

GNU Compiler Collection Internals

Richard M. Stallman

Last updated 28 December 2002

for GCC 3.3.4

For GCC Version 3.3.4

Published by the Free Software Foundation
59 Temple Place—Suite 330
Boston, MA 02111-1307, USA
Last printed April, 1998.
Printed copies are available for \$50 each.

Copyright © 1988, 1989, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, 2003, 2004 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with the Invariant Sections being “GNU General Public License” and “Funding Free Software”, the Front-Cover texts being (a) (see below), and with the Back-Cover Texts being (b) (see below). A copy of the license is included in the section entitled “GNU Free Documentation License”.

(a) The FSF’s Front-Cover Text is:

A GNU Manual

(b) The FSF’s Back-Cover Text is:

You have freedom to copy and modify this GNU Manual, like GNU software. Copies published by the Free Software Foundation raise funds for GNU development.

Short Contents

Introduction	1
1 Contributing to GCC Development	3
2 GCC and Portability	5
3 Interfacing to GCC Output	7
4 Language Front Ends in GCC	9
5 Source Tree Structure and Build System	11
6 Passes and Files of the Compiler	31
7 Trees: The intermediate representation used by the C and C++ front ends	39
8 RTL Representation	71
9 Machine Descriptions	117
10 Target Description Macros and Functions	193
11 Host Configuration Headers	327
12 Makefile Fragments	331
13 <code>collect2</code>	335
14 Standard Header File Directories	337
15 Memory Management and Type Information	339
Funding Free Software	345
The GNU Project and GNU/Linux	347
GNU GENERAL PUBLIC LICENSE	349
GNU Free Documentation License	355
Contributors to GCC	363
Option Index	375
Index	377

Table of Contents

Introduction	1
1 Contributing to GCC Development	3
2 GCC and Portability	5
3 Interfacing to GCC Output	7
4 Language Front Ends in GCC	9
5 Source Tree Structure and Build System ...	11
5.1 Configure Terms and History	11
5.2 Top Level Source Directory	11
5.3 The ‘gcc’ Subdirectory	12
5.3.1 Subdirectories of ‘gcc’	12
5.3.2 Configuration in the ‘gcc’ Directory	13
5.3.2.1 Scripts Used by ‘configure’	13
5.3.2.2 The ‘config.gcc’ File	14
5.3.2.3 Files Created by configure	14
5.3.3 Build System in the ‘gcc’ Directory	15
5.3.4 Makefile Targets	15
5.3.5 Library Source Files and Headers under the ‘gcc’ Directory	16
5.3.6 Headers Installed by GCC	16
5.3.7 Building Documentation	17
5.3.7.1 Texinfo Manuals	17
5.3.7.2 Man Page Generation	18
5.3.7.3 Miscellaneous Documentation	19
5.3.8 Anatomy of a Language Front End	19
5.3.8.1 The Front End ‘ <i>language</i> ’ Directory ...	20
5.3.8.2 The Front End ‘config-lang.in’ File ..	22
5.3.9 Anatomy of a Target Back End	23
5.4 Test Suites	24
5.4.1 Idioms Used in Test Suite Code	24
5.4.2 C Language Test Suites	25
5.4.3 The Java library test suites	27
5.4.4 Support for testing gcov	28
5.4.5 Support for testing profile-directed optimizations	28
5.4.6 Support for testing binary compatibility	29

6	Passes and Files of the Compiler	31
7	Trees: The intermediate representation used by the C and C++ front ends	39
7.1	Deficiencies	39
7.2	Overview	39
7.2.1	Trees	40
7.2.2	Identifiers	40
7.2.3	Containers	41
7.3	Types	41
7.4	Scopes	46
7.4.1	Namespaces	46
7.4.2	Classes	47
7.5	Declarations	49
7.6	Functions	51
7.6.1	Function Basics	52
7.6.2	Function Bodies	55
7.6.2.1	Statements	55
7.7	Attributes in trees	60
7.8	Expressions	61
8	RTL Representation	71
8.1	RTL Object Types	71
8.2	RTL Classes and Formats	72
8.3	Access to Operands	74
8.4	Flags in an RTL Expression	75
8.5	Machine Modes	80
8.6	Constant Expression Types	84
8.7	Registers and Memory	86
8.8	RTL Expressions for Arithmetic	90
8.9	Comparison Operations	93
8.10	Bit-Fields	94
8.11	Vector Operations	95
8.12	Conversions	96
8.13	Declarations	97
8.14	Side Effect Expressions	97
8.15	Embedded Side-Effects on Addresses	102
8.16	Assembler Instructions as Expressions	103
8.17	Insns	104
8.18	RTL Representation of Function-Call Insns	113
8.19	Structure Sharing Assumptions	113
8.20	Reading RTL	114

9	Machine Descriptions	117
9.1	Overview of How the Machine Description is Used	117
9.2	Everything about Instruction Patterns	117
9.3	Example of <code>define_insn</code>	118
9.4	RTL Template	119
9.5	Output Templates and Operand Substitution	123
9.6	C Statements for Assembler Output	124
9.7	Operand Constraints	126
9.7.1	Simple Constraints	126
9.7.2	Multiple Alternative Constraints	130
9.7.3	Register Class Preferences	131
9.7.4	Constraint Modifier Characters	131
9.7.5	Constraints for Particular Machines	132
9.8	Standard Pattern Names For Generation	143
9.9	When the Order of Patterns Matters	158
9.10	Interdependence of Patterns	158
9.11	Defining Jump Instruction Patterns	160
9.12	Defining Looping Instruction Patterns	161
9.13	Canonicalization of Instructions	163
9.14	Defining RTL Sequences for Code Generation	164
9.15	Defining How to Split Instructions	167
9.16	Including Patterns in Machine Descriptions	170
9.16.1	RTL Generation Tool Options for Directory Search	171
9.17	Machine-Specific Peephole Optimizers	171
9.17.1	RTL to Text Peephole Optimizers	171
9.17.2	RTL to RTL Peephole Optimizers	173
9.18	Instruction Attributes	175
9.18.1	Defining Attributes and their Values	175
9.18.2	Attribute Expressions	176
9.18.3	Assigning Attribute Values to Insns	178
9.18.4	Example of Attribute Specifications	179
9.18.5	Computing the Length of an Insn	180
9.18.6	Constant Attributes	181
9.18.7	Delay Slot Scheduling	182
9.18.8	Specifying processor pipeline description	183
9.18.8.1	Specifying Function Units	184
9.18.8.2	Describing instruction pipeline characteristics	186
9.18.8.3	Drawbacks of the old pipeline description	190
9.19	Conditional Execution	190
9.20	Constant Definitions	191

10	Target Description Macros and Functions	193
10.1	The Global <code>targetm</code> Variable	193
10.2	Controlling the Compilation Driver, ‘gcc’	193
10.3	Run-time Target Specification	200
10.4	Defining data structures for per-function information	204
10.5	Storage Layout	204
10.6	Layout of Source Language Data Types	213
10.7	Target Character Escape Sequences	217
10.8	Register Usage	217
10.8.1	Basic Characteristics of Registers	218
10.8.2	Order of Allocation of Registers	220
10.8.3	How Values Fit in Registers	220
10.8.4	Handling Leaf Functions	222
10.8.5	Registers That Form a Stack	223
10.9	Register Classes	223
10.10	Stack Layout and Calling Conventions	230
10.10.1	Basic Stack Layout	230
10.10.2	Exception Handling Support	233
10.10.3	Specifying How Stack Checking is Done	234
10.10.4	Registers That Address the Stack Frame	236
10.10.5	Eliminating Frame Pointer and Arg Pointer	238
10.10.6	Passing Function Arguments on the Stack	239
10.10.7	Passing Arguments in Registers	241
10.10.8	How Scalar Function Values Are Returned	245
10.10.9	How Large Values Are Returned	246
10.10.10	Caller-Saves Register Allocation	248
10.10.11	Function Entry and Exit	248
10.10.12	Generating Code for Profiling	252
10.10.13	Permitting tail calls	253
10.11	Implementing the Varargs Macros	253
10.12	Trampolines for Nested Functions	256
10.13	Implicit Calls to Library Routines	258
10.14	Addressing Modes	260
10.15	Condition Code Status	264
10.16	Describing Relative Costs of Operations	267
10.17	Adjusting the Instruction Scheduler	271
10.18	Dividing the Output into Sections (Texts, Data, ...)	275
10.19	Position Independent Code	278
10.20	Defining the Output Assembler Language	279
10.20.1	The Overall Framework of an Assembler File	279
10.20.2	Output of Data	281
10.20.3	Output of Uninitialized Variables	284
10.20.4	Output and Generation of Labels	285
10.20.5	How Initialization Functions Are Handled	292
10.20.6	Macros Controlling Initialization Routines	293
10.20.7	Output of Assembler Instructions	296
10.20.8	Output of Dispatch Tables	299

10.20.9	Assembler Commands for Exception Regions . . .	300
10.20.10	Assembler Commands for Alignment	301
10.21	Controlling Debugging Information Format	303
10.21.1	Macros Affecting All Debugging Formats	303
10.21.2	Specific Options for DBX Output	304
10.21.3	Open-Ended Hooks for DBX Format	306
10.21.4	File Names in DBX Format	307
10.21.5	Macros for SDB and DWARF Output	308
10.21.6	Macros for VMS Debug Format	310
10.22	Cross Compilation and Floating Point	310
10.23	Mode Switching Instructions	311
10.24	Defining target-specific uses of <code>__attribute__</code>	313
10.25	Defining coprocessor specifics for MIPS targets.	314
10.26	Miscellaneous Parameters	314
11	Host Configuration Headers	327
12	Makefile Fragments	331
12.1	Target Makefile Fragments	331
12.2	Host Makefile Fragments	333
13	collect2	335
14	Standard Header File Directories	337
15	Memory Management and Type Information	
	339
15.1	The Inside of a <code>GTY()</code>	339
15.2	Marking Roots for the Garbage Collector	342
15.3	Source Files Containing Type Information	342
	Funding Free Software	345
	The GNU Project and GNU/Linux	347
	GNU GENERAL PUBLIC LICENSE	349
	Preamble	349
	TERMS AND CONDITIONS FOR COPYING,	
	DISTRIBUTION AND MODIFICATION	349
	How to Apply These Terms to Your New Programs	354
	GNU Free Documentation License	355
	ADDENDUM: How to use this License for your documents	361

Contributors to GCC	363
Option Index	375
Index	377

Introduction

This manual documents the internals of the GNU compilers, including how to port them to new targets and some information about how to write front ends for new languages. It corresponds to GCC version 3.3.4. The use of the GNU compilers is documented in a separate manual. See section “Introduction” in *Using the GNU Compiler Collection (GCC)*.

This manual is mainly a reference manual rather than a tutorial. It discusses how to contribute to GCC (see Chapter 1 [Contributing], page 3), the characteristics of the machines supported by GCC as hosts and targets (see Chapter 2 [Portability], page 5), how GCC relates to the ABIs on such systems (see Chapter 3 [Interface], page 7), and the characteristics of the languages for which GCC front ends are written (see Chapter 4 [Languages], page 9). It then describes the GCC source tree structure and build system, some of the interfaces to GCC front ends, and how support for a target system is implemented in GCC.

Additional tutorial information is linked to from <http://gcc.gnu.org/readings.html>.

1 Contributing to GCC Development

If you would like to help pretest GCC releases to assure they work well, our current development sources are available by CVS (see <http://gcc.gnu.org/cvs.html>). Source and binary snapshots are also available for FTP; see <http://gcc.gnu.org/snapshots.html>.

If you would like to work on improvements to GCC, please read the advice at these URLs:

<http://gcc.gnu.org/contribute.html>
<http://gcc.gnu.org/contributewhy.html>

for information on how to make useful contributions and avoid duplication of effort. Suggested projects are listed at <http://gcc.gnu.org/projects/>.

2 GCC and Portability

GCC itself aims to be portable to any machine where `int` is at least a 32-bit type. It aims to target machines with a flat (non-segmented) byte addressed data address space (the code address space can be separate). Target ABIs may have 8, 16, 32 or 64-bit `int` type. `char` can be wider than 8 bits.

GCC gets most of the information about the target machine from a machine description which gives an algebraic formula for each of the machine's instructions. This is a very clean way to describe the target. But when the compiler needs information that is difficult to express in this fashion, I have not hesitated to define an ad-hoc parameter to the machine description. The purpose of portability is to reduce the total work needed on the compiler; it was not of interest for its own sake.

GCC does not contain machine dependent code, but it does contain code that depends on machine parameters such as endianness (whether the most significant byte has the highest or lowest address of the bytes in a word) and the availability of autoincrement addressing. In the RTL-generation pass, it is often necessary to have multiple strategies for generating code for a particular kind of syntax tree, strategies that are usable for different combinations of parameters. Often I have not tried to address all possible cases, but only the common ones or only the ones that I have encountered. As a result, a new target may require additional strategies. You will know if this happens because the compiler will call `abort`. Fortunately, the new strategies can be added in a machine-independent fashion, and will affect only the target machines that need them.

3 Interfacing to GCC Output

GCC is normally configured to use the same function calling convention normally in use on the target system. This is done with the machine-description macros described (see Chapter 10 [Target Macros], page 193).

However, returning of structure and union values is done differently on some target machines. As a result, functions compiled with PCC returning such types cannot be called from code compiled with GCC, and vice versa. This does not cause trouble often because few Unix library routines return structures or unions.

GCC code returns structures and unions that are 1, 2, 4 or 8 bytes long in the same registers used for `int` or `double` return values. (GCC typically allocates variables of such types in registers also.) Structures and unions of other sizes are returned by storing them into an address passed by the caller (usually in a register). 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. This is slower than the method used by GCC, and fails to be reentrant.

On some target machines, such as RISC machines and the 80386, the standard system convention is to pass to the subroutine the address of where to return the value. On these machines, GCC has been configured to be compatible with the standard compiler, when this method is used. It may not be compatible for structures of 1, 2, 4 or 8 bytes.

GCC uses the system's standard convention for passing arguments. On some machines, the first few arguments are passed in registers; in others, all are passed on the stack. It would be possible to use registers for argument passing on any machine, and this would probably result in a significant speedup. But the result would be complete incompatibility with code that follows the standard convention. So this change is practical only if you are switching to GCC as the sole C compiler for the system. We may implement register argument passing on certain machines once we have a complete GNU system so that we can compile the libraries with GCC.

On some machines (particularly the SPARC), certain types of arguments are passed "by invisible reference". This means that the value is stored in memory, and the address of the memory location is passed to the subroutine.

If you use `longjmp`, beware of automatic variables. ISO C says that automatic variables that are not declared `volatile` have undefined values after a `longjmp`. And this is all GCC promises to do, because it is very difficult to restore register variables correctly, and one of GCC's features is that it can put variables in registers without your asking it to.

If you want a variable to be unaltered by `longjmp`, and you don't want to write `volatile` because old C compilers don't accept it, just take the address of the variable. If a variable's address is ever taken, even if just to compute it and ignore it, then the variable cannot go in a register:

```
{
    int careful;
    &careful;
```

```
    ...  
}
```

Code compiled with GCC may call certain library routines. Most of them handle arithmetic for which there are no instructions. This includes multiply and divide on some machines, and floating point operations on any machine for which floating point support is disabled with `'-msoft-float'`. Some standard parts of the C library, such as `memcpy` or `memcpy`, are also called automatically. The usual function call interface is used for calling the library routines.

Some of these routines can be defined in mostly machine-independent C; they appear in `'libgcc2.c'`. Others must be hand-written in assembly language for each processor. Wherever they are defined, they are compiled into the support library, `'libgcc.a'`, which is automatically searched when you link programs with GCC.

4 Language Front Ends in GCC

The interface to front ends for languages in GCC, and in particular the `tree` structure (see Chapter 7 [Trees], page 39), was initially designed for C, and many aspects of it are still somewhat biased towards C and C-like languages. It is, however, reasonably well suited to other procedural languages, and front ends for many such languages have been written for GCC.

Writing a compiler as a front end for GCC, rather than compiling directly to assembler or generating C code which is then compiled by GCC, has several advantages:

- GCC front ends benefit from the support for many different target machines already present in GCC.
- GCC front ends benefit from all the optimizations in GCC. Some of these, such as alias analysis, may work better when GCC is compiling directly from source code than when it is compiling from generated C code.
- Better debugging information is generated when compiling directly from source code than when going via intermediate generated C code.

Because of the advantages of writing a compiler as a GCC front end, GCC front ends have also been created for languages very different from those for which GCC was designed, such as the declarative logic/functional language Mercury. For these reasons, it may also be useful to implement compilers created for specialized purposes (for example, as part of a research project) as GCC front ends.

5 Source Tree Structure and Build System

This chapter describes the structure of the GCC source tree, and how GCC is built. The user documentation for building and installing GCC is in a separate manual (<http://gcc.gnu.org/install/>), with which it is presumed that you are familiar.

5.1 Configure Terms and History

The configure and build process has a long and colorful history, and can be confusing to anyone who doesn't know why things are the way they are. While there are other documents which describe the configuration process in detail, here are a few things that everyone working on GCC should know.

There are three system names that the build knows about: the machine you are building on (*build*), the machine that you are building for (*host*), and the machine that GCC will produce code for (*target*). When you configure GCC, you specify these with '`--build=`', '`--host=`', and '`--target=`'.

Specifying the host without specifying the build should be avoided, as `configure` may (and once did) assume that the host you specify is also the build, which may not be true.

If build, host, and target are all the same, this is called a *native*. If build and host are the same but target is different, this is called a *cross*. If build, host, and target are all different this is called a *canadian* (for obscure reasons dealing with Canada's political party and the background of the person working on the build at that time). If host and target are the same, but build is different, you are using a cross-compiler to build a native for a different system. Some people call this a *host-x-host*, *crossed native*, or *cross-built native*. If build and target are the same, but host is different, you are using a cross compiler to build a cross compiler that produces code for the machine you're building on. This is rare, so there is no common way of describing it (although I propose calling it a *crossback*).

If build and host are the same, the GCC you are building will also be used to build the target libraries (like `libstdc++`). If build and host are different, you must have already build and installed a cross compiler that will be used to build the target libraries (if you configured with '`--target=foo-bar`', this compiler will be called `foo-bar-gcc`).

In the case of target libraries, the machine you're building for is the machine you specified with '`--target`'. So, build is the machine you're building on (no change there), host is the machine you're building for (the target libraries are built for the target, so host is the target you specified), and target doesn't apply (because you're not building a compiler, you're building libraries). The configure/make process will adjust these variables as needed. It also sets `$with_cross_host` to the original '`--host`' value in case you need it.

The `libiberty` support library is built up to three times: once for the host, once for the target (even if they are the same), and once for the build if build and host are different. This allows it to be used by all programs which are generated in the course of the build process.

5.2 Top Level Source Directory

The top level source directory in a GCC distribution contains several files and directories that are shared with other software distributions such as that of GNU Binutils. It also contains several subdirectories that contain parts of GCC and its runtime libraries:

- ‘boehm-gc’ The Boehm conservative garbage collector, used as part of the Java runtime library.
- ‘contrib’ Contributed scripts that may be found useful in conjunction with GCC. One of these, ‘contrib/texi2pod.pl’, is used to generate man pages from Texinfo manuals as part of the GCC build process.
- ‘fastjar’ An implementation of the `jar` command, used with the Java front end.
- ‘gcc’ The main sources of GCC itself (except for runtime libraries), including optimizers, support for different target architectures, language front ends, and test suites. See Section 5.3 [The ‘gcc’ Subdirectory], page 12, for details.
- ‘include’ Headers for the `libiberty` library.
- ‘libf2c’ The Fortran runtime library.
- ‘libffi’ The `libffi` library, used as part of the Java runtime library.
- ‘libiberty’ The `libiberty` library, used for portability and for some generally useful data structures and algorithms. See section “Introduction” in GNU *libiberty*, for more information about this library.
- ‘libjava’ The Java runtime library.
- ‘libobjc’ The Objective-C runtime library.
- ‘libstdc++-v3’ The C++ runtime library.
- ‘maintainer-scripts’ Scripts used by the `gccadmin` account on `gcc.gnu.org`.
- ‘zlib’ The `zlib` compression library, used by the Java front end and as part of the Java runtime library.

The build system in the top level directory, including how recursion into subdirectories works and how building runtime libraries for multilibs is handled, is documented in a separate manual, included with GNU Binutils. See section “GNU configure and build system” in *The GNU configure and build system*, for details.

5.3 The ‘gcc’ Subdirectory

The ‘gcc’ directory contains many files that are part of the C sources of GCC, other files used as part of the configuration and build process, and subdirectories including documentation and a test suite. The files that are sources of GCC are documented in a separate chapter. See Chapter 6 [Passes and Files of the Compiler], page 31.

5.3.1 Subdirectories of ‘gcc’

The ‘gcc’ directory contains the following subdirectories:

- ‘*language*’ Subdirectories for various languages. Directories containing a file ‘*config-lang.in*’ are language subdirectories. The contents of the subdirectories ‘*cp*’ (for C++) and ‘*objc*’ (for Objective-C) are documented in this manual (see Chapter 6 [Passes and Files of the Compiler], page 31); those for other languages are not. See Section 5.3.8 [Anatomy of a Language Front End], page 19, for details of the files in these directories.
- ‘*config*’ Configuration files for supported architectures and operating systems. See Section 5.3.9 [Anatomy of a Target Back End], page 23, for details of the files in this directory.
- ‘*doc*’ Texinfo documentation for GCC, together with automatically generated man pages and support for converting the installation manual to HTML. See Section 5.3.7 [Documentation], page 17.
- ‘*fixinc*’ The support for fixing system headers to work with GCC. See ‘*fixinc/README*’ for more information. The headers fixed by this mechanism are installed in ‘*libsubdir/include*’. Along with those headers, ‘*README-fixinc*’ is also installed, as ‘*libsubdir/include/README*’.
- ‘*ginclude*’ System headers installed by GCC, mainly those required by the C standard of freestanding implementations. See Section 5.3.6 [Headers Installed by GCC], page 17, for details of when these and other headers are installed.
- ‘*intl*’ GNU *libintl*, from GNU *gettext*, for systems which do not include it in *libc*. Properly, this directory should be at top level, parallel to the ‘*gcc*’ directory.
- ‘*po*’ Message catalogs with translations of messages produced by GCC into various languages, ‘*language.po*’. This directory also contains ‘*gcc.pot*’, the template for these message catalogues, ‘*exgettext*’, a wrapper around *gettext* to extract the messages from the GCC sources and create ‘*gcc.pot*’, which is run by ‘*make gcc.pot*’, and ‘*EXCLUDES*’, a list of files from which messages should not be extracted.
- ‘*testsuite*’ The GCC test suites (except for those for runtime libraries). See Section 5.4 [Test Suites], page 24.

5.3.2 Configuration in the ‘gcc’ Directory

The ‘gcc’ directory is configured with an Autoconf-generated script ‘*configure*’. The ‘*configure*’ script is generated from ‘*configure.in*’ and ‘*aclocal.m4*’. From the files ‘*configure.in*’ and ‘*acconfig.h*’, Autoheader generates the file ‘*config.in*’. The file ‘*cstamp-h.in*’ is used as a timestamp.

5.3.2.1 Scripts Used by ‘configure’

‘configure’ uses some other scripts to help in its work:

- The standard GNU ‘`config.sub`’ and ‘`config.guess`’ files, kept in the top level directory, are used. **FIXME:** when is the ‘`config.guess`’ file in the ‘`gcc`’ directory (that just calls the top level one) used?
- The file ‘`config.gcc`’ is used to handle configuration specific to the particular build, host or target machine. (In general, this should only be used for features that cannot reasonably be tested in Autoconf feature tests.) See Section 5.3.2.2 [The ‘`config.gcc`’ File], page 14, for details of the contents of this file.
- Each language subdirectory has a file ‘`language/config-lang.in`’ that is used for front-end-specific configuration. See Section 5.3.8.2 [The Front End ‘`config-lang.in`’ File], page 22, for details of this file.
- A helper script ‘`configure.frag`’ is used as part of creating the output of ‘configure’.

5.3.2.2 The ‘config.gcc’ File

FIXME: document the contents of this file, and what variables should be set to control build, host and target configuration.

5.3.2.3 Files Created by configure

Here we spell out what files will be set up by ‘configure’ in the ‘`gcc`’ directory. Some other files are created as temporary files in the configuration process, and are not used in the subsequent build; these are not documented.

- ‘`Makefile`’ is constructed from ‘`Makefile.in`’, together with the host and target fragments (see Chapter 12 [Makefile Fragments], page 331) ‘`t-target`’ and ‘`x-host`’ from ‘`config`’, if any, and language Makefile fragments ‘`language/Make-lang.in`’.
- ‘`auto-host.h`’ contains information about the host machine determined by ‘configure’. If the host machine is different from the build machine, then ‘`auto-build.h`’ is also created, containing such information about the build machine.
- ‘`config.status`’ is a script that may be run to recreate the current configuration.
- ‘`configargs.h`’ is a header containing details of the arguments passed to ‘configure’ to configure GCC, and of the thread model used.
- ‘`cstamp-h`’ is used as a timestamp.
- ‘`fixinc/Makefile`’ is constructed from ‘`fixinc/Makefile.in`’.
- ‘`gccbug`’, a script for reporting bugs in GCC, is constructed from ‘`gccbug.in`’.
- ‘`intl/Makefile`’ is constructed from ‘`intl/Makefile.in`’.
- ‘`mklibgcc`’, a shell script to create a Makefile to build libgcc, is constructed from ‘`mklibgcc.in`’.
- If a language ‘`config-lang.in`’ file (see Section 5.3.8.2 [The Front End ‘`config-lang.in`’ File], page 22) sets `outputs`, then the files listed in `outputs` there are also generated.

The following configuration headers are created from the Makefile, using `mkconfig.sh`, rather than directly by `configure`. `config.h`, `hconfig.h` and `tconfig.h` all contain the `xm-machine.h` header, if any, appropriate to the host, build and target machines respectively, the configuration headers for the target, and some definitions; for the host and build machines, these include the autoconfigured headers generated by `configure`. The other configuration headers are determined by `config.gcc`. They also contain the typedefs for `rtx`, `rtvec` and `tree`.

- `config.h`, for use in programs that run on the host machine.
- `hconfig.h`, for use in programs that run on the build machine.
- `tconfig.h`, for use in programs and libraries for the target machine.
- `tm_p.h`, which includes the header `machine-protos.h` that contains prototypes for functions in the target `.c` file. FIXME: why is such a separate header necessary?

5.3.3 Build System in the `gcc` Directory

FIXME: describe the build system, including what is built in what stages. Also list the various source files that are used in the build process but aren't source files of GCC itself and so aren't documented below (see Chapter 6 [Passes], page 31).

5.3.4 Makefile Targets

<code>all</code>	This is the default target. Depending on what your build/host/target configuration is, it coordinates all the things that need to be built.
<code>doc</code>	Produce info-formatted documentation. Also, <code>make dvi</code> is available for DVI-formatted documentation, and <code>make generated-manpages</code> to generate man pages.
<code>mostlyclean</code>	Delete the files made while building the compiler.
<code>clean</code>	That, and all the other files built by <code>make all</code> .
<code>distclean</code>	That, and all the files created by <code>configure</code> .
<code>extraclean</code>	That, and any temporary or intermediate files, like emacs backup files.
<code>maintainer-clean</code>	Distclean plus any file that can be generated from other files. Note that additional tools may be required beyond what is normally needed to build gcc.
<code>install</code>	Installs gcc.
<code>uninstall</code>	Deletes installed files.
<code>check</code>	Run the testsuite. This creates a <code>testsuite</code> subdirectory that has various <code>.sum</code> and <code>.log</code> files containing the results of the testing. You can run subsets with, for example, <code>make check-gcc</code> . You can specify specific tests by setting

RUNTESTFLAGS to be the name of the `.exp` file, optionally followed by (for some tests) an equals and a file wildcard, like:

```
make check-gcc RUNTESTFLAGS="execute.exp=19980413-*
```

Note that running the testsuite may require additional tools be installed, such as TCL or dejagnu.

bootstrap

Builds `gcc` three times—once with the native compiler, once with the native-built compiler it just built, and once with the compiler it built the second time. In theory, the last two should produce the same results, which `make compare` can check. Each step of this process is called a “stage”, and the results of each stage N ($N = 1 \dots 3$) are copied to a subdirectory `stageN/`.

bootstrap-lean

Like `bootstrap`, except that the various stages are removed once they’re no longer needed. This saves disk space.

bubblestrap

Once bootstrapped, this incrementally rebuilds each of the three stages, one at a time. It does this by “bubbling” the stages up from their subdirectories, rebuilding them, and copying them back to their subdirectories. This will allow you to, for example, quickly rebuild a bootstrapped compiler after changing the sources, without having to do a full bootstrap.

quickstrap

Rebuilds the most recently built stage. Since each stage requires special invocation, using this target means you don’t have to keep track of which stage you’re on or what invocation that stage needs.

cleanstrap

Removed everything (`make clean`) and rebuilds (`make bootstrap`).

stageN ($N = 1 \dots 4$)

For each stage, moves the appropriate files to the `stageN` subdirectory.

unstageN ($N = 1 \dots 4$)

Undoes the corresponding `stageN`.

restageN ($N = 1 \dots 4$)

Undoes the corresponding `stageN` and rebuilds it with the appropriate flags.

compare

Compares the results of stages 2 and 3. This ensures that the compiler is running properly, since it should produce the same object files regardless of how it itself was compiled.

5.3.5 Library Source Files and Headers under the ‘gcc’ Directory

FIXME: list here, with explanation, all the C source files and headers under the `gcc` directory that aren’t built into the GCC executable but rather are part of runtime libraries and object files, such as `crtstuff.c` and `unwind-dw2.c`. See Section 5.3.6 [Headers Installed by GCC], page 17, for more information about the `ginclude` directory.

5.3.6 Headers Installed by GCC

In general, GCC expects the system C library to provide most of the headers to be used with it. However, GCC will fix those headers if necessary to make them work with GCC, and will install some headers required of freestanding implementations. These headers are installed in `'libsubdir/include'`. Headers for non-C runtime libraries are also installed by GCC; these are not documented here. (FIXME: document them somewhere.)

Several of the headers GCC installs are in the `'ginclude'` directory. These headers, `'iso646.h'`, `'stdarg.h'`, `'stdbool.h'`, and `'stddef.h'`, are installed in `'libsubdir/include'`, unless the target Makefile fragment (see Section 12.1 [Target Fragment], page 331) overrides this by setting `USER_H`.

In addition to these headers and those generated by fixing system headers to work with GCC, some other headers may also be installed in `'libsubdir/include'`. `'config.gcc'` may set `extra_headers`; this specifies additional headers under `'config'` to be installed on some systems.

GCC installs its own version of `<float.h>`, from `'ginclude/float.h'`. This is done to cope with command-line options that change the representation of floating point numbers.

GCC also installs its own version of `<limits.h>`; this is generated from `'glimits.h'`, together with `'limitx.h'` and `'limity.h'` if the system also has its own version of `<limits.h>`. (GCC provides its own header because it is required of ISO C freestanding implementations, but needs to include the system header from its own header as well because other standards such as POSIX specify additional values to be defined in `<limits.h>`.) The system's `<limits.h>` header is used via `'libsubdir/include/syslimits.h'`, which is copied from `'gsyslimits.h'` if it does not need fixing to work with GCC; if it needs fixing, `'syslimits.h'` is the fixed copy.

5.3.7 Building Documentation

The main GCC documentation is in the form of manuals in Texinfo format. These are installed in Info format, and DVI versions may be generated by `'make dvi'`. In addition, some man pages are generated from the Texinfo manuals, there are some other text files with miscellaneous documentation, and runtime libraries have their own documentation outside the `'gcc'` directory. (FIXME: document the documentation for runtime libraries somewhere.)

5.3.7.1 Texinfo Manuals

The manuals for GCC as a whole, and the C and C++ front ends, are in files `'doc/*.texi'`. Other front ends have their own manuals in files `'language/*.texi'`. Common files `'doc/include/*.texi'` are provided which may be included in multiple manuals; the following files are in `'doc/include'`:

`'fdl.texi'`

The GNU Free Documentation License.

`'funding.texi'`

The section "Funding Free Software".

`'gcc-common.texi'`

Common definitions for manuals.

`'gpl.texi'`

The GNU General Public License.

`'texinfo.tex'`

A copy of `'texinfo.tex'` known to work with the GCC manuals.

DVI formatted manuals are generated by `'make dvi'`, which uses `texi2dvi` (via the Makefile macro `$(TEXI2DVI)`). Info manuals are generated by `'make info'` (which is run as part of a bootstrap); this generates the manuals in the source directory, using `makeinfo` via the Makefile macro `$(MAKEINFO)`, and they are included in release distributions.

Manuals are also provided on the GCC web site, in both HTML and PostScript forms. This is done via the script `'maintainer-scripts/update_web_docs'`. Each manual to be provided online must be listed in the definition of `MANUALS` in that file; a file `'name.texi'` must only appear once in the source tree, and the output manual must have the same name as the source file. (However, other Texinfo files, included in manuals but not themselves the root files of manuals, may have names that appear more than once in the source tree.) The manual file `'name.texi'` should only include other files in its own directory or in `'doc/include'`. HTML manuals will be generated by `'makeinfo --html'` and PostScript manuals by `texi2dvi` and `dvips`. All Texinfo files that are parts of manuals must be checked into CVS, even if they are generated files, for the generation of online manuals to work.

The installation manual, `'doc/install.texi'`, is also provided on the GCC web site. The HTML version is generated by the script `'doc/install.texi2html'`.

5.3.7.2 Man Page Generation

Because of user demand, in addition to full Texinfo manuals, man pages are provided which contain extracts from those manuals. These man pages are generated from the Texinfo manuals using `'contrib/texi2pod.pl'` and `pod2man`. (The man page for `g++`, `'cp/g++.1'`, just contains a `'.so'` reference to `'gcc.1'`, but all the other man pages are generated from Texinfo manuals.)

Because many systems may not have the necessary tools installed to generate the man pages, they are only generated if the `'configure'` script detects that recent enough tools are installed, and the Makefiles allow generating man pages to fail without aborting the build. Man pages are also included in release distributions. They are generated in the source directory.

Magic comments in Texinfo files starting `'@c man'` control what parts of a Texinfo file go into a man page. Only a subset of Texinfo is supported by `'texi2pod.pl'`, and it may be necessary to add support for more Texinfo features to this script when generating new man pages. To improve the man page output, some special Texinfo macros are provided in `'doc/include/gcc-common.texi'` which `'texi2pod.pl'` understands:

`@gcctabopt`

Use in the form `'@table @gcctabopt'` for tables of options, where for printed output the effect of `'@code'` is better than that of `'@option'` but for man page output a different effect is wanted.

@gccoptlist

Use for summary lists of options in manuals.

@gol

Use at the end of each line inside ‘@gccoptlist’. This is necessary to avoid problems with differences in how the ‘@gccoptlist’ macro is handled by different Texinfo formatters.

FIXME: describe the ‘texi2pod.pl’ input language and magic comments in more detail.

5.3.7.3 Miscellaneous Documentation

In addition to the formal documentation that is installed by GCC, there are several other text files with miscellaneous documentation:

‘ABOUT-GCC-NLS’

Notes on GCC’s Native Language Support. FIXME: this should be part of this manual rather than a separate file.

‘ABOUT-NLS’

Notes on the Free Translation Project.

‘COPYING’ The GNU General Public License.**‘COPYING.LIB’**

The GNU Lesser General Public License.

‘*ChangeLog*’**‘*/ChangeLog*’**

Change log files for various parts of GCC.

‘LANGUAGES’

Details of a few changes to the GCC front-end interface. FIXME: the information in this file should be part of general documentation of the front-end interface in this manual.

‘ONEWS’

Information about new features in old versions of GCC. (For recent versions, the information is on the GCC web site.)

‘README.Portability’

Information about portability issues when writing code in GCC. FIXME: why isn’t this part of this manual or of the GCC Coding Conventions?

‘SERVICE’ A pointer to the GNU Service Directory.

FIXME: document such files in subdirectories, at least ‘config’, ‘cp’, ‘objc’, ‘testsuite’.

5.3.8 Anatomy of a Language Front End

A front end for a language in GCC has the following parts:

- A directory ‘*language*’ under ‘gcc’ containing source files for that front end. See Section 5.3.8.1 [The Front End ‘*language*’ Directory], page 21, for details.
- A mention of the language in the list of supported languages in ‘gcc/doc/install.texi’.

- Details of contributors to that front end in `'gcc/doc/contrib.texi'`. If the details are in that front end's own manual then there should be a link to that manual's list in `'contrib.texi'`.
- Information about support for that language in `'gcc/doc/frontends.texi'`.
- Information about standards for that language, and the front end's support for them, in `'gcc/doc/standards.texi'`. This may be a link to such information in the front end's own manual.
- Details of source file suffixes for that language and `'-x lang'` options supported, in `'gcc/doc/invoke.texi'`.
- Entries in `default_compilers` in `'gcc.c'` for source file suffixes for that language.
- Preferably test suites, which may be under `'gcc/testsuite'` or runtime library directories. FIXME: document somewhere how to write test suite harnesses.
- Probably a runtime library for the language, outside the `'gcc'` directory. FIXME: document this further.
- Details of the directories of any runtime libraries in `'gcc/doc/sourcebuild.texi'`.

If the front end is added to the official GCC CVS repository, the following are also necessary:

- At least one Bugzilla component for bugs in that front end and runtime libraries. This category needs to be mentioned in `'gcc/gccbug.in'`, as well as being added to the Bugzilla database.
- Normally, one or more maintainers of that front end listed in `'MAINTAINERS'`.
- Mentions on the GCC web site in `'index.html'` and `'frontends.html'`, with any relevant links on `'readings.html'`. (Front ends that are not an official part of GCC may also be listed on `'frontends.html'`, with relevant links.)
- A news item on `'index.html'`, and possibly an announcement on the `gcc-announce@gcc.gnu.org` mailing list.
- The front end's manuals should be mentioned in `'maintainer-scripts/update_web_docs'` (see Section 5.3.7.1 [Texinfo Manuals], page 17) and the online manuals should be linked to from `'onlinedocs/index.html'`.
- Any old releases or CVS repositories of the front end, before its inclusion in GCC, should be made available on the GCC FTP site `ftp://gcc.gnu.org/pub/gcc/old-releases/`.
- The release and snapshot script `'maintainer-scripts/gcc_release'` should be updated to generate appropriate tarballs for this front end. The associated `'maintainer-scripts/snapshot-README'` and `'maintainer-scripts/snapshot-index.html'` files should be updated to list the tarballs and diffs for this front end.
- If this front end includes its own version files that include the current date, `'maintainer-scripts/update_version'` should be updated accordingly.
- `'CVSROOT/modules'` in the GCC CVS repository should be updated.

5.3.8.1 The Front End ‘*language*’ Directory

A front end ‘*language*’ directory contains the source files of that front end (but not of any runtime libraries, which should be outside the ‘*gcc*’ directory). This includes documentation, and possibly some subsidiary programs build alongside the front end. Certain files are special and other parts of the compiler depend on their names:

‘*config-lang.in*’

This file is required in all language subdirectories. See Section 5.3.8.2 [The Front End ‘*config-lang.in*’ File], page 22, for details of its contents

‘*Make-lang.in*’

This file is required in all language subdirectories. It contains targets *lang.hook* (where *lang* is the setting of *language* in ‘*config-lang.in*’) for the following values of *hook*, and any other Makefile rules required to build those targets (which may if necessary use other Makefiles specified in *outputs* in ‘*config-lang.in*’, although this is deprecated).

all.build

all.cross

start.encap

rest.encap

FIXME: exactly what goes in each of these targets?

info Build info documentation for the front end, in the source directory. This target is only called by ‘*make bootstrap*’ if a suitable version of *makeinfo* is available, so does not need to check for this, and should fail if an error occurs.

dvi Build DVI documentation for the front end, in the build directory. This should be done using $\$(\text{TEXI2DVI})$, with appropriate ‘*-I*’ arguments pointing to directories of included files.

generated-manpages

Build generated man pages for the front end from Texinfo manuals (see Section 5.3.7.2 [Man Page Generation], page 18), in the source directory. This target is only called if the necessary tools are available, but should ignore errors so as not to stop the build if errors occur; man pages are optional and the tools involved may be installed in a broken way.

install-normal

FIXME: what is this target for?

install-common

Install everything that is part of the front end, apart from the compiler executables listed in *compilers* in ‘*config-lang.in*’ that are installed in ‘*libsubdir*’ by the main ‘*Makefile*’.

install-info

Install info documentation for the front end, if it is present in the source directory. (It may not be present if a suitable version of

`makeinfo` was not installed.) This target should run the command `install-info` to update the info directory, but should ignore errors when running that command.

`install-man`

Install man pages for the front end. This target should ignore errors.

`uninstall`

Uninstall files installed by installing the compiler. This is currently documented not to be supported, so the hook need not do anything.

`mostlyclean`

`clean`

`distclean`

`extraclean`

`maintainer-clean`

Except for `extraclean`, the language parts of the standard GNU `*clean` targets. See section “Standard Targets for Users” in *GNU Coding Standards*, for details of the standard targets. `extraclean` does `distclean` and also deletes anything likely to be found in the source directory that shouldn’t be in the distribution. For GCC, `maintainer-clean` should delete all generated files in the source directory that are not checked into CVS, but should not delete anything checked into CVS.

`stage1`

`stage2`

`stage3`

`stage4`

Move to the stage directory files not included in `stagestuff` in `config-lang.in` or otherwise moved by the main `Makefile`.

`‘lang-options.h’`

This file provides entries for `documented_lang_options` in `toplev.c` describing command-line options the front end accepts for `‘--help’` output.

`‘lang-specs.h’`

This file provides entries for `default_compilers` in `gcc.c` which override the default of giving an error that a compiler for that language is not installed.

`‘language-tree.def’`

This file, which need not exist, defines any language-specific tree codes.

5.3.8.2 The Front End `‘config-lang.in’` File

Each language subdirectory contains a `‘config-lang.in’` file. In addition the main directory contains `‘c-config-lang.in’`, which contains limited information for the C language. This file is a shell script that may define some variables describing the language:

`language` This definition must be present, and gives the name of the language for some purposes such as arguments to `‘--enable-languages’`.

lang_requires

If defined, this variable lists (space-separated) language front ends other than C that this front end requires to be enabled (with the names given being their language settings). For example, the Java front end depends on the C++ front end, so sets `lang_requires=c++`.

target_libs

If defined, this variable lists (space-separated) targets in the top level `Makefile` to build the runtime libraries for this language, such as `target-libobjc`.

lang_dirs

If defined, this variable lists (space-separated) top level directories (parallel to `gcc`), apart from the runtime libraries, that should not be configured if this front end is not built.

build_by_default

If defined to `no`, this language front end is not built unless enabled in a `--enable-languages` argument. Otherwise, front ends are built by default, subject to any special logic in `configure.in` (as is present to disable the Ada front end if the Ada compiler is not already installed).

boot_language

If defined to `yes`, this front end is built in stage 1 of the bootstrap. This is only relevant to front ends written in their own languages.

compilers

If defined, a space-separated list of compiler executables that should be installed in `libsubdir`. The names here will each end with `\$(exeext)`.

stagemstuff

If defined, a space-separated list of files that should be moved to the `stagen` directories in each stage of bootstrap.

outputs

If defined, a space-separated list of files that should be generated by `configure` substituting values in them. This mechanism can be used to create a file `language/Makefile` from `language/Makefile.in`, but this is deprecated, building everything from the single `gcc/Makefile` is preferred.

gtfiles

If defined, a space-separated list of files that should be scanned by `genctype.c` to generate the garbage collection tables and routines for this language. This excludes the files that are common to all front ends. See Chapter 15 [Type Information], page 339.

5.3.9 Anatomy of a Target Back End

A back end for a target architecture in GCC has the following parts:

- A directory `machine` under `gcc/config`, containing a machine description `machine.md` file (see Chapter 9 [Machine Descriptions], page 117), header files `machine.h` and `machine-protos.h` and a source file `machine.c` (see Chapter 10 [Target Description Macros and Functions], page 193), possibly a target Makefile fragment `t-machine` (see Section 12.1 [The Target Makefile Fragment], page 331),

and maybe some other files. The names of these files may be changed from the defaults given by explicit specifications in ‘`config.gcc`’.

- If necessary, a file ‘`machine-modes.def`’ in the ‘`machine`’ directory, containing additional machine modes to represent condition codes. See Section 10.15 [Condition Code], page 264, for further details.
- Entries in ‘`config.gcc`’ (see Section 5.3.2.2 [The ‘`config.gcc`’ File], page 14) for the systems with this target architecture.
- Documentation in ‘`gcc/doc/invoke.texi`’ for any command-line options supported by this target (see Section 10.3 [Run-time Target Specification], page 200). This means both entries in the summary table of options and details of the individual options.
- Documentation in ‘`gcc/doc/extend.texi`’ for any target-specific attributes supported (see Section 10.24 [Defining target-specific uses of `__attribute__`], page 313), including where the same attribute is already supported on some targets, which are enumerated in the manual.
- Documentation in ‘`gcc/doc/extend.texi`’ for any target-specific pragmas supported.
- Documentation in ‘`gcc/doc/extend.texi`’ of any target-specific built-in functions supported.
- Documentation in ‘`gcc/doc/md.texi`’ of any target-specific constraint letters (see Section 9.7.5 [Constraints for Particular Machines], page 132).
- A note in ‘`gcc/doc/contrib.texi`’ under the person or people who contributed the target support.
- Entries in ‘`gcc/doc/install.texi`’ for all target triplets supported with this target architecture, giving details of any special notes about installation for this target, or saying that there are no special notes if there are none.
- Possibly other support outside the ‘`gcc`’ directory for runtime libraries. FIXME: reference docs for this. The `libstdc++` porting manual needs to be installed as info for this to work, or to be a chapter of this manual.

