 35
 - pedantic 35
 - pedantic-errors 36
 - W 30, 36
 - w 36
 - Waggregate-return 37
 - Wall 30, 36
 - Wbad-function-cast 37, 42
 - Wcast-align 37
 - Wcast-qual 37
 - Wchar-subscripts 37
 - Wcomment 37
 - Wconversion 37
 - Weffc++ 30
 - Werror 37
 - Werror-implicit-function-declaration 38
 - Wfloat-equal 38
 - Wformat 37
 - Wid-clash-len 37
 - Wimplicit 37

- Wimplicit-function-declaration 37
- Wimplicit-int 37
- Winline 38
- Wlarger-than-len 38
- Wlong-long 38
- Wmain 38
- Wmissing-declarations 38
- Wmissing-noreturn 38
- Wmissing-prototypes 42
- Wmultichar 38
- Wnested-externs 42
- Wno-import 38–39
- Wno-non-template-friend 30, 43
- Wnon-virtual-dtor 30
- Wno-pmf-conversions 30
- Wold-style-cast 30, 44
- Woverloaded-virtual 30, 44
- Wpacked 39
- Wpadded 39
- Wparentheses 39
- Wpointer-arith 40
- Wredundant-decls 40
- Wreorder 30, 43
- Wreturn-type 40
- Wshadow 40
- Wsign-compare 40
- Wsign-promo 31
- Wstrict-prototypes 43
- Wswitch 40
- Wsynth 31, 44
- Wtemplate-debugging 43
- Wtraditional 43
- Wtrigraphs 40
- Wundef 40
- Wuninitialized 40
- Wunknown-pragmas 41
- Wunreachable-code 41
- Wunused 41
- Wunused-function 42
- Wunused-label 42
- Wunused-parameter 42
- Wunused-value 42
- Wunused-variable 42
- Wwrite-strings 43
- warnings 44, 265, 361
- Wbad-function-cast 37, 42
- Wcast-align 37
- Wcast-qual 37
- Wchar-subscripts 37
- Wcomment 37, 265
- Wconversion 37, 358
- Wctor-dtor-privacy 30
- we32k-*-* 318
- weak 172
- Web support site iii
- Wefc++ 30
- Werror 37
- Werror-implicit-function-declaration 38
- Wfloat-equal 38
- Wformat 37
- whitespace 221, 238, 240, 242, 261
- Wid-clash-len 37
- Wimplicit 37
- Wimplicit-function-declaration 37
- Wimplicit-int 37
- Win32 API thread support 303
- Winline 38
- with-gnu-as 295
- with-gnu-ld 295
- with-stabs 295
- WI 69
- Wlarger-than-len 38
- Wlong-long 38
- Wmain 38
- Wmissing-declarations 38
- Wmissing-noreturn 38
- Wmissing-prototypes 42
- Wmultichar 38
- Wnested-externs 42
- Wno-deprecated 30
- Wno-import 38–39
- Wno-non-template-friend 30, 43
- Wnon-virtual-dtor 30
- Wno-pmf-conversions 30
- Wold-style-cast 30, 44
- word 171
- Woverloaded-virtual 30, 44
- Wp 64
- Wpacked 39
- Wpadded 39
- Wparentheses 39
- Wpointer-arith 40, 158
- Wredundant-decls 40
- Wreorder 30, 43
- Wreturn-type 40
- Wshadow 40
- Wsign-compare 40
- Wsign-promo 31
- Wstrict-prototypes 43
- Wswitch 40
- Wsynth 31, 44
- Wtemplate-debugging 43
- Wtraditional 43, 265
- Wtrigraphs 40, 265
- Wundef 40
- Wuninitialized 40
- Wunknown-pragmas 30, 41–42
- Wunreachable-code 41
- Wunused 41
- Wunused-function 42
- Wunused-label 42
- Wunused-parameter 42
- Wunused-value 42
- Wunused-variable 42
- Wwrite-strings 43

X

X11 header files 349

x386 options

-mno-ieee-fp 102

x86 options

-m386 102

-m486 102

-maccumulate-outgoing-args 105

-malign-double 103

-malign-functions= 104

-malign-jumps= 104

-malign-loops= 104

-march= 102

-mcpu= 102

-mieee-fp 102

-minline-all-stringops 105

-mno-align-double 103

-mno-align-stringops 105

-mno-fancy-math-387 102

-mno-fp-ret-in-387 102

-mno-svr3-shlib 103

-mno-wide-multiply 103

-mpentium 102

-mpentiumpro 102

-mpreferred-stack-boundary= 104

-mpush-args 105

-mreg-alloc= 103

-mregparm= 104

-mrtd 103

-msoft-float 102

-msvr3-shlib 103

-mthreads 105

-mwide-multiply 103

XCOFF 50

XCOFF format 51

-Xlinker 69

Y

-Ym 125

-YP 125

system V options

-mno-strict-align 97

Z

zero-length arrays 156