If the back end is added to the official GCC CVS repository, the following are also necessary:

- An entry for the target architecture in ‘`readings.html`’ on the GCC web site, with any relevant links.
- A news item about the contribution of support for that target architecture, in ‘`index.html`’ on the GCC web site.
- Normally, one or more maintainers of that target listed in ‘`MAINTAINERS`’. Some existing architectures may be unmaintained, but it would be unusual to add support for a target that does not have a maintainer when support is added.

5.4 Test Suites

GCC contains several test suites to help maintain compiler quality. Most of the runtime libraries and language front ends in GCC have test suites. Currently only the C language test suites are documented here; FIXME: document the others.

5.4.1 Idioms Used in Test Suite Code

In the ‘`gcc.c-torture`’ test suites, test cases are commonly named after the date on which they were added. This allows people to tell at a glance whether a test failure is because of a recently found bug that has not yet been fixed, or whether it may be a regression. In other test suites, more descriptive names are used. In general C test cases have a trailing ‘`-n.c`’, starting with ‘`-1.c`’, in case other test cases with similar names are added later.

Test cases should use `abort ()` to indicate failure and `exit (0)` for success; on some targets these may be redefined to indicate failure and success in other ways.

In the ‘`gcc.dg`’ test suite, it is often necessary to test that an error is indeed a hard error and not just a warning—for example, where it is a constraint violation in the C standard, which must become an error with ‘`-pedantic-errors`’. The following idiom, where the first line shown is line *line* of the file and the line that generates the error, is used for this:

```
/* { dg-bogus "warning" "warning in place of error" } */
/* { dg-error "regexp" "message" { target *-*-* } line } */
```

It may be necessary to check that an expression is an integer constant expression and has a certain value. To check that *E* has value *V*, an idiom similar to the following is used:

```
char x[(E) == (V) ? 1 : -1];
```

In ‘`gcc.dg`’ tests, `__typeof__` is sometimes used to make assertions about the types of expressions. See, for example, ‘`gcc.dg/c99-condexpr-1.c`’. The more subtle uses depend on the exact rules for the types of conditional expressions in the C standard; see, for example, ‘`gcc.dg/c99-intconst-1.c`’.

It is useful to be able to test that optimizations are being made properly. This cannot be done in all cases, but it can be done where the optimization will lead to code being optimized away (for example, where flow analysis or alias analysis should show that certain code cannot be called) or to functions not being called because they have been expanded as built-in functions. Such tests go in ‘`gcc.c-torture/execute`’. Where code should be optimized away, a call to a nonexistent function such as `link_failure ()` may be inserted; a definition

```
#ifndef __OPTIMIZE__
void
link_failure (void)
{
    abort ();
}
#endif
```

will also be needed so that linking still succeeds when the test is run without optimization. When all calls to a built-in function should have been optimized and no calls to the non-built-in version of the function should remain, that function may be defined as `static` to call `abort ()` (although redeclaring a function as `static` may not work on all targets).

All testcases must be portable. Target-specific testcases must have appropriate code to avoid causing failures on unsupported systems; unfortunately, the mechanisms for this differ by directory.

FIXME: discuss non-C test suites here.

5.4.2 C Language Test Suites

GCC contains the following C language test suites, in the ‘gcc/testsuite’ directory:

‘gcc.dg’ This contains tests of particular features of the C compiler, using the more modern ‘dg’ harness. Correctness tests for various compiler features should go here if possible.

Magic comments determine whether the file is preprocessed, compiled, linked or run. In these tests, error and warning message texts are compared against expected texts or regular expressions given in comments. These tests are run with the options ‘-ansi -pedantic’ unless other options are given in the test. Except as noted below they are not run with multiple optimization options.

‘gcc.dg/cpp’

This subdirectory contains tests of the preprocessor.

‘gcc.dg/debug’

This subdirectory contains tests for debug formats. Tests in this subdirectory are run for each debug format that the compiler supports.

‘gcc.dg/format’

This subdirectory contains tests of the ‘-Wformat’ format checking. Tests in this directory are run with and without ‘-DWIDE’.

‘gcc.dg/noncompile’

This subdirectory contains tests of code that should not compile and does not need any special compilation options. They are run with multiple optimization options, since sometimes invalid code crashes the compiler with optimization.

‘gcc.dg/special’

FIXME: describe this.

‘gcc.c-torture’

This contains particular code fragments which have historically broken easily. These tests are run with multiple optimization options, so tests for features which only break at some optimization levels belong here. This also contains tests to check that certain optimizations occur. It might be worthwhile to separate the correctness tests cleanly from the code quality tests, but it hasn’t been done yet.

‘gcc.c-torture/compat’

FIXME: describe this.

This directory should probably not be used for new tests.

‘gcc.c-torture/compile’

This test suite contains test cases that should compile, but do not need to link or run. These test cases are compiled with several different combinations of optimization options. All warnings are disabled for these test cases, so this directory is not suitable if you wish to test for the presence or absence of compiler warnings. While special options can be set, and tests disabled on specific platforms, by the use of ‘.x’ files, mostly these test cases should not

contain platform dependencies. **FIXME:** discuss how defines such as `NO_LABEL_VALUES` and `STACK_SIZE` are used.

`'gcc.c-torture/execute'`

This test suite contains test cases that should compile, link and run; otherwise the same comments as for `'gcc.c-torture/compile'` apply.

`'gcc.c-torture/execute/ieee'`

This contains tests which are specific to IEEE floating point.

`'gcc.c-torture/unordered'`

FIXME: describe this.

This directory should probably not be used for new tests.

`'gcc.c-torture/misc-tests'`

This directory contains C tests that require special handling. Some of these tests have individual expect files, and others share special-purpose expect files:

`'bprob*.c'`

Test `'-fbranch-probabilities'` using `'bprob.exp'`, which in turn uses the generic, language-independent framework (see Section 5.4.5 [Support for testing profile-directed optimizations], page 29).

`'dg-*.c'` Test the testsuite itself using `'dg-test.exp'`.

`'gcov*.c'` Test `gcov` output using `'gcov.exp'`, which in turn uses the language-independent support (see Section 5.4.4 [Support for testing gcov], page 28).

`'i386-pf-*.c'`

Test i386-specific support for data prefetch using `'i386-prefetch.exp'`.

FIXME: merge in `'testsuite/README.gcc'` and discuss the format of test cases and magic comments more.

5.4.3 The Java library test suites.

Runtime tests are executed via `'make check'` in the `'target/libjava/testsuite'` directory in the build tree. Additional runtime tests can be checked into this testsuite.

Regression testing of the core packages in libgcj is also covered by the Mauve test suite. The Mauve Project develops tests for the Java Class Libraries. These tests are run as part of libgcj testing by placing the Mauve tree within the libjava testsuite sources at `'libjava/testsuite/libjava.mauve/mauve'`, or by specifying the location of that tree when invoking `'make'`, as in `'make MAUVEDIR=~ /mauve check'`.

To detect regressions, a mechanism in `'mauve.exp'` compares the failures for a test run against the list of expected failures in `'libjava/testsuite/libjava.mauve/xfails'` from the source hierarchy. Update this file when adding new failing tests to Mauve, or when fixing bugs in libgcj that had caused Mauve test failures.

The Jacks project provides a test suite for Java compilers that can be used to test changes that affect the GCJ front end. This test suite is run as part of Java testing by placing the Jacks tree within the the libjava testsuite sources at `'libjava/testsuite/libjava.jacks/jacks'`.

We encourage developers to contribute test cases to Mauve and Jacks.

5.4.4 Support for testing gcov

Language-independent support for testing gcov, and for checking that branch profiling produces expected values, is provided by the expect file `'gcov.exp'`. gcov tests also rely on procedures in `'gcc.dg.exp'` to compile and run the test program. A typical gcov test contains the following DejaGNU commands within comments:

```
{ dg-options "-fprofile-arcs -ftest-coverage" }
{ dg-do run { target native } }
{ dg-final { run-gcov sourcefile } }
```

Checks of gcov output can include line counts, branch percentages, and call return percentages. All of these checks are requested via commands that appear in comments in the test's source file. Commands to check line counts are processed by default. Commands to check branch percentages and call return percentages are processed if there is a file with the same basename as the source file and a suffix `'.x'` that contains a line `set gcov_verify_branches 1` or `set gcov_verify_calls 1`, respectively.

A line count command appears within a comment on the source line that is expected to get the specified count and has the form `count(cnt)`. A test should only check line counts for lines that will get the same count for any architecture.

Commands to check branch percentages (`branch`) and call return percentages (`returns`) are very similar to each other. A beginning command appears on or before the first of a range of lines that will report the percentage, and the ending command follows that range of lines. The beginning command can include a list of percentages, all of which are expected to be found within the range. A range is terminated by the next command of the same kind. A command `branch(end)` or `returns(end)` marks the end of a range without starting a new one. For example:

```
if (i > 10 && j > i && j < 20) /* branch(27 50 75) */
                               /* branch(end) */
    foo (i, j);
```

For a call return percentage, the value specified is the percentage of calls reported to return. For a branch percentage, the value is either the expected percentage or 100 minus that value, since the direction of a branch can differ depending on the target or the optimization level.

Not all branches and calls need to be checked. A test should not check for branches that might be optimized away or replaced with predicated instructions. Don't check for calls inserted by the compiler or ones that might be inlined or optimized away.

A single test can check for combinations of line counts, branch percentages, and call return percentages. The command to check a line count must appear on the line that will report that count, but commands to check branch percentages and call return percentages can bracket the lines that report them.

5.4.5 Support for testing profile-directed optimizations

The file `'profopt.exp'` provides language-independent support for checking correct execution of a test built with profile-directed optimization. This testing requires that a test program be built and executed twice. The first time it is compiled to generate profile data, and the second time it is compiled to use the data that was generated during the first execution. The second execution is to verify that the test produces the expected results.

To check that the optimization actually generated better code, a test can be built and run a third time with normal optimizations to verify that the performance is better with the profile-directed optimizations. `'profopt.exp'` has the beginnings of this kind of support.

`'profopt.exp'` provides generic support for profile-directed optimizations. Each set of tests that uses it provides information about a specific optimization:

```

tool          tool being tested, e.g., gcc
profile_option
              options used to generate profile data
feedback_option
              options used to optimize using that profile data
prof_ext      suffix of profile data files
PROFOPT_OPTIONS
              list of options with which to run each test, similar to the lists for torture tests

```

5.4.6 Support for testing binary compatibility

The file `'compat.exp'` provides language-independent support for binary compatibility testing. It supports testing interoperability of two compilers that follow the same ABI, or of multiple sets of compiler options that should not affect binary compatibility. It is intended to be used for test suites that complement ABI test suites.

A test supported by this framework has three parts, each in a separate source file: a main program and two pieces that interact with each other to split up the functionality being tested.

```

'testname_main.suffix'
    Contains the main program, which calls a function in file 'testname_x.suffix'.
'testname_x.suffix'
    Contains at least one call to a function in 'testname_y.suffix'.
'testname_y.suffix'
    Shares data with, or gets arguments from, 'testname_x.suffix'.

```

Within each test, the main program and one functional piece are compiled by the GCC under test. The other piece can be compiled by an alternate compiler. If no alternate compiler is specified, then all three source files are all compiled by the GCC under test. It's also possible to specify a pair of lists of compiler options, one list for each compiler, so that each test will be compiled with each pair of options.

`'compat.exp'` defines default pairs of compiler options. These can be overridden by defining the environment variable `COMPAT_OPTIONS` as:

```
COMPAT_OPTIONS="[list [list {tst1} {alt1}]
...[list {tstn} {altn}]]"
```

where *tsti* and *alti* are lists of options, with *tsti* used by the compiler under test and *alti* used by the alternate compiler. For example, with `[list [list {-g -O0} {-O3}] [list {-fpic} {-fPIC -O2}]]`, the test is first built with `-g -O0` by the compiler under test and with `-O3` by the alternate compiler. The test is built a second time using `-fpic` by the compiler under test and `-fPIC -O2` by the alternate compiler.

An alternate compiler is specified by defining an environment variable; for C++ define `ALT_CXX_UNDER_TEST` to be the full pathname of an installed compiler. That will be written to the `'site.exp'` file used by DejaGNU. The default is to build each test with the compiler under test using the first of each pair of compiler options from `COMPAT_OPTIONS`. When `ALT_CXX_UNDER_TEST` is `same`, each test is built using the compiler under test but with combinations of the options from `COMPAT_OPTIONS`.

To run only the C++ compatibility suite using the compiler under test and another version of GCC using specific compiler options, do the following from `'objdir/gcc'`:

```
rm site.exp
make -k \
  ALT_CXX_UNDER_TEST=${alt_prefix}/bin/g++ \
  COMPAT_OPTIONS="lists as shown above" \
  check-c++ \
  RUNTESTFLAGS="compat.exp"
```

A test that fails when the source files are compiled with different compilers, but passes when the files are compiled with the same compiler, demonstrates incompatibility of the generated code or runtime support. A test that fails for the alternate compiler but passes for the compiler under test probably tests for a bug that was fixed in the compiler under test but is present in the alternate compiler.

6 Passes and Files of the Compiler

The overall control structure of the compiler is in ‘`toplev.c`’. This file is responsible for initialization, decoding arguments, opening and closing files, and sequencing the passes.

The parsing pass is invoked only once, to parse the entire input. A high level tree representation is then generated from the input, one function at a time. This tree code is then transformed into RTL intermediate code, and processed. The files involved in transforming the trees into RTL are ‘`expr.c`’, ‘`expmed.c`’, and ‘`stmt.c`’. The order of trees that are processed, is not necessarily the same order they are generated from the input, due to deferred inlining, and other considerations.

Each time the parsing pass reads a complete function definition or top-level declaration, it calls either the function `rest_of_compilation`, or the function `rest_of_decl_compilation` in ‘`toplev.c`’, which are responsible for all further processing necessary, ending with output of the assembler language. All other compiler passes run, in sequence, within `rest_of_compilation`. When that function returns from compiling a function definition, the storage used for that function definition’s compilation is entirely freed, unless it is an inline function, or was deferred for some reason (this can occur in templates, for example). (see section “An Inline Function is As Fast As a Macro” in *Using the GNU Compiler Collection (GCC)*).

Here is a list of all the passes of the compiler and their source files. Also included is a description of where debugging dumps can be requested with ‘`-d`’ options.

- Parsing. This pass reads the entire text of a function definition, constructing a high level tree representation. (Because of the semantic analysis that takes place during this pass, it does more than is formally considered to be parsing.)

The tree representation does not entirely follow C syntax, because it is intended to support other languages as well.

Language-specific data type analysis is also done in this pass, and every tree node that represents an expression has a data type attached. Variables are represented as declaration nodes.

The language-independent source files for parsing are ‘`tree.c`’, ‘`fold-const.c`’, and ‘`stor-layout.c`’. There are also header files ‘`tree.h`’ and ‘`tree.def`’ which define the format of the tree representation.

C preprocessing, for language front ends, that want or require it, is performed by `cpplib`, which is covered in separate documentation. In particular, the internals are covered in See section “Cpplib internals” in *Cpplib Internals*.

The source files to parse C are ‘`c-convert.c`’, ‘`c-decl.c`’, ‘`c-errors.c`’, ‘`c-lang.c`’, ‘`c-objc-common.c`’, ‘`c-parse.in`’, ‘`c-aux-info.c`’, and ‘`c-typeck.c`’, along with a header file ‘`c-tree.h`’ and some files shared with Objective-C and C++.

The source files for parsing C++ are in ‘`cp/`’. They are ‘`parse.y`’, ‘`class.c`’, ‘`cvt.c`’, ‘`decl.c`’, ‘`decl2.c`’, ‘`except.c`’, ‘`expr.c`’, ‘`init.c`’, ‘`lex.c`’, ‘`method.c`’, ‘`ptree.c`’, ‘`search.c`’, ‘`spew.c`’, ‘`semantics.c`’, ‘`tree.c`’, ‘`typeck2.c`’, and ‘`typeck.c`’, along with header files ‘`cp-tree.def`’, ‘`cp-tree.h`’, and ‘`decl.h`’.

The special source files for parsing Objective-C are in ‘`objc/`’. They are ‘`objc-act.c`’, ‘`objc-tree.def`’, and ‘`objc-act.h`’. Certain C-specific files are used for this as well.

The files `'c-common.c'`, `'c-common.def'`, `'c-format.c'`, `'c-opts.c'`, `'c-pragma.c'`, `'c-semantic.c'`, and `'c-lex.c'`, along with header files `'c-common.h'`, `'c-dump.h'`, and `'c-pragma.h'`, are also used for all of the above languages.

- Tree optimization. This is the optimization of the tree representation, before converting into RTL code.

Currently, the main optimization performed here is tree-based inlining. This is implemented in `'tree-inline.c'` and used by both C and C++. Note that tree based inlining turns off rtx based inlining (since it's more powerful, it would be a waste of time to do rtx based inlining in addition).

Constant folding and some arithmetic simplifications are also done during this pass, on the tree representation. The routines that perform these tasks are located in `'fold-const.c'`.

- RTL generation. This is the conversion of syntax tree into RTL code.

This is where the bulk of target-parameter-dependent code is found, since often it is necessary for strategies to apply only when certain standard kinds of instructions are available. The purpose of named instruction patterns is to provide this information to the RTL generation pass.

Optimization is done in this pass for `if`-conditions that are comparisons, boolean operations or conditional expressions. Tail recursion is detected at this time also. Decisions are made about how best to arrange loops and how to output `switch` statements.

The source files for RTL generation include `'stmt.c'`, `'calls.c'`, `'expr.c'`, `'expflow.c'`, `'expmed.c'`, `'function.c'`, `'optabs.c'` and `'emit-rtl.c'`. Also, the file `'insn-emit.c'`, generated from the machine description by the program `genemit`, is used in this pass. The header file `'expr.h'` is used for communication within this pass.

The header files `'insn-flags.h'` and `'insn-codes.h'`, generated from the machine description by the programs `genflags` and `gencodes`, tell this pass which standard names are available for use and which patterns correspond to them.

Aside from debugging information output, none of the following passes refers to the tree structure representation of the function (only part of which is saved).

The decision of whether the function can and should be expanded inline in its subsequent callers is made at the end of rtl generation. The function must meet certain criteria, currently related to the size of the function and the types and number of parameters it has. Note that this function may contain loops, recursive calls to itself (tail-recursive functions can be inlined!), `gotos`, in short, all constructs supported by GCC. The file `'integrate.c'` contains the code to save a function's rtl for later inlining and to inline that rtl when the function is called. The header file `'integrate.h'` is also used for this purpose.

The option `'-dr'` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `'rtl'` to the input file name.

- Sibling call optimization. This pass performs tail recursion elimination, and tail and sibling call optimizations. The purpose of these optimizations is to reduce the overhead of function calls, whenever possible.

The source file of this pass is `'sibcall.c'`

The option `-di` causes a debugging dump of the RTL code after this pass is run. This dump file's name is made by appending `.sibling` to the input file name.

- **Jump optimization.** This pass simplifies jumps to the following instruction, jumps across jumps, and jumps to jumps. It deletes unreferenced labels and unreachable code, except that unreachable code that contains a loop is not recognized as unreachable in this pass. (Such loops are deleted later in the basic block analysis.) It also converts some code originally written with jumps into sequences of instructions that directly set values from the results of comparisons, if the machine has such instructions.

Jump optimization is performed two or three times. The first time is immediately following RTL generation. The second time is after CSE, but only if CSE says repeated jump optimization is needed. The last time is right before the final pass. That time, cross-jumping and deletion of no-op move instructions are done together with the optimizations described above.

The source file of this pass is `jump.c`.

The option `-dj` causes a debugging dump of the RTL code after this pass is run for the first time. This dump file's name is made by appending `.jump` to the input file name.

- **Register scan.** This pass finds the first and last use of each register, as a guide for common subexpression elimination. Its source is in `regclass.c`.
- **Jump threading.** This pass detects a condition jump that branches to an identical or inverse test. Such jumps can be `threaded` through the second conditional test. The source code for this pass is in `jump.c`. This optimization is only performed if `-fthread-jumps` is enabled.
- **Static Single Assignment (SSA) based optimization passes.** The SSA conversion passes (to/from) are turned on by the `-fssa` option (it is also done automatically if you enable an SSA optimization pass). These passes utilize a form called Static Single Assignment. In SSA form, each variable (pseudo register) is only set once, giving you def-use and use-def chains for free, and enabling a lot more optimization passes to be run in linear time. Conversion to and from SSA form is handled by functions in `ssa.c`.

The option `-de` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.ssa` to the input file name.

- **SSA Conditional Constant Propagation.** Turned on by the `-fssa-ccp` option. This pass performs conditional constant propagation to simplify instructions including conditional branches. This pass is more aggressive than the constant propagation done by the CSE and GCSE passes, but operates in linear time.

The option `-dW` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.ssaccp` to the input file name.

- **SSA Aggressive Dead Code Elimination.** Turned on by the `-fssa-dce` option. This pass performs elimination of code considered unnecessary because it has no externally visible effects on the program. It operates in linear time.

The option `-dX` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.ssadce` to the input file name.

- **Common subexpression elimination.** This pass also does constant propagation. Its source files are `cse.c`, and `cselib.c`. If constant propagation causes conditional

jumps to become unconditional or to become no-ops, jump optimization is run again when CSE is finished.

The option `-ds` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.cse` to the input file name.

- Global common subexpression elimination. This pass performs two different types of GCSE depending on whether you are optimizing for size or not (LCM based GCSE tends to increase code size for a gain in speed, while Morel-Renvoise based GCSE does not). When optimizing for size, GCSE is done using Morel-Renvoise Partial Redundancy Elimination, with the exception that it does not try to move invariants out of loops—that is left to the loop optimization pass. If MR PRE GCSE is done, code hoisting (aka unification) is also done, as well as load motion. If you are optimizing for speed, LCM (lazy code motion) based GCSE is done. LCM is based on the work of Knoop, Ruting, and Steffen. LCM based GCSE also does loop invariant code motion. We also perform load and store motion when optimizing for speed. Regardless of which type of GCSE is used, the GCSE pass also performs global constant and copy propagation.

The source file for this pass is `gcse.c`, and the LCM routines are in `lcm.c`.

The option `-dG` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.gcse` to the input file name.

- Loop optimization. This pass moves constant expressions out of loops, and optionally does strength-reduction and loop unrolling as well. Its source files are `loop.c` and `unroll.c`, plus the header `loop.h` used for communication between them. Loop unrolling uses some functions in `integrate.c` and the header `integrate.h`. Loop dependency analysis routines are contained in `dependence.c`.

The option `-dL` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.loop` to the input file name.

- If `-frerun-cse-after-loop` was enabled, a second common subexpression elimination pass is performed after the loop optimization pass. Jump threading is also done again at this time if it was specified.

The option `-dt` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.cse2` to the input file name.

- Data flow analysis (`flow.c`). This pass divides the program into basic blocks (and in the process deletes unreachable loops); then it computes which pseudo-registers are live at each point in the program, and makes the first instruction that uses a value point at the instruction that computed the value.

This pass also deletes computations whose results are never used, and combines memory references with add or subtract instructions to make autoincrement or autodecrement addressing.

The option `-df` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.flow` to the input file name. If stupid register allocation is in use, this dump file reflects the full results of such allocation.

- Instruction combination (`combine.c`). This pass attempts to combine groups of two or three instructions that are related by data flow into single instructions. It combines the RTL expressions for the instructions by substitution, simplifies the result using algebra, and then attempts to match the result against the machine description.

The option `-dc` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.combine` to the input file name.

- If-conversion is a transformation that transforms control dependencies into data dependencies (IE it transforms conditional code into a single control stream). It is implemented in the file `ifcvt.c`.

The option `-dE` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.ce` to the input file name.

- Register movement (`regmove.c`). This pass looks for cases where matching constraints would force an instruction to need a reload, and this reload would be a register-to-register move. It then attempts to change the registers used by the instruction to avoid the move instruction.

The option `-dN` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.regmove` to the input file name.

- Instruction scheduling (`sched.c`). This pass looks for instructions whose output will not be available by the time that it is used in subsequent instructions. (Memory loads and floating point instructions often have this behavior on RISC machines). It re-orders instructions within a basic block to try to separate the definition and use of items that otherwise would cause pipeline stalls.

Instruction scheduling is performed twice. The first time is immediately after instruction combination and the second is immediately after reload.

The option `-ds` causes a debugging dump of the RTL code after this pass is run for the first time. The dump file's name is made by appending `.sched` to the input file name.

- Register allocation. These passes make sure that all occurrences of pseudo registers are eliminated, either by allocating them to a hard register, replacing them by an equivalent expression (e.g. a constant) or by placing them on the stack. This is done in several subpasses:

- Register class preferencing. The RTL code is scanned to find out which register class is best for each pseudo register. The source file is `regclass.c`.
- Local register allocation (`local-alloc.c`). This pass allocates hard registers to pseudo registers that are used only within one basic block. Because the basic block is linear, it can use fast and powerful techniques to do a very good job.

The option `-dl` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.lreg` to the input file name.

- Global register allocation (`global.c`). This pass allocates hard registers for the remaining pseudo registers (those whose life spans are not contained in one basic block).
- Graph coloring register allocator. The files `ra.c`, `ra-build.c`, `ra-colorize.c`, `ra-debug.c`, `ra-rewrite.c` together with the header `ra.h` contain another register allocator, which is used when the option `-fnew-ra` is given. In that case it is run instead of the above mentioned local and global register allocation passes, and the option `-dl` causes a debugging dump of its work.
- Reloading. This pass renumbers pseudo registers with the hardware registers numbers they were allocated. Pseudo registers that did not get hard registers are re-

placed with stack slots. Then it finds instructions that are invalid because a value has failed to end up in a register, or has ended up in a register of the wrong kind. It fixes up these instructions by reloading the problematical values temporarily into registers. Additional instructions are generated to do the copying.

The reload pass also optionally eliminates the frame pointer and inserts instructions to save and restore call-clobbered registers around calls.

Source files are `reload.c` and `reload1.c`, plus the header `reload.h` used for communication between them.

The option `-dg` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.greg` to the input file name.

- Instruction scheduling is repeated here to try to avoid pipeline stalls due to memory loads generated for spilled pseudo registers.

The option `-dR` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.sched2` to the input file name.

- Basic block reordering. This pass implements profile guided code positioning. If profile information is not available, various types of static analysis are performed to make the predictions normally coming from the profile feedback (IE execution frequency, branch probability, etc). It is implemented in the file `bb-reorder.c`, and the various prediction routines are in `predict.c`.

The option `-dB` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.bbro` to the input file name.

- Delayed branch scheduling. This optional pass attempts to find instructions that can go into the delay slots of other instructions, usually jumps and calls. The source file name is `reorg.c`.

The option `-dd` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.dbr` to the input file name.

- Branch shortening. On many RISC machines, branch instructions have a limited range. Thus, longer sequences of instructions must be used for long branches. In this pass, the compiler figures out what how far each instruction will be from each other instruction, and therefore whether the usual instructions, or the longer sequences, must be used for each branch.

- Conversion from usage of some hard registers to usage of a register stack may be done at this point. Currently, this is supported only for the floating-point registers of the Intel 80387 coprocessor. The source file name is `reg-stack.c`.

The options `-dk` causes a debugging dump of the RTL code after this pass. This dump file's name is made by appending `.stack` to the input file name.

- Final. This pass outputs the assembler code for the function. It is also responsible for identifying spurious test and compare instructions. Machine-specific peephole optimizations are performed at the same time. The function entry and exit sequences are generated directly as assembler code in this pass; they never exist as RTL.

The source files are `final.c` plus `insn-output.c`; the latter is generated automatically from the machine description by the tool `genoutput`. The header file `conditions.h` is used for communication between these files.

- Debugging information output. This is run after final because it must output the stack slot offsets for pseudo registers that did not get hard registers. Source files are ‘dbxout.c’ for DBX symbol table format, ‘sdbout.c’ for SDB symbol table format, ‘dwarfout.c’ for DWARF symbol table format, files ‘dwarf2out.c’ and ‘dwarf2asm.c’ for DWARF2 symbol table format, and ‘vmsdbgout.c’ for VMS debug symbol table format.

Some additional files are used by all or many passes:

- Every pass uses ‘machmode.def’ and ‘machmode.h’ which define the machine modes.
- Several passes use ‘real.h’, which defines the default representation of floating point constants and how to operate on them.
- All the passes that work with RTL use the header files ‘rtl.h’ and ‘rtl.def’, and subroutines in file ‘rtl.c’. The tools `gen*` also use these files to read and work with the machine description RTL.
- All the tools that read the machine description use support routines found in ‘gensupport.c’, ‘errors.c’, and ‘read-rtl.c’.
- Several passes refer to the header file ‘insn-config.h’ which contains a few parameters (C macro definitions) generated automatically from the machine description RTL by the tool `genconfig`.
- Several passes use the instruction recognizer, which consists of ‘recog.c’ and ‘recog.h’, plus the files ‘insn-recog.c’ and ‘insn-extract.c’ that are generated automatically from the machine description by the tools ‘genrekog’ and ‘genextract’.
- Several passes use the header files ‘regs.h’ which defines the information recorded about pseudo register usage, and ‘basic-block.h’ which defines the information recorded about basic blocks.
- ‘hard-reg-set.h’ defines the type `HARD_REG_SET`, a bit-vector with a bit for each hard register, and some macros to manipulate it. This type is just `int` if the machine has few enough hard registers; otherwise it is an array of `int` and some of the macros expand into loops.
- Several passes use instruction attributes. A definition of the attributes defined for a particular machine is in file ‘insn-attr.h’, which is generated from the machine description by the program ‘genattr’. The file ‘insn-attrtab.c’ contains subroutines to obtain the attribute values for insns and information about processor pipeline characteristics for the instruction scheduler. It is generated from the machine description by the program ‘genattrtab’.

7 Trees: The intermediate representation used by the C and C++ front ends

This chapter documents the internal representation used by GCC to represent C and C++ source programs. When presented with a C or C++ source program, GCC parses the program, performs semantic analysis (including the generation of error messages), and then produces the internal representation described here. This representation contains a complete representation for the entire translation unit provided as input to the front end. This representation is then typically processed by a code-generator in order to produce machine code, but could also be used in the creation of source browsers, intelligent editors, automatic documentation generators, interpreters, and any other programs needing the ability to process C or C++ code.

This chapter explains the internal representation. In particular, it documents the internal representation for C and C++ source constructs, and the macros, functions, and variables that can be used to access these constructs. The C++ representation is largely a superset of the representation used in the C front end. There is only one construct used in C that does not appear in the C++ front end and that is the GNU “nested function” extension. Many of the macros documented here do not apply in C because the corresponding language constructs do not appear in C.

If you are developing a “back end”, be it is a code-generator or some other tool, that uses this representation, you may occasionally find that you need to ask questions not easily answered by the functions and macros available here. If that situation occurs, it is quite likely that GCC already supports the functionality you desire, but that the interface is simply not documented here. In that case, you should ask the GCC maintainers (via mail to gcc@gcc.gnu.org) about documenting the functionality you require. Similarly, if you find yourself writing functions that do not deal directly with your back end, but instead might be useful to other people using the GCC front end, you should submit your patches for inclusion in GCC.

7.1 Deficiencies

There are many places in which this document is incomplet and incorrekt. It is, as of yet, only *preliminary* documentation.

7.2 Overview

The central data structure used by the internal representation is the **tree**. These nodes, while all of the C type **tree**, are of many varieties. A **tree** is a pointer type, but the object to which it points may be of a variety of types. From this point forward, we will refer to trees in ordinary type, rather than in **this font**, except when talking about the actual C type **tree**.

You can tell what kind of node a particular tree is by using the `TREE_CODE` macro. Many, many macros take trees as input and return trees as output. However, most macros require a certain kind of tree node as input. In other words, there is a type-system for trees, but it is not reflected in the C type-system.

For safety, it is useful to configure GCC with ‘`--enable-checking`’. Although this results in a significant performance penalty (since all tree types are checked at run-time), and is therefore inappropriate in a release version, it is extremely helpful during the development process.

Many macros behave as predicates. Many, although not all, of these predicates end in ‘`_P`’. Do not rely on the result type of these macros being of any particular type. You may, however, rely on the fact that the type can be compared to 0, so that statements like

```
if (TEST_P (t) && !TEST_P (y))
    x = 1;
```

and

```
int i = (TEST_P (t) != 0);
```

are legal. Macros that return `int` values now may be changed to return `tree` values, or other pointers in the future. Even those that continue to return `int` may return multiple nonzero codes where previously they returned only zero and one. Therefore, you should not write code like

```
if (TEST_P (t) == 1)
```

as this code is not guaranteed to work correctly in the future.

You should not take the address of values returned by the macros or functions described here. In particular, no guarantee is given that the values are lvalues.

In general, the names of macros are all in uppercase, while the names of functions are entirely in lower case. There are rare exceptions to this rule. You should assume that any macro or function whose name is made up entirely of uppercase letters may evaluate its arguments more than once. You may assume that a macro or function whose name is made up entirely of lowercase letters will evaluate its arguments only once.

The `error_mark_node` is a special tree. Its tree code is `ERROR_MARK`, but since there is only ever one node with that code, the usual practice is to compare the tree against `error_mark_node`. (This test is just a test for pointer equality.) If an error has occurred during front-end processing the flag `errorcount` will be set. If the front end has encountered code it cannot handle, it will issue a message to the user and set `sorrycount`. When these flags are set, any macro or function which normally returns a tree of a particular kind may instead return the `error_mark_node`. Thus, if you intend to do any processing of erroneous code, you must be prepared to deal with the `error_mark_node`.

Occasionally, a particular tree slot (like an operand to an expression, or a particular field in a declaration) will be referred to as “reserved for the back end.” These slots are used to store RTL when the tree is converted to RTL for use by the GCC back end. However, if that process is not taking place (e.g., if the front end is being hooked up to an intelligent editor), then those slots may be used by the back end presently in use.

If you encounter situations that do not match this documentation, such as tree nodes of types not mentioned here, or macros documented to return entities of a particular kind that instead return entities of some different kind, you have found a bug, either in the front end or in the documentation. Please report these bugs as you would any other bug.

7.2.1 Trees

This section is not here yet.

7.2.2 Identifiers

An `IDENTIFIER_NODE` represents a slightly more general concept than the standard C or C++ concept of identifier. In particular, an `IDENTIFIER_NODE` may contain a '\$', or other extraordinary characters.

There are never two distinct `IDENTIFIER_NODES` representing the same identifier. Therefore, you may use pointer equality to compare `IDENTIFIER_NODES`, rather than using a routine like `strcmp`.

You can use the following macros to access identifiers:

`IDENTIFIER_POINTER`

The string represented by the identifier, represented as a `char*`. This string is always NUL-terminated, and contains no embedded NUL characters.

`IDENTIFIER_LENGTH`

The length of the string returned by `IDENTIFIER_POINTER`, not including the trailing NUL. This value of `IDENTIFIER_LENGTH (x)` is always the same as `strlen (IDENTIFIER_POINTER (x))`.

`IDENTIFIER_OPNAME_P`

This predicate holds if the identifier represents the name of an overloaded operator. In this case, you should not depend on the contents of either the `IDENTIFIER_POINTER` or the `IDENTIFIER_LENGTH`.

`IDENTIFIER_TYPENAME_P`

This predicate holds if the identifier represents the name of a user-defined conversion operator. In this case, the `TREE_TYPE` of the `IDENTIFIER_NODE` holds the type to which the conversion operator converts.

7.2.3 Containers

Two common container data structures can be represented directly with tree nodes. A `TREE_LIST` is a singly linked list containing two trees per node. These are the `TREE_PURPOSE` and `TREE_VALUE` of each node. (Often, the `TREE_PURPOSE` contains some kind of tag, or additional information, while the `TREE_VALUE` contains the majority of the payload. In other cases, the `TREE_PURPOSE` is simply `NULL_TREE`, while in still others both the `TREE_PURPOSE` and `TREE_VALUE` are of equal stature.) Given one `TREE_LIST` node, the next node is found by following the `TREE_CHAIN`. If the `TREE_CHAIN` is `NULL_TREE`, then you have reached the end of the list.

A `TREE_VEC` is a simple vector. The `TREE_VEC_LENGTH` is an integer (not a tree) giving the number of nodes in the vector. The nodes themselves are accessed using the `TREE_VEC_ELT` macro, which takes two arguments. The first is the `TREE_VEC` in question; the second is an integer indicating which element in the vector is desired. The elements are indexed from zero.

7.3 Types

All types have corresponding tree nodes. However, you should not assume that there is exactly one tree node corresponding to each type. There are often several nodes each of which correspond to the same type.

For the most part, different kinds of types have different tree codes. (For example, pointer types use a `POINTER_TYPE` code while arrays use an `ARRAY_TYPE` code.) However, pointers to member functions use the `RECORD_TYPE` code. Therefore, when writing a `switch` statement that depends on the code associated with a particular type, you should take care to handle pointers to member functions under the `RECORD_TYPE` case label.

In C++, an array type is not qualified; rather the type of the array elements is qualified. This situation is reflected in the intermediate representation. The macros described here will always examine the qualification of the underlying element type when applied to an array type. (If the element type is itself an array, then the recursion continues until a non-array type is found, and the qualification of this type is examined.) So, for example, `CP_TYPE_CONST_P` will hold of the type `const int () [7]`, denoting an array of seven `ints`.

The following functions and macros deal with cv-qualification of types:

`CP_TYPE_QUALS`

This macro returns the set of type qualifiers applied to this type. This value is `TYPE_UNQUALIFIED` if no qualifiers have been applied. The `TYPE_QUAL_CONST` bit is set if the type is `const`-qualified. The `TYPE_QUAL_VOLATILE` bit is set if the type is `volatile`-qualified. The `TYPE_QUAL_RESTRICT` bit is set if the type is `restrict`-qualified.

`CP_TYPE_CONST_P`

This macro holds if the type is `const`-qualified.

`CP_TYPE_VOLATILE_P`

This macro holds if the type is `volatile`-qualified.

`CP_TYPE_RESTRICT_P`

This macro holds if the type is `restrict`-qualified.

`CP_TYPE_CONST_NON_VOLATILE_P`

This predicate holds for a type that is `const`-qualified, but *not* `volatile`-qualified; other cv-qualifiers are ignored as well: only the `const`-ness is tested.

`TYPE_MAIN_VARIANT`

This macro returns the unqualified version of a type. It may be applied to an unqualified type, but it is not always the identity function in that case.

A few other macros and functions are usable with all types:

`TYPE_SIZE`

The number of bits required to represent the type, represented as an `INTEGER_CST`. For an incomplete type, `TYPE_SIZE` will be `NULL_TREE`.

`TYPE_ALIGN`

The alignment of the type, in bits, represented as an `int`.

TYPE_NAME

This macro returns a declaration (in the form of a `TYPE_DECL`) for the type. (Note this macro does *not* return a `IDENTIFIER_NODE`, as you might expect, given its name!) You can look at the `DECL_NAME` of the `TYPE_DECL` to obtain the actual name of the type. The `TYPE_NAME` will be `NULL_TREE` for a type that is not a built-in type, the result of a `typedef`, or a named class type.

CP_INTEGRAL_TYPE

This predicate holds if the type is an integral type. Notice that in C++, enumerations are *not* integral types.

ARITHMETIC_TYPE_P

This predicate holds if the type is an integral type (in the C++ sense) or a floating point type.

CLASS_TYPE_P

This predicate holds for a class-type.

TYPE_BUILT_IN

This predicate holds for a built-in type.

TYPE_PTRMEM_P

This predicate holds if the type is a pointer to data member.

TYPE_PTR_P

This predicate holds if the type is a pointer type, and the pointee is not a data member.

TYPE_PTRFN_P

This predicate holds for a pointer to function type.

TYPE_PTROB_P

This predicate holds for a pointer to object type. Note however that it does not hold for the generic pointer to object type `void *`. You may use `TYPE_PTROBV_P` to test for a pointer to object type as well as `void *`.

same_type_p

This predicate takes two types as input, and holds if they are the same type. For example, if one type is a `typedef` for the other, or both are `typedefs` for the same type. This predicate also holds if the two trees given as input are simply copies of one another; i.e., there is no difference between them at the source level, but, for whatever reason, a duplicate has been made in the representation. You should never use `==` (pointer equality) to compare types; always use `same_type_p` instead.

Detailed below are the various kinds of types, and the macros that can be used to access them. Although other kinds of types are used elsewhere in G++, the types described here are the only ones that you will encounter while examining the intermediate representation.

VOID_TYPE

Used to represent the `void` type.

INTEGER_TYPE

Used to represent the various integral types, including `char`, `short`, `int`, `long`, and `long long`. This code is not used for enumeration types, nor for the `bool`

type. Note that GCC's `CHAR_TYPE` node is *not* used to represent `char`. The `TYPE_PRECISION` is the number of bits used in the representation, represented as an `unsigned int`. (Note that in the general case this is not the same value as `TYPE_SIZE`; suppose that there were a 24-bit integer type, but that alignment requirements for the ABI required 32-bit alignment. Then, `TYPE_SIZE` would be an `INTEGER_CST` for 32, while `TYPE_PRECISION` would be 24.) The integer type is unsigned if `TREE_UNSIGNED` holds; otherwise, it is signed.

The `TYPE_MIN_VALUE` is an `INTEGER_CST` for the smallest integer that may be represented by this type. Similarly, the `TYPE_MAX_VALUE` is an `INTEGER_CST` for the largest integer that may be represented by this type.

`REAL_TYPE`

Used to represent the `float`, `double`, and `long double` types. The number of bits in the floating-point representation is given by `TYPE_PRECISION`, as in the `INTEGER_TYPE` case.

`COMPLEX_TYPE`

Used to represent GCC built-in `__complex__` data types. The `TREE_TYPE` is the type of the real and imaginary parts.

`ENUMERAL_TYPE`

Used to represent an enumeration type. The `TYPE_PRECISION` gives (as an `int`), the number of bits used to represent the type. If there are no negative enumeration constants, `TREE_UNSIGNED` will hold. The minimum and maximum enumeration constants may be obtained with `TYPE_MIN_VALUE` and `TYPE_MAX_VALUE`, respectively; each of these macros returns an `INTEGER_CST`.

The actual enumeration constants themselves may be obtained by looking at the `TYPE_VALUES`. This macro will return a `TREE_LIST`, containing the constants. The `TREE_PURPOSE` of each node will be an `IDENTIFIER_NODE` giving the name of the constant; the `TREE_VALUE` will be an `INTEGER_CST` giving the value assigned to that constant. These constants will appear in the order in which they were declared. The `TREE_TYPE` of each of these constants will be the type of enumeration type itself.

`BOOLEAN_TYPE`

Used to represent the `bool` type.

`POINTER_TYPE`

Used to represent pointer types, and pointer to data member types. The `TREE_TYPE` gives the type to which this type points. If the type is a pointer to data member type, then `TYPE_PTRMEM_P` will hold. For a pointer to data member type of the form `'T X : *'`, `TYPE_PTRMEM_CLASS_TYPE` will be the type X, while `TYPE_PTRMEM_POINTED_TO_TYPE` will be the type T.

`REFERENCE_TYPE`

Used to represent reference types. The `TREE_TYPE` gives the type to which this type refers.

`FUNCTION_TYPE`

Used to represent the type of non-member functions and of static member functions. The `TREE_TYPE` gives the return type of the function. The `TYPE_`

`ARG_TYPES` are a `TREE_LIST` of the argument types. The `TREE_VALUE` of each node in this list is the type of the corresponding argument; the `TREE_PURPOSE` is an expression for the default argument value, if any. If the last node in the list is `void_list_node` (a `TREE_LIST` node whose `TREE_VALUE` is the `void_type_node`), then functions of this type do not take variable arguments. Otherwise, they do take a variable number of arguments.

Note that in C (but not in C++) a function declared like `void f()` is an unprototyped function taking a variable number of arguments; the `TYPE_ARG_TYPES` of such a function will be `NULL`.

`METHOD_TYPE`

Used to represent the type of a non-static member function. Like a `FUNCTION_TYPE`, the return type is given by the `TREE_TYPE`. The type of `*this`, i.e., the class of which functions of this type are a member, is given by the `TYPE_METHOD_BASETYPE`. The `TYPE_ARG_TYPES` is the parameter list, as for a `FUNCTION_TYPE`, and includes the `this` argument.

`ARRAY_TYPE`

Used to represent array types. The `TREE_TYPE` gives the type of the elements in the array. If the array-bound is present in the type, the `TYPE_DOMAIN` is an `INTEGER_TYPE` whose `TYPE_MIN_VALUE` and `TYPE_MAX_VALUE` will be the lower and upper bounds of the array, respectively. The `TYPE_MIN_VALUE` will always be an `INTEGER_CST` for zero, while the `TYPE_MAX_VALUE` will be one less than the number of elements in the array, i.e., the highest value which may be used to index an element in the array.

`RECORD_TYPE`

Used to represent `struct` and `class` types, as well as pointers to member functions and similar constructs in other languages. `TYPE_FIELDS` contains the items contained in this type, each of which can be a `FIELD_DECL`, `VAR_DECL`, `CONST_DECL`, or `TYPE_DECL`. You may not make any assumptions about the ordering of the fields in the type or whether one or more of them overlap. If `TYPE_PTRMEMFUNC_P` holds, then this type is a pointer-to-member type. In that case, the `TYPE_PTRMEMFUNC_FN_TYPE` is a `POINTER_TYPE` pointing to a `METHOD_TYPE`. The `METHOD_TYPE` is the type of a function pointed to by the pointer-to-member function. If `TYPE_PTRMEMFUNC_P` does not hold, this type is a class type. For more information, see Section 7.4.2 [Classes], page 48.

`UNION_TYPE`

Used to represent `union` types. Similar to `RECORD_TYPE` except that all `FIELD_DECL` nodes in `TYPE_FIELD` start at bit position zero.

`QUAL_UNION_TYPE`

Used to represent part of a variant record in Ada. Similar to `UNION_TYPE` except that each `FIELD_DECL` has a `DECL_QUALIFIER` field, which contains a boolean expression that indicates whether the field is present in the object. The type will only have one field, so each field's `DECL_QUALIFIER` is only evaluated if none of the expressions in the previous fields in `TYPE_FIELDS` are nonzero. Normally these expressions will reference a field in the outer object using a `PLACEHOLDER_EXPR`.

UNKNOWN_TYPE

This node is used to represent a type the knowledge of which is insufficient for a sound processing.

OFFSET_TYPE

This node is used to represent a data member; for example a pointer-to-data-member is represented by a **POINTER_TYPE** whose **TREE_TYPE** is an **OFFSET_TYPE**. For a data member **X::m** the **TYPE_OFFSET_BASETYPE** is **X** and the **TREE_TYPE** is the type of **m**.

TYPENAME_TYPE

Used to represent a construct of the form `typename T::A`. The **TYPE_CONTEXT** is **T**; the **TYPE_NAME** is an **IDENTIFIER_NODE** for **A**. If the type is specified via a template-id, then **TYPENAME_TYPE_FULLNAME** yields a **TEMPLATE_ID_EXPR**. The **TREE_TYPE** is non-NULL if the node is implicitly generated in support for the implicit typename extension; in which case the **TREE_TYPE** is a type node for the base-class.

TYPEOF_TYPE

Used to represent the `__typeof__` extension. The **TYPE_FIELDS** is the expression the type of which is being represented.

There are variables whose values represent some of the basic types. These include:

void_type_node

A node for `void`.

integer_type_node

A node for `int`.

unsigned_type_node.

A node for `unsigned int`.

char_type_node.

A node for `char`.

It may sometimes be useful to compare one of these variables with a type in hand, using `same_type_p`.

7.4 Scopes

The root of the entire intermediate representation is the variable `global_namespace`. This is the namespace specified with `::` in C++ source code. All other namespaces, types, variables, functions, and so forth can be found starting with this namespace.

Besides namespaces, the other high-level scoping construct in C++ is the class. (Throughout this manual the term *class* is used to mean the types referred to in the ANSI/ISO C++ Standard as classes; these include types defined with the `class`, `struct`, and `union` keywords.)

7.4.1 Namespaces

A namespace is represented by a `NAMESPACE_DECL` node.

However, except for the fact that it is distinguished as the root of the representation, the global namespace is no different from any other namespace. Thus, in what follows, we describe namespaces generally, rather than the global namespace in particular.

The following macros and functions can be used on a `NAMESPACE_DECL`:

`DECL_NAME`

This macro is used to obtain the `IDENTIFIER_NODE` corresponding to the unqualified name of the name of the namespace (see Section 7.2.2 [Identifiers], page 41). The name of the global namespace is `::`, even though in C++ the global namespace is unnamed. However, you should use comparison with `global_namespace`, rather than `DECL_NAME` to determine whether or not a namespace is the global one. An unnamed namespace will have a `DECL_NAME` equal to `anonymous_namespace_name`. Within a single translation unit, all unnamed namespaces will have the same name.

`DECL_CONTEXT`

This macro returns the enclosing namespace. The `DECL_CONTEXT` for the `global_namespace` is `NULL_TREE`.

`DECL_NAMESPACE_ALIAS`

If this declaration is for a namespace alias, then `DECL_NAMESPACE_ALIAS` is the namespace for which this one is an alias.

Do not attempt to use `cp_namespace_decls` for a namespace which is an alias. Instead, follow `DECL_NAMESPACE_ALIAS` links until you reach an ordinary, non-alias, namespace, and call `cp_namespace_decls` there.

`DECL_NAMESPACE_STD_P`

This predicate holds if the namespace is the special `::std` namespace.

`cp_namespace_decls`

This function will return the declarations contained in the namespace, including types, overloaded functions, other namespaces, and so forth. If there are no declarations, this function will return `NULL_TREE`. The declarations are connected through their `TREE_CHAIN` fields.

Although most entries on this list will be declarations, `TREE_LIST` nodes may also appear. In this case, the `TREE_VALUE` will be an `OVERLOAD`. The value of the `TREE_PURPOSE` is unspecified; back ends should ignore this value. As with the other kinds of declarations returned by `cp_namespace_decls`, the `TREE_CHAIN` will point to the next declaration in this list.

For more information on the kinds of declarations that can occur on this list, See Section 7.5 [Declarations], page 49. Some declarations will not appear on this list. In particular, no `FIELD_DECL`, `LABEL_DECL`, or `PARAM_DECL` nodes will appear here.

This function cannot be used with namespaces that have `DECL_NAMESPACE_ALIAS` set.

7.4.2 Classes

A class type is represented by either a `RECORD_TYPE` or a `UNION_TYPE`. A class declared with the `union` tag is represented by a `UNION_TYPE`, while classes declared with either the `struct` or the `class` tag are represented by `RECORD_TYPE`s. You can use the `CLASSTYPE_DECLARED_CLASS` macro to discern whether or not a particular type is a `class` as opposed to a `struct`. This macro will be true only for classes declared with the `class` tag.

Almost all non-function members are available on the `TYPE_FIELDS` list. Given one member, the next can be found by following the `TREE_CHAIN`. You should not depend in any way on the order in which fields appear on this list. All nodes on this list will be ‘DECL’ nodes. A `FIELD_DECL` is used to represent a non-static data member, a `VAR_DECL` is used to represent a static data member, and a `TYPE_DECL` is used to represent a type. Note that the `CONST_DECL` for an enumeration constant will appear on this list, if the enumeration type was declared in the class. (Of course, the `TYPE_DECL` for the enumeration type will appear here as well.) There are no entries for base classes on this list. In particular, there is no `FIELD_DECL` for the “base-class portion” of an object.

The `TYPE_VFIELD` is a compiler-generated field used to point to virtual function tables. It may or may not appear on the `TYPE_FIELDS` list. However, back ends should handle the `TYPE_VFIELD` just like all the entries on the `TYPE_FIELDS` list.

The function members are available on the `TYPE_METHODS` list. Again, subsequent members are found by following the `TREE_CHAIN` field. If a function is overloaded, each of the overloaded functions appears; no `OVERLOAD` nodes appear on the `TYPE_METHODS` list. Implicitly declared functions (including default constructors, copy constructors, assignment operators, and destructors) will appear on this list as well.

Every class has an associated *binfo*, which can be obtained with `TYPE_BINFO`. Binfos are used to represent base-classes. The binfo given by `TYPE_BINFO` is the degenerate case, whereby every class is considered to be its own base-class. The base classes for a particular binfo can be obtained with `BINFO_BASETYPES`. These base-classes are themselves binfos. The class type associated with a binfo is given by `BINFO_TYPE`. It is always the case that `BINFO_TYPE (TYPE_BINFO (x))` is the same type as `x`, up to qualifiers. However, it is not always the case that `TYPE_BINFO (BINFO_TYPE (y))` is always the same binfo as `y`. The reason is that if `y` is a binfo representing a base-class `B` of a derived class `D`, then `BINFO_TYPE (y)` will be `B`, and `TYPE_BINFO (BINFO_TYPE (y))` will be `B` as its own base-class, rather than as a base-class of `D`.

The `BINFO_BASETYPES` is a `TREE_VEC` (see Section 7.2.3 [Containers], page 41). Base types appear in left-to-right order in this vector. You can tell whether or `public`, `protected`, or `private` inheritance was used by using the `TREE_VIA_PUBLIC`, `TREE_VIA_PROTECTED`, and `TREE_VIA_PRIVATE` macros. Each of these macros takes a `BINFO` and is true if and only if the indicated kind of inheritance was used. If `TREE_VIA_VIRTUAL` holds of a binfo, then its `BINFO_TYPE` was inherited from virtually.

The following macros can be used on a tree node representing a class-type.

`LOCAL_CLASS_P`

This predicate holds if the class is local class *i.e.* declared inside a function body.

TYPE_POLYMORPHIC_P

This predicate holds if the class has at least one virtual function (declared or inherited).

TYPE_HAS_DEFAULT_CONSTRUCTOR

This predicate holds whenever its argument represents a class-type with default constructor.

CLASSTYPE_HAS_MUTABLE**TYPE_HAS_MUTABLE_P**

These predicates hold for a class-type having a mutable data member.

CLASSTYPE_NON_POD_P

This predicate holds only for class-types that are not PODs.

TYPE_HAS_NEW_OPERATOR

This predicate holds for a class-type that defines `operator new`.

TYPE_HAS_ARRAY_NEW_OPERATOR

This predicate holds for a class-type for which `operator new[]` is defined.

TYPE_OVERLOADS_CALL_EXPR

This predicate holds for class-type for which the function call `operator()` is overloaded.

TYPE_OVERLOADS_ARRAY_REF

This predicate holds for a class-type that overloads `operator[]`

TYPE_OVERLOADS_ARROW

This predicate holds for a class-type for which `operator->` is overloaded.

7.5 Declarations

This section covers the various kinds of declarations that appear in the internal representation, except for declarations of functions (represented by `FUNCTION_DECL` nodes), which are described in Section 7.6 [Functions], page 52.

Some macros can be used with any kind of declaration. These include:

DECL_NAME

This macro returns an `IDENTIFIER_NODE` giving the name of the entity.

TREE_TYPE

This macro returns the type of the entity declared.

DECL_SOURCE_FILE

This macro returns the name of the file in which the entity was declared, as a `char*`. For an entity declared implicitly by the compiler (like `__builtin_memcpy`), this will be the string "`<internal>`".

DECL_SOURCE_LINE

This macro returns the line number at which the entity was declared, as an `int`.

DECL_ARTIFICIAL

This predicate holds if the declaration was implicitly generated by the compiler. For example, this predicate will hold of an implicitly declared member function, or of the `TYPE_DECL` implicitly generated for a class type. Recall that in C++ code like:

```
struct S {};
```

is roughly equivalent to C code like:

```
struct S {};  
typedef struct S S;
```

The implicitly generated `typedef` declaration is represented by a `TYPE_DECL` for which `DECL_ARTIFICIAL` holds.

DECL_NAMESPACE_SCOPE_P

This predicate holds if the entity was declared at a namespace scope.

DECL_CLASS_SCOPE_P

This predicate holds if the entity was declared at a class scope.

DECL_FUNCTION_SCOPE_P

This predicate holds if the entity was declared inside a function body.

The various kinds of declarations include:

LABEL_DECL

These nodes are used to represent labels in function bodies. For more information, see Section 7.6 [Functions], page 52. These nodes only appear in block scopes.

CONST_DECL

These nodes are used to represent enumeration constants. The value of the constant is given by `DECL_INITIAL` which will be an `INTEGER_CST` with the same type as the `TREE_TYPE` of the `CONST_DECL`, i.e., an `ENUMERAL_TYPE`.

RESULT_DECL

These nodes represent the value returned by a function. When a value is assigned to a `RESULT_DECL`, that indicates that the value should be returned, via bitwise copy, by the function. You can use `DECL_SIZE` and `DECL_ALIGN` on a `RESULT_DECL`, just as with a `VAR_DECL`.

TYPE_DECL

These nodes represent `typedef` declarations. The `TREE_TYPE` is the type declared to have the name given by `DECL_NAME`. In some cases, there is no associated name.

VAR_DECL

These nodes represent variables with namespace or block scope, as well as static data members. The `DECL_SIZE` and `DECL_ALIGN` are analogous to `TYPE_SIZE` and `TYPE_ALIGN`. For a declaration, you should always use the `DECL_SIZE` and `DECL_ALIGN` rather than the `TYPE_SIZE` and `TYPE_ALIGN` given by the `TREE_TYPE`, since special attributes may have been applied to the variable to give it a particular size and alignment. You may use the predicates `DECL_THIS_STATIC`

or `DECL_THIS_EXTERN` to test whether the storage class specifiers `static` or `extern` were used to declare a variable.

If this variable is initialized (but does not require a constructor), the `DECL_INITIAL` will be an expression for the initializer. The initializer should be evaluated, and a bitwise copy into the variable performed. If the `DECL_INITIAL` is the `error_mark_node`, there is an initializer, but it is given by an explicit statement later in the code; no bitwise copy is required.

GCC provides an extension that allows either automatic variables, or global variables, to be placed in particular registers. This extension is being used for a particular `VAR_DECL` if `DECL_REGISTER` holds for the `VAR_DECL`, and if `DECL_ASSEMBLER_NAME` is not equal to `DECL_NAME`. In that case, `DECL_ASSEMBLER_NAME` is the name of the register into which the variable will be placed.

PARAM_DECL

Used to represent a parameter to a function. Treat these nodes similarly to `VAR_DECL` nodes. These nodes only appear in the `DECL_ARGUMENTS` for a `FUNCTION_DECL`.

The `DECL_ARG_TYPE` for a `PARAM_DECL` is the type that will actually be used when a value is passed to this function. It may be a wider type than the `TREE_TYPE` of the parameter; for example, the ordinary type might be `short` while the `DECL_ARG_TYPE` is `int`.

FIELD_DECL

These nodes represent non-static data members. The `DECL_SIZE` and `DECL_ALIGN` behave as for `VAR_DECL` nodes. The `DECL_FIELD_BITPOS` gives the first bit used for this field, as an `INTEGER_CST`. These values are indexed from zero, where zero indicates the first bit in the object.

If `DECL_C_BIT_FIELD` holds, this field is a bit-field.

NAMESPACE_DECL

See Section 7.4.1 [Namespaces], page 47.

TEMPLATE_DECL

These nodes are used to represent class, function, and variable (static data member) templates. The `DECL_TEMPLATE_SPECIALIZATIONS` are a `TREE_LIST`. The `TREE_VALUE` of each node in the list is a `TEMPLATE_DECL`s or `FUNCTION_DECL`s representing specializations (including instantiations) of this template. Back ends can safely ignore `TEMPLATE_DECL`s, but should examine `FUNCTION_DECL` nodes on the specializations list just as they would ordinary `FUNCTION_DECL` nodes.

For a class template, the `DECL_TEMPLATE_INSTANTIATIONS` list contains the instantiations. The `TREE_VALUE` of each node is an instantiation of the class. The `DECL_TEMPLATE_SPECIALIZATIONS` contains partial specializations of the class.

USING_DECL

Back ends can safely ignore these nodes.

7.6 Functions

A function is represented by a `FUNCTION_DECL` node. A set of overloaded functions is sometimes represented by a `OVERLOAD` node.

An `OVERLOAD` node is not a declaration, so none of the ‘`DECL_`’ macros should be used on an `OVERLOAD`. An `OVERLOAD` node is similar to a `TREE_LIST`. Use `OVL_CURRENT` to get the function associated with an `OVERLOAD` node; use `OVL_NEXT` to get the next `OVERLOAD` node in the list of overloaded functions. The macros `OVL_CURRENT` and `OVL_NEXT` are actually polymorphic; you can use them to work with `FUNCTION_DECL` nodes as well as with overloads. In the case of a `FUNCTION_DECL`, `OVL_CURRENT` will always return the function itself, and `OVL_NEXT` will always be `NULL_TREE`.

To determine the scope of a function, you can use the `DECL_REAL_CONTEXT` macro. This macro will return the class (either a `RECORD_TYPE` or a `UNION_TYPE`) or namespace (a `NAMESPACE_DECL`) of which the function is a member. For a virtual function, this macro returns the class in which the function was actually defined, not the base class in which the virtual declaration occurred. If a friend function is defined in a class scope, the `DECL_CLASS_CONTEXT` macro can be used to determine the class in which it was defined. For example, in

```
class C { friend void f() {} };
```

the `DECL_REAL_CONTEXT` for `f` will be the `global_namespace`, but the `DECL_CLASS_CONTEXT` will be the `RECORD_TYPE` for `C`.

The `DECL_REAL_CONTEXT` and `DECL_CLASS_CONTEXT` are not available in C; instead you should simply use `DECL_CONTEXT`. In C, the `DECL_CONTEXT` for a function maybe another function. This representation indicates that the GNU nested function extension is in use. For details on the semantics of nested functions, see the GCC Manual. The nested function can refer to local variables in its containing function. Such references are not explicitly marked in the tree structure; back ends must look at the `DECL_CONTEXT` for the referenced `VAR_DECL`. If the `DECL_CONTEXT` for the referenced `VAR_DECL` is not the same as the function currently being processed, and neither `DECL_EXTERNAL` nor `DECL_STATIC` hold, then the reference is to a local variable in a containing function, and the back end must take appropriate action.

7.6.1 Function Basics

The following macros and functions can be used on a `FUNCTION_DECL`:

`DECL_MAIN_P`

This predicate holds for a function that is the program entry point `::code`.

`DECL_NAME`

This macro returns the unqualified name of the function, as an `IDENTIFIER_NODE`. For an instantiation of a function template, the `DECL_NAME` is the unqualified name of the template, not something like `f<int>`. The value of `DECL_NAME` is undefined when used on a constructor, destructor, overloaded operator, or type-conversion operator, or any function that is implicitly generated by the compiler. See below for macros that can be used to distinguish these cases.

DECL_ASSEMBLER_NAME

This macro returns the mangled name of the function, also an IDENTIFIER_NODE. This name does not contain leading underscores on systems that prefix all identifiers with underscores. The mangled name is computed in the same way on all platforms; if special processing is required to deal with the object file format used on a particular platform, it is the responsibility of the back end to perform those modifications. (Of course, the back end should not modify DECL_ASSEMBLER_NAME itself.)

DECL_EXTERNAL

This predicate holds if the function is undefined.

TREE_PUBLIC

This predicate holds if the function has external linkage.

DECL_LOCAL_FUNCTION_P

This predicate holds if the function was declared at block scope, even though it has a global scope.

DECL_ANTICIPATED

This predicate holds if the function is a built-in function but its prototype is not yet explicitly declared.

DECL_EXTERN_C_FUNCTION_P

This predicate holds if the function is declared as an ‘extern "C"’ function.

DECL_LINKONCE_P

This macro holds if multiple copies of this function may be emitted in various translation units. It is the responsibility of the linker to merge the various copies. Template instantiations are the most common example of functions for which DECL_LINKONCE_P holds; G++ instantiates needed templates in all translation units which require them, and then relies on the linker to remove duplicate instantiations.

FIXME: This macro is not yet implemented.

DECL_FUNCTION_MEMBER_P

This macro holds if the function is a member of a class, rather than a member of a namespace.

DECL_STATIC_FUNCTION_P

This predicate holds if the function is a static member function.

DECL_NONSTATIC_MEMBER_FUNCTION_P

This macro holds for a non-static member function.

DECL_CONST_MEMFUNC_P

This predicate holds for a `const`-member function.

DECL_VOLATILE_MEMFUNC_P

This predicate holds for a `volatile`-member function.

DECL_CONSTRUCTOR_P

This macro holds if the function is a constructor.

DECL_NONCONVERTING_P

This predicate holds if the constructor is a non-converting constructor.

DECL_COMPLETE_CONSTRUCTOR_P

This predicate holds for a function which is a constructor for an object of a complete type.

DECL_BASE_CONSTRUCTOR_P

This predicate holds for a function which is a constructor for a base class sub-object.

DECL_COPY_CONSTRUCTOR_P

This predicate holds for a function which is a copy-constructor.

DECL_DESTRUCTOR_P

This macro holds if the function is a destructor.

DECL_COMPLETE_DESTRUCTOR_P

This predicate holds if the function is the destructor for an object a complete type.

DECL_OVERLOADED_OPERATOR_P

This macro holds if the function is an overloaded operator.

DECL_CONV_FN_P

This macro holds if the function is a type-conversion operator.

DECL_GLOBAL_CTOR_P

This predicate holds if the function is a file-scope initialization function.

DECL_GLOBAL_DTOR_P

This predicate holds if the function is a file-scope finalization function.

DECL_THUNK_P

This predicate holds if the function is a thunk.

These functions represent stub code that adjusts the `this` pointer and then jumps to another function. When the jumped-to function returns, control is transferred directly to the caller, without returning to the thunk. The first parameter to the thunk is always the `this` pointer; the thunk should add `THUNK_DELTA` to this value. (The `THUNK_DELTA` is an `int`, not an `INTEGER_CST`.)

Then, if `THUNK_VCALL_OFFSET` (an `INTEGER_CST`) is nonzero the adjusted `this` pointer must be adjusted again. The complete calculation is given by the following pseudo-code:

```

    this += THUNK_DELTA
    if (THUNK_VCALL_OFFSET)
        this += *((ptrdiff_t **) this)[THUNK_VCALL_OFFSET]

```

Finally, the thunk should jump to the location given by `DECL_INITIAL`; this will always be an expression for the address of a function.

DECL_NON_THUNK_FUNCTION_P

This predicate holds if the function is *not* a thunk function.

GLOBAL_INIT_PRIORITY

If either `DECL_GLOBAL_CTOR_P` or `DECL_GLOBAL_DTOR_P` holds, then this gives the initialization priority for the function. The linker will arrange that all functions for which `DECL_GLOBAL_CTOR_P` holds are run in increasing order of priority before `main` is called. When the program exits, all functions for which `DECL_GLOBAL_DTOR_P` holds are run in the reverse order.

DECL_ARTIFICIAL

This macro holds if the function was implicitly generated by the compiler, rather than explicitly declared. In addition to implicitly generated class member functions, this macro holds for the special functions created to implement static initialization and destruction, to compute run-time type information, and so forth.

DECL_ARGUMENTS

This macro returns the `PARAM_DECL` for the first argument to the function. Subsequent `PARAM_DECL` nodes can be obtained by following the `TREE_CHAIN` links.

DECL_RESULT

This macro returns the `RESULT_DECL` for the function.

TREE_TYPE

This macro returns the `FUNCTION_TYPE` or `METHOD_TYPE` for the function.

TYPE_RAISES_EXCEPTIONS

This macro returns the list of exceptions that a (member-)function can raise. The returned list, if non `NULL`, is comprised of nodes whose `TREE_VALUE` represents a type.

TYPE_NOTHROW_P

This predicate holds when the exception-specification of its arguments is of the form `'()'`.

DECL_ARRAY_DELETE_OPERATOR_P

This predicate holds if the function is an overloaded operator `delete[]`.

7.6.2 Function Bodies

A function that has a definition in the current translation unit will have a non-`NULL` `DECL_INITIAL`. However, back ends should not make use of the particular value given by `DECL_INITIAL`.

The `DECL_SAVED_TREE` macro will give the complete body of the function. This node will usually be a `COMPOUND_STMT` representing the outermost block of the function, but it may also be a `TRY_BLOCK`, a `RETURN_INIT`, or any other valid statement.

7.6.2.1 Statements

There are tree nodes corresponding to all of the source-level statement constructs. These are enumerated here, together with a list of the various macros that can be used to obtain information about them. There are a few macros that can be used with all statements:

STMT_LINENO

This macro returns the line number for the statement. If the statement spans multiple lines, this value will be the number of the first line on which the statement occurs. Although we mention `CASE_LABEL` below as if it were a statement, they do not allow the use of `STMT_LINENO`. There is no way to obtain the line number for a `CASE_LABEL`.

Statements do not contain information about the file from which they came; that information is implicit in the `FUNCTION_DECL` from which the statements originate.

STMT_IS_FULL_EXPR_P

In C++, statements normally constitute “full expressions”; temporaries created during a statement are destroyed when the statement is complete. However, G++ sometimes represents expressions by statements; these statements will not have `STMT_IS_FULL_EXPR_P` set. Temporaries created during such statements should be destroyed when the innermost enclosing statement with `STMT_IS_FULL_EXPR_P` set is exited.

Here is the list of the various statement nodes, and the macros used to access them. This documentation describes the use of these nodes in non-template functions (including instantiations of template functions). In template functions, the same nodes are used, but sometimes in slightly different ways.

Many of the statements have substatements. For example, a `while` loop will have a body, which is itself a statement. If the substatement is `NULL_TREE`, it is considered equivalent to a statement consisting of a single `;`, i.e., an expression statement in which the expression has been omitted. A substatement may in fact be a list of statements, connected via their `TREE_CHAINS`. So, you should always process the statement tree by looping over substatements, like this:

```
void process_stmt (stmt)
  tree stmt;
{
  while (stmt)
  {
    switch (TREE_CODE (stmt))
    {
      case IF_STMT:
        process_stmt (THEN_CLAUSE (stmt));
        /* More processing here. */
        break;

      ...
    }

    stmt = TREE_CHAIN (stmt);
  }
}
```

In other words, while the `then` clause of an `if` statement in C++ can be only one statement (although that one statement may be a compound statement), the intermediate representation will sometimes use several statements chained together.

ASM_STMT

Used to represent an inline assembly statement. For an inline assembly statement like:

```
asm ("mov x, y");
```

The `ASM_STRING` macro will return a `STRING_CST` node for "mov x, y". If the original statement made use of the extended-assembly syntax, then `ASM_OUTPUTS`, `ASM_INPUTS`, and `ASM_CLOBBERS` will be the outputs, inputs, and clobbers for the statement, represented as `STRING_CST` nodes. The extended-assembly syntax looks like:

```
asm ("fsinx %1,%0" : "=f" (result) : "f" (angle));
```

The first string is the `ASM_STRING`, containing the instruction template. The next two strings are the output and inputs, respectively; this statement has no clobbers. As this example indicates, "plain" assembly statements are merely a special case of extended assembly statements; they have no cv-qualifiers, outputs, inputs, or clobbers. All of the strings will be NUL-terminated, and will contain no embedded NUL-characters.

If the assembly statement is declared `volatile`, or if the statement was not an extended assembly statement, and is therefore implicitly volatile, then the predicate `ASM_VOLATILE_P` will hold of the `ASM_STMT`.

BREAK_STMT

Used to represent a `break` statement. There are no additional fields.

CASE_LABEL

Use to represent a `case` label, range of `case` labels, or a `default` label. If `CASE_LOW` is `NULL_TREE`, then this is a `default` label. Otherwise, if `CASE_HIGH` is `NULL_TREE`, then this is an ordinary `case` label. In this case, `CASE_LOW` is an expression giving the value of the label. Both `CASE_LOW` and `CASE_HIGH` are `INTEGER_CST` nodes. These values will have the same type as the condition expression in the switch statement.

Otherwise, if both `CASE_LOW` and `CASE_HIGH` are defined, the statement is a range of case labels. Such statements originate with the extension that allows users to write things of the form:

```
case 2 ... 5:
```

The first value will be `CASE_LOW`, while the second will be `CASE_HIGH`.

CLEANUP_STMT

Used to represent an action that should take place upon exit from the enclosing scope. Typically, these actions are calls to destructors for local objects, but back ends cannot rely on this fact. If these nodes are in fact representing such destructors, `CLEANUP_DECL` will be the `VAR_DECL` destroyed. Otherwise, `CLEANUP_DECL` will be `NULL_TREE`. In any case, the `CLEANUP_EXPR` is the expression to execute. The cleanups executed on exit from a scope should be run

in the reverse order of the order in which the associated `CLEANUP_STMTs` were encountered.

`COMPOUND_STMT`

Used to represent a brace-enclosed block. The first substatement is given by `COMPOUND_BODY`. Subsequent substatements are found by following the `TREE_CHAIN` link from one substatement to the next. The `COMPOUND_BODY` will be `NULL_TREE` if there are no substatements.

`CONTINUE_STMT`

Used to represent a `continue` statement. There are no additional fields.

`CTOR_STMT`

Used to mark the beginning (if `CTOR_BEGIN_P` holds) or end (if `CTOR_END_P` holds) of the main body of a constructor. See also `SUBOBJECT` for more information on how to use these nodes.

`DECL_STMT`

Used to represent a local declaration. The `DECL_STMT_DECL` macro can be used to obtain the entity declared. This declaration may be a `LABEL_DECL`, indicating that the label declared is a local label. (As an extension, GCC allows the declaration of labels with scope.) In C, this declaration may be a `FUNCTION_DECL`, indicating the use of the GCC nested function extension. For more information, see Section 7.6 [Functions], page 52.

`DO_STMT`

Used to represent a `do` loop. The body of the loop is given by `DO_BODY` while the termination condition for the loop is given by `DO_COND`. The condition for a `do`-statement is always an expression.

`EMPTY_CLASS_EXPR`

Used to represent a temporary object of a class with no data whose address is never taken. (All such objects are interchangeable.) The `TREE_TYPE` represents the type of the object.

`EXPR_STMT`

Used to represent an expression statement. Use `EXPR_STMT_EXPR` to obtain the expression.

`FILE_STMT`

Used to record a change in filename within the body of a function. Use `FILE_STMT_FILENAME` to obtain the new filename.

`FOR_STMT`

Used to represent a `for` statement. The `FOR_INIT_STMT` is the initialization statement for the loop. The `FOR_COND` is the termination condition. The `FOR_EXPR` is the expression executed right before the `FOR_COND` on each loop iteration; often, this expression increments a counter. The body of the loop is given by `FOR_BODY`. Note that `FOR_INIT_STMT` and `FOR_BODY` return statements, while `FOR_COND` and `FOR_EXPR` return expressions.

GOTO_STMT

Used to represent a `goto` statement. The `GOTO_DESTINATION` will usually be a `LABEL_DECL`. However, if the “computed goto” extension has been used, the `GOTO_DESTINATION` will be an arbitrary expression indicating the destination. This expression will always have pointer type. Additionally the `GOTO_FAKE_P` flag is set whenever the `goto` statement does not come from source code, but it is generated implicitly by the compiler. This is used for branch prediction.

HANDLER

Used to represent a C++ `catch` block. The `HANDLER_TYPE` is the type of exception that will be caught by this handler; it is equal (by pointer equality) to `NULL` if this handler is for all types. `HANDLER_PARMS` is the `DECL_STMT` for the catch parameter, and `HANDLER_BODY` is the `COMPOUND_STMT` for the block itself.

IF_STMT

Used to represent an `if` statement. The `IF_COND` is the expression.

If the condition is a `TREE_LIST`, then the `TREE_PURPOSE` is a statement (usually a `DECL_STMT`). Each time the condition is evaluated, the statement should be executed. Then, the `TREE_VALUE` should be used as the conditional expression itself. This representation is used to handle C++ code like this:

```
if (int i = 7) ...
```

where there is a new local variable (or variables) declared within the condition.

The `THEN_CLAUSE` represents the statement given by the `then` condition, while the `ELSE_CLAUSE` represents the statement given by the `else` condition.

LABEL_STMT

Used to represent a label. The `LABEL_DECL` declared by this statement can be obtained with the `LABEL_STMT_LABEL` macro. The `IDENTIFIER_NODE` giving the name of the label can be obtained from the `LABEL_DECL` with `DECL_NAME`.

RETURN_INIT

If the function uses the G++ “named return value” extension, meaning that the function has been defined like:

```
S f(int) return s {...}
```

then there will be a `RETURN_INIT`. There is never a named returned value for a constructor. The first argument to the `RETURN_INIT` is the name of the object returned; the second argument is the initializer for the object. The object is initialized when the `RETURN_INIT` is encountered. The object referred to is the actual object returned; this extension is a manual way of doing the “return-value optimization.” Therefore, the object must actually be constructed in the place where the object will be returned.

RETURN_STMT

Used to represent a `return` statement. The `RETURN_EXPR` is the expression returned; it will be `NULL_TREE` if the statement was just

```
return;
```

SCOPE_STMT

A scope-statement represents the beginning or end of a scope. If `SCOPE_BEGIN_P` holds, this statement represents the beginning of a scope; if `SCOPE_END_P` holds this statement represents the end of a scope. On exit from a scope, all cleanups from `CLEANUP_STMTs` occurring in the scope must be run, in reverse order to the order in which they were encountered. If `SCOPE_NULLIFIED_P` or `SCOPE_NO_CLEANUPS_P` holds of the scope, back ends should behave as if the `SCOPE_STMT` were not present at all.

SUBOBJECT

In a constructor, these nodes are used to mark the point at which a subobject of `this` is fully constructed. If, after this point, an exception is thrown before a `CTOR_STMT` with `CTOR_END_P` set is encountered, the `SUBOBJECT_CLEANUP` must be executed. The cleanups must be executed in the reverse order in which they appear.

SWITCH_STMT

Used to represent a `switch` statement. The `SWITCH_COND` is the expression on which the switch is occurring. See the documentation for an `IF_STMT` for more information on the representation used for the condition. The `SWITCH_BODY` is the body of the switch statement. The `SWITCH_TYPE` is the original type of switch expression as given in the source, before any compiler conversions.

TRY_BLOCK

Used to represent a `try` block. The body of the try block is given by `TRY_STMTs`. Each of the catch blocks is a `HANDLER` node. The first handler is given by `TRY_HANDLERS`. Subsequent handlers are obtained by following the `TREE_CHAIN` link from one handler to the next. The body of the handler is given by `HANDLER_BODY`.

If `CLEANUP_P` holds of the `TRY_BLOCK`, then the `TRY_HANDLERS` will not be a `HANDLER` node. Instead, it will be an expression that should be executed if an exception is thrown in the try block. It must rethrow the exception after executing that code. And, if an exception is thrown while the expression is executing, `terminate` must be called.

USING_STMT

Used to represent a `using` directive. The namespace is given by `USING_STMT_NAMESPACE`, which will be a `NAMESPACE_DECL`. This node is needed inside template functions, to implement using directives during instantiation.

WHILE_STMT

Used to represent a `while` loop. The `WHILE_COND` is the termination condition for the loop. See the documentation for an `IF_STMT` for more information on the representation used for the condition.

The `WHILE_BODY` is the body of the loop.

7.7 Attributes in trees

Attributes, as specified using the `__attribute__` keyword, are represented internally as a `TREE_LIST`. The `TREE_PURPOSE` is the name of the attribute, as an `IDENTIFIER_NODE`. The `TREE_VALUE` is a `TREE_LIST` of the arguments of the attribute, if any, or `NULL_TREE` if there are no arguments; the arguments are stored as the `TREE_VALUE` of successive entries in the list, and may be identifiers or expressions. The `TREE_CHAIN` of the attribute is the next attribute in a list of attributes applying to the same declaration or type, or `NULL_TREE` if there are no further attributes in the list.

Attributes may be attached to declarations and to types; these attributes may be accessed with the following macros. All attributes are stored in this way, and many also cause other changes to the declaration or type or to other internal compiler data structures.

`tree DECL_ATTRIBUTES (tree decl)` [Tree Macro]

This macro returns the attributes on the declaration *decl*.

`tree TYPE_ATTRIBUTES (tree type)` [Tree Macro]

This macro returns the attributes on the type *type*.

7.8 Expressions

The internal representation for expressions is for the most part quite straightforward. However, there are a few facts that one must bear in mind. In particular, the expression “tree” is actually a directed acyclic graph. (For example there may be many references to the integer constant zero throughout the source program; many of these will be represented by the same expression node.) You should not rely on certain kinds of node being shared, nor should rely on certain kinds of nodes being unshared.

The following macros can be used with all expression nodes:

`TREE_TYPE`

Returns the type of the expression. This value may not be precisely the same type that would be given the expression in the original program.

In what follows, some nodes that one might expect to always have type `bool` are documented to have either integral or boolean type. At some point in the future, the C front end may also make use of this same intermediate representation, and at this point these nodes will certainly have integral type. The previous sentence is not meant to imply that the C++ front end does not or will not give these nodes integral type.

Below, we list the various kinds of expression nodes. Except where noted otherwise, the operands to an expression are accessed using the `TREE_OPERAND` macro. For example, to access the first operand to a binary plus expression `expr`, use:

```
TREE_OPERAND (expr, 0)
```

As this example indicates, the operands are zero-indexed.

The table below begins with constants, moves on to unary expressions, then proceeds to binary expressions, and concludes with various other kinds of expressions:

INTEGER_CST

These nodes represent integer constants. Note that the type of these constants is obtained with `TREE_TYPE`; they are not always of type `int`. In particular, `char` constants are represented with `INTEGER_CST` nodes. The value of the integer constant `e` is given by

$$\begin{aligned} & ((\text{TREE_INT_CST_HIGH } (e) \ll \text{HOST_BITS_PER_WIDE_INT}) \\ & + \text{TREE_INST_CST_LOW } (e)) \end{aligned}$$

`HOST_BITS_PER_WIDE_INT` is at least thirty-two on all platforms. Both `TREE_INT_CST_HIGH` and `TREE_INT_CST_LOW` return a `HOST_WIDE_INT`. The value of an `INTEGER_CST` is interpreted as a signed or unsigned quantity depending on the type of the constant. In general, the expression given above will overflow, so it should not be used to calculate the value of the constant.

The variable `integer_zero_node` is an integer constant with value zero. Similarly, `integer_one_node` is an integer constant with value one. The `size_zero_node` and `size_one_node` variables are analogous, but have type `size_t` rather than `int`.

The function `tree_int_cst_lt` is a predicate which holds if its first argument is less than its second. Both constants are assumed to have the same signedness (i.e., either both should be signed or both should be unsigned.) The full width of the constant is used when doing the comparison; the usual rules about promotions and conversions are ignored. Similarly, `tree_int_cst_equal` holds if the two constants are equal. The `tree_int_cst_sgn` function returns the sign of a constant. The value is 1, 0, or -1 according on whether the constant is greater than, equal to, or less than zero. Again, the signedness of the constant's type is taken into account; an unsigned constant is never less than zero, no matter what its bit-pattern.

REAL_CST

FIXME: Talk about how to obtain representations of this constant, do comparisons, and so forth.

COMPLEX_CST

These nodes are used to represent complex number constants, that is a `__complex__` whose parts are constant nodes. The `TREE_REALPART` and `TREE_IMAGPART` return the real and the imaginary parts respectively.

VECTOR_CST

These nodes are used to represent vector constants, whose parts are constant nodes. Each individual constant node is either an integer or a double constant node. The first operand is a `TREE_LIST` of the constant nodes and is accessed through `TREE_VECTOR_CST_ELTS`.

STRING_CST

These nodes represent string-constants. The `TREE_STRING_LENGTH` returns the length of the string, as an `int`. The `TREE_STRING_POINTER` is a `char*` containing the string itself. The string may not be NUL-terminated, and it may contain embedded NUL characters. Therefore, the `TREE_STRING_LENGTH` includes the trailing NUL if it is present.

For wide string constants, the `TREE_STRING_LENGTH` is the number of bytes in the string, and the `TREE_STRING_POINTER` points to an array of the bytes of the string, as represented on the target system (that is, as integers in the target endianness). Wide and non-wide string constants are distinguished only by the `TREE_TYPE` of the `STRING_CST`.

FIXME: The formats of string constants are not well-defined when the target system bytes are not the same width as host system bytes.

`PTRMEM_CST`

These nodes are used to represent pointer-to-member constants. The `PTRMEM_CST_CLASS` is the class type (either a `RECORD_TYPE` or `UNION_TYPE` within which the pointer points), and the `PTRMEM_CST_MEMBER` is the declaration for the pointed to object. Note that the `DECL_CONTEXT` for the `PTRMEM_CST_MEMBER` is in general different from the `PTRMEM_CST_CLASS`. For example, given:

```
struct B { int i; };
struct D : public B {};
int D::*dp = &D::i;
```

The `PTRMEM_CST_CLASS` for `&D::i` is `D`, even though the `DECL_CONTEXT` for the `PTRMEM_CST_MEMBER` is `B`, since `B::i` is a member of `B`, not `D`.

`VAR_DECL`

These nodes represent variables, including static data members. For more information, see Section 7.5 [Declarations], page 49.

`NEGATE_EXPR`

These nodes represent unary negation of the single operand, for both integer and floating-point types. The type of negation can be determined by looking at the type of the expression.

`BIT_NOT_EXPR`

These nodes represent bitwise complement, and will always have integral type. The only operand is the value to be complemented.

`TRUTH_NOT_EXPR`

These nodes represent logical negation, and will always have integral (or boolean) type. The operand is the value being negated.

`PREDECREMENT_EXPR`

`PREINCREMENT_EXPR`

`POSTDECREMENT_EXPR`

`POSTINCREMENT_EXPR`

These nodes represent increment and decrement expressions. The value of the single operand is computed, and the operand incremented or decremented. In the case of `PREDECREMENT_EXPR` and `PREINCREMENT_EXPR`, the value of the expression is the value resulting after the increment or decrement; in the case of `POSTDECREMENT_EXPR` and `POSTINCREMENT_EXPR` is the value before the increment or decrement occurs. The type of the operand, like that of the result, will be either integral, boolean, or floating-point.

ADDR_EXPR

These nodes are used to represent the address of an object. (These expressions will always have pointer or reference type.) The operand may be another expression, or it may be a declaration.

As an extension, GCC allows users to take the address of a label. In this case, the operand of the ADDR_EXPR will be a LABEL_DECL. The type of such an expression is void*.

If the object addressed is not an lvalue, a temporary is created, and the address of the temporary is used.

INDIRECT_REF

These nodes are used to represent the object pointed to by a pointer. The operand is the pointer being dereferenced; it will always have pointer or reference type.

FIX_TRUNC_EXPR

These nodes represent conversion of a floating-point value to an integer. The single operand will have a floating-point type, while the complete expression will have an integral (or boolean) type. The operand is rounded towards zero.

FLOAT_EXPR

These nodes represent conversion of an integral (or boolean) value to a floating-point value. The single operand will have integral type, while the complete expression will have a floating-point type.

FIXME: How is the operand supposed to be rounded? Is this dependent on '-mieee'?

COMPLEX_EXPR

These nodes are used to represent complex numbers constructed from two expressions of the same (integer or real) type. The first operand is the real part and the second operand is the imaginary part.

CONJ_EXPR

These nodes represent the conjugate of their operand.

REALPART_EXPR**IMAGPART_EXPR**

These nodes represent respectively the real and the imaginary parts of complex numbers (their sole argument).

NON_LVALUE_EXPR

These nodes indicate that their one and only operand is not an lvalue. A backend can treat these identically to the single operand.

NOP_EXPR

These nodes are used to represent conversions that do not require any code-generation. For example, conversion of a char* to an int* does not require any code be generated; such a conversion is represented by a NOP_EXPR. The single operand is the expression to be converted. The conversion from a pointer to a reference is also represented with a NOP_EXPR.

CONVERT_EXPR

These nodes are similar to `NOP_EXPRs`, but are used in those situations where code may need to be generated. For example, if an `int*` is converted to an `int` code may need to be generated on some platforms. These nodes are never used for C++-specific conversions, like conversions between pointers to different classes in an inheritance hierarchy. Any adjustments that need to be made in such cases are always indicated explicitly. Similarly, a user-defined conversion is never represented by a `CONVERT_EXPR`; instead, the function calls are made explicit.

THROW_EXPR

These nodes represent `throw` expressions. The single operand is an expression for the code that should be executed to throw the exception. However, there is one implicit action not represented in that expression; namely the call to `__throw`. This function takes no arguments. If `setjmp/longjmp` exceptions are used, the function `__sjthrow` is called instead. The normal GCC back end uses the function `emit_throw` to generate this code; you can examine this function to see what needs to be done.

LSHIFT_EXPR**RSHIFT_EXPR**

These nodes represent left and right shifts, respectively. The first operand is the value to shift; it will always be of integral type. The second operand is an expression for the number of bits by which to shift. Right shift should be treated as arithmetic, i.e., the high-order bits should be zero-filled when the expression has unsigned type and filled with the sign bit when the expression has signed type. Note that the result is undefined if the second operand is larger than the first operand's type size.

BIT_IOR_EXPR**BIT_XOR_EXPR****BIT_AND_EXPR**

These nodes represent bitwise inclusive or, bitwise exclusive or, and bitwise and, respectively. Both operands will always have integral type.

TRUTH_ANDIF_EXPR**TRUTH_ORIF_EXPR**

These nodes represent logical and and logical or, respectively. These operators are not strict; i.e., the second operand is evaluated only if the value of the expression is not determined by evaluation of the first operand. The type of the operands, and the result type, is always of boolean or integral type.

TRUTH_AND_EXPR**TRUTH_OR_EXPR****TRUTH_XOR_EXPR**

These nodes represent logical and, logical or, and logical exclusive or. They are strict; both arguments are always evaluated. There are no corresponding oper-

ators in C or C++, but the front end will sometimes generate these expressions anyhow, if it can tell that strictness does not matter.

PLUS_EXPR
MINUS_EXPR
MULT_EXPR
TRUNC_DIV_EXPR
TRUNC_MOD_EXPR
RDIV_EXPR

These nodes represent various binary arithmetic operations. Respectively, these operations are addition, subtraction (of the second operand from the first), multiplication, integer division, integer remainder, and floating-point division. The operands to the first three of these may have either integral or floating type, but there will never be case in which one operand is of floating type and the other is of integral type.

The result of a TRUNC_DIV_EXPR is always rounded towards zero. The TRUNC_MOD_EXPR of two operands *a* and *b* is always $a - (a/b)*b$ where the division is as if computed by a TRUNC_DIV_EXPR.

ARRAY_REF

These nodes represent array accesses. The first operand is the array; the second is the index. To calculate the address of the memory accessed, you must scale the index by the size of the type of the array elements. The type of these expressions must be the type of a component of the array.

ARRAY_RANGE_REF

These nodes represent access to a range (or “slice”) of an array. The operands are the same as that for ARRAY_REF and have the same meanings. The type of these expressions must be an array whose component type is the same as that of the first operand. The range of that array type determines the amount of data these expressions access.

EXACT_DIV_EXPR

Document.

LT_EXPR
LE_EXPR
GT_EXPR
GE_EXPR
EQ_EXPR
NE_EXPR

These nodes represent the less than, less than or equal to, greater than, greater than or equal to, equal, and not equal comparison operators. The first and second operand with either be both of integral type or both of floating type. The result type of these expressions will always be of integral or boolean type.

MODIFY_EXPR

These nodes represent assignment. The left-hand side is the first operand; the right-hand side is the second operand. The left-hand side will be a VAR_DECL, INDIRECT_REF, COMPONENT_REF, or other lvalue.

These nodes are used to represent not only assignment with '=' but also compound assignments (like '+='), by reduction to '=' assignment. In other words, the representation for 'i += 3' looks just like that for 'i = i + 3'.

INIT_EXPR

These nodes are just like `MODIFY_EXPR`, but are used only when a variable is initialized, rather than assigned to subsequently.

COMPONENT_REF

These nodes represent non-static data member accesses. The first operand is the object (rather than a pointer to it); the second operand is the `FIELD_DECL` for the data member.

COMPOUND_EXPR

These nodes represent comma-expressions. The first operand is an expression whose value is computed and thrown away prior to the evaluation of the second operand. The value of the entire expression is the value of the second operand.

COND_EXPR

These nodes represent `?:` expressions. The first operand is of boolean or integral type. If it evaluates to a nonzero value, the second operand should be evaluated, and returned as the value of the expression. Otherwise, the third operand is evaluated, and returned as the value of the expression.

The second operand must have the same type as the entire expression, unless it unconditionally throws an exception or calls a noreturn function, in which case it should have void type. The same constraints apply to the third operand. This allows array bounds checks to be represented conveniently as `(i >= 0 && i < 10) ? i : abort()`.

As a GNU extension, the C language front-ends allow the second operand of the `?:` operator may be omitted in the source. For example, `x ? : 3` is equivalent to `x ? x : 3`, assuming that `x` is an expression without side-effects. In the tree representation, however, the second operand is always present, possibly protected by `SAVE_EXPR` if the first argument does cause side-effects.

CALL_EXPR

These nodes are used to represent calls to functions, including non-static member functions. The first operand is a pointer to the function to call; it is always an expression whose type is a `POINTER_TYPE`. The second argument is a `TREE_LIST`. The arguments to the call appear left-to-right in the list. The `TREE_VALUE` of each list node contains the expression corresponding to that argument. (The value of `TREE_PURPOSE` for these nodes is unspecified, and should be ignored.) For non-static member functions, there will be an operand corresponding to the `this` pointer. There will always be expressions corresponding to all of the arguments, even if the function is declared with default arguments and some arguments are not explicitly provided at the call sites.

STMT_EXPR

These nodes are used to represent GCC's statement-expression extension. The statement-expression extension allows code like this:

```
int f() { return ({ int j; j = 3; j + 7; }); }
```

In other words, an sequence of statements may occur where a single expression would normally appear. The `STMT_EXPR` node represents such an expression. The `STMT_EXPR_STMT` gives the statement contained in the expression; this is always a `COMPOUND_STMT`. The value of the expression is the value of the last sub-statement in the `COMPOUND_STMT`. More precisely, the value is the value computed by the last `EXPR_STMT` in the outermost scope of the `COMPOUND_STMT`. For example, in:

```
{ 3; }
```

the value is 3 while in:

```
{ if (x) { 3; } }
```

(represented by a nested `COMPOUND_STMT`), there is no value. If the `STMT_EXPR` does not yield a value, it's type will be `void`.

`BIND_EXPR`

These nodes represent local blocks. The first operand is a list of temporary variables, connected via their `TREE_CHAIN` field. These will never require cleanups. The scope of these variables is just the body of the `BIND_EXPR`. The body of the `BIND_EXPR` is the second operand.

`LOOP_EXPR`

These nodes represent “infinite” loops. The `LOOP_EXPR_BODY` represents the body of the loop. It should be executed forever, unless an `EXIT_EXPR` is encountered.

`EXIT_EXPR`

These nodes represent conditional exits from the nearest enclosing `LOOP_EXPR`. The single operand is the condition; if it is nonzero, then the loop should be exited. An `EXIT_EXPR` will only appear within a `LOOP_EXPR`.

`CLEANUP_POINT_EXPR`

These nodes represent full-expressions. The single operand is an expression to evaluate. Any destructor calls engendered by the creation of temporaries during the evaluation of that expression should be performed immediately after the expression is evaluated.

`CONSTRUCTOR`

These nodes represent the brace-enclosed initializers for a structure or array. The first operand is reserved for use by the back end. The second operand is a `TREE_LIST`. If the `TREE_TYPE` of the `CONSTRUCTOR` is a `RECORD_TYPE` or `UNION_TYPE`, then the `TREE_PURPOSE` of each node in the `TREE_LIST` will be a `FIELD_DECL` and the `TREE_VALUE` of each node will be the expression used to initialize that field. You should not depend on the fields appearing in any particular order, nor should you assume that all fields will be represented. Unrepresented fields may be assigned any value.

If the `TREE_TYPE` of the `CONSTRUCTOR` is an `ARRAY_TYPE`, then the `TREE_PURPOSE` of each element in the `TREE_LIST` will be an `INTEGER_CST`. This constant indicates which element of the array (indexed from zero) is being assigned to;

again, the `TREE_VALUE` is the corresponding initializer. If the `TREE_PURPOSE` is `NULL_TREE`, then the initializer is for the next available array element.

Conceptually, before any initialization is done, the entire area of storage is initialized to zero.

`COMPOUND_LITERAL_EXPR`

These nodes represent ISO C99 compound literals. The `COMPOUND_LITERAL_EXPR_DECL_STMT` is a `DECL_STMT` containing an anonymous `VAR_DECL` for the unnamed object represented by the compound literal; the `DECL_INITIAL` of that `VAR_DECL` is a `CONSTRUCTOR` representing the brace-enclosed list of initializers in the compound literal. That anonymous `VAR_DECL` can also be accessed directly by the `COMPOUND_LITERAL_EXPR_DECL` macro.

`SAVE_EXPR`

A `SAVE_EXPR` represents an expression (possibly involving side-effects) that is used more than once. The side-effects should occur only the first time the expression is evaluated. Subsequent uses should just reuse the computed value. The first operand to the `SAVE_EXPR` is the expression to evaluate. The side-effects should be executed where the `SAVE_EXPR` is first encountered in a depth-first preorder traversal of the expression tree.

`TARGET_EXPR`

A `TARGET_EXPR` represents a temporary object. The first operand is a `VAR_DECL` for the temporary variable. The second operand is the initializer for the temporary. The initializer is evaluated, and copied (bitwise) into the temporary. Often, a `TARGET_EXPR` occurs on the right-hand side of an assignment, or as the second operand to a comma-expression which is itself the right-hand side of an assignment, etc. In this case, we say that the `TARGET_EXPR` is “normal”; otherwise, we say it is “orphaned”. For a normal `TARGET_EXPR` the temporary variable should be treated as an alias for the left-hand side of the assignment, rather than as a new temporary variable.

The third operand to the `TARGET_EXPR`, if present, is a cleanup-expression (i.e., destructor call) for the temporary. If this expression is orphaned, then this expression must be executed when the statement containing this expression is complete. These cleanups must always be executed in the order opposite to that in which they were encountered. Note that if a temporary is created on one branch of a conditional operator (i.e., in the second or third operand to a `COND_EXPR`), the cleanup must be run only if that branch is actually executed. See `STMT_IS_FULL_EXPR_P` for more information about running these cleanups.

`AGGR_INIT_EXPR`

An `AGGR_INIT_EXPR` represents the initialization as the return value of a function call, or as the result of a constructor. An `AGGR_INIT_EXPR` will only appear as the second operand of a `TARGET_EXPR`. The first operand to the `AGGR_INIT_EXPR` is the address of a function to call, just as in a `CALL_EXPR`. The second operand are the arguments to pass that function, as a `TREE_LIST`, again in a manner similar to that of a `CALL_EXPR`. The value of the expression is that returned by the function.

If `AGGR_INIT_VIA_CTOR_P` holds of the `AGGR_INIT_EXPR`, then the initialization is via a constructor call. The address of the third operand of the `AGGR_INIT_EXPR`, which is always a `VAR_DECL`, is taken, and this value replaces the first argument in the argument list. In this case, the value of the expression is the `VAR_DECL` given by the third operand to the `AGGR_INIT_EXPR`; constructors do not return a value.

`VTABLE_REF`

A `VTABLE_REF` indicates that the interior expression computes a value that is a vtable entry. It is used with `-fvtable-gc` to track the reference through to front end to the middle end, at which point we transform this to a `REG_VTABLE_REF` note, which survives the balance of code generation.

The first operand is the expression that computes the vtable reference. The second operand is the `VAR_DECL` of the vtable. The third operand is an `INTEGER_CST` of the byte offset into the vtable.

`VA_ARG_EXPR`

This node is used to implement support for the C/C++ variable argument-list mechanism. It represents expressions like `va_arg (ap, type)`. Its `TREE_TYPE` yields the tree representation for `type` and its sole argument yields the representation for `ap`.

8 RTL Representation

Most of the work of the compiler is done on an intermediate representation called register transfer language. In this language, the instructions to be output are described, pretty much one by one, in an algebraic form that describes what the instruction does.

RTL is inspired by Lisp lists. It has both an internal form, made up of structures that point at other structures, and a textual form that is used in the machine description and in printed debugging dumps. The textual form uses nested parentheses to indicate the pointers in the internal form.

8.1 RTL Object Types

RTL uses five kinds of objects: expressions, integers, wide integers, strings and vectors. Expressions are the most important ones. An RTL expression (“RTX”, for short) is a C structure, but it is usually referred to with a pointer; a type that is given the typedef name `rtx`.

An integer is simply an `int`; their written form uses decimal digits. A wide integer is an integral object whose type is `HOST_WIDE_INT`; their written form uses decimal digits.

A string is a sequence of characters. In core it is represented as a `char *` in usual C fashion, and it is written in C syntax as well. However, strings in RTL may never be null. If you write an empty string in a machine description, it is represented in core as a null pointer rather than as a pointer to a null character. In certain contexts, these null pointers instead of strings are valid. Within RTL code, strings are most commonly found inside `symbol_ref` expressions, but they appear in other contexts in the RTL expressions that make up machine descriptions.

In a machine description, strings are normally written with double quotes, as you would in C. However, strings in machine descriptions may extend over many lines, which is invalid C, and adjacent string constants are not concatenated as they are in C. Any string constant may be surrounded with a single set of parentheses. Sometimes this makes the machine description easier to read.

There is also a special syntax for strings, which can be useful when C code is embedded in a machine description. Wherever a string can appear, it is also valid to write a C-style brace block. The entire brace block, including the outermost pair of braces, is considered to be the string constant. Double quote characters inside the braces are not special. Therefore, if you write string constants in the C code, you need not escape each quote character with a backslash.

A vector contains an arbitrary number of pointers to expressions. The number of elements in the vector is explicitly present in the vector. The written form of a vector consists of square brackets (`[...]`) surrounding the elements, in sequence and with whitespace separating them. Vectors of length zero are not created; null pointers are used instead.

Expressions are classified by *expression codes* (also called RTX codes). The expression code is a name defined in `rtl.def`, which is also (in upper case) a C enumeration constant. The possible expression codes and their meanings are machine-independent. The code of an RTX can be extracted with the macro `GET_CODE (x)` and altered with `PUT_CODE (x, newcode)`.

The expression code determines how many operands the expression contains, and what kinds of objects they are. In RTL, unlike Lisp, you cannot tell by looking at an operand what kind of object it is. Instead, you must know from its context—from the expression code of the containing expression. For example, in an expression of code `subreg`, the first operand is to be regarded as an expression and the second operand as an integer. In an expression of code `plus`, there are two operands, both of which are to be regarded as expressions. In a `symbol_ref` expression, there is one operand, which is to be regarded as a string.

Expressions are written as parentheses containing the name of the expression type, its flags and machine mode if any, and then the operands of the expression (separated by spaces).

Expression code names in the ‘`md`’ file are written in lower case, but when they appear in C code they are written in upper case. In this manual, they are shown as follows: `const_int`.

In a few contexts a null pointer is valid where an expression is normally wanted. The written form of this is `(nil)`.

8.2 RTL Classes and Formats

The various expression codes are divided into several *classes*, which are represented by single characters. You can determine the class of an RTX code with the macro `GET_RTX_CLASS` (*code*). Currently, ‘`rtx.def`’ defines these classes:

- o An RTX code that represents an actual object, such as a register (`REG`) or a memory location (`MEM`, `SYMBOL_REF`). Constants and basic transforms on objects (`ADDRESSOF`, `HIGH`, `LO_SUM`) are also included. Note that `SUBREG` and `STRICT_LOW_PART` are not in this class, but in class `x`.
- < An RTX code for a comparison, such as `NE` or `LT`.
- 1 An RTX code for a unary arithmetic operation, such as `NEG`, `NOT`, or `ABS`. This category also includes value extension (sign or zero) and conversions between integer and floating point.
- c An RTX code for a commutative binary operation, such as `PLUS` or `AND`. `NE` and `EQ` are comparisons, so they have class `<`.
- 2 An RTX code for a non-commutative binary operation, such as `MINUS`, `DIV`, or `ASHIFTRT`.
- b An RTX code for a bit-field operation. Currently only `ZERO_EXTRACT` and `SIGN_EXTRACT`. These have three inputs and are lvalues (so they can be used for insertion as well). See Section 8.10 [Bit-Fields], page 95.
- 3 An RTX code for other three input operations. Currently only `IF_THEN_ELSE`.
- i An RTX code for an entire instruction: `INSN`, `JUMP_INSN`, and `CALL_INSN`. See Section 8.17 [Insns], page 104.
- m An RTX code for something that matches in insns, such as `MATCH_DUP`. These only occur in machine descriptions.

- a An RTX code for an auto-increment addressing mode, such as `POST_INC`.
- x All other RTX codes. This category includes the remaining codes used only in machine descriptions (`DEFINE_*`, etc.). It also includes all the codes describing side effects (`SET`, `USE`, `CLOBBER`, etc.) and the non-insns that may appear on an insn chain, such as `NOTE`, `BARRIER`, and `CODE_LABEL`.

For each expression code, `rtl.def` specifies the number of contained objects and their kinds using a sequence of characters called the *format* of the expression code. For example, the format of `subreg` is `'ei'`.

These are the most commonly used format characters:

- e An expression (actually a pointer to an expression).
- i An integer.
- w A wide integer.
- s A string.
- E A vector of expressions.

A few other format characters are used occasionally:

- u 'u' is equivalent to 'e' except that it is printed differently in debugging dumps. It is used for pointers to insns.
- n 'n' is equivalent to 'i' except that it is printed differently in debugging dumps. It is used for the line number or code number of a `note` insn.
- S 'S' indicates a string which is optional. In the RTL objects in core, 'S' is equivalent to 's', but when the object is read, from an `'md'` file, the string value of this operand may be omitted. An omitted string is taken to be the null string.
- V 'V' indicates a vector which is optional. In the RTL objects in core, 'V' is equivalent to 'E', but when the object is read from an `'md'` file, the vector value of this operand may be omitted. An omitted vector is effectively the same as a vector of no elements.
- B 'B' indicates a pointer to basic block structure.
- O 'O' means a slot whose contents do not fit any normal category. 'O' slots are not printed at all in dumps, and are often used in special ways by small parts of the compiler.

There are macros to get the number of operands and the format of an expression code:

`GET_RTX_LENGTH (code)`

Number of operands of an RTX of code *code*.

`GET_RTX_FORMAT (code)`

The format of an RTX of code *code*, as a C string.

Some classes of RTX codes always have the same format. For example, it is safe to assume that all comparison operations have format `ee`.

- 1 All codes of this class have format `e`.
- <
- c
- 2 All codes of these classes have format `ee`.
- b
- 3 All codes of these classes have format `eee`.
- i All codes of this class have formats that begin with `iuueiee`. See Section 8.17 [Insn], page 104. Note that not all RTL objects linked onto an insn chain are of class `i`.
- o
- m
- x You can make no assumptions about the format of these codes.

8.3 Access to Operands

Operands of expressions are accessed using the macros `XEXP`, `XINT`, `XWINT` and `XSTR`. Each of these macros takes two arguments: an expression-pointer (RTL) and an operand number (counting from zero). Thus,

```
XEXP (x, 2)
```

accesses operand 2 of expression `x`, as an expression.

```
XINT (x, 2)
```

accesses the same operand as an integer. `XSTR`, used in the same fashion, would access it as a string.

Any operand can be accessed as an integer, as an expression or as a string. You must choose the correct method of access for the kind of value actually stored in the operand. You would do this based on the expression code of the containing expression. That is also how you would know how many operands there are.

For example, if `x` is a `subreg` expression, you know that it has two operands which can be correctly accessed as `XEXP (x, 0)` and `XINT (x, 1)`. If you did `XINT (x, 0)`, you would get the address of the expression operand but cast as an integer; that might occasionally be useful, but it would be cleaner to write `(int) XEXP (x, 0)`. `XEXP (x, 1)` would also compile without error, and would return the second, integer operand cast as an expression pointer, which would probably result in a crash when accessed. Nothing stops you from writing `XEXP (x, 28)` either, but this will access memory past the end of the expression with unpredictable results.

Access to operands which are vectors is more complicated. You can use the macro `XVEC` to get the vector-pointer itself, or the macros `XVECEXP` and `XVECLEN` to access the elements and length of a vector.

```
XVEC (exp, idx)
```

Access the vector-pointer which is operand number `idx` in `exp`.

```
XVECLEN (exp, idx)
```

Access the length (number of elements) in the vector which is in operand number `idx` in `exp`. This value is an `int`.

XVECEXP (*exp*, *idx*, *eltnum*)

Access element number *eltnum* in the vector which is in operand number *idx* in *exp*. This value is an RTX.

It is up to you to make sure that *eltnum* is not negative and is less than **XVECLEN** (*exp*, *idx*).

All the macros defined in this section expand into lvalues and therefore can be used to assign the operands, lengths and vector elements as well as to access them.

8.4 Flags in an RTL Expression

RTL expressions contain several flags (one-bit bit-fields) that are used in certain types of expression. Most often they are accessed with the following macros, which expand into lvalues.

CONSTANT_POOL_ADDRESS_P (*x*)

Nonzero in a **symbol_ref** if it refers to part of the current function's constant pool. For most targets these addresses are in a **.rodata** section entirely separate from the function, but for some targets the addresses are close to the beginning of the function. In either case GCC assumes these addresses can be addressed directly, perhaps with the help of base registers. Stored in the **unchanging** field and printed as **'/u'**.

CONST_OR_PURE_CALL_P (*x*)

In a **call_insn**, **note**, or an **expr_list** for notes, indicates that the insn represents a call to a const or pure function. Stored in the **unchanging** field and printed as **'/u'**.

INSN_ANNULLED_BRANCH_P (*x*)

In a **jump_insn**, **call_insn**, or **insn** indicates that the branch is an annulling one. See the discussion under **sequence** below. Stored in the **unchanging** field and printed as **'/u'**.

INSN_DEAD_CODE_P (*x*)

In an **insn** during the dead-code elimination pass, nonzero if the insn is dead. Stored in the **in_struct** field and printed as **'/s'**.

INSN_DELETED_P (*x*)

In an **insn**, **call_insn**, **jump_insn**, **code_label**, **barrier**, or **note**, nonzero if the insn has been deleted. Stored in the **volatile** field and printed as **'/v'**.

INSN_FROM_TARGET_P (*x*)

In an **insn** or **jump_insn** or **call_insn** in a delay slot of a branch, indicates that the insn is from the target of the branch. If the branch insn has **INSN_ANNULLED_BRANCH_P** set, this insn will only be executed if the branch is taken. For annulled branches with **INSN_FROM_TARGET_P** clear, the insn will be executed only if the branch is not taken. When **INSN_ANNULLED_BRANCH_P** is not set, this insn will always be executed. Stored in the **in_struct** field and printed as **'/s'**.

LABEL_OUTSIDE_LOOP_P (x)

In `label_ref` expressions, nonzero if this is a reference to a label that is outside the innermost loop containing the reference to the label. Stored in the `in_struct` field and printed as `‘/s’`.

LABEL_PRESERVE_P (x)

In a `code_label` or `note`, indicates that the label is referenced by code or data not visible to the RTL of a given function. Labels referenced by a non-local goto will have this bit set. Stored in the `in_struct` field and printed as `‘/s’`.

LABEL_REF_NONLOCAL_P (x)

In `label_ref` and `reg_label` expressions, nonzero if this is a reference to a non-local label. Stored in the `volatile` field and printed as `‘/v’`.

MEM_IN_STRUCT_P (x)

In `mem` expressions, nonzero for reference to an entire structure, union or array, or to a component of one. Zero for references to a scalar variable or through a pointer to a scalar. If both this flag and `MEM_SCALAR_P` are clear, then we don't know whether this `mem` is in a structure or not. Both flags should never be simultaneously set. Stored in the `in_struct` field and printed as `‘/s’`.

MEM_KEEP_ALIAS_SET_P (x)

In `mem` expressions, 1 if we should keep the alias set for this `mem` unchanged when we access a component. Set to 1, for example, when we are already in a non-addressable component of an aggregate. Stored in the `jump` field and printed as `‘/j’`.

MEM_SCALAR_P (x)

In `mem` expressions, nonzero for reference to a scalar known not to be a member of a structure, union, or array. Zero for such references and for indirections through pointers, even pointers pointing to scalar types. If both this flag and `MEM_IN_STRUCT_P` are clear, then we don't know whether this `mem` is in a structure or not. Both flags should never be simultaneously set. Stored in the `frame_related` field and printed as `‘/f’`.

MEM_VOLATILE_P (x)

In `mem`, `asm_operands`, and `asm_input` expressions, nonzero for volatile memory references. Stored in the `volatile` field and printed as `‘/v’`.

REG_FUNCTION_VALUE_P (x)

Nonzero in a `reg` if it is the place in which this function's value is going to be returned. (This happens only in a hard register.) Stored in the `integrated` field and printed as `‘/i’`.

REG_LOOP_TEST_P (x)

In `reg` expressions, nonzero if this register's entire life is contained in the exit test code for some loop. Stored in the `in_struct` field and printed as `‘/s’`.

REG_POINTER (x)

Nonzero in a `reg` if the register holds a pointer. Stored in the `frame_related` field and printed as `‘/f’`.

REG_USERVAR_P (x)

In a `reg`, nonzero if it corresponds to a variable present in the user's source code. Zero for temporaries generated internally by the compiler. Stored in the `volatile` field and printed as `'v'`.

The same hard register may be used also for collecting the values of functions called by this one, but `REG_FUNCTION_VALUE_P` is zero in this kind of use.

RTX_FRAME_RELATED_P (x)

Nonzero in an `insn`, `call_insn`, `jump_insn`, `barrier`, or `set` which is part of a function prologue and sets the stack pointer, sets the frame pointer, or saves a register. This flag should also be set on an instruction that sets up a temporary register to use in place of the frame pointer. Stored in the `frame_related` field and printed as `'f'`.

In particular, on RISC targets where there are limits on the sizes of immediate constants, it is sometimes impossible to reach the register save area directly from the stack pointer. In that case, a temporary register is used that is near enough to the register save area, and the Canonical Frame Address, i.e., DWARF2's logical frame pointer, register must (temporarily) be changed to be this temporary register. So, the instruction that sets this temporary register must be marked as `RTX_FRAME_RELATED_P`.

If the marked instruction is overly complex (defined in terms of what `dwarf2out_frame_debug_expr` can handle), you will also have to create a `REG_FRAME_RELATED_EXPR` note and attach it to the instruction. This note should contain a simple expression of the computation performed by this instruction, i.e., one that `dwarf2out_frame_debug_expr` can handle.

This flag is required for exception handling support on targets with RTL prologues.

RTX_INTEGRATED_P (x)

Nonzero in an `insn`, `call_insn`, `jump_insn`, `barrier`, `code_label`, `insn_list`, `const`, or `note` if it resulted from an in-line function call. Stored in the `integrated` field and printed as `'i'`.

RTX_UNCHANGING_P (x)

Nonzero in a `reg`, `mem`, or `concat` if the register or memory is set at most once, anywhere. This does not mean that it is function invariant.

GCC uses this flag to determine whether two references conflict. As implemented by `true_dependence` in `'alias.c'` for memory references, unchanging memory can't conflict with non-unchanging memory; a non-unchanging read can conflict with a non-unchanging write; an unchanging read can conflict with an unchanging write (since there may be a single store to this address to initialize it); and an unchanging store can conflict with a non-unchanging read. This means we must make conservative assumptions when choosing the value of this flag for a memory reference to an object containing both unchanging and non-unchanging fields: we must set the flag when writing to the object and clear it when reading from the object.

Stored in the `unchanging` field and printed as `'u'`.

SCHED_GROUP_P (x)

During instruction scheduling, in an `insn`, `call_insn` or `jump_insn`, indicates that the previous `insn` must be scheduled together with this `insn`. This is used to ensure that certain groups of instructions will not be split up by the instruction scheduling pass, for example, `use` `insns` before a `call_insn` may not be separated from the `call_insn`. Stored in the `in_struct` field and printed as `‘/s’`.

SET_IS_RETURN_P (x)

For a `set`, nonzero if it is for a return. Stored in the `jump` field and printed as `‘/j’`.

SIBLING_CALL_P (x)

For a `call_insn`, nonzero if the `insn` is a sibling call. Stored in the `jump` field and printed as `‘/j’`.

STRING_POOL_ADDRESS_P (x)

For a `symbol_ref` expression, nonzero if it addresses this function’s string constant pool. Stored in the `frame_related` field and printed as `‘/f’`.

SUBREG_PROMOTED_UNSIGNED_P (x)

Returns a value greater than zero for a `subreg` that has `SUBREG_PROMOTED_VAR_P` nonzero if the object being referenced is kept zero-extended, zero if it is kept sign-extended, and less than zero if it is extended some other way via the `ptr_extend` instruction. Stored in the `unchanging` field and `volatil` field, printed as `‘/u’` and `‘/v’`. This macro may only be used to get the value it may not be used to change the value. Use `SUBREG_PROMOTED_UNSIGNED_SET` to change the value.

SUBREG_PROMOTED_UNSIGNED_SET (x)

Set the `unchanging` and `volatil` fields in a `subreg` to reflect zero, sign, or other extension. If `volatil` is zero, then `unchanging` as nonzero means zero extension and as zero means sign extension. If `volatil` is nonzero then some other type of extension was done via the `ptr_extend` instruction.

SUBREG_PROMOTED_VAR_P (x)

Nonzero in a `subreg` if it was made when accessing an object that was promoted to a wider mode in accord with the `PROMOTED_MODE` machine description macro (see Section 10.5 [Storage Layout], page 205). In this case, the mode of the `subreg` is the declared mode of the object and the mode of `SUBREG_REG` is the mode of the register that holds the object. Promoted variables are always either sign- or zero-extended to the wider mode on every assignment. Stored in the `in_struct` field and printed as `‘/s’`.

SYMBOL_REF_USED (x)

In a `symbol_ref`, indicates that `x` has been used. This is normally only used to ensure that `x` is only declared external once. Stored in the `used` field.

SYMBOL_REF_WEAK (x)

In a `symbol_ref`, indicates that `x` has been declared weak. Stored in the `integrated` field and printed as `‘/i’`.

SYMBOL_REF_FLAG (x)

In a `symbol_ref`, this is used as a flag for machine-specific purposes. Stored in the `volatile` field and printed as `‘/v’`.

Most uses of `SYMBOL_REF_FLAG` are historic and may be subsumed by `SYMBOL_REF_FLAGS`. Certainly use of `SYMBOL_REF_FLAGS` is mandatory if the target requires more than one bit of storage.

These are the fields to which the above macros refer:

call This flag is currently unused.
In an RTL dump, this flag is represented as `‘/c’`.

frame_related

In an `insn` or `set` expression, 1 means that it is part of a function prologue and sets the stack pointer, sets the frame pointer, saves a register, or sets up a temporary register to use in place of the frame pointer.

In `reg` expressions, 1 means that the register holds a pointer.

In `symbol_ref` expressions, 1 means that the reference addresses this function’s string constant pool.

In `mem` expressions, 1 means that the reference is to a scalar.

In an RTL dump, this flag is represented as `‘/f’`.

in_struct

In `mem` expressions, it is 1 if the memory datum referred to is all or part of a structure or array; 0 if it is (or might be) a scalar variable. A reference through a C pointer has 0 because the pointer might point to a scalar variable. This information allows the compiler to determine something about possible cases of aliasing.

In `reg` expressions, it is 1 if the register has its entire life contained within the test expression of some loop.

In `subreg` expressions, 1 means that the `subreg` is accessing an object that has had its mode promoted from a wider mode.

In `label_ref` expressions, 1 means that the referenced label is outside the innermost loop containing the `insn` in which the `label_ref` was found.

In `code_label` expressions, it is 1 if the label may never be deleted. This is used for labels which are the target of non-local gotos. Such a label that would have been deleted is replaced with a `note` of type `NOTE_INSN_DELETED_LABEL`.

In an `insn` during dead-code elimination, 1 means that the `insn` is dead code.

In an `insn` or `jump_insn` during reorg for an `insn` in the delay slot of a branch, 1 means that this `insn` is from the target of the branch.

In an `insn` during instruction scheduling, 1 means that this `insn` must be scheduled as part of a group together with the previous `insn`.

In an RTL dump, this flag is represented as `‘/s’`.

integrated

In an `insn`, `insn_list`, or `const`, 1 means the RTL was produced by procedure integration.

- In `reg` expressions, 1 means the register contains the value to be returned by the current function. On machines that pass parameters in registers, the same register number may be used for parameters as well, but this flag is not set on such uses.
- In `symbol_ref` expressions, 1 means the referenced symbol is weak.
- In an RTL dump, this flag is represented as `‘/i’`.
- jump** In a `mem` expression, 1 means we should keep the alias set for this mem unchanged when we access a component.
- In a `set`, 1 means it is for a return.
- In a `call_insn`, 1 means it is a sibling call.
- In an RTL dump, this flag is represented as `‘/j’`.
- unchanging** In `reg` and `mem` expressions, 1 means that the value of the expression never changes.
- In `subreg` expressions, it is 1 if the `subreg` references an unsigned object whose mode has been promoted to a wider mode.
- In an `insn` or `jump_insn` in the delay slot of a branch instruction, 1 means an annulling branch should be used.
- In a `symbol_ref` expression, 1 means that this symbol addresses something in the per-function constant pool.
- In a `call_insn`, `note`, or an `expr_list` of notes, 1 means that this instruction is a call to a const or pure function.
- In an RTL dump, this flag is represented as `‘/u’`.
- used** This flag is used directly (without an access macro) at the end of RTL generation for a function, to count the number of times an expression appears in insns. Expressions that appear more than once are copied, according to the rules for shared structure (see Section 8.19 [Sharing], page 114).
- For a `reg`, it is used directly (without an access macro) by the leaf register renumbering code to ensure that each register is only renumbered once.
- In a `symbol_ref`, it indicates that an external declaration for the symbol has already been written.
- volatil** In a `mem`, `asm_operands`, or `asm_input` expression, it is 1 if the memory reference is volatile. Volatile memory references may not be deleted, reordered or combined.
- In a `symbol_ref` expression, it is used for machine-specific purposes.
- In a `reg` expression, it is 1 if the value is a user-level variable. 0 indicates an internal compiler temporary.
- In an `insn`, 1 means the `insn` has been deleted.
- In `label_ref` and `reg_label` expressions, 1 means a reference to a non-local label.
- In an RTL dump, this flag is represented as `‘/v’`.

8.5 Machine Modes

A machine mode describes a size of data object and the representation used for it. In the C code, machine modes are represented by an enumeration type, `enum machine_mode`, defined in `'machmode.def'`. Each RTL expression has room for a machine mode and so do certain kinds of tree expressions (declarations and types, to be precise).

In debugging dumps and machine descriptions, the machine mode of an RTL expression is written after the expression code with a colon to separate them. The letters 'mode' which appear at the end of each machine mode name are omitted. For example, `(reg:SI 38)` is a `reg` expression with machine mode `SImode`. If the mode is `VOIDmode`, it is not written at all.

Here is a table of machine modes. The term "byte" below refers to an object of `BITS_PER_UNIT` bits (see Section 10.5 [Storage Layout], page 205).

<code>BImode</code>	"Bit" mode represents a single bit, for predicate registers.
<code>QImode</code>	"Quarter-Integer" mode represents a single byte treated as an integer.
<code>HImode</code>	"Half-Integer" mode represents a two-byte integer.
<code>PSImode</code>	"Partial Single Integer" mode represents an integer which occupies four bytes but which doesn't really use all four. On some machines, this is the right mode to use for pointers.
<code>SImode</code>	"Single Integer" mode represents a four-byte integer.
<code>PDImode</code>	"Partial Double Integer" mode represents an integer which occupies eight bytes but which doesn't really use all eight. On some machines, this is the right mode to use for certain pointers.
<code>DImode</code>	"Double Integer" mode represents an eight-byte integer.
<code>TImode</code>	"Tetra Integer" (?) mode represents a sixteen-byte integer.
<code>OImode</code>	"Octa Integer" (?) mode represents a thirty-two-byte integer.
<code>QFmode</code>	"Quarter-Floating" mode represents a quarter-precision (single byte) floating point number.
<code>HFmode</code>	"Half-Floating" mode represents a half-precision (two byte) floating point number.
<code>TQFmode</code>	"Three-Quarter-Floating" (?) mode represents a three-quarter-precision (three byte) floating point number.
<code>SFmode</code>	"Single Floating" mode represents a four byte floating point number. In the common case, of a processor with IEEE arithmetic and 8-bit bytes, this is a single-precision IEEE floating point number; it can also be used for double-precision (on processors with 16-bit bytes) and single-precision VAX and IBM types.
<code>DFmode</code>	"Double Floating" mode represents an eight byte floating point number. In the common case, of a processor with IEEE arithmetic and 8-bit bytes, this is a double-precision IEEE floating point number.

- XFmode** “Extended Floating” mode represents a twelve byte floating point number. This mode is used for IEEE extended floating point. On some systems not all bits within these bytes will actually be used.
- TFmode** “Tetra Floating” mode represents a sixteen byte floating point number. This gets used for both the 96-bit extended IEEE floating-point types padded to 128 bits, and true 128-bit extended IEEE floating-point types.
- CCmode** “Condition Code” mode represents the value of a condition code, which is a machine-specific set of bits used to represent the result of a comparison operation. Other machine-specific modes may also be used for the condition code. These modes are not used on machines that use `cc0` (see see Section 10.15 [Condition Code], page 264).
- BLKmode** “Block” mode represents values that are aggregates to which none of the other modes apply. In RTL, only memory references can have this mode, and only if they appear in string-move or vector instructions. On machines which have no such instructions, **BLKmode** will not appear in RTL.
- VOIDmode** Void mode means the absence of a mode or an unspecified mode. For example, RTL expressions of code `const_int` have mode **VOIDmode** because they can be taken to have whatever mode the context requires. In debugging dumps of RTL, **VOIDmode** is expressed by the absence of any mode.

QCmode, HCmode, SCmode, DCmode, XCmode, TCmode

These modes stand for a complex number represented as a pair of floating point values. The floating point values are in **QFmode**, **HFmode**, **SFmode**, **DFmode**, **XFmode**, and **TFmode**, respectively.

CQImode, CHImode, CSImode, CDImode, CTImode, COImode

These modes stand for a complex number represented as a pair of integer values. The integer values are in **QImode**, **HImode**, **SImode**, **DImode**, **TImode**, and **OImode**, respectively.

The machine description defines **Pmode** as a C macro which expands into the machine mode used for addresses. Normally this is the mode whose size is `BITS_PER_WORD`, **SImode** on 32-bit machines.

The only modes which a machine description *must* support are **QImode**, and the modes corresponding to `BITS_PER_WORD`, `FLOAT_TYPE_SIZE` and `DOUBLE_TYPE_SIZE`. The compiler will attempt to use **DImode** for 8-byte structures and unions, but this can be prevented by overriding the definition of `MAX_FIXED_MODE_SIZE`. Alternatively, you can have the compiler use **TImode** for 16-byte structures and unions. Likewise, you can arrange for the C type `short int` to avoid using **HImode**.

Very few explicit references to machine modes remain in the compiler and these few references will soon be removed. Instead, the machine modes are divided into mode classes. These are represented by the enumeration type `enum mode_class` defined in ‘`machmode.h`’. The possible mode classes are:

- MODE_INT** Integer modes. By default these are **BImode**, **QImode**, **HImode**, **SImode**, **DImode**, **TImode**, and **OImode**.

MODE_PARTIAL_INT

The “partial integer” modes, PQImode, PHImode, PSImode and PDImode.

MODE_FLOAT

Floating point modes. By default these are QFmode, HFmode, TQFmode, SFmode, DFmode, XFmode and TFmode.

MODE_COMPLEX_INT

Complex integer modes. (These are not currently implemented).

MODE_COMPLEX_FLOAT

Complex floating point modes. By default these are QCmode, HCmode, SCmode, DCmode, XCmode, and TCmode.

MODE_FUNCTION

Algol or Pascal function variables including a static chain. (These are not currently implemented).

MODE_CC

Modes representing condition code values. These are CCmode plus any modes listed in the EXTRA_CC_MODES macro. See Section 9.11 [Jump Patterns], page 160, also see Section 10.15 [Condition Code], page 264.

MODE_RANDOM

This is a catchall mode class for modes which don't fit into the above classes. Currently VOIDmode and BLKmode are in MODE_RANDOM.

Here are some C macros that relate to machine modes:

GET_MODE (x)

Returns the machine mode of the RTX *x*.

PUT_MODE (x, newmode)

Alters the machine mode of the RTX *x* to be *newmode*.

NUM_MACHINE_MODES

Stands for the number of machine modes available on the target machine. This is one greater than the largest numeric value of any machine mode.

GET_MODE_NAME (m)

Returns the name of mode *m* as a string.

GET_MODE_CLASS (m)

Returns the mode class of mode *m*.

GET_MODE_WIDER_MODE (m)

Returns the next wider natural mode. For example, the expression GET_MODE_WIDER_MODE (QImode) returns HImode.

GET_MODE_SIZE (m)

Returns the size in bytes of a datum of mode *m*.

GET_MODE_BITSIZE (m)

Returns the size in bits of a datum of mode *m*.

GET_MODE_MASK (*m*)

Returns a bitmask containing 1 for all bits in a word that fit within mode *m*. This macro can only be used for modes whose bitsize is less than or equal to `HOST_BITS_PER_INT`.

GET_MODE_ALIGNMENT (*m*)

Return the required alignment, in bits, for an object of mode *m*.

GET_MODE_UNIT_SIZE (*m*)

Returns the size in bytes of the subunits of a datum of mode *m*. This is the same as `GET_MODE_SIZE` except in the case of complex modes. For them, the unit size is the size of the real or imaginary part.

GET_MODE_NUNITS (*m*)

Returns the number of units contained in a mode, i.e., `GET_MODE_SIZE` divided by `GET_MODE_UNIT_SIZE`.

GET_CLASS_NARROWEST_MODE (*c*)

Returns the narrowest mode in mode class *c*.

The global variables `byte_mode` and `word_mode` contain modes whose classes are `MODE_INT` and whose bitsizes are either `BITS_PER_UNIT` or `BITS_PER_WORD`, respectively. On 32-bit machines, these are `QImode` and `SImode`, respectively.

8.6 Constant Expression Types

The simplest RTL expressions are those that represent constant values.

(*const_int i*)

This type of expression represents the integer value *i*. *i* is customarily accessed with the macro `INTVAL` as in `INTVAL (exp)`, which is equivalent to `XWINT (exp, 0)`.

There is only one expression object for the integer value zero; it is the value of the variable `const0_rtx`. Likewise, the only expression for integer value one is found in `const1_rtx`, the only expression for integer value two is found in `const2_rtx`, and the only expression for integer value negative one is found in `constm1_rtx`. Any attempt to create an expression of code `const_int` and value zero, one, two or negative one will return `const0_rtx`, `const1_rtx`, `const2_rtx` or `constm1_rtx` as appropriate.

Similarly, there is only one object for the integer whose value is `STORE_FLAG_VALUE`. It is found in `const_true_rtx`. If `STORE_FLAG_VALUE` is one, `const_true_rtx` and `const1_rtx` will point to the same object. If `STORE_FLAG_VALUE` is `-1`, `const_true_rtx` and `constm1_rtx` will point to the same object.

(*const_double:m addr i0 i1 ...*)

Represents either a floating-point constant of mode *m* or an integer constant too large to fit into `HOST_BITS_PER_WIDE_INT` bits but small enough to fit within twice that number of bits (GCC does not provide a mechanism to represent even larger constants). In the latter case, *m* will be `VOIDmode`.

`(const_vector:m [x0 x1 ...])`

Represents a vector constant. The square brackets stand for the vector containing the constant elements. *x0*, *x1* and so on are the `const_int` or `const_double` elements.

The number of units in a `const_vector` is obtained with the macro `CONST_VECTOR_NUNITS` as in `CONST_VECTOR_NUNITS (v)`.

Individual elements in a vector constant are accessed with the macro `CONST_VECTOR_ELT` as in `CONST_VECTOR_ELT (v, n)` where *v* is the vector constant and *n* is the element desired.

addr is used to contain the `mem` expression that corresponds to the location in memory that at which the constant can be found. If it has not been allocated a memory location, but is on the chain of all `const_double` expressions in this compilation (maintained using an undisplayed field), *addr* contains `const0_rtx`. If it is not on the chain, *addr* contains `cc0_rtx`. *addr* is customarily accessed with the macro `CONST_DOUBLE_MEM` and the chain field via `CONST_DOUBLE_CHAIN`.

If *m* is `VOIDmode`, the bits of the value are stored in *i0* and *i1*. *i0* is customarily accessed with the macro `CONST_DOUBLE_LOW` and *i1* with `CONST_DOUBLE_HIGH`.

If the constant is floating point (regardless of its precision), then the number of integers used to store the value depends on the size of `REAL_VALUE_TYPE` (see Section 10.22 [Floating Point], page 310). The integers represent a floating point number, but not precisely in the target machine's or host machine's floating point format. To convert them to the precise bit pattern used by the target machine, use the macro `REAL_VALUE_TO_TARGET_DOUBLE` and friends (see Section 10.20.2 [Data Output], page 281).

The macro `CONST0_RTX (mode)` refers to an expression with value 0 in mode *mode*. If mode *mode* is of mode class `MODE_INT`, it returns `const0_rtx`. If mode *mode* is of mode class `MODE_FLOAT`, it returns a `CONST_DOUBLE` expression in mode *mode*. Otherwise, it returns a `CONST_VECTOR` expression in mode *mode*. Similarly, the macro `CONST1_RTX (mode)` refers to an expression with value 1 in mode *mode* and similarly for `CONST2_RTX`. The `CONST1_RTX` and `CONST2_RTX` macros are undefined for vector modes.

`(const_string str)`

Represents a constant string with value *str*. Currently this is used only for insn attributes (see Section 9.18 [Insn Attributes], page 175) since constant strings in C are placed in memory.

`(symbol_ref:mode symbol)`

Represents the value of an assembler label for data. *symbol* is a string that describes the name of the assembler label. If it starts with a '*', the label is the rest of *symbol* not including the '*'. Otherwise, the label is *symbol*, usually prefixed with '_

The `symbol_ref` contains a mode, which is usually `Pmode`. Usually that is the only mode for which a symbol is directly valid.

(`label_ref label`)

Represents the value of an assembler label for code. It contains one operand, an expression, which must be a `code_label` or a `note` of type `NOTE_INSN_DELETED_LABEL` that appears in the instruction sequence to identify the place where the label should go.

The reason for using a distinct expression type for code label references is so that jump optimization can distinguish them.

(`const:m exp`)

Represents a constant that is the result of an assembly-time arithmetic computation. The operand, `exp`, is an expression that contains only constants (`const_int`, `symbol_ref` and `label_ref` expressions) combined with `plus` and `minus`. However, not all combinations are valid, since the assembler cannot do arbitrary arithmetic on relocatable symbols.

`m` should be `Pmode`.

(`high:m exp`)

Represents the high-order bits of `exp`, usually a `symbol_ref`. The number of bits is machine-dependent and is normally the number of bits specified in an instruction that initializes the high order bits of a register. It is used with `lo_sum` to represent the typical two-instruction sequence used in RISC machines to reference a global memory location.

`m` should be `Pmode`.

8.7 Registers and Memory

Here are the RTL expression types for describing access to machine registers and to main memory.

(`reg:m n`)

For small values of the integer `n` (those that are less than `FIRST_PSEUDO_REGISTER`), this stands for a reference to machine register number `n`: a *hard register*. For larger values of `n`, it stands for a temporary value or *pseudo register*. The compiler's strategy is to generate code assuming an unlimited number of such pseudo registers, and later convert them into hard registers or into memory references.

`m` is the machine mode of the reference. It is necessary because machines can generally refer to each register in more than one mode. For example, a register may contain a full word but there may be instructions to refer to it as a half word or as a single byte, as well as instructions to refer to it as a floating point number of various precisions.

Even for a register that the machine can access in only one mode, the mode must always be specified.

The symbol `FIRST_PSEUDO_REGISTER` is defined by the machine description, since the number of hard registers on the machine is an invariant characteristic of the machine. Note, however, that not all of the machine registers must be

general registers. All the machine registers that can be used for storage of data are given hard register numbers, even those that can be used only in certain instructions or can hold only certain types of data.

A hard register may be accessed in various modes throughout one function, but each pseudo register is given a natural mode and is accessed only in that mode. When it is necessary to describe an access to a pseudo register using a nonnatural mode, a **subreg** expression is used.

A **reg** expression with a machine mode that specifies more than one word of data may actually stand for several consecutive registers. If in addition the register number specifies a hardware register, then it actually represents several consecutive hardware registers starting with the specified one.

Each pseudo register number used in a function's RTL code is represented by a unique **reg** expression.

Some pseudo register numbers, those within the range of **FIRST_VIRTUAL_REGISTER** to **LAST_VIRTUAL_REGISTER** only appear during the RTL generation phase and are eliminated before the optimization phases. These represent locations in the stack frame that cannot be determined until RTL generation for the function has been completed. The following virtual register numbers are defined:

VIRTUAL_INCOMING_ARGS_REGNUM

This points to the first word of the incoming arguments passed on the stack. Normally these arguments are placed there by the caller, but the callee may have pushed some arguments that were previously passed in registers.

When RTL generation is complete, this virtual register is replaced by the sum of the register given by **ARG_POINTER_REGNUM** and the value of **FIRST_PARM_OFFSET**.

VIRTUAL_STACK_VARS_REGNUM

If **FRAME_GROWS_DOWNWARD** is defined, this points to immediately above the first variable on the stack. Otherwise, it points to the first variable on the stack.

VIRTUAL_STACK_VARS_REGNUM is replaced with the sum of the register given by **FRAME_POINTER_REGNUM** and the value **STARTING_FRAME_OFFSET**.

VIRTUAL_STACK_DYNAMIC_REGNUM

This points to the location of dynamically allocated memory on the stack immediately after the stack pointer has been adjusted by the amount of memory desired.

This virtual register is replaced by the sum of the register given by **STACK_POINTER_REGNUM** and the value **STACK_DYNAMIC_OFFSET**.

VIRTUAL_OUTGOING_ARGS_REGNUM

This points to the location in the stack at which outgoing arguments should be written when the stack is pre-pushed (arguments pushed using push insns should always use **STACK_POINTER_REGNUM**).

This virtual register is replaced by the sum of the register given by `STACK_POINTER_REGNUM` and the value `STACK_POINTER_OFFSET`.

(`subreg:m reg bytenum`)

`subreg` expressions are used to refer to a register in a machine mode other than its natural one, or to refer to one register of a multi-part `reg` that actually refers to several registers.

Each pseudo-register has a natural mode. If it is necessary to operate on it in a different mode—for example, to perform a fullword move instruction on a pseudo-register that contains a single byte—the pseudo-register must be enclosed in a `subreg`. In such a case, *bytenum* is zero.

Usually *m* is at least as narrow as the mode of *reg*, in which case it is restricting consideration to only the bits of *reg* that are in *m*.

Sometimes *m* is wider than the mode of *reg*. These `subreg` expressions are often called *paradoxical*. They are used in cases where we want to refer to an object in a wider mode but do not care what value the additional bits have. The reload pass ensures that paradoxical references are only made to hard registers.

The other use of `subreg` is to extract the individual registers of a multi-register value. Machine modes such as `DImode` and `TImode` can indicate values longer than a word, values which usually require two or more consecutive registers. To access one of the registers, use a `subreg` with mode `SImode` and a *bytenum* offset that says which register.

Storing in a non-paradoxical `subreg` has undefined results for bits belonging to the same word as the `subreg`. This laxity makes it easier to generate efficient code for such instructions. To represent an instruction that preserves all the bits outside of those in the `subreg`, use `strict_low_part` around the `subreg`. The compilation parameter `WORDS_BIG_ENDIAN`, if set to 1, says that byte number zero is part of the most significant word; otherwise, it is part of the least significant word.

The compilation parameter `BYTES_BIG_ENDIAN`, if set to 1, says that byte number zero is the most significant byte within a word; otherwise, it is the least significant byte within a word.

On a few targets, `FLOAT_WORDS_BIG_ENDIAN` disagrees with `WORDS_BIG_ENDIAN`. However, most parts of the compiler treat floating point values as if they had the same endianness as integer values. This works because they handle them solely as a collection of integer values, with no particular numerical value. Only `real.c` and the runtime libraries care about `FLOAT_WORDS_BIG_ENDIAN`.

Between the combiner pass and the reload pass, it is possible to have a paradoxical `subreg` which contains a `mem` instead of a `reg` as its first operand. After the reload pass, it is also possible to have a non-paradoxical `subreg` which contains a `mem`; this usually occurs when the `mem` is a stack slot which replaced a pseudo register.

Note that it is not valid to access a `DFmode` value in `SFmode` using a `subreg`. On some machines the most significant part of a `DFmode` value does not have the same format as a single-precision floating value.

It is also not valid to access a single word of a multi-word value in a hard register when less registers can hold the value than would be expected from its size. For example, some 32-bit machines have floating-point registers that can hold an entire `DFmode` value. If register 10 were such a register (`subreg:SI (reg:DF 10) 1`) would be invalid because there is no way to convert that reference to a single machine register. The reload pass prevents `subreg` expressions such as these from being formed.

The first operand of a `subreg` expression is customarily accessed with the `SUBREG_REG` macro and the second operand is customarily accessed with the `SUBREG_BYTE` macro.

`(scratch:m)`

This represents a scratch register that will be required for the execution of a single instruction and not used subsequently. It is converted into a `reg` by either the local register allocator or the reload pass.

`scratch` is usually present inside a `clobber` operation (see Section 8.14 [Side Effects], page 97).

`(cc0)`

This refers to the machine's condition code register. It has no operands and may not have a machine mode. There are two ways to use it:

- To stand for a complete set of condition code flags. This is best on most machines, where each comparison sets the entire series of flags.

With this technique, `(cc0)` may be validly used in only two contexts: as the destination of an assignment (in test and compare instructions) and in comparison operators comparing against zero (`const_int` with value zero; that is to say, `const0_rtx`).

- To stand for a single flag that is the result of a single condition. This is useful on machines that have only a single flag bit, and in which comparison instructions must specify the condition to test.

With this technique, `(cc0)` may be validly used in only two contexts: as the destination of an assignment (in test and compare instructions) where the source is a comparison operator, and as the first operand of `if_then_else` (in a conditional branch).

There is only one expression object of code `cc0`; it is the value of the variable `cc0_rtx`. Any attempt to create an expression of code `cc0` will return `cc0_rtx`.

Instructions can set the condition code implicitly. On many machines, nearly all instructions set the condition code based on the value that they compute or store. It is not necessary to record these actions explicitly in the RTL because the machine description includes a prescription for recognizing the instructions that do so (by means of the macro `NOTICE_UPDATE_CC`). See Section 10.15 [Condition Code], page 264. Only instructions whose sole purpose is to set the condition code, and instructions that use the condition code, need mention `(cc0)`.

On some machines, the condition code register is given a register number and a `reg` is used instead of `(cc0)`. This is usually the preferable approach if only a small subset of instructions modify the condition code. Other machines store

condition codes in general registers; in such cases a pseudo register should be used.

Some machines, such as the SPARC and RS/6000, have two sets of arithmetic instructions, one that sets and one that does not set the condition code. This is best handled by normally generating the instruction that does not set the condition code, and making a pattern that both performs the arithmetic and sets the condition code register (which would not be `cc0`) in this case). For examples, search for `'addcc'` and `'andcc'` in `'sparc.md'`.

(`pc`) This represents the machine's program counter. It has no operands and may not have a machine mode. (`pc`) may be validly used only in certain specific contexts in jump instructions.

There is only one expression object of code `pc`; it is the value of the variable `pc_rtx`. Any attempt to create an expression of code `pc` will return `pc_rtx`.

All instructions that do not jump alter the program counter implicitly by incrementing it, but there is no need to mention this in the RTL.

(`mem:m addr alias`)

This RTX represents a reference to main memory at an address represented by the expression `addr`. `m` specifies how large a unit of memory is accessed. `alias` specifies an alias set for the reference. In general two items are in different alias sets if they cannot reference the same memory address.

The construct (`mem:BLK (scratch)`) is considered to alias all other memories. Thus it may be used as a memory barrier in epilogue stack deallocation patterns.

(`addressof:m reg`)

This RTX represents a request for the address of register `reg`. Its mode is always `Pmode`. If there are any `addressof` expressions left in the function after CSE, `reg` is forced into the stack and the `addressof` expression is replaced with a `plus` expression for the address of its stack slot.

8.8 RTL Expressions for Arithmetic

Unless otherwise specified, all the operands of arithmetic expressions must be valid for mode `m`. An operand is valid for mode `m` if it has mode `m`, or if it is a `const_int` or `const_double` and `m` is a mode of class `MODE_INT`.

For commutative binary operations, constants should be placed in the second operand.

(`plus:m x y`)

Represents the sum of the values represented by `x` and `y` carried out in machine mode `m`.

(`lo_sum:m x y`)

Like `plus`, except that it represents that sum of `x` and the low-order bits of `y`. The number of low order bits is machine-dependent but is normally the number of bits in a `Pmode` item minus the number of bits set by the `high` code (see Section 8.6 [Constants], page 84).

`m` should be `Pmode`.

`(minus:m x y)`

Like `plus` but represents subtraction.

`(ss_plus:m x y)`

Like `plus`, but using signed saturation in case of an overflow.

`(us_plus:m x y)`

Like `plus`, but using unsigned saturation in case of an overflow.

`(ss_minus:m x y)`

Like `minus`, but using signed saturation in case of an overflow.

`(us_minus:m x y)`

Like `minus`, but using unsigned saturation in case of an overflow.

`(compare:m x y)`

Represents the result of subtracting `y` from `x` for purposes of comparison. The result is computed without overflow, as if with infinite precision.

Of course, machines can't really subtract with infinite precision. However, they can pretend to do so when only the sign of the result will be used, which is the case when the result is stored in the condition code. And that is the *only* way this kind of expression may validly be used: as a value to be stored in the condition codes, either `(cc0)` or a register. See Section 8.9 [Comparisons], page 93.

The mode `m` is not related to the modes of `x` and `y`, but instead is the mode of the condition code value. If `(cc0)` is used, it is `VOIDmode`. Otherwise it is some mode in class `MODE_CC`, often `CCmode`. See Section 10.15 [Condition Code], page 264. If `m` is `VOIDmode` or `CCmode`, the operation returns sufficient information (in an unspecified format) so that any comparison operator can be applied to the result of the `COMPARE` operation. For other modes in class `MODE_CC`, the operation only returns a subset of this information.

Normally, `x` and `y` must have the same mode. Otherwise, `compare` is valid only if the mode of `x` is in class `MODE_INT` and `y` is a `const_int` or `const_double` with mode `VOIDmode`. The mode of `x` determines what mode the comparison is to be done in; thus it must not be `VOIDmode`.

If one of the operands is a constant, it should be placed in the second operand and the comparison code adjusted as appropriate.

A `compare` specifying two `VOIDmode` constants is not valid since there is no way to know in what mode the comparison is to be performed; the comparison must either be folded during the compilation or the first operand must be loaded into a register while its mode is still known.

`(neg:m x)`

Represents the negation (subtraction from zero) of the value represented by `x`, carried out in mode `m`.

`(mult:m x y)`

Represents the signed product of the values represented by `x` and `y` carried out in machine mode `m`.

Some machines support a multiplication that generates a product wider than the operands. Write the pattern for this as

```
(mult:m (sign_extend:m x) (sign_extend:m y))
```

where *m* is wider than the modes of *x* and *y*, which need not be the same.

For unsigned widening multiplication, use the same idiom, but with `zero_extend` instead of `sign_extend`.

```
(div:m x y)
```

Represents the quotient in signed division of *x* by *y*, carried out in machine mode *m*. If *m* is a floating point mode, it represents the exact quotient; otherwise, the integerized quotient.

Some machines have division instructions in which the operands and quotient widths are not all the same; you should represent such instructions using `truncate` and `sign_extend` as in,

```
(truncate:m1 (div:m2 x (sign_extend:m2 y)))
```

```
(udiv:m x y)
```

Like `div` but represents unsigned division.

```
(mod:m x y)
```

```
(umod:m x y)
```

Like `div` and `udiv` but represent the remainder instead of the quotient.

```
(smin:m x y)
```

```
(smax:m x y)
```

Represents the smaller (for `smin`) or larger (for `smax`) of *x* and *y*, interpreted as signed integers in mode *m*.

```
(umin:m x y)
```

```
(umax:m x y)
```

Like `smin` and `smax`, but the values are interpreted as unsigned integers.

```
(not:m x)
```

Represents the bitwise complement of the value represented by *x*, carried out in mode *m*, which must be a fixed-point machine mode.

```
(and:m x y)
```

Represents the bitwise logical-and of the values represented by *x* and *y*, carried out in machine mode *m*, which must be a fixed-point machine mode.

```
(ior:m x y)
```

Represents the bitwise inclusive-or of the values represented by *x* and *y*, carried out in machine mode *m*, which must be a fixed-point mode.

```
(xor:m x y)
```

Represents the bitwise exclusive-or of the values represented by *x* and *y*, carried out in machine mode *m*, which must be a fixed-point mode.

```
(ashift:m x c)
```

Represents the result of arithmetically shifting *x* left by *c* places. *x* have mode *m*, a fixed-point machine mode. *c* be a fixed-point mode or be a constant

with mode `VOIDmode`; which mode is determined by the mode called for in the machine description entry for the left-shift instruction. For example, on the VAX, the mode of `c` is `QImode` regardless of `m`.

`(lshiftrt:m x c)`

`(ashiftrt:m x c)`

Like `ashift` but for right shift. Unlike the case for left shift, these two operations are distinct.

`(rotate:m x c)`

`(rotatert:m x c)`

Similar but represent left and right rotate. If `c` is a constant, use `rotate`.

`(abs:m x)`

Represents the absolute value of `x`, computed in mode `m`.

`(sqrt:m x)`

Represents the square root of `x`, computed in mode `m`. Most often `m` will be a floating point mode.

`(ffs:m x)`

Represents one plus the index of the least significant 1-bit in `x`, represented as an integer of mode `m`. (The value is zero if `x` is zero.) The mode of `x` need not be `m`; depending on the target machine, various mode combinations may be valid.

8.9 Comparison Operations

Comparison operators test a relation on two operands and are considered to represent a machine-dependent nonzero value described by, but not necessarily equal to, `STORE_FLAG_VALUE` (see Section 10.26 [Misc], page 314) if the relation holds, or zero if it does not, for comparison operators whose results have a ‘`MODE.INT`’ mode, and `FLOAT_STORE_FLAG_VALUE` (see Section 10.26 [Misc], page 314) if the relation holds, or zero if it does not, for comparison operators that return floating-point values. The mode of the comparison operation is independent of the mode of the data being compared. If the comparison operation is being tested (e.g., the first operand of an `if_then_else`), the mode must be `VOIDmode`.

There are two ways that comparison operations may be used. The comparison operators may be used to compare the condition codes (`cc0`) against zero, as in `(eq (cc0) (const_int 0))`. Such a construct actually refers to the result of the preceding instruction in which the condition codes were set. The instruction setting the condition code must be adjacent to the instruction using the condition code; only `note` insns may separate them.

Alternatively, a comparison operation may directly compare two data objects. The mode of the comparison is determined by the operands; they must both be valid for a common machine mode. A comparison with both operands constant would be invalid as the machine mode could not be deduced from it, but such a comparison should never exist in RTL due to constant folding.

In the example above, if (`cc0`) were last set to `(compare x y)`, the comparison operation is identical to `(eq x y)`. Usually only one style of comparisons is supported on a particular

machine, but the combine pass will try to merge the operations to produce the `eq` shown in case it exists in the context of the particular `insn` involved.

Inequality comparisons come in two flavors, signed and unsigned. Thus, there are distinct expression codes `gt` and `gtu` for signed and unsigned greater-than. These can produce different results for the same pair of integer values: for example, 1 is signed greater-than `-1` but not unsigned greater-than, because `-1` when regarded as unsigned is actually `0xffffffff` which is greater than 1.

The signed comparisons are also used for floating point values. Floating point comparisons are distinguished by the machine modes of the operands.

`(eq:m x y)`
 STORE_FLAG_VALUE if the values represented by `x` and `y` are equal, otherwise 0.

`(ne:m x y)`
 STORE_FLAG_VALUE if the values represented by `x` and `y` are not equal, otherwise 0.

`(gt:m x y)`
 STORE_FLAG_VALUE if the `x` is greater than `y`. If they are fixed-point, the comparison is done in a signed sense.

`(gtu:m x y)`
 Like `gt` but does unsigned comparison, on fixed-point numbers only.

`(lt:m x y)`
`(ltu:m x y)`
 Like `gt` and `gtu` but test for “less than”.

`(ge:m x y)`
`(geu:m x y)`
 Like `gt` and `gtu` but test for “greater than or equal”.

`(le:m x y)`
`(leu:m x y)`
 Like `gt` and `gtu` but test for “less than or equal”.

`(if_then_else cond then else)`
 This is not a comparison operation but is listed here because it is always used in conjunction with a comparison operation. To be precise, `cond` is a comparison expression. This expression represents a choice, according to `cond`, between the value represented by `then` and the one represented by `else`.

On most machines, `if_then_else` expressions are valid only to express conditional jumps.

`(cond [test1 value1 test2 value2 ...] default)`
 Similar to `if_then_else`, but more general. Each of `test1`, `test2`, ... is performed in turn. The result of this expression is the `value` corresponding to the first nonzero test, or `default` if none of the tests are nonzero expressions.

This is currently not valid for instruction patterns and is supported only for `insn` attributes. See Section 9.18 [Insn Attributes], page 175.

8.10 Bit-Fields

Special expression codes exist to represent bit-field instructions. These types of expressions are lvalues in RTL; they may appear on the left side of an assignment, indicating insertion of a value into the specified bit-field.

`(sign_extract:m loc size pos)`

This represents a reference to a sign-extended bit-field contained or starting in *loc* (a memory or register reference). The bit-field is *size* bits wide and starts at bit *pos*. The compilation option `BITS_BIG_ENDIAN` says which end of the memory unit *pos* counts from.

If *loc* is in memory, its mode must be a single-byte integer mode. If *loc* is in a register, the mode to use is specified by the operand of the `insv` or `extv` pattern (see Section 9.8 [Standard Names], page 144) and is usually a full-word integer mode, which is the default if none is specified.

The mode of *pos* is machine-specific and is also specified in the `insv` or `extv` pattern.

The mode *m* is the same as the mode that would be used for *loc* if it were a register.

`(zero_extract:m loc size pos)`

Like `sign_extract` but refers to an unsigned or zero-extended bit-field. The same sequence of bits are extracted, but they are filled to an entire word with zeros instead of by sign-extension.

8.11 Vector Operations

All normal RTL expressions can be used with vector modes; they are interpreted as operating on each part of the vector independently. Additionally, there are a few new expressions to describe specific vector operations.

`(vec_merge:m vec1 vec2 items)`

This describes a merge operation between two vectors. The result is a vector of mode *m*; its elements are selected from either *vec1* or *vec2*. Which elements are selected is described by *items*, which is a bit mask represented by a `const_int`; a zero bit indicates the corresponding element in the result vector is taken from *vec2* while a set bit indicates it is taken from *vec1*.

`(vec_select:m vec1 selection)`

This describes an operation that selects parts of a vector. *vec1* is the source vector, *selection* is a `parallel` that contains a `const_int` for each of the subparts of the result vector, giving the number of the source subpart that should be stored into it.

`(vec_concat:m vec1 vec2)`

Describes a vector concat operation. The result is a concatenation of the vectors *vec1* and *vec2*; its length is the sum of the lengths of the two inputs.

`(vec_const:m subparts)`

This describes a constant vector. *subparts* is a `parallel` that contains a constant for each of the subparts of the vector.

`(vec_duplicate:m vec)`

This operation converts a small vector into a larger one by duplicating the input values. The output vector mode must have the same submodes as the input vector mode, and the number of output parts must be an integer multiple of the number of input parts.

8.12 Conversions

All conversions between machine modes must be represented by explicit conversion operations. For example, an expression which is the sum of a byte and a full word cannot be written as `(plus:SI (reg:QI 34) (reg:SI 80))` because the `plus` operation requires two operands of the same machine mode. Therefore, the byte-sized operand is enclosed in a conversion operation, as in

```
(plus:SI (sign_extend:SI (reg:QI 34)) (reg:SI 80))
```

The conversion operation is not a mere placeholder, because there may be more than one way of converting from a given starting mode to the desired final mode. The conversion operation code says how to do it.

For all conversion operations, *x* must not be `VOIDmode` because the mode in which to do the conversion would not be known. The conversion must either be done at compile-time or *x* must be placed into a register.

`(sign_extend:m x)`

Represents the result of sign-extending the value *x* to machine mode *m*. *m* must be a fixed-point mode and *x* a fixed-point value of a mode narrower than *m*.

`(zero_extend:m x)`

Represents the result of zero-extending the value *x* to machine mode *m*. *m* must be a fixed-point mode and *x* a fixed-point value of a mode narrower than *m*.

`(float_extend:m x)`

Represents the result of extending the value *x* to machine mode *m*. *m* must be a floating point mode and *x* a floating point value of a mode narrower than *m*.

`(truncate:m x)`

Represents the result of truncating the value *x* to machine mode *m*. *m* must be a fixed-point mode and *x* a fixed-point value of a mode wider than *m*.

`(ss_truncate:m x)`

Represents the result of truncating the value *x* to machine mode *m*, using signed saturation in the case of overflow. Both *m* and the mode of *x* must be fixed-point modes.

`(us_truncate:m x)`

Represents the result of truncating the value x to machine mode m , using unsigned saturation in the case of overflow. Both m and the mode of x must be fixed-point modes.

`(float_truncate:m x)`

Represents the result of truncating the value x to machine mode m . m must be a floating point mode and x a floating point value of a mode wider than m .

`(float:m x)`

Represents the result of converting fixed point value x , regarded as signed, to floating point mode m .

`(unsigned_float:m x)`

Represents the result of converting fixed point value x , regarded as unsigned, to floating point mode m .

`(fix:m x)`

When m is a fixed point mode, represents the result of converting floating point value x to mode m , regarded as signed. How rounding is done is not specified, so this operation may be used validly in compiling C code only for integer-valued operands.

`(unsigned_fix:m x)`

Represents the result of converting floating point value x to fixed point mode m , regarded as unsigned. How rounding is done is not specified.

`(fix:m x)`

When m is a floating point mode, represents the result of converting floating point value x (valid for mode m) to an integer, still represented in floating point mode m , by rounding towards zero.

8.13 Declarations

Declaration expression codes do not represent arithmetic operations but rather state assertions about their operands.

`(strict_low_part (subreg:m (reg:n r) 0))`

This expression code is used in only one context: as the destination operand of a `set` expression. In addition, the operand of this expression must be a non-paradoxical `subreg` expression.

The presence of `strict_low_part` says that the part of the register which is meaningful in mode n , but is not part of mode m , is not to be altered. Normally, an assignment to such a subreg is allowed to have undefined effects on the rest of the register when m is less than a word.

8.14 Side Effect Expressions

The expression codes described so far represent values, not actions. But machine instructions never produce values; they are meaningful only for their side effects on the state of the machine. Special expression codes are used to represent side effects.

The body of an instruction is always one of these side effect codes; the codes described above, which represent values, appear only as the operands of these.

(set *lval* *x*)

Represents the action of storing the value of *x* into the place represented by *lval*. *lval* must be an expression representing a place that can be stored in: **reg** (or **subreg**, **strict_low_part** or **zero_extract**), **mem**, **pc**, **parallel**, or **cc0**.

If *lval* is a **reg**, **subreg** or **mem**, it has a machine mode; then *x* must be valid for that mode.

If *lval* is a **reg** whose machine mode is less than the full width of the register, then it means that the part of the register specified by the machine mode is given the specified value and the rest of the register receives an undefined value. Likewise, if *lval* is a **subreg** whose machine mode is narrower than the mode of the register, the rest of the register can be changed in an undefined way.

If *lval* is a **strict_low_part** or **zero_extract** of a **subreg**, then the part of the register specified by the machine mode of the **subreg** is given the value *x* and the rest of the register is not changed.

If *lval* is **cc0**, it has no machine mode, and *x* may be either a **compare** expression or a value that may have any mode. The latter case represents a “test” instruction. The expression **(set (cc0) (reg:m n))** is equivalent to **(set (cc0) (compare (reg:m n) (const_int 0)))**. Use the former expression to save space during the compilation.

If *lval* is a **parallel**, it is used to represent the case of a function returning a structure in multiple registers. Each element of the **parallel** is an **expr_list** whose first operand is a **reg** and whose second operand is a **const_int** representing the offset (in bytes) into the structure at which the data in that register corresponds. The first element may be null to indicate that the structure is also passed partly in memory.

If *lval* is **pc**, we have a jump instruction, and the possibilities for *x* are very limited. It may be a **label_ref** expression (unconditional jump). It may be an **if_then_else** (conditional jump), in which case either the second or the third operand must be **pc** (for the case which does not jump) and the other of the two must be a **label_ref** (for the case which does jump). *x* may also be a **mem** or **(plus:SI (pc) y)**, where *y* may be a **reg** or a **mem**; these unusual patterns are used to represent jumps through branch tables.

If *lval* is neither **cc0** nor **pc**, the mode of *lval* must not be **VOIDmode** and the mode of *x* must be valid for the mode of *lval*.

lval is customarily accessed with the **SET_DEST** macro and *x* with the **SET_SRC** macro.

(return) As the sole expression in a pattern, represents a return from the current function, on machines where this can be done with one instruction, such as VAXen. On machines where a multi-instruction “epilogue” must be executed in order to return from the function, returning is done by jumping to a label which precedes the epilogue, and the **return** expression code is never used.

Inside an `if_then_else` expression, represents the value to be placed in `pc` to return to the caller.

Note that an `insn` pattern of `(return)` is logically equivalent to `(set (pc) (return))`, but the latter form is never used.

`(call function nargs)`

Represents a function call. *function* is a `mem` expression whose address is the address of the function to be called. *nargs* is an expression which can be used for two purposes: on some machines it represents the number of bytes of stack argument; on others, it represents the number of argument registers.

Each machine has a standard machine mode which *function* must have. The machine description defines macro `FUNCTION_MODE` to expand into the requisite mode name. The purpose of this mode is to specify what kind of addressing is allowed, on machines where the allowed kinds of addressing depend on the machine mode being addressed.

`(clobber x)`

Represents the storing or possible storing of an unpredictable, undescribed value into `x`, which must be a `reg`, `scratch`, `parallel` or `mem` expression.

One place this is used is in string instructions that store standard values into particular hard registers. It may not be worth the trouble to describe the values that are stored, but it is essential to inform the compiler that the registers will be altered, lest it attempt to keep data in them across the string instruction.

If `x` is `(mem:BLK (const_int 0))` or `(mem:BLK (scratch))`, it means that all memory locations must be presumed clobbered. If `x` is a `parallel`, it has the same meaning as a `parallel` in a `set` expression.

Note that the machine description classifies certain hard registers as “call-clobbered”. All function call instructions are assumed by default to clobber these registers, so there is no need to use `clobber` expressions to indicate this fact. Also, each function call is assumed to have the potential to alter any memory location, unless the function is declared `const`.

If the last group of expressions in a `parallel` are each a `clobber` expression whose arguments are `reg` or `match_scratch` (see Section 9.4 [RTL Template], page 119) expressions, the combiner phase can add the appropriate `clobber` expressions to an `insn` it has constructed when doing so will cause a pattern to be matched.

This feature can be used, for example, on a machine that whose multiply and add instructions don't use an MQ register but which has an add-accumulate instruction that does clobber the MQ register. Similarly, a combined instruction might require a temporary register while the constituent instructions might not.

When a `clobber` expression for a register appears inside a `parallel` with other side effects, the register allocator guarantees that the register is unoccupied both before and after that `insn`. However, the reload phase may allocate a register used for one of the inputs unless the ‘&’ constraint is specified for the selected alternative (see Section 9.7.4 [Modifiers], page 131). You can clobber either a specific hard register, a pseudo register, or a `scratch` expression; in

the latter two cases, GCC will allocate a hard register that is available there for use as a temporary.

For instructions that require a temporary register, you should use `scratch` instead of a pseudo-register because this will allow the combiner phase to add the `clobber` when required. You do this by coding `(clobber (match_scratch ...))`. If you do clobber a pseudo register, use one which appears nowhere else—generate a new one each time. Otherwise, you may confuse CSE.

There is one other known use for clobbering a pseudo register in a `parallel`: when one of the input operands of the `insn` is also clobbered by the `insn`. In this case, using the same pseudo register in the `clobber` and elsewhere in the `insn` produces the expected results.

(use x) Represents the use of the value of `x`. It indicates that the value in `x` at this point in the program is needed, even though it may not be apparent why this is so. Therefore, the compiler will not attempt to delete previous instructions whose only effect is to store a value in `x`. `x` must be a `reg` expression.

In some situations, it may be tempting to add a `use` of a register in a `parallel` to describe a situation where the value of a special register will modify the behavior of the instruction. An hypothetical example might be a pattern for an addition that can either wrap around or use saturating addition depending on the value of a special control register:

```
(parallel [(set (reg:SI 2) (unspec:SI [(reg:SI 3)
                                       (reg:SI 4)] 0))
          (use (reg:SI 1))])
```

This will not work, several of the optimizers only look at expressions locally; it is very likely that if you have multiple `insns` with identical inputs to the `unspec`, they will be optimized away even if register 1 changes in between.

This means that `use` can *only* be used to describe that the register is live. You should think twice before adding `use` statements, more often you will want to use `unspec` instead. The `use` RTX is most commonly useful to describe that a fixed register is implicitly used in an `insn`. It is also safe to use in patterns where the compiler knows for other reasons that the result of the whole pattern is variable, such as `'movstrm'` or `'call'` patterns.

During the reload phase, an `insn` that has a `use` as pattern can carry a `reg_equal` note. These `use` `insns` will be deleted before the reload phase exits.

During the delayed branch scheduling phase, `x` may be an `insn`. This indicates that `x` previously was located at this place in the code and its data dependencies need to be taken into account. These `use` `insns` will be deleted before the delayed branch scheduling phase exits.

(parallel [x0 x1 ...])

Represents several side effects performed in parallel. The square brackets stand for a vector; the operand of `parallel` is a vector of expressions. `x0`, `x1` and so on are individual side effect expressions—expressions of code `set`, `call`, `return`, `clobber` or `use`.

“In parallel” means that first all the values used in the individual side-effects are computed, and second all the actual side-effects are performed. For example,

```
(parallel [(set (reg:SI 1) (mem:SI (reg:SI 1)))
           (set (mem:SI (reg:SI 1)) (reg:SI 1))])
```

says unambiguously that the values of hard register 1 and the memory location addressed by it are interchanged. In both places where `(reg:SI 1)` appears as a memory address it refers to the value in register 1 *before* the execution of the `insn`.

It follows that it is *incorrect* to use `parallel` and expect the result of one `set` to be available for the next one. For example, people sometimes attempt to represent a jump-if-zero instruction this way:

```
(parallel [(set (cc0) (reg:SI 34))
           (set (pc) (if_then_else
                  (eq (cc0) (const_int 0))
                  (label_ref ...)
                  (pc)))])
```

But this is incorrect, because it says that the jump condition depends on the condition code value *before* this instruction, not on the new value that is set by this instruction.

Peephole optimization, which takes place together with final assembly code output, can produce `insns` whose patterns consist of a `parallel` whose elements are the operands needed to output the resulting assembler code—often `reg`, `mem` or constant expressions. This would not be well-formed RTL at any other stage in compilation, but it is ok then because no further optimization remains to be done. However, the definition of the macro `NOTICE_UPDATE_CC`, if any, must deal with such `insns` if you define any peephole optimizations.

`(cond_exec [cond expr])`

Represents a conditionally executed expression. The `expr` is executed only if the `cond` is nonzero. The `cond` expression must not have side-effects, but the `expr` may very well have side-effects.

`(sequence [insns ...])`

Represents a sequence of `insns`. Each of the `insns` that appears in the vector is suitable for appearing in the chain of `insns`, so it must be an `insn`, `jump_insn`, `call_insn`, `code_label`, `barrier` or `note`.

A `sequence` RTX is never placed in an actual `insn` during RTL generation. It represents the sequence of `insns` that result from a `define_expand` *before* those `insns` are passed to `emit_insn` to insert them in the chain of `insns`. When actually inserted, the individual sub-`insns` are separated out and the `sequence` is forgotten.

After delay-slot scheduling is completed, an `insn` and all the `insns` that reside in its delay slots are grouped together into a `sequence`. The `insn` requiring the delay slot is the first `insn` in the vector; subsequent `insns` are to be placed in the delay slot.

`INSN_ANNULLED_BRANCH_P` is set on an `insn` in a delay slot to indicate that a branch `insn` should be used that will conditionally annul the effect of the `insns` in the delay slots. In such a case, `INSN_FROM_TARGET_P` indicates that the `insn`

is from the target of the branch and should be executed only if the branch is taken; otherwise the *insn* should be executed only if the branch is not taken. See Section 9.18.7 [Delay Slots], page 182.

These expression codes appear in place of a side effect, as the body of an *insn*, though strictly speaking they do not always describe side effects as such:

(asm_input *s*)

Represents literal assembler code as described by the string *s*.

(unspec [*operands ...*] *index*)

(unspec_volatile [*operands ...*] *index*)

Represents a machine-specific operation on *operands*. *index* selects between multiple machine-specific operations. `unspec_volatile` is used for volatile operations and operations that may trap; `unspec` is used for other operations.

These codes may appear inside a `pattern` of an *insn*, inside a `parallel`, or inside an expression.

(addr_vec:*m* [*lr0 lr1 ...*])

Represents a table of jump addresses. The vector elements *lr0*, etc., are `label_ref` expressions. The mode *m* specifies how much space is given to each address; normally *m* would be `Pmode`.

(addr_diff_vec:*m* *base* [*lr0 lr1 ...*] *min max flags*)

Represents a table of jump addresses expressed as offsets from *base*. The vector elements *lr0*, etc., are `label_ref` expressions and so is *base*. The mode *m* specifies how much space is given to each address-difference. *min* and *max* are set up by branch shortening and hold a label with a minimum and a maximum address, respectively. *flags* indicates the relative position of *base*, *min* and *max* to the containing *insn* and of *min* and *max* to *base*. See `rtl.def` for details.

(prefetch:*m* *addr rw locality*)

Represents prefetch of memory at address *addr*. Operand *rw* is 1 if the prefetch is for data to be written, 0 otherwise; targets that do not support write prefetches should treat this as a normal prefetch. Operand *locality* specifies the amount of temporal locality; 0 if there is none or 1, 2, or 3 for increasing levels of temporal locality; targets that do not support locality hints should ignore this.

This *insn* is used to minimize cache-miss latency by moving data into a cache before it is accessed. It should use only non-faulting data prefetch instructions.

8.15 Embedded Side-Effects on Addresses

Six special side-effect expression codes appear as memory addresses.

(pre_dec:*m* *x*)

Represents the side effect of decrementing *x* by a standard amount and represents also the value that *x* has after being decremented. *x* must be a `reg` or `mem`, but most machines allow only a `reg`. *m* must be the machine mode for

pointers on the machine in use. The amount x is decremented by is the length in bytes of the machine mode of the containing memory reference of which this expression serves as the address. Here is an example of its use:

```
(mem:DF (pre_dec:SI (reg:SI 39)))
```

This says to decrement pseudo register 39 by the length of a DFmode value and use the result to address a DFmode value.

```
(pre_inc:m x)
```

Similar, but specifies incrementing x instead of decrementing it.

```
(post_dec:m x)
```

Represents the same side effect as `pre_dec` but a different value. The value represented here is the value x has *before* being decremented.

```
(post_inc:m x)
```

Similar, but specifies incrementing x instead of decrementing it.

```
(post_modify:m x y)
```

Represents the side effect of setting x to y and represents x before x is modified. x must be a `reg` or `mem`, but most machines allow only a `reg`. m must be the machine mode for pointers on the machine in use.

The expression y must be one of three forms:

```
(plus:m x z), (minus:m x z), or (plus:m x i),
```

where z is an index register and i is a constant.

Here is an example of its use:

```
(mem:SF (post_modify:SI (reg:SI 42) (plus (reg:SI 42)
                                           (reg:SI 48))))
```

This says to modify pseudo register 42 by adding the contents of pseudo register 48 to it, after the use of what ever 42 points to.

```
(pre_modify:m x expr)
```

Similar except side effects happen before the use.

These embedded side effect expressions must be used with care. Instruction patterns may not use them. Until the ‘flow’ pass of the compiler, they may occur only to represent pushes onto the stack. The ‘flow’ pass finds cases where registers are incremented or decremented in one instruction and used as an address shortly before or after; these cases are then transformed to use pre- or post-increment or -decrement.

If a register used as the operand of these expressions is used in another address in an `insn`, the original value of the register is used. Uses of the register outside of an address are not permitted within the same `insn` as a use in an embedded side effect expression because such `insns` behave differently on different machines and hence must be treated as ambiguous and disallowed.

An instruction that can be represented with an embedded side effect could also be represented using `parallel` containing an additional `set` to describe how the address register is altered. This is not done because machines that allow these operations at all typically allow them wherever a memory address is called for. Describing them as additional parallel stores would require doubling the number of entries in the machine description.

8.16 Assembler Instructions as Expressions

The RTX code `asm_operands` represents a value produced by a user-specified assembler instruction. It is used to represent an `asm` statement with arguments. An `asm` statement with a single output operand, like this:

```
asm ("foo %1,%2,%0" : "=a" (outputvar) : "g" (x + y), "di" (*z));
```

is represented using a single `asm_operands` RTX which represents the value that is stored in `outputvar`:

```
(set rtx-for-outputvar
  (asm_operands "foo %1,%2,%0" "a" 0
    [rtx-for-addition-result rtx-for-*z]
    [(asm_input:m1 "g")
     (asm_input:m2 "di")]))
```

Here the operands of the `asm_operands` RTX are the assembler template string, the output-operand's constraint, the index-number of the output operand among the output operands specified, a vector of input operand RTX's, and a vector of input-operand modes and constraints. The mode `m1` is the mode of the sum `x+y`; `m2` is that of `*z`.

When an `asm` statement has multiple output values, its `insn` has several such `set` RTX's inside of a `parallel`. Each `set` contains a `asm_operands`; all of these share the same assembler template and vectors, but each contains the constraint for the respective output operand. They are also distinguished by the output-operand index number, which is 0, 1, ... for successive output operands.

8.17 Insns

The RTL representation of the code for a function is a doubly-linked chain of objects called *insns*. Insns are expressions with special codes that are used for no other purpose. Some insns are actual instructions; others represent dispatch tables for `switch` statements; others represent labels to jump to or various sorts of declarative information.

In addition to its own specific data, each `insn` must have a unique id-number that distinguishes it from all other insns in the current function (after delayed branch scheduling, copies of an `insn` with the same id-number may be present in multiple places in a function, but these copies will always be identical and will only appear inside a `sequence`), and chain pointers to the preceding and following insns. These three fields occupy the same position in every `insn`, independent of the expression code of the `insn`. They could be accessed with `XEXP` and `XINT`, but instead three special macros are always used:

`INSN_UID` (*i*)

Accesses the unique id of `insn i`.

`PREV_INSN` (*i*)

Accesses the chain pointer to the `insn` preceding *i*. If *i* is the first `insn`, this is a null pointer.

`NEXT_INSN` (*i*)

Accesses the chain pointer to the `insn` following *i*. If *i* is the last `insn`, this is a null pointer.

The first *insn* in the chain is obtained by calling `get_insns`; the last *insn* is the result of calling `get_last_insn`. Within the chain delimited by these *insns*, the `NEXT_INSN` and `PREV_INSN` pointers must always correspond: if *insn* is not the first *insn*,

```
NEXT_INSN (PREV_INSN (insn)) == insn
```

is always true and if *insn* is not the last *insn*,

```
PREV_INSN (NEXT_INSN (insn)) == insn
```

is always true.

After delay slot scheduling, some of the *insns* in the chain might be **sequence** expressions, which contain a vector of *insns*. The value of `NEXT_INSN` in all but the last of these *insns* is the next *insn* in the vector; the value of `NEXT_INSN` of the last *insn* in the vector is the same as the value of `NEXT_INSN` for the **sequence** in which it is contained. Similar rules apply for `PREV_INSN`.

This means that the above invariants are not necessarily true for *insns* inside **sequence** expressions. Specifically, if *insn* is the first *insn* in a **sequence**, `NEXT_INSN (PREV_INSN (insn))` is the *insn* containing the **sequence** expression, as is the value of `PREV_INSN (NEXT_INSN (insn))` if *insn* is the last *insn* in the **sequence** expression. You can use these expressions to find the containing **sequence** expression.

Every *insn* has one of the following six expression codes:

insn The expression code `insn` is used for instructions that do not jump and do not do function calls. **sequence** expressions are always contained in *insns* with code `insn` even if one of those *insns* should jump or do function calls.

Insns with code `insn` have four additional fields beyond the three mandatory ones listed above. These four are described in a table below.

jump_insn

The expression code `jump_insn` is used for instructions that may jump (or, more generally, may contain `label_ref` expressions). If there is an instruction to return from the current function, it is recorded as a `jump_insn`.

`jump_insn` *insns* have the same extra fields as `insn` *insns*, accessed in the same way and in addition contain a field `JUMP_LABEL` which is defined once jump optimization has completed.

For simple conditional and unconditional jumps, this field contains the `code_label` to which this *insn* will (possibly conditionally) branch. In a more complex jump, `JUMP_LABEL` records one of the labels that the *insn* refers to; the only way to find the others is to scan the entire body of the *insn*. In an `addr_vec`, `JUMP_LABEL` is `NULL_RTX`.

Return *insns* count as jumps, but since they do not refer to any labels, their `JUMP_LABEL` is `NULL_RTX`.

call_insn

The expression code `call_insn` is used for instructions that may do function calls. It is important to distinguish these instructions because they imply that certain registers and memory locations may be altered unpredictably.

`call_insn` *insns* have the same extra fields as `insn` *insns*, accessed in the same way and in addition contain a field `CALL_INSN_FUNCTION_USAGE`, which contains

a list (chain of `expr_list` expressions) containing `use` and `clobber` expressions that denote hard registers and MEMs used or clobbered by the called function.

A `MEM` generally points to a stack slots in which arguments passed to the libcall by reference (see Section 10.10.7 [Register Arguments], page 242) are stored. If the argument is caller-copied (see Section 10.10.7 [Register Arguments], page 242), the stack slot will be mentioned in `CLOBBER` and `USE` entries; if it's callee-copied, only a `USE` will appear, and the `MEM` may point to addresses that are not stack slots. These `MEMs` are used only in libcalls, because, unlike regular function calls, `CONST_CALLs` (which libcalls generally are, see Section 8.4 [Flags], page 75) aren't assumed to read and write all memory, so flow would consider the stores dead and remove them. Note that, since a libcall must never return values in memory (see Section 10.10.9 [Aggregate Return], page 247), there will never be a `CLOBBER` for a memory address holding a return value.

`CLOBBERed` registers in this list augment registers specified in `CALL_USED_REGISTERS` (see Section 10.8.1 [Register Basics], page 218).

`code_label`

A `code_label` insn represents a label that a jump insn can jump to. It contains two special fields of data in addition to the three standard ones. `CODE_LABEL_NUMBER` is used to hold the *label number*, a number that identifies this label uniquely among all the labels in the compilation (not just in the current function). Ultimately, the label is represented in the assembler output as an assembler label, usually of the form '`Ln`' where *n* is the label number.

When a `code_label` appears in an RTL expression, it normally appears within a `label_ref` which represents the address of the label, as a number.

Besides as a `code_label`, a label can also be represented as a `note` of type `NOTE_INSN_DELETED_LABEL`.

The field `LABEL_NUSES` is only defined once the jump optimization phase is completed. It contains the number of times this label is referenced in the current function.

The field `LABEL_KIND` differentiates four different types of labels: `LABEL_NORMAL`, `LABEL_STATIC_ENTRY`, `LABEL_GLOBAL_ENTRY`, and `LABEL_WEAK_ENTRY`. The only labels that do not have type `LABEL_NORMAL` are *alternate entry points* to the current function. These may be static (visible only in the containing translation unit), global (exposed to all translation units), or weak (global, but can be overridden by another symbol with the same name).

Much of the compiler treats all four kinds of label identically. Some of it needs to know whether or not a label is an alternate entry point; for this purpose, the macro `LABEL_ALT_ENTRY_P` is provided. It is equivalent to testing whether '`LABEL_KIND (label) == LABEL_NORMAL`'. The only place that cares about the distinction between static, global, and weak alternate entry points, besides the front-end code that creates them, is the function `output_alternate_entry_point`, in '`final.c`'.

To set the kind of a label, use the `SET_LABEL_KIND` macro.

`barrier`

Barriers are placed in the instruction stream when control cannot flow past them. They are placed after unconditional jump instructions to indicate that

the jumps are unconditional and after calls to `volatile` functions, which do not return (e.g., `exit`). They contain no information beyond the three standard fields.

note `note` insns are used to represent additional debugging and declarative information. They contain two nonstandard fields, an integer which is accessed with the macro `NOTE_LINE_NUMBER` and a string accessed with `NOTE_SOURCE_FILE`. If `NOTE_LINE_NUMBER` is positive, the note represents the position of a source line and `NOTE_SOURCE_FILE` is the source file name that the line came from. These notes control generation of line number data in the assembler output. Otherwise, `NOTE_LINE_NUMBER` is not really a line number but a code with one of the following values (and `NOTE_SOURCE_FILE` must contain a null pointer):

`NOTE_INSN_DELETED`

Such a note is completely ignorable. Some passes of the compiler delete insns by altering them into notes of this kind.

`NOTE_INSN_DELETED_LABEL`

This marks what used to be a `code_label`, but was not used for other purposes than taking its address and was transformed to mark that no code jumps to it.

`NOTE_INSN_BLOCK_BEG`

`NOTE_INSN_BLOCK_END`

These types of notes indicate the position of the beginning and end of a level of scoping of variable names. They control the output of debugging information.

`NOTE_INSN_EH_REGION_BEG`

`NOTE_INSN_EH_REGION_END`

These types of notes indicate the position of the beginning and end of a level of scoping for exception handling. `NOTE_BLOCK_NUMBER` identifies which `CODE_LABEL` or `note` of type `NOTE_INSN_DELETED_LABEL` is associated with the given region.

`NOTE_INSN_LOOP_BEG`

`NOTE_INSN_LOOP_END`

These types of notes indicate the position of the beginning and end of a `while` or `for` loop. They enable the loop optimizer to find loops quickly.

`NOTE_INSN_LOOP_CONT`

Appears at the place in a loop that `continue` statements jump to.

`NOTE_INSN_LOOP_VTOP`

This note indicates the place in a loop where the exit test begins for those loops in which the exit test has been duplicated. This position becomes another virtual start of the loop when considering loop invariants.

NOTE_INSN_FUNCTION_END

Appears near the end of the function body, just before the label that **return** statements jump to (on machine where a single instruction does not suffice for returning). This note may be deleted by jump optimization.

NOTE_INSN_SETJMP

Appears following each call to **setjmp** or a related function.

These codes are printed symbolically when they appear in debugging dumps.

The machine mode of an **insn** is normally **VOIDmode**, but some phases use the mode for various purposes.

The common subexpression elimination pass sets the mode of an **insn** to **QImode** when it is the first **insn** in a block that has already been processed.

The second Haifa scheduling pass, for targets that can multiple issue, sets the mode of an **insn** to **TImode** when it is believed that the instruction begins an issue group. That is, when the instruction cannot issue simultaneously with the previous. This may be relied on by later passes, in particular machine-dependent reorg.

Here is a table of the extra fields of **insn**, **jump_insn** and **call_insn** insns:

PATTERN (*i*)

An expression for the side effect performed by this **insn**. This must be one of the following codes: **set**, **call**, **use**, **clobber**, **return**, **asm_input**, **asm_output**, **addr_vec**, **addr_diff_vec**, **trap_if**, **unspec**, **unspec_volatile**, **parallel**, **cond_exec**, or **sequence**. If it is a **parallel**, each element of the **parallel** must be one these codes, except that **parallel** expressions cannot be nested and **addr_vec** and **addr_diff_vec** are not permitted inside a **parallel** expression.

INSN_CODE (*i*)

An integer that says which pattern in the machine description matches this **insn**, or **-1** if the matching has not yet been attempted.

Such matching is never attempted and this field remains **-1** on an **insn** whose pattern consists of a single **use**, **clobber**, **asm_input**, **addr_vec** or **addr_diff_vec** expression.

Matching is also never attempted on insns that result from an **asm** statement. These contain at least one **asm_operands** expression. The function **asm_noperands** returns a non-negative value for such insns.

In the debugging output, this field is printed as a number followed by a symbolic representation that locates the pattern in the 'md' file as some small positive or negative offset from a named pattern.

LOG_LINKS (*i*)

A list (chain of **insn_list** expressions) giving information about dependencies between instructions within a basic block. Neither a jump nor a label may come between the related insns.

REG_NOTES (*i*)

A list (chain of `expr_list` and `insn_list` expressions) giving miscellaneous information about the `insn`. It is often information pertaining to the registers used in this `insn`.

The `LOG_LINKS` field of an `insn` is a chain of `insn_list` expressions. Each of these has two operands: the first is an `insn`, and the second is another `insn_list` expression (the next one in the chain). The last `insn_list` in the chain has a null pointer as second operand. The significant thing about the chain is which `insns` appear in it (as first operands of `insn_list` expressions). Their order is not significant.

This list is originally set up by the flow analysis pass; it is a null pointer until then. Flow only adds links for those data dependencies which can be used for instruction combination. For each `insn`, the flow analysis pass adds a link to `insns` which store into registers values that are used for the first time in this `insn`. The instruction scheduling pass adds extra links so that every dependence will be represented. Links represent data dependencies, antidependencies and output dependencies; the machine mode of the link distinguishes these three types: antidependencies have mode `REG_DEP_ANTI`, output dependencies have mode `REG_DEP_OUTPUT`, and data dependencies have mode `VOIDmode`.

The `REG_NOTES` field of an `insn` is a chain similar to the `LOG_LINKS` field but it includes `expr_list` expressions in addition to `insn_list` expressions. There are several kinds of register notes, which are distinguished by the machine mode, which in a register note is really understood as being an `enum reg_note`. The first operand `op` of the note is data whose meaning depends on the kind of note.

The macro `REG_NOTE_KIND (x)` returns the kind of register note. Its counterpart, the macro `PUT_REG_NOTE_KIND (x, newkind)` sets the register note type of `x` to be `newkind`.

Register notes are of three classes: They may say something about an input to an `insn`, they may say something about an output of an `insn`, or they may create a linkage between two `insns`. There are also a set of values that are only used in `LOG_LINKS`.

These register notes annotate inputs to an `insn`:

REG_DEAD The value in `op` dies in this `insn`; that is to say, altering the value immediately after this `insn` would not affect the future behavior of the program.

It does not follow that the register `op` has no useful value after this `insn` since `op` is not necessarily modified by this `insn`. Rather, no subsequent instruction uses the contents of `op`.

REG_UNUSED

The register `op` being set by this `insn` will not be used in a subsequent `insn`. This differs from a `REG_DEAD` note, which indicates that the value in an input will not be used subsequently. These two notes are independent; both may be present for the same register.

REG_INC The register `op` is incremented (or decremented; at this level there is no distinction) by an embedded side effect inside this `insn`. This means it appears in a `post_inc`, `pre_inc`, `post_dec` or `pre_dec` expression.

REG_NONNEG

The register *op* is known to have a nonnegative value when this insn is reached. This is used so that decrement and branch until zero instructions, such as the m68k *dbra*, can be matched.

The **REG_NONNEG** note is added to insns only if the machine description has a ‘*decrement_and_branch_until_zero*’ pattern.

REG_NO_CONFLICT

This insn does not cause a conflict between *op* and the item being set by this insn even though it might appear that it does. In other words, if the destination register and *op* could otherwise be assigned the same register, this insn does not prevent that assignment.

Insns with this note are usually part of a block that begins with a **clobber** insn specifying a multi-word pseudo register (which will be the output of the block), a group of insns that each set one word of the value and have the **REG_NO_CONFLICT** note attached, and a final insn that copies the output to itself with an attached **REG_EQUAL** note giving the expression being computed. This block is encapsulated with **REG_LIBCALL** and **REG_RETVAL** notes on the first and last insns, respectively.

REG_LABEL

This insn uses *op*, a *code_label* or a note of type **NOTE_INSN_DELETED_LABEL**, but is not a **jump_insn**, or it is a **jump_insn** that required the label to be held in a register. The presence of this note allows jump optimization to be aware that *op* is, in fact, being used, and flow optimization to build an accurate flow graph.

The following notes describe attributes of outputs of an insn:

REG_EQUIV**REG_EQUAL**

This note is only valid on an insn that sets only one register and indicates that that register will be equal to *op* at run time; the scope of this equivalence differs between the two types of notes. The value which the insn explicitly copies into the register may look different from *op*, but they will be equal at run time. If the output of the single **set** is a **strict_low_part** expression, the note refers to the register that is contained in **SUBREG_REG** of the **subreg** expression.

For **REG_EQUIV**, the register is equivalent to *op* throughout the entire function, and could validly be replaced in all its occurrences by *op*. (“Validly” here refers to the data flow of the program; simple replacement may make some insns invalid.) For example, when a constant is loaded into a register that is never assigned any other value, this kind of note is used.

When a parameter is copied into a pseudo-register at entry to a function, a note of this kind records that the register is equivalent to the stack slot where the parameter was passed. Although in this case the register may be set by other insns, it is still valid to replace the register by the stack slot throughout the function.

A `REG_EQUIV` note is also used on an instruction which copies a register parameter into a pseudo-register at entry to a function, if there is a stack slot where that parameter could be stored. Although other insns may set the pseudo-register, it is valid for the compiler to replace the pseudo-register by stack slot throughout the function, provided the compiler ensures that the stack slot is properly initialized by making the replacement in the initial copy instruction as well. This is used on machines for which the calling convention allocates stack space for register parameters. See `REG_PARM_STACK_SPACE` in Section 10.10.6 [Stack Arguments], page 239.

In the case of `REG_EQUAL`, the register that is set by this insn will be equal to *op* at run time at the end of this insn but not necessarily elsewhere in the function. In this case, *op* is typically an arithmetic expression. For example, when a sequence of insns such as a library call is used to perform an arithmetic operation, this kind of note is attached to the insn that produces or copies the final value.

These two notes are used in different ways by the compiler passes. `REG_EQUAL` is used by passes prior to register allocation (such as common subexpression elimination and loop optimization) to tell them how to think of that value. `REG_EQUIV` notes are used by register allocation to indicate that there is an available substitute expression (either a constant or a `mem` expression for the location of a parameter on the stack) that may be used in place of a register if insufficient registers are available.

Except for stack homes for parameters, which are indicated by a `REG_EQUIV` note and are not useful to the early optimization passes and pseudo registers that are equivalent to a memory location throughout their entire life, which is not detected until later in the compilation, all equivalences are initially indicated by an attached `REG_EQUAL` note. In the early stages of register allocation, a `REG_EQUAL` note is changed into a `REG_EQUIV` note if *op* is a constant and the insn represents the only set of its destination register.

Thus, compiler passes prior to register allocation need only check for `REG_EQUAL` notes and passes subsequent to register allocation need only check for `REG_EQUIV` notes.

`REG_WAS_0`

The single output of this insn contained zero before this insn. *op* is the insn that set it to zero. You can rely on this note if it is present and *op* has not been deleted or turned into a `note`; its absence implies nothing.

These notes describe linkages between insns. They occur in pairs: one insn has one of a pair of notes that points to a second insn, which has the inverse note pointing back to the first insn.

`REG_RETVAL`

This insn copies the value of a multi-insn sequence (for example, a library call), and *op* is the first insn of the sequence (for a library call, the first insn that was generated to set up the arguments for the library call).

Loop optimization uses this note to treat such a sequence as a single operation for code motion purposes and flow analysis uses this note to delete such sequences whose results are dead.

A `REG_EQUAL` note will also usually be attached to this insn to provide the expression being computed by the sequence.

These notes will be deleted after reload, since they are no longer accurate or useful.

`REG_LIBCALL`

This is the inverse of `REG_RETVAL`: it is placed on the first insn of a multi-insn sequence, and it points to the last one.

These notes are deleted after reload, since they are no longer useful or accurate.

`REG_CC_SETTER`

`REG_CC_USER`

On machines that use `cc0`, the insns which set and use `cc0` set and use `cc0` are adjacent. However, when branch delay slot filling is done, this may no longer be true. In this case a `REG_CC_USER` note will be placed on the insn setting `cc0` to point to the insn using `cc0` and a `REG_CC_SETTER` note will be placed on the insn using `cc0` to point to the insn setting `cc0`.

These values are only used in the `LOG_LINKS` field, and indicate the type of dependency that each link represents. Links which indicate a data dependence (a read after write dependence) do not use any code, they simply have mode `VOIDmode`, and are printed without any descriptive text.

`REG_DEP_ANTI`

This indicates an anti dependence (a write after read dependence).

`REG_DEP_OUTPUT`

This indicates an output dependence (a write after write dependence).

These notes describe information gathered from gcov profile data. They are stored in the `REG_NOTES` field of an insn as an `expr_list`.

`REG_BR_PROB`

This is used to specify the ratio of branches to non-branches of a branch insn according to the profile data. The value is stored as a value between 0 and `REG_BR_PROB_BASE`; larger values indicate a higher probability that the branch will be taken.

`REG_BR_PRED`

These notes are found in `JUMP` insns after delayed branch scheduling has taken place. They indicate both the direction and the likelihood of the `JUMP`. The format is a bitmask of `ATTR_FLAG_*` values.

`REG_FRAME_RELATED_EXPR`

This is used on an `RTX_FRAME_RELATED_P` insn wherein the attached expression is used in place of the actual insn pattern. This is done in cases where the pattern is either complex or misleading.

For convenience, the machine mode in an `insn_list` or `expr_list` is printed using these symbolic codes in debugging dumps.

The only difference between the expression codes `insn_list` and `expr_list` is that the first operand of an `insn_list` is assumed to be an `insn` and is printed in debugging dumps as the `insn`'s unique id; the first operand of an `expr_list` is printed in the ordinary way as an expression.

8.18 RTL Representation of Function-Call Insns

Insns that call subroutines have the RTL expression code `call_insn`. These insns must satisfy special rules, and their bodies must use a special RTL expression code, `call`.

A `call` expression has two operands, as follows:

```
(call (mem:fm addr) nbytes)
```

Here `nbytes` is an operand that represents the number of bytes of argument data being passed to the subroutine, `fm` is a machine mode (which must equal as the definition of the `FUNCTION_MODE` macro in the machine description) and `addr` represents the address of the subroutine.

For a subroutine that returns no value, the `call` expression as shown above is the entire body of the `insn`, except that the `insn` might also contain `use` or `clobber` expressions.

For a subroutine that returns a value whose mode is not `BLKmode`, the value is returned in a hard register. If this register's number is `r`, then the body of the `call` `insn` looks like this:

```
(set (reg:m r)
      (call (mem:fm addr) nbytes))
```

This RTL expression makes it clear (to the optimizer passes) that the appropriate register receives a useful value in this `insn`.

When a subroutine returns a `BLKmode` value, it is handled by passing to the subroutine the address of a place to store the value. So the `call` `insn` itself does not “return” any value, and it has the same RTL form as a `call` that returns nothing.

On some machines, the `call` instruction itself clobbers some register, for example to contain the return address. `call_insn` insns on these machines should have a body which is a `parallel` that contains both the `call` expression and `clobber` expressions that indicate which registers are destroyed. Similarly, if the `call` instruction requires some register other than the stack pointer that is not explicitly mentioned in its RTL, a `use` subexpression should mention that register.

Functions that are called are assumed to modify all registers listed in the configuration macro `CALL_USED_REGISTERS` (see Section 10.8.1 [Register Basics], page 218) and, with the exception of `const` functions and library calls, to modify all of memory.

Insns containing just `use` expressions directly precede the `call_insn` `insn` to indicate which registers contain inputs to the function. Similarly, if registers other than those in `CALL_USED_REGISTERS` are clobbered by the called function, insns containing a single `clobber` follow immediately after the `call` to indicate which registers.

8.19 Structure Sharing Assumptions

The compiler assumes that certain kinds of RTL expressions are unique; there do not exist two distinct objects representing the same value. In other cases, it makes an opposite assumption: that no RTL expression object of a certain kind appears in more than one place in the containing structure.

These assumptions refer to a single function; except for the RTL objects that describe global variables and external functions, and a few standard objects such as small integer constants, no RTL objects are common to two functions.

- Each pseudo-register has only a single `reg` object to represent it, and therefore only a single machine mode.
- For any symbolic label, there is only one `symbol_ref` object referring to it.
- All `const_int` expressions with equal values are shared.
- There is only one `pc` expression.
- There is only one `cc0` expression.
- There is only one `const_double` expression with value 0 for each floating point mode. Likewise for values 1 and 2.
- There is only one `const_vector` expression with value 0 for each vector mode, be it an integer or a double constant vector.
- No `label_ref` or `scratch` appears in more than one place in the RTL structure; in other words, it is safe to do a tree-walk of all the insns in the function and assume that each time a `label_ref` or `scratch` is seen it is distinct from all others that are seen.
- Only one `mem` object is normally created for each static variable or stack slot, so these objects are frequently shared in all the places they appear. However, separate but equal objects for these variables are occasionally made.
- When a single `asm` statement has multiple output operands, a distinct `asm_operands` expression is made for each output operand. However, these all share the vector which contains the sequence of input operands. This sharing is used later on to test whether two `asm_operands` expressions come from the same statement, so all optimizations must carefully preserve the sharing if they copy the vector at all.
- No RTL object appears in more than one place in the RTL structure except as described above. Many passes of the compiler rely on this by assuming that they can modify RTL objects in place without unwanted side-effects on other insns.
- During initial RTL generation, shared structure is freely introduced. After all the RTL for a function has been generated, all shared structure is copied by `unshare_all_rtl` in `'emit-rtl.c'`, after which the above rules are guaranteed to be followed.
- During the combiner pass, shared structure within an insn can exist temporarily. However, the shared structure is copied before the combiner is finished with the insn. This is done by calling `copy_rtx_if_shared`, which is a subroutine of `unshare_all_rtl`.

8.20 Reading RTL

To read an RTL object from a file, call `read_rtx`. It takes one argument, a stdio stream, and returns a single RTL object. This routine is defined in `'read-rtl.c'`. It is not available

in the compiler itself, only the various programs that generate the compiler back end from the machine description.

People frequently have the idea of using RTL stored as text in a file as an interface between a language front end and the bulk of GCC. This idea is not feasible.

GCC was designed to use RTL internally only. Correct RTL for a given program is very dependent on the particular target machine. And the RTL does not contain all the information about the program.

The proper way to interface GCC to a new language front end is with the “tree” data structure, described in the files ‘`tree.h`’ and ‘`tree.def`’. The documentation for this structure (see Chapter 7 [Trees], page 39) is incomplete.

9 Machine Descriptions

A machine description has two parts: a file of instruction patterns (`.md` file) and a C header file of macro definitions.

The `.md` file for a target machine contains a pattern for each instruction that the target machine supports (or at least each instruction that is worth telling the compiler about). It may also contain comments. A semicolon causes the rest of the line to be a comment, unless the semicolon is inside a quoted string.

See the next chapter for information on the C header file.

9.1 Overview of How the Machine Description is Used

There are three main conversions that happen in the compiler:

1. The front end reads the source code and builds a parse tree.
2. The parse tree is used to generate an RTL insn list based on named instruction patterns.
3. The insn list is matched against the RTL templates to produce assembler code.

For the generate pass, only the names of the insns matter, from either a named `define_insn` or a `define_expand`. The compiler will choose the pattern with the right name and apply the operands according to the documentation later in this chapter, without regard for the RTL template or operand constraints. Note that the names the compiler looks for are hard-coded in the compiler—it will ignore unnamed patterns and patterns with names it doesn't know about, but if you don't provide a named pattern it needs, it will abort.

If a `define_insn` is used, the template given is inserted into the insn list. If a `define_expand` is used, one of three things happens, based on the condition logic. The condition logic may manually create new insns for the insn list, say via `emit_insn()`, and invoke `DONE`. For certain named patterns, it may invoke `FAIL` to tell the compiler to use an alternate way of performing that task. If it invokes neither `DONE` nor `FAIL`, the template given in the pattern is inserted, as if the `define_expand` were a `define_insn`.

Once the insn list is generated, various optimization passes convert, replace, and rearrange the insns in the insn list. This is where the `define_split` and `define_peephole` patterns get used, for example.

Finally, the insn list's RTL is matched up with the RTL templates in the `define_insn` patterns, and those patterns are used to emit the final assembly code. For this purpose, each named `define_insn` acts like it's unnamed, since the names are ignored.

9.2 Everything about Instruction Patterns

Each instruction pattern contains an incomplete RTL expression, with pieces to be filled in later, operand constraints that restrict how the pieces can be filled in, and an output pattern or C code to generate the assembler output, all wrapped up in a `define_insn` expression.

A `define_insn` is an RTL expression containing four or five operands:

1. An optional name. The presence of a name indicate that this instruction pattern can perform a certain standard job for the RTL-generation pass of the compiler. This pass knows certain names and will use the instruction patterns with those names, if the names are defined in the machine description.

The absence of a name is indicated by writing an empty string where the name should go. Nameless instruction patterns are never used for generating RTL code, but they may permit several simpler insns to be combined later on.

Names that are not thus known and used in RTL-generation have no effect; they are equivalent to no name at all.

For the purpose of debugging the compiler, you may also specify a name beginning with the '*' character. Such a name is used only for identifying the instruction in RTL dumps; it is entirely equivalent to having a nameless pattern for all other purposes.

2. The *RTL template* (see Section 9.4 [RTL Template], page 119) is a vector of incomplete RTL expressions which show what the instruction should look like. It is incomplete because it may contain `match_operand`, `match_operator`, and `match_dup` expressions that stand for operands of the instruction.

If the vector has only one element, that element is the template for the instruction pattern. If the vector has multiple elements, then the instruction pattern is a **parallel** expression containing the elements described.

3. A condition. This is a string which contains a C expression that is the final test to decide whether an insn body matches this pattern.

For a named pattern, the condition (if present) may not depend on the data in the insn being matched, but only the target-machine-type flags. The compiler needs to test these conditions during initialization in order to learn exactly which named instructions are available in a particular run.

For nameless patterns, the condition is applied only when matching an individual insn, and only after the insn has matched the pattern's recognition template. The insn's operands may be found in the vector `operands`. For an insn where the condition has once matched, it can't be used to control register allocation, for example by excluding certain hard registers or hard register combinations.

4. The *output template*: a string that says how to output matching insns as assembler code. '%' in this string specifies where to substitute the value of an operand. See Section 9.5 [Output Template], page 123.

When simple substitution isn't general enough, you can specify a piece of C code to compute the output. See Section 9.6 [Output Statement], page 124.

5. Optionally, a vector containing the values of attributes for insns matching this pattern. See Section 9.18 [Insn Attributes], page 175.

9.3 Example of `define_insn`

Here is an actual example of an instruction pattern, for the 68000/68020.

```
(define_insn "tstsi"
  [(set (cc0)
        (match_operand:SI 0 "general_operand" "rm"))]
```

```

""
"*
{
  if (TARGET_68020 || ! ADDRESS_REG_P (operands[0]))
    return \"tstl %0\";
  return \"cpl #0,%0\";
})

```

This can also be written using braced strings:

```

(define_insn "tstsi"
  [(set (cc0)
        (match_operand:SI 0 "general_operand" "rm"))]
  ""
  {
    if (TARGET_68020 || ! ADDRESS_REG_P (operands[0]))
      return "tstl %0";
    return "cpl #0,%0";
  })

```

This is an instruction that sets the condition codes based on the value of a general operand. It has no condition, so any insn whose RTL description has the form shown may be handled according to this pattern. The name ‘tstsi’ means “test a SImode value” and tells the RTL generation pass that, when it is necessary to test such a value, an insn to do so can be constructed using this pattern.

The output control string is a piece of C code which chooses which output template to return based on the kind of operand and the specific type of CPU for which code is being generated.

‘rm’ is an operand constraint. Its meaning is explained below.

9.4 RTL Template

The RTL template is used to define which insns match the particular pattern and how to find their operands. For named patterns, the RTL template also says how to construct an insn from specified operands.

Construction involves substituting specified operands into a copy of the template. Matching involves determining the values that serve as the operands in the insn being matched. Both of these activities are controlled by special expression types that direct matching and substitution of the operands.

(match_operand:m n predicate constraint)

This expression is a placeholder for operand number *n* of the insn. When constructing an insn, operand number *n* will be substituted at this point. When matching an insn, whatever appears at this position in the insn will be taken as operand number *n*; but it must satisfy *predicate* or this instruction pattern will not match at all.

Operand numbers must be chosen consecutively counting from zero in each instruction pattern. There may be only one *match_operand* expression in the

pattern for each operand number. Usually operands are numbered in the order of appearance in `match_operand` expressions. In the case of a `define_expand`, any operand numbers used only in `match_dup` expressions have higher values than all other operand numbers.

predicate is a string that is the name of a C function that accepts two arguments, an expression and a machine mode. During matching, the function will be called with the putative operand as the expression and *m* as the mode argument (if *m* is not specified, `VOIDmode` will be used, which normally causes *predicate* to accept any mode). If it returns zero, this instruction pattern fails to match. *predicate* may be an empty string; then it means no test is to be done on the operand, so anything which occurs in this position is valid.

Most of the time, *predicate* will reject modes other than *m*—but not always. For example, the predicate `address_operand` uses *m* as the mode of memory ref that the address should be valid for. Many predicates accept `const_int` nodes even though their mode is `VOIDmode`.

constraint controls reloading and the choice of the best register class to use for a value, as explained later (see Section 9.7 [Constraints], page 126).

People are often unclear on the difference between the constraint and the predicate. The predicate helps decide whether a given insn matches the pattern. The constraint plays no role in this decision; instead, it controls various decisions in the case of an insn which does match.

On CISC machines, the most common *predicate* is `"general_operand"`. This function checks that the putative operand is either a constant, a register or a memory reference, and that it is valid for mode *m*.

For an operand that must be a register, *predicate* should be `"register_operand"`. Using `"general_operand"` would be valid, since the reload pass would copy any non-register operands through registers, but this would make GCC do extra work, it would prevent invariant operands (such as constant) from being removed from loops, and it would prevent the register allocator from doing the best possible job. On RISC machines, it is usually most efficient to allow *predicate* to accept only objects that the constraints allow.

For an operand that must be a constant, you must be sure to either use `"immediate_operand"` for *predicate*, or make the instruction pattern's extra condition require a constant, or both. You cannot expect the constraints to do this work! If the constraints allow only constants, but the predicate allows something else, the compiler will crash when that case arises.

`(match_scratch:m n constraint)`

This expression is also a placeholder for operand number *n* and indicates that operand must be a `scratch` or `reg` expression.

When matching patterns, this is equivalent to

```
(match_operand:m n "scratch_operand" pred)
```

but, when generating RTL, it produces a `(scratch:m)` expression.

If the last few expressions in a `parallel` are `clobber` expressions whose operands are either a hard register or `match_scratch`, the combiner can add or delete them when necessary. See Section 8.14 [Side Effects], page 97.

`(match_dup n)`

This expression is also a placeholder for operand number *n*. It is used when the operand needs to appear more than once in the insn.

In construction, `match_dup` acts just like `match_operand`: the operand is substituted into the insn being constructed. But in matching, `match_dup` behaves differently. It assumes that operand number *n* has already been determined by a `match_operand` appearing earlier in the recognition template, and it matches only an identical-looking expression.

Note that `match_dup` should not be used to tell the compiler that a particular register is being used for two operands (example: `add` that adds one register to another; the second register is both an input operand and the output operand). Use a matching constraint (see Section 9.7.1 [Simple Constraints], page 126) for those. `match_dup` is for the cases where one operand is used in two places in the template, such as an instruction that computes both a quotient and a remainder, where the opcode takes two input operands but the RTL template has to refer to each of those twice; once for the quotient pattern and once for the remainder pattern.

`(match_operator:m n predicate [operands...])`

This pattern is a kind of placeholder for a variable RTL expression code.

When constructing an insn, it stands for an RTL expression whose expression code is taken from that of operand *n*, and whose operands are constructed from the patterns *operands*.

When matching an expression, it matches an expression if the function *predicate* returns nonzero on that expression *and* the patterns *operands* match the operands of the expression.

Suppose that the function `commutative_operator` is defined as follows, to match any expression whose operator is one of the commutative arithmetic operators of RTL and whose mode is *mode*:

```
int
commutative_operator (x, mode)
  rtx x;
  enum machine_mode mode;
{
  enum rtx_code code = GET_CODE (x);
  if (GET_MODE (x) != mode)
    return 0;
  return (GET_RTX_CLASS (code) == 'c'
          || code == EQ || code == NE);
}
```

Then the following pattern will match any RTL expression consisting of a commutative operator applied to two general operands:

```
(match_operator:SI 3 "commutative_operator"
 [(match_operand:SI 1 "general_operand" "g")
  (match_operand:SI 2 "general_operand" "g")])
```

Here the vector *[operands...]* contains two patterns because the expressions to be matched all contain two operands.

When this pattern does match, the two operands of the commutative operator are recorded as operands 1 and 2 of the insn. (This is done by the two instances

of `match_operand`.) Operand 3 of the `insn` will be the entire commutative expression: use `GET_CODE (operands[3])` to see which commutative operator was used.

The machine mode `m` of `match_operator` works like that of `match_operand`: it is passed as the second argument to the predicate function, and that function is solely responsible for deciding whether the expression to be matched “has” that mode.

When constructing an `insn`, argument 3 of the gen-function will specify the operation (i.e. the expression code) for the expression to be made. It should be an RTL expression, whose expression code is copied into a new expression whose operands are arguments 1 and 2 of the gen-function. The subexpressions of argument 3 are not used; only its expression code matters.

When `match_operator` is used in a pattern for matching an `insn`, it usually best if the operand number of the `match_operator` is higher than that of the actual operands of the `insn`. This improves register allocation because the register allocator often looks at operands 1 and 2 of `insns` to see if it can do register tying.

There is no way to specify constraints in `match_operator`. The operand of the `insn` which corresponds to the `match_operator` never has any constraints because it is never reloaded as a whole. However, if parts of its `operands` are matched by `match_operand` patterns, those parts may have constraints of their own.

`(match_op_dup:m n [operands...])`

Like `match_dup`, except that it applies to operators instead of operands. When constructing an `insn`, operand number `n` will be substituted at this point. But in matching, `match_op_dup` behaves differently. It assumes that operand number `n` has already been determined by a `match_operator` appearing earlier in the recognition template, and it matches only an identical-looking expression.

`(match_parallel n predicate [subpat...])`

This pattern is a placeholder for an `insn` that consists of a `parallel` expression with a variable number of elements. This expression should only appear at the top level of an `insn` pattern.

When constructing an `insn`, operand number `n` will be substituted at this point. When matching an `insn`, it matches if the body of the `insn` is a `parallel` expression with at least as many elements as the vector of `subpat` expressions in the `match_parallel`, if each `subpat` matches the corresponding element of the `parallel`, and the function `predicate` returns nonzero on the `parallel` that is the body of the `insn`. It is the responsibility of the predicate to validate elements of the `parallel` beyond those listed in the `match_parallel`.

A typical use of `match_parallel` is to match load and store multiple expressions, which can contain a variable number of elements in a `parallel`. For example,

```
(define_insn ""
  [(match_parallel 0 "load_multiple_operation"
    [(set (match_operand:SI 1 "gpc_reg_operand" "=r")
```

```

        (match_operand:SI 2 "memory_operand" "m")
      (use (reg:SI 179))
      (clobber (reg:SI 179))]]]
    ""
    "loadm 0,0,%1,%2")

```

This example comes from ‘a29k.md’. The function `load_multiple_operation` is defined in ‘a29k.c’ and checks that subsequent elements in the `parallel` are the same as the `set` in the pattern, except that they are referencing subsequent registers and memory locations.

An insn that matches this pattern might look like:

```

(parallel
 [(set (reg:SI 20) (mem:SI (reg:SI 100)))
  (use (reg:SI 179))
  (clobber (reg:SI 179))
  (set (reg:SI 21)
       (mem:SI (plus:SI (reg:SI 100)
                        (const_int 4))))
  (set (reg:SI 22)
       (mem:SI (plus:SI (reg:SI 100)
                        (const_int 8))))])

```

`(match_par_dup n [subpat...])`

Like `match_op_dup`, but for `match_parallel` instead of `match_operator`.

`(match_insn predicate)`

Match a complete insn. Unlike the other `match_*` recognizers, `match_insn` does not take an operand number.

The machine mode *m* of `match_insn` works like that of `match_operand`: it is passed as the second argument to the predicate function, and that function is solely responsible for deciding whether the expression to be matched “has” that mode.

`(match_insn2 n predicate)`

Match a complete insn.

The machine mode *m* of `match_insn2` works like that of `match_operand`: it is passed as the second argument to the predicate function, and that function is solely responsible for deciding whether the expression to be matched “has” that mode.

9.5 Output Templates and Operand Substitution

The *output template* is a string which specifies how to output the assembler code for an instruction pattern. Most of the template is a fixed string which is output literally. The character ‘%’ is used to specify where to substitute an operand; it can also be used to identify places where different variants of the assembler require different syntax.

In the simplest case, a ‘%’ followed by a digit *n* says to output operand *n* at that point in the string.

‘%’ followed by a letter and a digit says to output an operand in an alternate fashion. Four letters have standard, built-in meanings described below. The machine description macro `PRINT_OPERAND` can define additional letters with nonstandard meanings.

`%cdigit` can be used to substitute an operand that is a constant value without the syntax that normally indicates an immediate operand.

`%ndigit` is like `%cdigit` except that the value of the constant is negated before printing.

`%adigit` can be used to substitute an operand as if it were a memory reference, with the actual operand treated as the address. This may be useful when outputting a “load address” instruction, because often the assembler syntax for such an instruction requires you to write the operand as if it were a memory reference.

`%ldigit` is used to substitute a `label_ref` into a jump instruction.

`%=` outputs a number which is unique to each instruction in the entire compilation. This is useful for making local labels to be referred to more than once in a single template that generates multiple assembler instructions.

`%` followed by a punctuation character specifies a substitution that does not use an operand. Only one case is standard: `%%` outputs a `%` into the assembler code. Other nonstandard cases can be defined in the `PRINT_OPERAND` macro. You must also define which punctuation characters are valid with the `PRINT_OPERAND_PUNCT_VALID_P` macro.

The template may generate multiple assembler instructions. Write the text for the instructions, with `\;` between them.

When the RTL contains two operands which are required by constraint to match each other, the output template must refer only to the lower-numbered operand. Matching operands are not always identical, and the rest of the compiler arranges to put the proper RTL expression for printing into the lower-numbered operand.

One use of nonstandard letters or punctuation following `%` is to distinguish between different assembler languages for the same machine; for example, Motorola syntax versus MIT syntax for the 68000. Motorola syntax requires periods in most opcode names, while MIT syntax does not. For example, the opcode `move1` in MIT syntax is `move.1` in Motorola syntax. The same file of patterns is used for both kinds of output syntax, but the character sequence `%.` is used in each place where Motorola syntax wants a period. The `PRINT_OPERAND` macro for Motorola syntax defines the sequence to output a period; the macro for MIT syntax defines it to do nothing.

As a special case, a template consisting of the single character `#` instructs the compiler to first split the `insn`, and then output the resulting instructions separately. This helps eliminate redundancy in the output templates. If you have a `define_insn` that needs to emit multiple assembler instructions, and there is an matching `define_split` already defined, then you can simply use `#` as the output template instead of writing an output template that emits the multiple assembler instructions.

If the macro `ASSEMBLER_DIALECT` is defined, you can use construct of the form `{option0|option1|option2}` in the templates. These describe multiple variants of assembler language syntax. See Section 10.20.7 [Instruction Output], page 296.

9.6 C Statements for Assembler Output

Often a single fixed template string cannot produce correct and efficient assembler code for all the cases that are recognized by a single instruction pattern. For example, the opcodes

may depend on the kinds of operands; or some unfortunate combinations of operands may require extra machine instructions.

If the output control string starts with a '@', then it is actually a series of templates, each on a separate line. (Blank lines and leading spaces and tabs are ignored.) The templates correspond to the pattern's constraint alternatives (see Section 9.7.2 [Multi-Alternative], page 130). For example, if a target machine has a two-address add instruction 'addr' to add into a register and another 'addm' to add a register to memory, you might write this pattern:

```
(define_insn "addsi3"
  [(set (match_operand:SI 0 "general_operand" "=r,m")
        (plus:SI (match_operand:SI 1 "general_operand" "0,0")
                 (match_operand:SI 2 "general_operand" "g,r")))]
  ""
  "@
  addr %2,%0
  addm %2,%0")
```

If the output control string starts with a '*', then it is not an output template but rather a piece of C program that should compute a template. It should execute a `return` statement to return the template-string you want. Most such templates use C string literals, which require doublequote characters to delimit them. To include these doublequote characters in the string, prefix each one with '\ '.

If the output control string is written as a brace block instead of a double-quoted string, it is automatically assumed to be C code. In that case, it is not necessary to put in a leading asterisk, or to escape the doublequotes surrounding C string literals.

The operands may be found in the array `operands`, whose C data type is `rtx []`.

It is very common to select different ways of generating assembler code based on whether an immediate operand is within a certain range. Be careful when doing this, because the result of `INTVAL` is an integer on the host machine. If the host machine has more bits in an `int` than the target machine has in the mode in which the constant will be used, then some of the bits you get from `INTVAL` will be superfluous. For proper results, you must carefully disregard the values of those bits.

It is possible to output an assembler instruction and then go on to output or compute more of them, using the subroutine `output_asm_insn`. This receives two arguments: a template-string and a vector of operands. The vector may be `operands`, or it may be another array of `rtx` that you declare locally and initialize yourself.

When an `insn` pattern has multiple alternatives in its constraints, often the appearance of the assembler code is determined mostly by which alternative was matched. When this is so, the C code can test the variable `which_alternative`, which is the ordinal number of the alternative that was actually satisfied (0 for the first, 1 for the second alternative, etc.).

For example, suppose there are two opcodes for storing zero, 'clrreg' for registers and 'clrmem' for memory locations. Here is how a pattern could use `which_alternative` to choose between them:

```
(define_insn ""
  [(set (match_operand:SI 0 "general_operand" "=r,m")
        (const_int 0))]
  ""
  {
```

```

return (which_alternative == 0
        ? "clrreg %0" : "clrmem %0");
})

```

The example above, where the assembler code to generate was *solely* determined by the alternative, could also have been specified as follows, having the output control string start with a '@':

```

(define_insn ""
  [(set (match_operand:SI 0 "general_operand" "=r,m")
        (const_int 0))]
  ""
  "@
  clrreg %0
  clrmem %0")

```

9.7 Operand Constraints

Each `match_operand` in an instruction pattern can specify a constraint for the type of operands allowed. 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.

9.7.1 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:

whitespace

Whitespace characters are ignored and can be inserted at any position except the first. This enables each alternative for different operands to be visually aligned in the machine description even if they have different number of constraints and modifiers.

'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.

Note that 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.
- ‘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` or `const_vector`) is allowed.
- ‘G’, ‘H’ ‘G’ and ‘H’ may be 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. We arrange for this to happen by defining the letter ‘K’ to mean “any integer outside the range -128 to 127 ”, and then specifying ‘Ks’ 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, even if it does not satisfy `general_operand`. This is normally used in the constraint of a `match_scratch` when certain alternatives will not actually require a scratch register.

‘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 number is allowed to be more than a single digit. If multiple digits are encountered consecutively, they are interpreted as a single decimal integer. There is scant chance for ambiguity, since to-date it has never been desirable that ‘10’ be interpreted as matching either operand 1 *or* operand 0. Should this be desired, one can use multiple alternatives instead.

This is called a *matching constraint* and what it really means is that the assembler has only a single operand that fills two roles considered separate in the RTL insn. For example, an add insn has two input operands and one output operand in the RTL, 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.

For operands to match in a particular case usually means that they are identical-looking RTL expressions. But in a few special cases specific kinds of dissimilarity are allowed. For example, `*x` as an input operand will match `*x++` as an output operand. For proper results in such cases, the output template should always use the output-operand’s number when printing the operand.

‘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.

other-letters

Other letters can be defined in machine-dependent fashion to stand for particular classes of registers or other arbitrary operand types. ‘d’, ‘a’ and ‘f’ are defined on the 68000/68020 to stand for data, address and floating point registers.

The machine description macro `REG_CLASS_FROM_LETTER` has first cut at the otherwise unused letters. If it evaluates to `NO_REGS`, then `EXTRA_CONSTRAINT` is evaluated.

A typical use for `EXTRA_CONSTRAINT` would be to distinguish certain types of memory references that affect other insn operands.

In order to have valid assembler code, each operand must satisfy its constraint. But a failure to do so does not prevent the pattern from applying to an insn. Instead, it directs the compiler to modify the code so that the constraint will be satisfied. Usually this is done by copying an operand into a register.

Contrast, therefore, the two instruction patterns that follow:

```
(define_insn ""
  [(set (match_operand:SI 0 "general_operand" "=r")
        (plus:SI (match_dup 0)
                  (match_operand:SI 1 "general_operand" "r")))]
  ""
  "...")
```

which has two operands, one of which must appear in two places, and

```
(define_insn ""
  [(set (match_operand:SI 0 "general_operand" "=r")
        (plus:SI (match_operand:SI 1 "general_operand" "0")
                  (match_operand:SI 2 "general_operand" "r")))]
  ""
  "...")
```

which has three operands, two of which are required by a constraint to be identical. If we are considering an `insn` of the form

```
(insn n prev next
  (set (reg:SI 3)
        (plus:SI (reg:SI 6) (reg:SI 109)))
  ...)
```

the first pattern would not apply at all, because this `insn` does not contain two identical subexpressions in the right place. The pattern would say, “That does not look like an add instruction; try other patterns.” The second pattern would say, “Yes, that’s an add instruction, but there is something wrong with it.” It would direct the reload pass of the compiler to generate additional `insns` to make the constraint true. The results might look like this:

```
(insn n2 prev n
  (set (reg:SI 3) (reg:SI 6))
  ...)

(insn n n2 next
  (set (reg:SI 3)
        (plus:SI (reg:SI 3) (reg:SI 109)))
  ...)
```

It is up to you to make sure that each operand, in each pattern, has constraints that can handle any RTL expression that could be present for that operand. (When multiple alternatives are in use, each pattern must, for each possible combination of operand expressions, have at least one alternative which can handle that combination of operands.) The constraints don’t need to *allow* any possible operand—when this is the case, they do not constrain—but they must at least point the way to reloading any possible operand so that it will fit.

- If the constraint accepts whatever operands the predicate permits, there is no problem: reloading is never necessary for this operand.

For example, an operand whose constraints permit everything except registers is safe provided its predicate rejects registers.

An operand whose predicate accepts only constant values is safe provided its constraints include the letter ‘i’. If any possible constant value is accepted, then nothing less than ‘i’ will do; if the predicate is more selective, then the constraints may also be more selective.

- Any operand expression can be reloaded by copying it into a register. So if an operand's constraints allow some kind of register, it is certain to be safe. It need not permit all classes of registers; the compiler knows how to copy a register into another register of the proper class in order to make an instruction valid.
- A nonoffsettable memory reference can be reloaded by copying the address into a register. So if the constraint uses the letter 'o', all memory references are taken care of.
- A constant operand can be reloaded by allocating space in memory to hold it as preinitialized data. Then the memory reference can be used in place of the constant. So if the constraint uses the letters 'o' or 'm', constant operands are not a problem.
- If the constraint permits a constant and a pseudo register used in an insn was not allocated to a hard register and is equivalent to a constant, the register will be replaced with the constant. If the predicate does not permit a constant and the insn is re-recognized for some reason, the compiler will crash. Thus the predicate must always recognize any objects allowed by the constraint.

If the operand's predicate can recognize registers, but the constraint does not permit them, it can make the compiler crash. When this operand happens to be a register, the reload pass will be stymied, because it does not know how to copy a register temporarily into memory.

If the predicate accepts a unary operator, the constraint applies to the operand. For example, the MIPS processor at ISA level 3 supports an instruction which adds two registers in `SI`mode to produce a `DI`mode result, but only if the registers are correctly sign extended. This predicate for the input operands accepts a `sign_extend` of an `SI`mode register. Write the constraint to indicate the type of register that is required for the operand of the `sign_extend`.

9.7.2 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. Here is how it is done for fullword logical-or on the 68000:

```
(define_insn "iorsi3"
  [(set (match_operand:SI 0 "general_operand" "=m,d")
        (ior:SI (match_operand:SI 1 "general_operand" "%0,0")
                 (match_operand:SI 2 "general_operand" "dKs,dmKs")))]
  ...)
```

The first alternative has 'm' (memory) for operand 0, '0' for operand 1 (meaning it must match operand 0), and 'dKs' for operand 2. The second alternative has 'd' (data register) for operand 0, '0' for operand 1, and 'dmKs' for operand 2. The '=' and '%' in the constraints

apply to all the alternatives; their meaning is explained in the next section (see Section 9.7.3 [Class Preferences], page 131).

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.

When an `insn` pattern has multiple alternatives in its constraints, often the appearance of the assembler code is determined mostly by which alternative was matched. When this is so, the C code for writing the assembler code can use the variable `which_alternative`, which is the ordinal number of the alternative that was actually satisfied (0 for the first, 1 for the second alternative, etc.). See Section 9.6 [Output Statement], page 124.

9.7.3 Register Class Preferences

The operand constraints have another function: they enable the compiler to decide which kind of hardware register a pseudo register is best allocated to. The compiler examines the constraints that apply to the `insns` that use the pseudo register, looking for the machine-dependent letters such as ‘d’ and ‘a’ that specify classes of registers. The pseudo register is put in whichever class gets the most “votes”. The constraint letters ‘g’ and ‘r’ also vote: they vote in favor of a general register. The machine description says which registers are considered general.

Of course, on some machines all registers are equivalent, and no register classes are defined. Then none of this complexity is relevant.

9.7.4 Constraint Modifier Characters

Here 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.
- If you specify ‘=’ or ‘+’ in a constraint, you put it in the first character of the constraint string.

‘&’ 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.

An input operand can be tied to an earlyclobber operand if its only use as an input occurs before the early result is written. Adding alternatives of this form often allows GCC to produce better code when only some of the inputs can be affected by the earlyclobber. See, for example, the ‘*mulsi3*’ insn of the ARM.

‘&’ 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. This is often used in patterns for addition instructions that really have only two operands: the result must go in one of the arguments. Here for example, is how the 68000 halfword-add instruction is defined:

```
(define_insn "addhi3"
  [(set (match_operand:HI 0 "general_operand" "=m,r")
        (plus:HI (match_operand:HI 1 "general_operand" "%0,0")
                  (match_operand:HI 2 "general_operand" "di,g")))]
  ...)
```

GCC can only handle one commutative pair in an asm; if you use more, the compiler may fail.

‘#’ 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.

‘*’ Says that the following character should be ignored when choosing register preferences. ‘*’ has no effect on the meaning of the constraint as a constraint, and no effect on reloading.

Here is an example: the 68000 has an instruction to sign-extend a halfword in a data register, and can also sign-extend a value by copying it into an address register. While either kind of register is acceptable, the constraints on an address-register destination are less strict, so it is best if register allocation makes an address register its goal. Therefore, ‘*’ is used so that the ‘d’ constraint letter (for data register) is ignored when computing register preferences.

```
(define_insn "extendhi2"
  [(set (match_operand:SI 0 "general_operand" ".*d,a")
        (sign_extend:SI
          (match_operand:HI 1 "general_operand" "0,g")))]
  ...)
```

9.7.5 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 Section 9.7.1 [Simple Constraints], page 126), and ‘I’, usually the letter indicating the most common immediate-constant format.

For each machine architecture, the ‘`config/machine/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 these 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, here is a summary of the machine-dependent constraints available on some particular machines.

ARM family—‘`arm.h`’

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 <code>asm</code> statements)
R	An item in the constant pool

S	A symbol in the text segment of the current file
<i>AVR family</i> —‘avr.h’	
l	Registers from r0 to r15
a	Registers from r16 to r23
d	Registers from r16 to r31
w	Registers from r24 to r31. These registers can be used in ‘adiw’ command
e	Pointer register (r26–r31)
b	Base pointer register (r28–r31)
q	Stack pointer register (SPH:SPL)
t	Temporary register r0
x	Register pair X (r27:r26)
y	Register pair Y (r29:r28)
z	Register pair Z (r31:r30)
I	Constant greater than -1 , less than 64
J	Constant greater than -64 , less than 1
K	Constant integer 2
L	Constant integer 0
M	Constant that fits in 8 bits
N	Constant integer -1
O	Constant integer 8, 16, or 24
P	Constant integer 1
G	A floating point constant 0.0
<i>IBM RS6000</i> —‘rs6000.h’	
b	Address base register
f	Floating point register
h	‘MQ’, ‘CTR’, or ‘LINK’ register
q	‘MQ’ register
c	‘CTR’ register
l	‘LINK’ register
x	‘CR’ register (condition register) number 0
y	‘CR’ register (condition register)
z	‘FPMEM’ stack memory for FPR-GPR transfers

I	Signed 16-bit constant
J	Unsigned 16-bit constant shifted left 16 bits (use ‘L’ instead for SImode constants)
K	Unsigned 16-bit constant
L	Signed 16-bit constant shifted left 16 bits
M	Constant larger than 31
N	Exact power of 2
O	Zero
P	Constant whose negation is a signed 16-bit constant
G	Floating point constant that can be loaded into a register with one instruction per word
Q	Memory operand that is an offset from a register (‘m’ is preferable for <code>asm</code> statements)
R	AIX TOC entry
S	Constant suitable as a 64-bit mask operand
T	Constant suitable as a 32-bit mask operand
U	System V Release 4 small data area reference

Intel 386—‘i386.h’

q	‘a’, b, c, or d register for the i386. For x86-64 it is equivalent to ‘r’ class. (for 8-bit instructions that do not use upper halves)
Q	‘a’, b, c, or d register. (for 8-bit instructions, that do use upper halves)
R	Legacy register—equivalent to <code>r</code> class in i386 mode. (for non-8-bit registers used together with 8-bit upper halves in a single instruction)
A	Specifies the ‘a’ or ‘d’ registers. This is primarily useful for 64-bit integer values (when in 32-bit mode) intended to be returned with the ‘d’ register holding the most significant bits and the ‘a’ register holding the least significant bits.
f	Floating point register
t	First (top of stack) floating point register
u	Second floating point register
a	‘a’ register
b	‘b’ register
c	‘c’ register
C	Specifies constant that can be easily constructed in SSE register without loading it from memory.

d	‘d’ register
D	‘di’ register
S	‘si’ register
x	‘xmm’ SSE register
y	MMX register
I	Constant in range 0 to 31 (for 32-bit shifts)
J	Constant in range 0 to 63 (for 64-bit shifts)
K	‘0xff’
L	‘0xffff’
M	0, 1, 2, or 3 (shifts for <code>lea</code> instruction)
N	Constant in range 0 to 255 (for <code>out</code> instruction)
Z	Constant in range 0 to 0xffffffff or symbolic reference known to fit specified range. (for using immediates in zero extending 32-bit to 64-bit x86-64 instructions)
e	Constant in range -2147483648 to 2147483647 or symbolic reference known to fit specified range. (for using immediates in 64-bit x86-64 instructions)
G	Standard 80387 floating point constant

Intel 960—‘i960.h’

f	Floating point register (fp0 to fp3)
l	Local register (r0 to r15)
b	Global register (g0 to g15)
d	Any local or global register
I	Integers from 0 to 31
J	0
K	Integers from -31 to 0
G	Floating point 0
H	Floating point 1

Intel IA-64—‘ia64.h’

a	General register r0 to r3 for <code>addl</code> instruction
b	Branch register
c	Predicate register (‘c’ as in “conditional”)
d	Application register residing in M-unit
e	Application register residing in I-unit

f	Floating-point register
m	Memory operand. Remember that 'm' allows postincrement and postdecrement which require printing with '%Pn' on IA-64. Use 'S' to disallow postincrement and postdecrement.
G	Floating-point constant 0.0 or 1.0
I	14-bit signed integer constant
J	22-bit signed integer constant
K	8-bit signed integer constant for logical instructions
L	8-bit adjusted signed integer constant for compare pseudo-ops
M	6-bit unsigned integer constant for shift counts
N	9-bit signed integer constant for load and store postincrements
O	The constant zero
P	0 or -1 for <code>dep</code> instruction
Q	Non-volatile memory for floating-point loads and stores
R	Integer constant in the range 1 to 4 for <code>shladd</code> instruction
S	Memory operand except postincrement and postdecrement

FRV—'frv.h'

a	Register in the class <code>ACC_REGS</code> (<code>acc0</code> to <code>acc7</code>).
b	Register in the class <code>EVEN_ACC_REGS</code> (<code>acc0</code> to <code>acc7</code>).
c	Register in the class <code>CC_REGS</code> (<code>fcc0</code> to <code>fcc3</code> and <code>icc0</code> to <code>icc3</code>).
d	Register in the class <code>GPR_REGS</code> (<code>gr0</code> to <code>gr63</code>).
e	Register in the class <code>EVEN_REGS</code> (<code>gr0</code> to <code>gr63</code>). Odd registers are excluded not in the class but through the use of a machine mode larger than 4 bytes.
f	Register in the class <code>FPR_REGS</code> (<code>fr0</code> to <code>fr63</code>).
h	Register in the class <code>FEVEN_REGS</code> (<code>fr0</code> to <code>fr63</code>). Odd registers are excluded not in the class but through the use of a machine mode larger than 4 bytes.
l	Register in the class <code>LR_REG</code> (the <code>lr</code> register).
q	Register in the class <code>QUAD_REGS</code> (<code>gr2</code> to <code>gr63</code>). Register numbers not divisible by 4 are excluded not in the class but through the use of a machine mode larger than 8 bytes.
t	Register in the class <code>ICC_REGS</code> (<code>icc0</code> to <code>icc3</code>).
u	Register in the class <code>FCC_REGS</code> (<code>fcc0</code> to <code>fcc3</code>).
v	Register in the class <code>ICR_REGS</code> (<code>cc4</code> to <code>cc7</code>).

w	Register in the class <code>FCR_REGS</code> (<code>cc0</code> to <code>cc3</code>).
x	Register in the class <code>QUAD_FPR_REGS</code> (<code>fr0</code> to <code>fr63</code>). Register numbers not divisible by 4 are excluded not in the class but through the use of a machine mode larger than 8 bytes.
z	Register in the class <code>SPR_REGS</code> (<code>1cr</code> and <code>1r</code>).
A	Register in the class <code>QUAD_ACC_REGS</code> (<code>acc0</code> to <code>acc7</code>).
B	Register in the class <code>ACCG_REGS</code> (<code>accg0</code> to <code>accg7</code>).
C	Register in the class <code>CR_REGS</code> (<code>cc0</code> to <code>cc7</code>).
G	Floating point constant zero
I	6-bit signed integer constant
J	10-bit signed integer constant
L	16-bit signed integer constant
M	16-bit unsigned integer constant
N	12-bit signed integer constant that is negative—i.e. in the range of <code>-2048</code> to <code>-1</code>
O	Constant zero
P	12-bit signed integer constant that is greater than zero—i.e. in the range of <code>1</code> to <code>2047</code> .

IP2K—‘ip2k.h’

a	‘DP’ or ‘IP’ registers (general address)
f	‘IP’ register
j	‘IPL’ register
k	‘IPH’ register
b	‘DP’ register
y	‘DPH’ register
z	‘DPL’ register
q	‘SP’ register
c	‘DP’ or ‘SP’ registers (offsettable address)
d	Non-pointer registers (not ‘SP’, ‘DP’, ‘IP’)
u	Non-SP registers (everything except ‘SP’)
R	Indirect thru ‘IP’ - Avoid this except for <code>QImode</code> , since we can’t access extra bytes
S	Indirect thru ‘SP’ or ‘DP’ with short displacement (<code>0..127</code>)
T	Data-section immediate value

I	Integers from -255 to -1
J	Integers from 0 to 7—valid bit number in a register
K	Integers from 0 to 127—valid displacement for addressing mode
L	Integers from 1 to 127
M	Integer -1
N	Integer 1
O	Zero
P	Integers from 0 to 255

MIPS—`'mips.h'`

d	General-purpose integer register
f	Floating-point register (if available)
h	'Hi' register
l	'Lo' register
x	'Hi' or 'Lo' register
y	General-purpose integer register
z	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 ('m' is preferable for <code>asm</code> statements)
R	Memory reference that can be loaded with one instruction ('m' is preferable for <code>asm</code> statements)
S	Memory reference in external OSF/rose PIC format ('m' is preferable for <code>asm</code> statements)

Motorola 680x0—`'m68k.h'`

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

Motorola 68HC11 & 68HC12 families—'m68hc11.h'

a	Register 'a'
b	Register 'b'
d	Register 'd'
q	An 8-bit register
t	Temporary soft register <code>..tmp</code>
u	A soft register <code>..d1</code> to <code>..d31</code>
w	Stack pointer register
x	Register 'x'
y	Register 'y'
z	Pseudo register 'z' (replaced by 'x' or 'y' at the end)
A	An address register: x, y or z
B	An address register: x or y
D	Register pair (x:d) to form a 32-bit value
L	Constants in the range -65536 to 65535
M	Constants whose 16-bit low part is zero
N	Constant integer 1 or -1
O	Constant integer 16
P	Constants in the range -8 to 2

SPARC—'sparc.h'

f	Floating-point register on the SPARC-V8 architecture and lower floating-point register on the SPARC-V9 architecture.
---	--

e	Floating-point register. It is equivalent to ‘f’ on the SPARC-V8 architecture and contains both lower and upper floating-point registers on the SPARC-V9 architecture.
c	Floating-point condition code register.
d	Lower floating-point register. It is only valid on the SPARC-V9 architecture when the Visual Instruction Set is available.
b	Floating-point register. It is only valid on the SPARC-V9 architecture when the Visual Instruction Set is available.
h	64-bit global or out register for the SPARC-V8+ architecture.
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)
L	A constant in the range supported by <code>movcc</code> instructions
M	A constant in the range supported by <code>movrcc</code> instructions
N	Same as ‘K’, except that it verifies that bits that are not in the lower 32-bit range are all zero. Must be used instead of ‘K’ for modes wider than <code>SImode</code>
O	The constant 4096
G	Floating-point zero
H	Signed 13-bit constant, sign-extended to 32 or 64 bits
Q	Floating-point constant whose integral representation can be moved into an integer register using a single <code>sethi</code> instruction
R	Floating-point constant whose integral representation can be moved into an integer register using a single <code>mov</code> instruction
S	Floating-point constant whose integral representation can be moved into an integer register using a <code>high/lo_sum</code> instruction sequence
T	Memory address aligned to an 8-byte boundary
U	Even register
W	Memory address for ‘e’ constraint registers.

TMS320C3x/C4x—‘c4x.h’

a	Auxiliary (address) register (ar0-ar7)
b	Stack pointer register (sp)
c	Standard (32-bit) precision integer register
f	Extended (40-bit) precision register (r0-r11)
k	Block count register (bk)

q	Extended (40-bit) precision low register (r0-r7)
t	Extended (40-bit) precision register (r0-r1)
u	Extended (40-bit) precision register (r2-r3)
v	Repeat count register (rc)
x	Index register (ir0-ir1)
y	Status (condition code) register (st)
z	Data page register (dp)
G	Floating-point zero
H	Immediate 16-bit floating-point constant
I	Signed 16-bit constant
J	Signed 8-bit constant
K	Signed 5-bit constant
L	Unsigned 16-bit constant
M	Unsigned 8-bit constant
N	Ones complement of unsigned 16-bit constant
O	High 16-bit constant (32-bit constant with 16 LSBs zero)
Q	Indirect memory reference with signed 8-bit or index register displacement
R	Indirect memory reference with unsigned 5-bit displacement
S	Indirect memory reference with 1 bit or index register displacement
T	Direct memory reference
U	Symbolic address

S/390 and zSeries—‘s390.h’

a	Address register (general purpose register except r0)
d	Data register (arbitrary general purpose register)
f	Floating-point register
I	Unsigned 8-bit constant (0–255)
J	Unsigned 12-bit constant (0–4095)
K	Signed 16-bit constant (–32768–32767)
L	Unsigned 16-bit constant (0–65535)
Q	Memory reference without index register
S	Symbolic constant suitable for use with the <code>larl</code> instruction

Xstormy16—‘stormy16.h’

a	Register r0.
---	--------------

b	Register r1.
c	Register r2.
d	Register r8.
e	Registers r0 through r7.
t	Registers r0 and r1.
y	The carry register.
z	Registers r8 and r9.
I	A constant between 0 and 3 inclusive.
J	A constant that has exactly one bit set.
K	A constant that has exactly one bit clear.
L	A constant between 0 and 255 inclusive.
M	A constant between -255 and 0 inclusive.
N	A constant between -3 and 0 inclusive.
O	A constant between 1 and 4 inclusive.
P	A constant between -4 and -1 inclusive.
Q	A memory reference that is a stack push.
R	A memory reference that is a stack pop.
S	A memory reference that refers to a constant address of known value.
T	The register indicated by Rx (not implemented yet).
U	A constant that is not between 2 and 15 inclusive.

Xtensa—`xtensa.h`'

a	General-purpose 32-bit register
b	One-bit boolean register
A	MAC16 40-bit accumulator register
I	Signed 12-bit integer constant, for use in MOVI instructions
J	Signed 8-bit integer constant, for use in ADDI instructions
K	Integer constant valid for BccI instructions
L	Unsigned constant valid for BccUI instructions

9.8 Standard Pattern Names For Generation

Here is a table of the instruction names that are meaningful in the RTL generation pass of the compiler. Giving one of these names to an instruction pattern tells the RTL generation pass that it can use the pattern to accomplish a certain task.

`'movm'` Here *m* stands for a two-letter machine mode name, in lower case. This instruction pattern moves data with that machine mode from operand 1 to operand 0. For example, `'movsi'` moves full-word data.

If operand 0 is a `subreg` with mode *m* of a register whose own mode is wider than *m*, the effect of this instruction is to store the specified value in the part of the register that corresponds to mode *m*. Bits outside of *m*, but which are within the same target word as the `subreg` are undefined. Bits which are outside the target word are left unchanged.

This class of patterns is special in several ways. First of all, each of these names up to and including full word size *must* be defined, because there is no other way to copy a datum from one place to another. If there are patterns accepting operands in larger modes, `'movm'` must be defined for integer modes of those sizes.

Second, these patterns are not used solely in the RTL generation pass. Even the reload pass can generate move insns to copy values from stack slots into temporary registers. When it does so, one of the operands is a hard register and the other is an operand that can need to be reloaded into a register.

Therefore, when given such a pair of operands, the pattern must generate RTL which needs no reloading and needs no temporary registers—no registers other than the operands. For example, if you support the pattern with a `define_expand`, then in such a case the `define_expand` mustn't call `force_reg` or any other such function which might generate new pseudo registers.

This requirement exists even for subword modes on a RISC machine where fetching those modes from memory normally requires several insns and some temporary registers.

During reload a memory reference with an invalid address may be passed as an operand. Such an address will be replaced with a valid address later in the reload pass. In this case, nothing may be done with the address except to use it as it stands. If it is copied, it will not be replaced with a valid address. No attempt should be made to make such an address into a valid address and no routine (such as `change_address`) that will do so may be called. Note that `general_operand` will fail when applied to such an address.

The global variable `reload_in_progress` (which must be explicitly declared if required) can be used to determine whether such special handling is required.

The variety of operands that have reloads depends on the rest of the machine description, but typically on a RISC machine these can only be pseudo registers that did not get hard registers, while on other machines explicit memory references will get optional reloads.

If a scratch register is required to move an object to or from memory, it can be allocated using `gen_reg_rtx` prior to life analysis.

If there are cases which need scratch registers during or after reload, you must define `SECONDARY_INPUT_RELOAD_CLASS` and/or `SECONDARY_OUTPUT_RELOAD_CLASS` to detect them, and provide patterns `'reload_inm'` or `'reload_outm'` to handle them. See Section 10.9 [Register Classes], page 223.

The global variable `no_new_pseudos` can be used to determine if it is unsafe to create new pseudo registers. If this variable is nonzero, then it is unsafe to call `gen_reg_rtx` to allocate a new pseudo.

The constraints on a `'movm'` must permit moving any hard register to any other hard register provided that `HARD_REGNO_MODE_OK` permits mode *m* in both registers and `REGISTER_MOVE_COST` applied to their classes returns a value of 2.

It is obligatory to support floating point `'movm'` instructions into and out of any registers that can hold fixed point values, because unions and structures (which have modes `SImode` or `DImode`) can be in those registers and they may have floating point members.

There may also be a need to support fixed point `'movm'` instructions in and out of floating point registers. Unfortunately, I have forgotten why this was so, and I don't know whether it is still true. If `HARD_REGNO_MODE_OK` rejects fixed point values in floating point registers, then the constraints of the fixed point `'movm'` instructions must be designed to avoid ever trying to reload into a floating point register.

`'reload_inm'`

`'reload_outm'`

Like `'movm'`, but used when a scratch register is required to move between operand 0 and operand 1. Operand 2 describes the scratch register. See the discussion of the `SECONDARY_RELOAD_CLASS` macro in see Section 10.9 [Register Classes], page 223.

There are special restrictions on the form of the `match_operands` used in these patterns. First, only the predicate for the reload operand is examined, i.e., `reload_in` examines operand 1, but not the predicates for operand 0 or 2. Second, there may be only one alternative in the constraints. Third, only a single register class letter may be used for the constraint; subsequent constraint letters are ignored. As a special exception, an empty constraint string matches the `ALL_REGS` register class. This may relieve ports of the burden of defining an `ALL_REGS` constraint letter just for these patterns.

`'movstrictm'`

Like `'movm'` except that if operand 0 is a `subreg` with mode *m* of a register whose natural mode is wider, the `'movstrictm'` instruction is guaranteed not to alter any of the register except the part which belongs to mode *m*.

`'load_multiple'`

Load several consecutive memory locations into consecutive registers. Operand 0 is the first of the consecutive registers, operand 1 is the first memory location, and operand 2 is a constant: the number of consecutive registers.

Define this only if the target machine really has such an instruction; do not define this if the most efficient way of loading consecutive registers from memory is to do them one at a time.

On some machines, there are restrictions as to which consecutive registers can be stored into memory, such as particular starting or ending register numbers or only a range of valid counts. For those machines, use a `define_expand` (see Section 9.14 [Expander Definitions], page 164) and make the pattern fail if the restrictions are not met.

Write the generated insn as a `parallel` with elements being a `set` of one register from the appropriate memory location (you may also need `use` or `clobber` elements). Use a `match_parallel` (see Section 9.4 [RTL Template], page 119) to recognize the insn. See ‘rs6000.md’ for examples of the use of this insn pattern.

‘store_multiple’

Similar to ‘load_multiple’, but store several consecutive registers into consecutive memory locations. Operand 0 is the first of the consecutive memory locations, operand 1 is the first register, and operand 2 is a constant: the number of consecutive registers.

‘pushm’

Output a push instruction. Operand 0 is value to push. Used only when `PUSH_ROUNDING` is defined. For historical reason, this pattern may be missing and in such case an `mov` expander is used instead, with a `MEM` expression forming the push operation. The `mov` expander method is deprecated.

‘addm3’

Add operand 2 and operand 1, storing the result in operand 0. All operands must have mode *m*. This can be used even on two-address machines, by means of constraints requiring operands 1 and 0 to be the same location.

‘subm3’, ‘mulm3’

‘divm3’, ‘udivm3’, ‘modm3’, ‘umodm3’

‘sminm3’, ‘smaxm3’, ‘uminm3’, ‘umaxm3’

‘andm3’, ‘iorm3’, ‘xorm3’

Similar, for other arithmetic operations.

‘minm3’, ‘maxm3’

Floating point min and max operations. If both operands are zeros, or if either operand is NaN, then it is unspecified which of the two operands is returned as the result.

‘mulhisi3’

Multiply operands 1 and 2, which have mode `HImode`, and store a `SImode` product in operand 0.

‘mulqihi3’, ‘mulsidi3’

Similar widening-multiplication instructions of other widths.

‘umulqihi3’, ‘umulhisi3’, ‘umulsidi3’

Similar widening-multiplication instructions that do unsigned multiplication.

‘smulm3_highpart’

Perform a signed multiplication of operands 1 and 2, which have mode *m*, and store the most significant half of the product in operand 0. The least significant half of the product is discarded.

`'umulm3_highpart'`

Similar, but the multiplication is unsigned.

`'divmodm4'`

Signed division that produces both a quotient and a remainder. Operand 1 is divided by operand 2 to produce a quotient stored in operand 0 and a remainder stored in operand 3.

For machines with an instruction that produces both a quotient and a remainder, provide a pattern for `'divmodm4'` but do not provide patterns for `'divm3'` and `'modm3'`. This allows optimization in the relatively common case when both the quotient and remainder are computed.

If an instruction that just produces a quotient or just a remainder exists and is more efficient than the instruction that produces both, write the output routine of `'divmodm4'` to call `find_reg_note` and look for a `REG_UNUSED` note on the quotient or remainder and generate the appropriate instruction.

`'udivmodm4'`

Similar, but does unsigned division.

`'ashlm3'`

Arithmetic-shift operand 1 left by a number of bits specified by operand 2, and store the result in operand 0. Here m is the mode of operand 0 and operand 1; operand 2's mode is specified by the instruction pattern, and the compiler will convert the operand to that mode before generating the instruction.

`'ashrm3'`, `'lshrm3'`, `'rotlm3'`, `'rotrm3'`

Other shift and rotate instructions, analogous to the `ashlm3` instructions.

`'negm2'`

Negate operand 1 and store the result in operand 0.

`'absm2'`

Store the absolute value of operand 1 into operand 0.

`'sqrtm2'`

Store the square root of operand 1 into operand 0.

The `sqrt` built-in function of C always uses the mode which corresponds to the C data type `double` and the `sqrtf` built-in function uses the mode which corresponds to the C data type `float`.

`'cosm2'`

Store the cosine of operand 1 into operand 0.

The `cos` built-in function of C always uses the mode which corresponds to the C data type `double` and the `cosf` built-in function uses the mode which corresponds to the C data type `float`.

`'sinm2'`

Store the sine of operand 1 into operand 0.

The `sin` built-in function of C always uses the mode which corresponds to the C data type `double` and the `sinf` built-in function uses the mode which corresponds to the C data type `float`.

`'expm2'`

Store the exponential of operand 1 into operand 0.

The `exp` built-in function of C always uses the mode which corresponds to the C data type `double` and the `expf` built-in function uses the mode which corresponds to the C data type `float`.

- ‘logm2’** Store the natural logarithm of operand 1 into operand 0.
The `log` built-in function of C always uses the mode which corresponds to the C data type `double` and the `logf` built-in function uses the mode which corresponds to the C data type `float`.
- ‘ffsm2’** Store into operand 0 one plus the index of the least significant 1-bit of operand 1. If operand 1 is zero, store zero. *m* is the mode of operand 0; operand 1’s mode is specified by the instruction pattern, and the compiler will convert the operand to that mode before generating the instruction.
The `ffs` built-in function of C always uses the mode which corresponds to the C data type `int`.
- ‘one_cmplm2’** Store the bitwise-complement of operand 1 into operand 0.
- ‘cmpm’** Compare operand 0 and operand 1, and set the condition codes. The RTL pattern should look like this:

```
(set (cc0) (compare (match_operand:m 0 ...)
                   (match_operand:m 1 ...)))
```
- ‘tstm’** Compare operand 0 against zero, and set the condition codes. The RTL pattern should look like this:

```
(set (cc0) (match_operand:m 0 ...))
```


‘tstm’ patterns should not be defined for machines that do not use `(cc0)`. Doing so would confuse the optimizer since it would no longer be clear which `set` operations were comparisons. The **‘cmpm’** patterns should be used instead.
- ‘movstrm’** Block move instruction. The addresses of the destination and source strings are the first two operands, and both are in mode `Pmode`.
The number of bytes to move is the third operand, in mode *m*. Usually, you specify `word_mode` for *m*. However, if you can generate better code knowing the range of valid lengths is smaller than those representable in a full word, you should provide a pattern with a mode corresponding to the range of values you can handle efficiently (e.g., `QImode` for values in the range 0–127; note we avoid numbers that appear negative) and also a pattern with `word_mode`.
The fourth operand is the known shared alignment of the source and destination, in the form of a `const_int` rtx. Thus, if the compiler knows that both source and destination are word-aligned, it may provide the value 4 for this operand.
Descriptions of multiple `movstrm` patterns can only be beneficial if the patterns for smaller modes have fewer restrictions on their first, second and fourth operands. Note that the mode *m* in `movstrm` does not impose any restriction on the mode of individually moved data units in the block.
These patterns need not give special consideration to the possibility that the source and destination strings might overlap.
- ‘clrstrm’** Block clear instruction. The addresses of the destination string is the first operand, in mode `Pmode`. The number of bytes to clear is the second operand, in mode *m*. See **‘movstrm’** for a discussion of the choice of mode.

The third operand is the known alignment of the destination, in the form of a `const_int` rtx. Thus, if the compiler knows that the destination is word-aligned, it may provide the value 4 for this operand.

The use for multiple `clrstrm` is as for `movstrm`.

- '`cmpstrm`' Block compare instruction, with five operands. Operand 0 is the output; it has mode *m*. The remaining four operands are like the operands of '`movstrm`'. The two memory blocks specified are compared byte by byte in lexicographic order. The effect of the instruction is to store a value in operand 0 whose sign indicates the result of the comparison.
- '`strlenm`' Compute the length of a string, with three operands. Operand 0 is the result (of mode *m*), operand 1 is a `mem` referring to the first character of the string, operand 2 is the character to search for (normally zero), and operand 3 is a constant describing the known alignment of the beginning of the string.
- '`floatmn2`' Convert signed integer operand 1 (valid for fixed point mode *m*) to floating point mode *n* and store in operand 0 (which has mode *n*).
- '`floatunsmn2`' Convert unsigned integer operand 1 (valid for fixed point mode *m*) to floating point mode *n* and store in operand 0 (which has mode *n*).
- '`fixmn2`' Convert operand 1 (valid for floating point mode *m*) to fixed point mode *n* as a signed number and store in operand 0 (which has mode *n*). This instruction's result is defined only when the value of operand 1 is an integer.
- '`fixunsmn2`' Convert operand 1 (valid for floating point mode *m*) to fixed point mode *n* as an unsigned number and store in operand 0 (which has mode *n*). This instruction's result is defined only when the value of operand 1 is an integer.
- '`ftruncm2`' Convert operand 1 (valid for floating point mode *m*) to an integer value, still represented in floating point mode *m*, and store it in operand 0 (valid for floating point mode *m*).
- '`fix_truncmn2`' Like '`fixmn2`' but works for any floating point value of mode *m* by converting the value to an integer.
- '`fixuns_truncmn2`' Like '`fixunsmn2`' but works for any floating point value of mode *m* by converting the value to an integer.
- '`truncmn2`' Truncate operand 1 (valid for mode *m*) to mode *n* and store in operand 0 (which has mode *n*). Both modes must be fixed point or both floating point.
- '`extendmn2`' Sign-extend operand 1 (valid for mode *m*) to mode *n* and store in operand 0 (which has mode *n*). Both modes must be fixed point or both floating point.

- `'zero_extendmn2'`
Zero-extend operand 1 (valid for mode *m*) to mode *n* and store in operand 0 (which has mode *n*). Both modes must be fixed point.
- `'extv'`
Extract a bit-field from operand 1 (a register or memory operand), where operand 2 specifies the width in bits and operand 3 the starting bit, and store it in operand 0. Operand 0 must have mode `word_mode`. Operand 1 may have mode `byte_mode` or `word_mode`; often `word_mode` is allowed only for registers. Operands 2 and 3 must be valid for `word_mode`.
The RTL generation pass generates this instruction only with constants for operands 2 and 3.
The bit-field value is sign-extended to a full word integer before it is stored in operand 0.
- `'extzv'`
Like `'extv'` except that the bit-field value is zero-extended.
- `'insv'`
Store operand 3 (which must be valid for `word_mode`) into a bit-field in operand 0, where operand 1 specifies the width in bits and operand 2 the starting bit. Operand 0 may have mode `byte_mode` or `word_mode`; often `word_mode` is allowed only for registers. Operands 1 and 2 must be valid for `word_mode`.
The RTL generation pass generates this instruction only with constants for operands 1 and 2.
- `'movmodecc'`
Conditionally move operand 2 or operand 3 into operand 0 according to the comparison in operand 1. If the comparison is true, operand 2 is moved into operand 0, otherwise operand 3 is moved.
The mode of the operands being compared need not be the same as the operands being moved. Some machines, `sparc64` for example, have instructions that conditionally move an integer value based on the floating point condition codes and vice versa.
If the machine does not have conditional move instructions, do not define these patterns.
- `'scond'`
Store zero or nonzero in the operand according to the condition codes. Value stored is nonzero iff the condition *cond* is true. *cond* is the name of a comparison operation expression code, such as `eq`, `lt` or `leu`.
You specify the mode that the operand must have when you write the `match_operand` expression. The compiler automatically sees which mode you have used and supplies an operand of that mode.
The value stored for a true condition must have 1 as its low bit, or else must be negative. Otherwise the instruction is not suitable and you should omit it from the machine description. You describe to the compiler exactly which value is stored by defining the macro `STORE_FLAG_VALUE` (see Section 10.26 [Misc], page 314). If a description cannot be found that can be used for all the `'scond'` patterns, you should omit those operations from the machine description.
These operations may fail, but should do so only in relatively uncommon cases; if they would fail for common cases involving integer comparisons, it is best to omit these patterns.

If these operations are omitted, the compiler will usually generate code that copies the constant one to the target and branches around an assignment of zero to the target. If this code is more efficient than the potential instructions used for the `'scond'` pattern followed by those required to convert the result into a 1 or a zero in `SI`mode, you should omit the `'scond'` operations from the machine description.

`'bcond'` Conditional branch instruction. Operand 0 is a `label_ref` that refers to the label to jump to. Jump if the condition codes meet condition `cond`.

Some machines do not follow the model assumed here where a comparison instruction is followed by a conditional branch instruction. In that case, the `'cmpm'` (and `'tstm'`) patterns should simply store the operands away and generate all the required insns in a `define_expand` (see Section 9.14 [Expander Definitions], page 164) for the conditional branch operations. All calls to expand `'bcond'` patterns are immediately preceded by calls to expand either a `'cmpm'` pattern or a `'tstm'` pattern.

Machines that use a pseudo register for the condition code value, or where the mode used for the comparison depends on the condition being tested, should also use the above mechanism. See Section 9.11 [Jump Patterns], page 160.

The above discussion also applies to the `'movmodecc'` and `'scond'` patterns.

`'jump'` A jump inside a function; an unconditional branch. Operand 0 is the `label_ref` of the label to jump to. This pattern name is mandatory on all machines.

`'call'` Subroutine call instruction returning no value. Operand 0 is the function to call; operand 1 is the number of bytes of arguments pushed as a `const_int`; operand 2 is the number of registers used as operands.

On most machines, operand 2 is not actually stored into the RTL pattern. It is supplied for the sake of some RISC machines which need to put this information into the assembler code; they can put it in the RTL instead of operand 1.

Operand 0 should be a `mem RTX` whose address is the address of the function. Note, however, that this address can be a `symbol_ref` expression even if it would not be a legitimate memory address on the target machine. If it is also not a valid argument for a call instruction, the pattern for this operation should be a `define_expand` (see Section 9.14 [Expander Definitions], page 164) that places the address into a register and uses that register in the call instruction.

`'call_value'`

Subroutine call instruction returning a value. Operand 0 is the hard register in which the value is returned. There are three more operands, the same as the three operands of the `'call'` instruction (but with numbers increased by one).

Subroutines that return `BLKmode` objects use the `'call'` insn.

`'call_pop'`, `'call_value_pop'`

Similar to `'call'` and `'call_value'`, except used if defined and if `RETURN_POPS_ARGS` is nonzero. They should emit a `parallel` that contains both the function call and a `set` to indicate the adjustment made to the frame pointer.

For machines where `RETURN_POPS_ARGS` can be nonzero, the use of these patterns increases the number of functions for which the frame pointer can be eliminated, if desired.

‘`untyped_call`’

Subroutine call instruction returning a value of any type. Operand 0 is the function to call; operand 1 is a memory location where the result of calling the function is to be stored; operand 2 is a `parallel` expression where each element is a `set` expression that indicates the saving of a function return value into the result block.

This instruction pattern should be defined to support `__builtin_apply` on machines where special instructions are needed to call a subroutine with arbitrary arguments or to save the value returned. This instruction pattern is required on machines that have multiple registers that can hold a return value (i.e. `FUNCTION_VALUE_REGNO_P` is true for more than one register).

‘`return`’

Subroutine return instruction. This instruction pattern name should be defined only if a single instruction can do all the work of returning from a function.

Like the ‘`movm`’ patterns, this pattern is also used after the RTL generation phase. In this case it is to support machines where multiple instructions are usually needed to return from a function, but some class of functions only requires one instruction to implement a return. Normally, the applicable functions are those which do not need to save any registers or allocate stack space.

For such machines, the condition specified in this pattern should only be true when `reload_completed` is nonzero and the function’s epilogue would only be a single instruction. For machines with register windows, the routine `leaf_function_p` may be used to determine if a register window push is required.

Machines that have conditional return instructions should define patterns such as

```
(define_insn ""
  [(set (pc)
        (if_then_else (match_operator
                       0 "comparison_operator"
                       [(cc0) (const_int 0)])
                       (return)
                       (pc)))]
  "condition"
  "...")
```

where *condition* would normally be the same condition specified on the named ‘`return`’ pattern.

‘`untyped_return`’

Untyped subroutine return instruction. This instruction pattern should be defined to support `__builtin_return` on machines where special instructions are needed to return a value of any type.

Operand 0 is a memory location where the result of calling a function with `__builtin_apply` is stored; operand 1 is a `parallel` expression where each element is a `set` expression that indicates the restoring of a function return value from the result block.

`'nop'` No-op instruction. This instruction pattern name should always be defined to output a no-op in assembler code. (`const_int 0`) will do as an RTL pattern.

`'indirect_jump'`
An instruction to jump to an address which is operand zero. This pattern name is mandatory on all machines.

`'casesi'` Instruction to jump through a dispatch table, including bounds checking. This instruction takes five operands:

1. The index to dispatch on, which has mode `SImode`.
2. The lower bound for indices in the table, an integer constant.
3. The total range of indices in the table—the largest index minus the smallest one (both inclusive).
4. A label that precedes the table itself.
5. A label to jump to if the index has a value outside the bounds. (If the machine-description macro `CASE_DROPS_THROUGH` is defined, then an out-of-bounds index drops through to the code following the jump table instead of jumping to this label. In that case, this label is not actually used by the `'casesi'` instruction, but it is always provided as an operand.)

The table is a `addr_vec` or `addr_diff_vec` inside of a `jump_insn`. The number of elements in the table is one plus the difference between the upper bound and the lower bound.

`'tablejump'`
Instruction to jump to a variable address. This is a low-level capability which can be used to implement a dispatch table when there is no `'casesi'` pattern.

This pattern requires two operands: the address or offset, and a label which should immediately precede the jump table. If the macro `CASE_VECTOR_PC_RELATIVE` evaluates to a nonzero value then the first operand is an offset which counts from the address of the table; otherwise, it is an absolute address to jump to. In either case, the first operand has mode `Pmode`.

The `'tablejump'` insn is always the last insn before the jump table it uses. Its assembler code normally has no need to use the second operand, but you should incorporate it in the RTL pattern so that the jump optimizer will not delete the table as unreachable code.

`'decrement_and_branch_until_zero'`
Conditional branch instruction that decrements a register and jumps if the register is nonzero. Operand 0 is the register to decrement and test; operand 1 is the label to jump to if the register is nonzero. See Section 9.12 [Looping Patterns], page 161.

This optional instruction pattern is only used by the combiner, typically for loops reversed by the loop optimizer when strength reduction is enabled.

`'doloop_end'`
Conditional branch instruction that decrements a register and jumps if the register is nonzero. This instruction takes five operands: Operand 0 is the

register to decrement and test; operand 1 is the number of loop iterations as a `const_int` or `const0_rtx` if this cannot be determined until run-time; operand 2 is the actual or estimated maximum number of iterations as a `const_int`; operand 3 is the number of enclosed loops as a `const_int` (an innermost loop has a value of 1); operand 4 is the label to jump to if the register is nonzero. See Section 9.12 [Looping Patterns], page 161.

This optional instruction pattern should be defined for machines with low-overhead looping instructions as the loop optimizer will try to modify suitable loops to utilize it. If nested low-overhead looping is not supported, use a `define_expand` (see Section 9.14 [Expander Definitions], page 164) and make the pattern fail if operand 3 is not `const1_rtx`. Similarly, if the actual or estimated maximum number of iterations is too large for this instruction, make it fail.

`'doloop_begin'`

Companion instruction to `doloop_end` required for machines that need to perform some initialization, such as loading special registers used by a low-overhead looping instruction. If initialization insns do not always need to be emitted, use a `define_expand` (see Section 9.14 [Expander Definitions], page 164) and make it fail.

`'canonicalize_funcptr_for_compare'`

Canonicalize the function pointer in operand 1 and store the result into operand 0.

Operand 0 is always a `reg` and has mode `Pmode`; operand 1 may be a `reg`, `mem`, `symbol_ref`, `const_int`, etc and also has mode `Pmode`.

Canonicalization of a function pointer usually involves computing the address of the function which would be called if the function pointer were used in an indirect call.

Only define this pattern if function pointers on the target machine can have different values but still call the same function when used in an indirect call.

`'save_stack_block'`

`'save_stack_function'`

`'save_stack_nonlocal'`

`'restore_stack_block'`

`'restore_stack_function'`

`'restore_stack_nonlocal'`

Most machines save and restore the stack pointer by copying it to or from an object of mode `Pmode`. Do not define these patterns on such machines.

Some machines require special handling for stack pointer saves and restores. On those machines, define the patterns corresponding to the non-standard cases by using a `define_expand` (see Section 9.14 [Expander Definitions], page 164) that produces the required insns. The three types of saves and restores are:

1. `'save_stack_block'` saves the stack pointer at the start of a block that allocates a variable-sized object, and `'restore_stack_block'` restores the stack pointer when the block is exited.

2. `'save_stack_function'` and `'restore_stack_function'` do a similar job for the outermost block of a function and are used when the function allocates variable-sized objects or calls `alloca`. Only the epilogue uses the restored stack pointer, allowing a simpler save or restore sequence on some machines.
3. `'save_stack_nonlocal'` is used in functions that contain labels branched to by nested functions. It saves the stack pointer in such a way that the inner function can use `'restore_stack_nonlocal'` to restore the stack pointer. The compiler generates code to restore the frame and argument pointer registers, but some machines require saving and restoring additional data such as register window information or stack backchains. Place insns in these patterns to save and restore any such required data.

When saving the stack pointer, operand 0 is the save area and operand 1 is the stack pointer. The mode used to allocate the save area defaults to `Pmode` but you can override that choice by defining the `STACK_SAVEAREA_MODE` macro (see Section 10.5 [Storage Layout], page 205). You must specify an integral mode, or `VOIDmode` if no save area is needed for a particular type of save (either because no save is needed or because a machine-specific save area can be used). Operand 0 is the stack pointer and operand 1 is the save area for restore operations. If `'save_stack_block'` is defined, operand 0 must not be `VOIDmode` since these saves can be arbitrarily nested.

A save area is a `mem` that is at a constant offset from `virtual_stack_vars_rtx` when the stack pointer is saved for use by nonlocal gotos and a `reg` in the other two cases.

`'allocate_stack'`

Subtract (or add if `STACK_GROWS_DOWNWARD` is undefined) operand 1 from the stack pointer to create space for dynamically allocated data.

Store the resultant pointer to this space into operand 0. If you are allocating space from the main stack, do this by emitting a move insn to copy `virtual_stack_dynamic_rtx` to operand 0. If you are allocating the space elsewhere, generate code to copy the location of the space to operand 0. In the latter case, you must ensure this space gets freed when the corresponding space on the main stack is free.

Do not define this pattern if all that must be done is the subtraction. Some machines require other operations such as stack probes or maintaining the back chain. Define this pattern to emit those operations in addition to updating the stack pointer.

`'probe'`

Some machines require instructions to be executed after space is allocated from the stack, for example to generate a reference at the bottom of the stack.

If you need to emit instructions before the stack has been adjusted, put them into the `'allocate_stack'` pattern. Otherwise, define this pattern to emit the required instructions.

No operands are provided.

`'check_stack'`

If stack checking cannot be done on your system by probing the stack with a load or store instruction (see Section 10.10.3 [Stack Checking], page 235), define this pattern to perform the needed check and signaling an error if the stack has overflowed. The single operand is the location in the stack furthest from the current stack pointer that you need to validate. Normally, on machines where this pattern is needed, you would obtain the stack limit from a global or thread-specific variable or register.

`'nonlocal_goto'`

Emit code to generate a non-local goto, e.g., a jump from one function to a label in an outer function. This pattern has four arguments, each representing a value to be used in the jump. The first argument is to be loaded into the frame pointer, the second is the address to branch to (code to dispatch to the actual label), the third is the address of a location where the stack is saved, and the last is the address of the label, to be placed in the location for the incoming static chain.

On most machines you need not define this pattern, since GCC will already generate the correct code, which is to load the frame pointer and static chain, restore the stack (using the `'restore_stack_nonlocal'` pattern, if defined), and jump indirectly to the dispatcher. You need only define this pattern if this code will not work on your machine.

`'nonlocal_goto_receiver'`

This pattern, if defined, contains code needed at the target of a nonlocal goto after the code already generated by GCC. You will not normally need to define this pattern. A typical reason why you might need this pattern is if some value, such as a pointer to a global table, must be restored when the frame pointer is restored. Note that a nonlocal goto only occurs within a unit-of-translation, so a global table pointer that is shared by all functions of a given module need not be restored. There are no arguments.

`'exception_receiver'`

This pattern, if defined, contains code needed at the site of an exception handler that isn't needed at the site of a nonlocal goto. You will not normally need to define this pattern. A typical reason why you might need this pattern is if some value, such as a pointer to a global table, must be restored after control flow is branched to the handler of an exception. There are no arguments.

`'builtin_setjmp_setup'`

This pattern, if defined, contains additional code needed to initialize the `jmp_buf`. You will not normally need to define this pattern. A typical reason why you might need this pattern is if some value, such as a pointer to a global table, must be restored. Though it is preferred that the pointer value be recalculated if possible (given the address of a label for instance). The single argument is a pointer to the `jmp_buf`. Note that the buffer is five words long and that the first three are normally used by the generic mechanism.

‘builtin_setjmp_receiver’

This pattern, if defined, contains code needed at the site of an built-in setjmp that isn't needed at the site of a nonlocal goto. You will not normally need to define this pattern. A typical reason why you might need this pattern is if some value, such as a pointer to a global table, must be restored. It takes one argument, which is the label to which builtin_longjmp transferred control; this pattern may be emitted at a small offset from that label.

‘builtin_longjmp’

This pattern, if defined, performs the entire action of the longjmp. You will not normally need to define this pattern unless you also define **builtin_setjmp_setup**. The single argument is a pointer to the **jmp_buf**.

‘eh_return’

This pattern, if defined, affects the way **__builtin_eh_return**, and thence the call frame exception handling library routines, are built. It is intended to handle non-trivial actions needed along the abnormal return path.

The pattern takes two arguments. The first is an offset to be applied to the stack pointer. It will have been copied to some appropriate location (typically **EH_RETURN_STACKADJ_RTX**) which will survive until after reload to when the normal epilogue is generated. The second argument is the address of the exception handler to which the function should return. This will normally need to be copied by the pattern to some special register or memory location.

This pattern only needs to be defined if call frame exception handling is to be used, and simple moves involving **EH_RETURN_STACKADJ_RTX** and **EH_RETURN_HANDLER_RTX** are not sufficient.

‘prologue’

This pattern, if defined, emits RTL for entry to a function. The function entry is responsible for setting up the stack frame, initializing the frame pointer register, saving callee saved registers, etc.

Using a prologue pattern is generally preferred over defining **TARGET_ASM_FUNCTION_PROLOGUE** to emit assembly code for the prologue.

The **prologue** pattern is particularly useful for targets which perform instruction scheduling.

‘epilogue’

This pattern emits RTL for exit from a function. The function exit is responsible for deallocating the stack frame, restoring callee saved registers and emitting the return instruction.

Using an epilogue pattern is generally preferred over defining **TARGET_ASM_FUNCTION_EPILOGUE** to emit assembly code for the epilogue.

The **epilogue** pattern is particularly useful for targets which perform instruction scheduling or which have delay slots for their return instruction.

‘sibcall_epilogue’

This pattern, if defined, emits RTL for exit from a function without the final branch back to the calling function. This pattern will be emitted before any sibling call (aka tail call) sites.

The `sibcall_epilogue` pattern must not clobber any arguments used for parameter passing or any stack slots for arguments passed to the current function.

‘trap’ This pattern, if defined, signals an error, typically by causing some kind of signal to be raised. Among other places, it is used by the Java front end to signal ‘invalid array index’ exceptions.

‘conditional_trap’

Conditional trap instruction. Operand 0 is a piece of RTL which performs a comparison. Operand 1 is the trap code, an integer.

A typical `conditional_trap` pattern looks like

```
(define_insn "conditional_trap"
  [(trap_if (match_operator 0 "trap_operator"
                [(cc0) (const_int 0)])
            (match_operand 1 "const_int_operand" "i"))]
  ""
  "...")
```

‘prefetch’

This pattern, if defined, emits code for a non-faulting data prefetch instruction. Operand 0 is the address of the memory to prefetch. Operand 1 is a constant 1 if the prefetch is preparing for a write to the memory address, or a constant 0 otherwise. Operand 2 is the expected degree of temporal locality of the data and is a value between 0 and 3, inclusive; 0 means that the data has no temporal locality, so it need not be left in the cache after the access; 3 means that the data has a high degree of temporal locality and should be left in all levels of cache possible; 1 and 2 mean, respectively, a low or moderate degree of temporal locality.

Targets that do not support write prefetches or locality hints can ignore the values of operands 1 and 2.

9.9 When the Order of Patterns Matters

Sometimes an `insn` can match more than one instruction pattern. Then the pattern that appears first in the machine description is the one used. Therefore, more specific patterns (patterns that will match fewer things) and faster instructions (those that will produce better code when they do match) should usually go first in the description.

In some cases the effect of ordering the patterns can be used to hide a pattern when it is not valid. For example, the 68000 has an instruction for converting a fullword to floating point and another for converting a byte to floating point. An instruction converting an integer to floating point could match either one. We put the pattern to convert the fullword first to make sure that one will be used rather than the other. (Otherwise a large integer might be generated as a single-byte immediate quantity, which would not work.) Instead of using this pattern ordering it would be possible to make the pattern for convert-a-byte smart enough to deal properly with any constant value.

9.10 Interdependence of Patterns

Every machine description must have a named pattern for each of the conditional branch names ‘*bcond*’. The recognition template must always have the form

```
(set (pc)
      (if_then_else (cond (cc0) (const_int 0))
                    (label_ref (match_operand 0 "" ""))
                    (pc)))
```

In addition, every machine description must have an anonymous pattern for each of the possible reverse-conditional branches. Their templates look like

```
(set (pc)
      (if_then_else (cond (cc0) (const_int 0))
                    (pc)
                    (label_ref (match_operand 0 "" ""))))
```

They are necessary because jump optimization can turn direct-conditional branches into reverse-conditional branches.

It is often convenient to use the `match_operator` construct to reduce the number of patterns that must be specified for branches. For example,

```
(define_insn ""
  [(set (pc)
        (if_then_else (match_operator 0 "comparison_operator"
                                [(cc0) (const_int 0)])
                      (pc)
                      (label_ref (match_operand 1 "" ""))))]
  "condition"
  "...")
```

In some cases machines support instructions identical except for the machine mode of one or more operands. For example, there may be “sign-extend halfword” and “sign-extend byte” instructions whose patterns are

```
(set (match_operand:SI 0 ...)
      (extend:SI (match_operand:HI 1 ...)))

(set (match_operand:SI 0 ...)
      (extend:SI (match_operand:QI 1 ...)))
```

Constant integers do not specify a machine mode, so an instruction to extend a constant value could match either pattern. The pattern it actually will match is the one that appears first in the file. For correct results, this must be the one for the widest possible mode (HI mode, here). If the pattern matches the QI mode instruction, the results will be incorrect if the constant value does not actually fit that mode.

Such instructions to extend constants are rarely generated because they are optimized away, but they do occasionally happen in nonoptimized compilations.

If a constraint in a pattern allows a constant, the reload pass may replace a register with a constant permitted by the constraint in some cases. Similarly for memory references. Because of this substitution, you should not provide separate patterns for increment and

decrement instructions. Instead, they should be generated from the same pattern that supports register-register add insns by examining the operands and generating the appropriate machine instruction.

9.11 Defining Jump Instruction Patterns

For most machines, GCC assumes that the machine has a condition code. A comparison insn sets the condition code, recording the results of both signed and unsigned comparison of the given operands. A separate branch insn tests the condition code and branches or not according its value. The branch insns come in distinct signed and unsigned flavors. Many common machines, such as the VAX, the 68000 and the 32000, work this way.

Some machines have distinct signed and unsigned compare instructions, and only one set of conditional branch instructions. The easiest way to handle these machines is to treat them just like the others until the final stage where assembly code is written. At this time, when outputting code for the compare instruction, peek ahead at the following branch using `next_cc0_user (insn)`. (The variable `insn` refers to the insn being output, in the output-writing code in an instruction pattern.) If the RTL says that is an unsigned branch, output an unsigned compare; otherwise output a signed compare. When the branch itself is output, you can treat signed and unsigned branches identically.

The reason you can do this is that GCC always generates a pair of consecutive RTL insns, possibly separated by `note` insns, one to set the condition code and one to test it, and keeps the pair inviolate until the end.

To go with this technique, you must define the machine-description macro `NOTICE_UPDATE_CC` to do `CC_STATUS_INIT`; in other words, no compare instruction is superfluous.

Some machines have compare-and-branch instructions and no condition code. A similar technique works for them. When it is time to “output” a compare instruction, record its operands in two static variables. When outputting the branch-on-condition-code instruction that follows, actually output a compare-and-branch instruction that uses the remembered operands.

It also works to define patterns for compare-and-branch instructions. In optimizing compilation, the pair of compare and branch instructions will be combined according to these patterns. But this does not happen if optimization is not requested. So you must use one of the solutions above in addition to any special patterns you define.

In many RISC machines, most instructions do not affect the condition code and there may not even be a separate condition code register. On these machines, the restriction that the definition and use of the condition code be adjacent insns is not necessary and can prevent important optimizations. For example, on the IBM RS/6000, there is a delay for taken branches unless the condition code register is set three instructions earlier than the conditional branch. The instruction scheduler cannot perform this optimization if it is not permitted to separate the definition and use of the condition code register.

On these machines, do not use `(cc0)`, but instead use a register to represent the condition code. If there is a specific condition code register in the machine, use a hard register. If the condition code or comparison result can be placed in any general register, or if there are multiple condition registers, use a pseudo register.

On some machines, the type of branch instruction generated may depend on the way the condition code was produced; for example, on the 68k and SPARC, setting the condition code directly from an add or subtract instruction does not clear the overflow bit the way that a test instruction does, so a different branch instruction must be used for some conditional branches. For machines that use (cc0), the set and use of the condition code must be adjacent (separated only by `note` insns) allowing flags in `cc_status` to be used. (See Section 10.15 [Condition Code], page 264.) Also, the comparison and branch insns can be located from each other by using the functions `prev_cc0_setter` and `next_cc0_user`.

However, this is not true on machines that do not use (cc0). On those machines, no assumptions can be made about the adjacency of the compare and branch insns and the above methods cannot be used. Instead, we use the machine mode of the condition code register to record different formats of the condition code register.

Registers used to store the condition code value should have a mode that is in class `MODE_CC`. Normally, it will be `CCmode`. If additional modes are required (as for the add example mentioned above in the SPARC), define the macro `EXTRA_CC_MODES` to list the additional modes required (see Section 10.15 [Condition Code], page 264). Also define `SELECT_CC_MODE` to choose a mode given an operand of a compare.

If it is known during RTL generation that a different mode will be required (for example, if the machine has separate compare instructions for signed and unsigned quantities, like most IBM processors), they can be specified at that time.

If the cases that require different modes would be made by instruction combination, the macro `SELECT_CC_MODE` determines which machine mode should be used for the comparison result. The patterns should be written using that mode. To support the case of the add on the SPARC discussed above, we have the pattern

```
(define_insn ""
  [(set (reg:CC_NOOV 0)
        (compare:CC_NOOV
          (plus:SI (match_operand:SI 0 "register_operand" "%r")
                  (match_operand:SI 1 "arith_operand" "rI"))
          (const_int 0)))]
  ""
  "...")
```

The `SELECT_CC_MODE` macro on the SPARC returns `CC_NOOVmode` for comparisons whose argument is a `plus`.

9.12 Defining Looping Instruction Patterns

Some machines have special jump instructions that can be utilized to make loops more efficient. A common example is the 68000 ‘`dbra`’ instruction which performs a decrement of a register and a branch if the result was greater than zero. Other machines, in particular digital signal processors (DSPs), have special block repeat instructions to provide low-overhead loop support. For example, the TI TMS320C3x/C4x DSPs have a block repeat instruction that loads special registers to mark the top and end of a loop and to count the number of loop iterations. This avoids the need for fetching and executing a ‘`dbra`’-like instruction and avoids pipeline stalls associated with the jump.

GCC has three special named patterns to support low overhead looping. They are ‘`decrement_and_branch_until_zero`’, ‘`doloop_begin`’, and ‘`doloop_end`’. The first pat-

tern, ‘decrement_and_branch_until_zero’, is not emitted during RTL generation but may be emitted during the instruction combination phase. This requires the assistance of the loop optimizer, using information collected during strength reduction, to reverse a loop to count down to zero. Some targets also require the loop optimizer to add a `REG_NONNEG` note to indicate that the iteration count is always positive. This is needed if the target performs a signed loop termination test. For example, the 68000 uses a pattern similar to the following for its `dbra` instruction:

```
(define_insn "decrement_and_branch_until_zero"
  [(set (pc)
    (if_then_else
      (ge (plus:SI (match_operand:SI 0 "general_operand" "+d*am")
        (const_int -1))
        (const_int 0))
      (label_ref (match_operand 1 "" ""))
      (pc)))
    (set (match_dup 0)
      (plus:SI (match_dup 0)
        (const_int -1)))]
  "find_reg_note (insn, REG_NONNEG, 0)"
  "...")
```

Note that since the `insn` is both a jump `insn` and has an output, it must deal with its own reloads, hence the ‘m’ constraints. Also note that since this `insn` is generated by the instruction combination phase combining two sequential `insns` together into an implicit parallel `insn`, the iteration counter needs to be biased by the same amount as the decrement operation, in this case `-1`. Note that the following similar pattern will not be matched by the combiner.

```
(define_insn "decrement_and_branch_until_zero"
  [(set (pc)
    (if_then_else
      (ge (match_operand:SI 0 "general_operand" "+d*am")
        (const_int 1))
      (label_ref (match_operand 1 "" ""))
      (pc)))
    (set (match_dup 0)
      (plus:SI (match_dup 0)
        (const_int -1)))]
  "find_reg_note (insn, REG_NONNEG, 0)"
  "...")
```

The other two special looping patterns, ‘`doloop_begin`’ and ‘`doloop_end`’, are emitted by the loop optimizer for certain well-behaved loops with a finite number of loop iterations using information collected during strength reduction.

The ‘`doloop_end`’ pattern describes the actual looping instruction (or the implicit looping operation) and the ‘`doloop_begin`’ pattern is an optional companion pattern that can be used for initialization needed for some low-overhead looping instructions.

Note that some machines require the actual looping instruction to be emitted at the top of the loop (e.g., the TMS320C3x/C4x DSPs). Emitting the true RTL for a looping instruction at the top of the loop can cause problems with flow analysis. So instead, a dummy `doloop` `insn` is emitted at the end of the loop. The machine dependent reorg pass checks for the presence of this `doloop` `insn` and then searches back to the top of the loop, where it inserts the true looping `insn` (provided there are no instructions in the loop which would cause problems). Any additional labels can be emitted at this point. In addition,

if the desired special iteration counter register was not allocated, this machine dependent reorg pass could emit a traditional compare and jump instruction pair.

The essential difference between the ‘`decrement_and_branch_until_zero`’ and the ‘`doloop_end`’ patterns is that the loop optimizer allocates an additional pseudo register for the latter as an iteration counter. This pseudo register cannot be used within the loop (i.e., general induction variables cannot be derived from it), however, in many cases the loop induction variable may become redundant and removed by the flow pass.

9.13 Canonicalization of Instructions

There are often cases where multiple RTL expressions could represent an operation performed by a single machine instruction. This situation is most commonly encountered with logical, branch, and multiply-accumulate instructions. In such cases, the compiler attempts to convert these multiple RTL expressions into a single canonical form to reduce the number of `insn` patterns required.

In addition to algebraic simplifications, following canonicalizations are performed:

- For commutative and comparison operators, a constant is always made the second operand. If a machine only supports a constant as the second operand, only patterns that match a constant in the second operand need be supplied.

For these operators, if only one operand is a `neg`, `not`, `mult`, `plus`, or `minus` expression, it will be the first operand.

- In combinations of `neg`, `mult`, `plus`, and `minus`, the `neg` operations (if any) will be moved inside the operations as far as possible. For instance, `(neg (mult A B))` is canonicalized as `(mult (neg A) B)`, but `(plus (mult (neg A) B) C)` is canonicalized as `(minus A (mult B C))`.
- For the `compare` operator, a constant is always the second operand on machines where `cc0` is used (see Section 9.11 [Jump Patterns], page 160). On other machines, there are rare cases where the compiler might want to construct a `compare` with a constant as the first operand. However, these cases are not common enough for it to be worthwhile to provide a pattern matching a constant as the first operand unless the machine actually has such an instruction.

An operand of `neg`, `not`, `mult`, `plus`, or `minus` is made the first operand under the same conditions as above.

- `(minus x (const_int n))` is converted to `(plus x (const_int -n))`.
- Within address computations (i.e., inside `mem`), a left shift is converted into the appropriate multiplication by a power of two.
- DeMorgan’s Law is used to move bitwise negation inside a bitwise logical-and or logical-or operation. If this results in only one operand being a `not` expression, it will be the first one.

A machine that has an instruction that performs a bitwise logical-and of one operand with the bitwise negation of the other should specify the pattern for that instruction as

```
(define_insn ""
  [(set (match_operand:m 0 ...)
```

```

        (and:m (not:m (match_operand:m 1 ...))
              (match_operand:m 2 ...)))]
    "... "
    "...")

```

Similarly, a pattern for a “NAND” instruction should be written

```

(define_insn ""
  [(set (match_operand:m 0 ...)
        (ior:m (not:m (match_operand:m 1 ...))
              (not:m (match_operand:m 2 ...))))])
    "... "
    "...")

```

In both cases, it is not necessary to include patterns for the many logically equivalent RTL expressions.

- The only possible RTL expressions involving both bitwise exclusive-or and bitwise negation are `(xor:m x y)` and `(not:m (xor:m x y))`.
- The sum of three items, one of which is a constant, will only appear in the form `(plus:m (plus:m x y) constant)`
- On machines that do not use `cc0`, `(compare x (const_int 0))` will be converted to `x`.
- Equality comparisons of a group of bits (usually a single bit) with zero will be written using `zero_extract` rather than the equivalent `and` or `sign_extract` operations.

9.14 Defining RTL Sequences for Code Generation

On some target machines, some standard pattern names for RTL generation cannot be handled with single `insn`, but a sequence of RTL `insns` can represent them. For these target machines, you can write a `define_expand` to specify how to generate the sequence of RTL.

A `define_expand` is an RTL expression that looks almost like a `define_insn`; but, unlike the latter, a `define_expand` is used only for RTL generation and it can produce more than one RTL `insn`.

A `define_expand` RTX has four operands:

- The name. Each `define_expand` must have a name, since the only use for it is to refer to it by name.
- The RTL template. This is a vector of RTL expressions representing a sequence of separate instructions. Unlike `define_insn`, there is no implicit surrounding `PARALLEL`.
- The condition, a string containing a C expression. This expression is used to express how the availability of this pattern depends on subclasses of target machine, selected by command-line options when GCC is run. This is just like the condition of a `define_insn` that has a standard name. Therefore, the condition (if present) may not depend on the data in the `insn` being matched, but only the target-machine-type flags. The compiler needs to test these conditions during initialization in order to learn exactly which named instructions are available in a particular run.
- The preparation statements, a string containing zero or more C statements which are to be executed before RTL code is generated from the RTL template.

Usually these statements prepare temporary registers for use as internal operands in the RTL template, but they can also generate RTL insns directly by calling routines such as `emit_insn`, etc. Any such insns precede the ones that come from the RTL template.

Every RTL insn emitted by a `define_expand` must match some `define_insn` in the machine description. Otherwise, the compiler will crash when trying to generate code for the insn or trying to optimize it.

The RTL template, in addition to controlling generation of RTL insns, also describes the operands that need to be specified when this pattern is used. In particular, it gives a predicate for each operand.

A true operand, which needs to be specified in order to generate RTL from the pattern, should be described with a `match_operand` in its first occurrence in the RTL template. This enters information on the operand's predicate into the tables that record such things. GCC uses the information to preload the operand into a register if that is required for valid RTL code. If the operand is referred to more than once, subsequent references should use `match_dup`.

The RTL template may also refer to internal "operands" which are temporary registers or labels used only within the sequence made by the `define_expand`. Internal operands are substituted into the RTL template with `match_dup`, never with `match_operand`. The values of the internal operands are not passed in as arguments by the compiler when it requests use of this pattern. Instead, they are computed within the pattern, in the preparation statements. These statements compute the values and store them into the appropriate elements of `operands` so that `match_dup` can find them.

There are two special macros defined for use in the preparation statements: `DONE` and `FAIL`. Use them with a following semicolon, as a statement.

DONE Use the `DONE` macro to end RTL generation for the pattern. The only RTL insns resulting from the pattern on this occasion will be those already emitted by explicit calls to `emit_insn` within the preparation statements; the RTL template will not be generated.

FAIL Make the pattern fail on this occasion. When a pattern fails, it means that the pattern was not truly available. The calling routines in the compiler will try other strategies for code generation using other patterns.

Failure is currently supported only for binary (addition, multiplication, shifting, etc.) and bit-field (`extv`, `extzv`, and `insv`) operations.

If the preparation falls through (invokes neither `DONE` nor `FAIL`), then the `define_expand` acts like a `define_insn` in that the RTL template is used to generate the insn.

The RTL template is not used for matching, only for generating the initial insn list. If the preparation statement always invokes `DONE` or `FAIL`, the RTL template may be reduced to a simple list of operands, such as this example:

```
(define_expand "addsi3"
  [(match_operand:SI 0 "register_operand" "")
   (match_operand:SI 1 "register_operand" "")
   (match_operand:SI 2 "register_operand" "")]
```



```

""
"operands[2]
  = force_reg (SImode, GEN_INT (65535)); ")

```

Note: If the `define_expand` is used to serve a standard binary or unary arithmetic operation or a bit-field operation, then the last insn it generates must not be a `code_label`, `barrier` or `note`. It must be an `insn`, `jump_insn` or `call_insn`. If you don't need a real insn at the end, emit an insn to copy the result of the operation into itself. Such an insn will generate no code, but it can avoid problems in the compiler.

9.15 Defining How to Split Instructions

There are two cases where you should specify how to split a pattern into multiple insns. On machines that have instructions requiring delay slots (see Section 9.18.7 [Delay Slots], page 182) or that have instructions whose output is not available for multiple cycles (see Section 9.18.8 [Processor pipeline description], page 183), the compiler phases that optimize these cases need to be able to move insns into one-instruction delay slots. However, some insns may generate more than one machine instruction. These insns cannot be placed into a delay slot.

Often you can rewrite the single insn as a list of individual insns, each corresponding to one machine instruction. The disadvantage of doing so is that it will cause the compilation to be slower and require more space. If the resulting insns are too complex, it may also suppress some optimizations. The compiler splits the insn if there is a reason to believe that it might improve instruction or delay slot scheduling.

The insn combiner phase also splits putative insns. If three insns are merged into one insn with a complex expression that cannot be matched by some `define_insn` pattern, the combiner phase attempts to split the complex pattern into two insns that are recognized. Usually it can break the complex pattern into two patterns by splitting out some subexpression. However, in some other cases, such as performing an addition of a large constant in two insns on a RISC machine, the way to split the addition into two insns is machine-dependent.

The `define_split` definition tells the compiler how to split a complex insn into several simpler insns. It looks like this:

```

(define_split
  [insn-pattern]
  "condition"
  [new-insn-pattern-1
   new-insn-pattern-2
   ...]
  "preparation-statements")

```

insn-pattern is a pattern that needs to be split and *condition* is the final condition to be tested, as in a `define_insn`. When an insn matching *insn-pattern* and satisfying *condition* is found, it is replaced in the insn list with the insns given by *new-insn-pattern-1*, *new-insn-pattern-2*, etc.

The *preparation-statements* are similar to those statements that are specified for `define_expand` (see Section 9.14 [Expander Definitions], page 164) and are executed before the new RTL is generated to prepare for the generated code or emit some insns whose pattern is not fixed. Unlike those in `define_expand`, however, these statements must not generate any

new pseudo-registers. Once reload has completed, they also must not allocate any space in the stack frame.

Patterns are matched against *insn-pattern* in two different circumstances. If an *insn* needs to be split for delay slot scheduling or *insn* scheduling, the *insn* is already known to be valid, which means that it must have been matched by some *define_insn* and, if *reload_completed* is nonzero, is known to satisfy the constraints of that *define_insn*. In that case, the new *insn* patterns must also be *insns* that are matched by some *define_insn* and, if *reload_completed* is nonzero, must also satisfy the constraints of those definitions.

As an example of this usage of *define_split*, consider the following example from 'a29k.md', which splits a *sign_extend* from HImode to SImode into a pair of shift *insns*:

```
(define_split
  [(set (match_operand:SI 0 "gen_reg_operand" "")
        (sign_extend:SI (match_operand:HI 1 "gen_reg_operand" "")))]
  ""
  [(set (match_dup 0)
        (ashift:SI (match_dup 1)
                   (const_int 16)))
   (set (match_dup 0)
        (ashiftrt:SI (match_dup 0)
                     (const_int 16)))]
  "
  { operands[1] = gen_lowpart (SImode, operands[1]); }")
```

When the combiner phase tries to split an *insn* pattern, it is always the case that the pattern is *not* matched by any *define_insn*. The combiner pass first tries to split a single *set* expression and then the same *set* expression inside a *parallel*, but followed by a *clobber* of a pseudo-reg to use as a scratch register. In these cases, the combiner expects exactly two new *insn* patterns to be generated. It will verify that these patterns match some *define_insn* definitions, so you need not do this test in the *define_split* (of course, there is no point in writing a *define_split* that will never produce *insns* that match).

Here is an example of this use of *define_split*, taken from 'rs6000.md':

```
(define_split
  [(set (match_operand:SI 0 "gen_reg_operand" "")
        (plus:SI (match_operand:SI 1 "gen_reg_operand" "")
                 (match_operand:SI 2 "non_add_cint_operand" "")))]
  ""
  [(set (match_dup 0) (plus:SI (match_dup 1) (match_dup 3)))
   (set (match_dup 0) (plus:SI (match_dup 0) (match_dup 4)))]
  "
  {
    int low = INTVAL (operands[2]) & 0xffff;
    int high = (unsigned) INTVAL (operands[2]) >> 16;

    if (low & 0x8000)
      high++, low |= 0xffff0000;

    operands[3] = GEN_INT (high << 16);
    operands[4] = GEN_INT (low);
  }")
```

Here the predicate *non_add_cint_operand* matches any *const_int* that is *not* a valid operand of a single *add* *insn*. The *add* with the smaller displacement is written so that it can be substituted into the address of a subsequent operation.

An example that uses a scratch register, from the same file, generates an equality comparison of a register and a large constant:

```
(define_split
  [(set (match_operand:CC 0 "cc_reg_operand" "")
        (compare:CC (match_operand:SI 1 "gen_reg_operand" "")
                    (match_operand:SI 2 "non_short_cint_operand" "")))
   (clobber (match_operand:SI 3 "gen_reg_operand" ""))]
  "find_single_use (operands[0], insn, 0)
  && (GET_CODE (*find_single_use (operands[0], insn, 0)) == EQ
      || GET_CODE (*find_single_use (operands[0], insn, 0)) == NE)"
  [(set (match_dup 3) (xor:SI (match_dup 1) (match_dup 4)))
   (set (match_dup 0) (compare:CC (match_dup 3) (match_dup 5)))]
  "
{
  /* Get the constant we are comparing against, C, and see what it
     looks like sign-extended to 16 bits. Then see what constant
     could be XOR'ed with C to get the sign-extended value. */

  int c = INTVAL (operands[2]);
  int sextc = (c << 16) >> 16;
  int xorv = c ^ sextc;

  operands[4] = GEN_INT (xorv);
  operands[5] = GEN_INT (sextc);
}");
```

To avoid confusion, don't write a single `define_split` that accepts some insns that match some `define_insn` as well as some insns that don't. Instead, write two separate `define_split` definitions, one for the insns that are valid and one for the insns that are not valid.

The splitter is allowed to split jump instructions into sequence of jumps or create new jumps in while splitting non-jump instructions. As the central flowgraph and branch prediction information needs to be updated, several restriction apply.

Splitting of jump instruction into sequence that over by another jump instruction is always valid, as compiler expect identical behavior of new jump. When new sequence contains multiple jump instructions or new labels, more assistance is needed. Splitter is required to create only unconditional jumps, or simple conditional jump instructions. Additionally it must attach a `REG_BR_PROB` note to each conditional jump. A global variable `split_branch_probability` hold the probability of original branch in case it was an simple conditional jump, `-1` otherwise. To simplify recomputing of edge frequencies, new sequence is required to have only forward jumps to the newly created labels.

For the common case where the pattern of a `define_split` exactly matches the pattern of a `define_insn`, use `define_insn_and_split`. It looks like this:

```
(define_insn_and_split
  [insn-pattern]
  "condition"
  "output-template"
  "split-condition"
  [new-insn-pattern-1
   new-insn-pattern-2
   ...]
  "preparation-statements"
  [insn-attributes])
```

insn-pattern, *condition*, *output-template*, and *insn-attributes* are used as in `define_insn`. The *new-insn-pattern* vector and the *preparation-statements* are used as in a `define_split`. The *split-condition* is also used as in `define_split`, with the additional behavior that if the condition starts with ‘&&’, the condition used for the split will be the constructed as a logical “and” of the split condition with the `insn` condition. For example, from `i386.md`:

```
(define_insn_and_split "zero_extendhi2_and"
  [(set (match_operand:SI 0 "register_operand" "=r")
        (zero_extend:SI (match_operand:HI 1 "register_operand" "0"))))
  (clobber (reg:CC 17))]
  "TARGET_ZERO_EXTEND_WITH_AND && !optimize_size"
  "#"
  "&& reload_completed"
  [(parallel [(set (match_dup 0)
                  (and:SI (match_dup 0) (const_int 65535)))
             (clobber (reg:CC 17))])]
  ""
  [(set_attr "type" "alu1")])
```

In this case, the actual split condition will be ‘`TARGET_ZERO_EXTEND_WITH_AND && !optimize_size && reload_completed`’.

The `define_insn_and_split` construction provides exactly the same functionality as two separate `define_insn` and `define_split` patterns. It exists for compactness, and as a maintenance tool to prevent having to ensure the two patterns’ templates match.

9.16 Including Patterns in Machine Descriptions.

The `include` pattern tells the compiler tools where to look for patterns that are in files other than in the file ‘`.md`’. This is used only at build time and there is no preprocessing allowed.

It looks like:

```
(include
  pathname)
```

For example:

```
(include "filestuff")
```

Where *pathname* is a string that specifies the location of the file, specifies the include file to be in ‘`gcc/config/target/filestuff`’. The directory ‘`gcc/config/target`’ is regarded as the default directory.

Machine descriptions may be split up into smaller more manageable subsections and placed into subdirectories.

By specifying:

```
(include "BOGUS/filestuff")
```

the include file is specified to be in ‘`gcc/config/target/BOGUS/filestuff`’. Specifying an absolute path for the include file such as;

```
(include "/u2/BOGUS/filestuff")
```

is permitted but is not encouraged.

9.16.1 RTL Generation Tool Options for Directory Search

The ‘`-Idir`’ option specifies directories to search for machine descriptions. For example:

```
genrecog -I/p1/abc/proc1 -I/p2/abcd/pro2 target.md
```

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 machine definition file, substituting your own version, since these directories are searched before the default machine description file directories. If you use more than one ‘`-I`’ option, the directories are scanned in left-to-right order; the standard default directory come after.

9.17 Machine-Specific Peephole Optimizers

In addition to instruction patterns the ‘`md`’ file may contain definitions of machine-specific peephole optimizations.

The combiner does not notice certain peephole optimizations when the data flow in the program does not suggest that it should try them. For example, sometimes two consecutive insns related in purpose can be combined even though the second one does not appear to use a register computed in the first one. A machine-specific peephole optimizer can detect such opportunities.

There are two forms of peephole definitions that may be used. The original `define_peephole` is run at assembly output time to match insns and substitute assembly text. Use of `define_peephole` is deprecated.

A newer `define_peephole2` matches insns and substitutes new insns. The `peephole2` pass is run after register allocation but before scheduling, which may result in much better code for targets that do scheduling.

9.17.1 RTL to Text Peephole Optimizers

A definition looks like this:

```
(define_peephole
  [insn-pattern-1
   insn-pattern-2
   ...]
  "condition"
  "template"
  "optional-insn-attributes")
```

The last string operand may be omitted if you are not using any machine-specific information in this machine description. If present, it must obey the same rules as in a `define_insn`.

In this skeleton, *insn-pattern-1* and so on are patterns to match consecutive insns. The optimization applies to a sequence of insns when *insn-pattern-1* matches the first one, *insn-pattern-2* matches the next, and so on.

Each of the insns matched by a peephole must also match a `define_insn`. Peepholes are checked only at the last stage just before code generation, and only optionally. Therefore, any insn which would match a peephole but no `define_insn` will cause a crash in code generation in an unoptimized compilation, or at various optimization stages.

The operands of the insns are matched with `match_operands`, `match_operator`, and `match_dup`, as usual. What is not usual is that the operand numbers apply to all the insn patterns in the definition. So, you can check for identical operands in two insns by using `match_operand` in one insn and `match_dup` in the other.

The operand constraints used in `match_operand` patterns do not have any direct effect on the applicability of the peephole, but they will be validated afterward, so make sure your constraints are general enough to apply whenever the peephole matches. If the peephole matches but the constraints are not satisfied, the compiler will crash.

It is safe to omit constraints in all the operands of the peephole; or you can write constraints which serve as a double-check on the criteria previously tested.

Once a sequence of insns matches the patterns, the *condition* is checked. This is a C expression which makes the final decision whether to perform the optimization (we do so if the expression is nonzero). If *condition* is omitted (in other words, the string is empty) then the optimization is applied to every sequence of insns that matches the patterns.

The defined peephole optimizations are applied after register allocation is complete. Therefore, the peephole definition can check which operands have ended up in which kinds of registers, just by looking at the operands.

The way to refer to the operands in *condition* is to write `operands[i]` for operand number *i* (as matched by `(match_operand i ...)`). Use the variable `insn` to refer to the last of the insns being matched; use `prev_active_insn` to find the preceding insns.

When optimizing computations with intermediate results, you can use *condition* to match only when the intermediate results are not used elsewhere. Use the C expression `dead_or_set_p (insn, op)`, where *insn* is the insn in which you expect the value to be used for the last time (from the value of `insn`, together with use of `prev_nonnote_insn`), and *op* is the intermediate value (from `operands[i]`).

Applying the optimization means replacing the sequence of insns with one new insn. The *template* controls ultimate output of assembler code for this combined insn. It works exactly like the template of a `define_insn`. Operand numbers in this template are the same ones used in matching the original sequence of insns.

The result of a defined peephole optimizer does not need to match any of the insn patterns in the machine description; it does not even have an opportunity to match them. The peephole optimizer definition itself serves as the insn pattern to control how the insn is output.

Defined peephole optimizers are run as assembler code is being output, so the insns they produce are never combined or rearranged in any way.

Here is an example, taken from the 68000 machine description:

```

(define_peephole
  [(set (reg:SI 15) (plus:SI (reg:SI 15) (const_int 4)))
   (set (match_operand:DF 0 "register_operand" "=f")
        (match_operand:DF 1 "register_operand" "ad"))]
  "FP_REG_P (operands[0]) && ! FP_REG_P (operands[1])"
  {
    rtx xoperands[2];
    xoperands[1] = gen_rtx (REG, SImode, REGNO (operands[1]) + 1);
  #ifdef MOTOROLA
    output_asm_insn ("move.l %1,(sp)", xoperands);
    output_asm_insn ("move.l %1,-(sp)", operands);
    return "fmove.d (sp)+,%0";
  #else
    output_asm_insn ("movel %1,sp@", xoperands);
    output_asm_insn ("movel %1,sp@-", operands);
    return "fmoved sp@+,%0";
  #endif
  })

```

The effect of this optimization is to change

```

jbsr _foobar
addq #4,sp
movel d1,sp@-
movel d0,sp@-
fmoved sp@+,fp0

```

into

```

jbsr _foobar
movel d1,sp@
movel d0,sp@-
fmoved sp@+,fp0

```

insn-pattern-1 and so on look *almost* like the second operand of `define_insn`. There is one important difference: the second operand of `define_insn` consists of one or more RTX's enclosed in square brackets. Usually, there is only one: then the same action can be written as an element of a `define_peephole`. But when there are multiple actions in a `define_insn`, they are implicitly enclosed in a `parallel`. Then you must explicitly write the `parallel`, and the square brackets within it, in the `define_peephole`. Thus, if an `insn` pattern looks like this,

```

(define_insn "divmodsi4"
  [(set (match_operand:SI 0 "general_operand" "=d")
        (div:SI (match_operand:SI 1 "general_operand" "0")
                (match_operand:SI 2 "general_operand" "dmsK")))
   (set (match_operand:SI 3 "general_operand" "=d")
        (mod:SI (match_dup 1) (match_dup 2)))]
  "TARGET_68020"
  "divs1%.1 %2,%3:%0")

```

then the way to mention this `insn` in a peephole is as follows:

```

(define_peephole
  [...
  (parallel
    [(set (match_operand:SI 0 "general_operand" "=d")
          (div:SI (match_operand:SI 1 "general_operand" "0")
                  (match_operand:SI 2 "general_operand" "dmsK")))
     (set (match_operand:SI 3 "general_operand" "=d")
          (mod:SI (match_dup 1) (match_dup 2)))]
    ...]
  ...))

```

9.17.2 RTL to RTL Peephole Optimizers

The `define_peephole2` definition tells the compiler how to substitute one sequence of instructions for another sequence, what additional scratch registers may be needed and what their lifetimes must be.

```
(define_peephole2
  [insn-pattern-1
   insn-pattern-2
   ...]
  "condition"
  [new-insn-pattern-1
   new-insn-pattern-2
   ...]
  "preparation-statements")
```

The definition is almost identical to `define_split` (see Section 9.15 [Insn Splitting], page 167) except that the pattern to match is not a single instruction, but a sequence of instructions.

It is possible to request additional scratch registers for use in the output template. If appropriate registers are not free, the pattern will simply not match.

Scratch registers are requested with a `match_scratch` pattern at the top level of the input pattern. The allocated register (initially) will be dead at the point requested within the original sequence. If the scratch is used at more than a single point, a `match_dup` pattern at the top level of the input pattern marks the last position in the input sequence at which the register must be available.

Here is an example from the IA-32 machine description:

```
(define_peephole2
  [(match_scratch:SI 2 "r")
   (parallel [(set (match_operand:SI 0 "register_operand" "")
                  (match_operator:SI 3 "arith_or_logical_operator"
                    [(match_dup 0)
                     (match_operand:SI 1 "memory_operand" "")]))
              (clobber (reg:CC 17))]])
  "! optimize_size && ! TARGET_READ_MODIFY"
  [(set (match_dup 2) (match_dup 1))
   (parallel [(set (match_dup 0)
                  (match_op_dup 3 [(match_dup 0) (match_dup 2)]))
              (clobber (reg:CC 17))]])
  "")
```

This pattern tries to split a load from its use in the hopes that we'll be able to schedule around the memory load latency. It allocates a single SImode register of class `GENERAL_REGS` ("`r`") that needs to be live only at the point just before the arithmetic.

A real example requiring extended scratch lifetimes is harder to come by, so here's a silly made-up example:

```
(define_peephole2
  [(match_scratch:SI 4 "r")
   (set (match_operand:SI 0 "" "") (match_operand:SI 1 "" ""))
   (set (match_operand:SI 2 "" "") (match_dup 1))
   (match_dup 4)
   (set (match_operand:SI 3 "" "") (match_dup 1))]
  "/* determine 1 does not overlap 0 and 2 */"
  [(set (match_dup 4) (match_dup 1))])
```

```
(set (match_dup 0) (match_dup 4))
(set (match_dup 2) (match_dup 4))]
(set (match_dup 3) (match_dup 4))]
"")
```

If we had not added the `(match_dup 4)` in the middle of the input sequence, it might have been the case that the register we chose at the beginning of the sequence is killed by the first or second `set`.

9.18 Instruction Attributes

In addition to describing the instruction supported by the target machine, the ‘`md`’ file also defines a group of *attributes* and a set of values for each. Every generated `insn` is assigned a value for each attribute. One possible attribute would be the effect that the `insn` has on the machine’s condition code. This attribute can then be used by `NOTICE_UPDATE_CC` to track the condition codes.

9.18.1 Defining Attributes and their Values

The `define_attr` expression is used to define each attribute required by the target machine. It looks like:

```
(define_attr name list-of-values default)
```

name is a string specifying the name of the attribute being defined.

list-of-values is either a string that specifies a comma-separated list of values that can be assigned to the attribute, or a null string to indicate that the attribute takes numeric values.

default is an attribute expression that gives the value of this attribute for `insns` that match patterns whose definition does not include an explicit value for this attribute. See Section 9.18.4 [Attr Example], page 180, for more information on the handling of defaults. See Section 9.18.6 [Constant Attributes], page 182, for information on attributes that do not depend on any particular `insn`.

For each defined attribute, a number of definitions are written to the ‘`insn-attr.h`’ file. For cases where an explicit set of values is specified for an attribute, the following are defined:

- A ‘`#define`’ is written for the symbol ‘`HAVE_ATTR_name`’.
- An enumeral class is defined for ‘`attr_name`’ with elements of the form ‘`upper_name_upper-value`’ where the attribute name and value are first converted to upper case.
- A function ‘`get_attr_name`’ is defined that is passed an `insn` and returns the attribute value for that `insn`.

For example, if the following is present in the ‘`md`’ file:

```
(define_attr "type" "branch,fp,load,store,arith" ...)
```

the following lines will be written to the file ‘`insn-attr.h`’.

```
#define HAVE_ATTR_type
enum attr_type {TYPE_BRANCH, TYPE_FP, TYPE_LOAD,
```

```

        TYPE_STORE, TYPE_ARITH};
extern enum attr_type get_attr_type ();

```

If the attribute takes numeric values, no `enum` type will be defined and the function to obtain the attribute's value will return `int`.

9.18.2 Attribute Expressions

RTL expressions used to define attributes use the codes described above plus a few specific to attribute definitions, to be discussed below. Attribute value expressions must have one of the following forms:

`(const_int i)`

The integer *i* specifies the value of a numeric attribute. *i* must be non-negative.

The value of a numeric attribute can be specified either with a `const_int`, or as an integer represented as a string in `const_string`, `eq_attr` (see below), `attr`, `symbol_ref`, simple arithmetic expressions, and `set_attr` overrides on specific instructions (see Section 9.18.3 [Tagging Insns], page 178).

`(const_string value)`

The string *value* specifies a constant attribute value. If *value* is specified as `"*"`, it means that the default value of the attribute is to be used for the insn containing this expression. `"*"` obviously cannot be used in the *default* expression of a `define_attr`.

If the attribute whose value is being specified is numeric, *value* must be a string containing a non-negative integer (normally `const_int` would be used in this case). Otherwise, it must contain one of the valid values for the attribute.

`(if_then_else test true-value false-value)`

test specifies an attribute test, whose format is defined below. The value of this expression is *true-value* if *test* is true, otherwise it is *false-value*.

`(cond [test1 value1 ...] default)`

The first operand of this expression is a vector containing an even number of expressions and consisting of pairs of *test* and *value* expressions. The value of the `cond` expression is that of the *value* corresponding to the first true *test* expression. If none of the *test* expressions are true, the value of the `cond` expression is that of the *default* expression.

test expressions can have one of the following forms:

`(const_int i)`

This test is true if *i* is nonzero and false otherwise.

`(not test)`

`(ior test1 test2)`

`(and test1 test2)`

These tests are true if the indicated logical function is true.

`(match_operand:m n pred constraints)`

This test is true if operand *n* of the insn whose attribute value is being determined has mode *m* (this part of the test is ignored if *m* is `VOIDmode`) and

the function specified by the string *pred* returns a nonzero value when passed operand *n* and mode *m* (this part of the test is ignored if *pred* is the null string).

The *constraints* operand is ignored and should be the null string.

```
(le arith1 arith2)
(leu arith1 arith2)
(lt arith1 arith2)
(ltu arith1 arith2)
(gt arith1 arith2)
(gtu arith1 arith2)
(ge arith1 arith2)
(geu arith1 arith2)
(ne arith1 arith2)
(eq arith1 arith2)
```

These tests are true if the indicated comparison of the two arithmetic expressions is true. Arithmetic expressions are formed with `plus`, `minus`, `mult`, `div`, `mod`, `abs`, `neg`, `and`, `ior`, `xor`, `not`, `ashift`, `lshiftrt`, and `ashiftrt` expressions.

`const_int` and `symbol_ref` are always valid terms (see Section 9.18.5 [Insn Lengths], page 180, for additional forms). `symbol_ref` is a string denoting a C expression that yields an `int` when evaluated by the `'get_attr_...'` routine. It should normally be a global variable.

```
(eq_attr name value)
```

name is a string specifying the name of an attribute.

value is a string that is either a valid value for attribute *name*, a comma-separated list of values, or `'!'` followed by a value or list. If *value* does not begin with a `'!'`, this test is true if the value of the *name* attribute of the current insn is in the list specified by *value*. If *value* begins with a `'!'`, this test is true if the attribute's value is *not* in the specified list.

For example,

```
(eq_attr "type" "load,store")
```

is equivalent to

```
(ior (eq_attr "type" "load") (eq_attr "type" "store"))
```

If *name* specifies an attribute of `'alternative'`, it refers to the value of the compiler variable `which_alternative` (see Section 9.6 [Output Statement], page 124) and the values must be small integers. For example,

```
(eq_attr "alternative" "2,3")
```

is equivalent to

```
(ior (eq (symbol_ref "which_alternative") (const_int 2))
      (eq (symbol_ref "which_alternative") (const_int 3)))
```

Note that, for most attributes, an `eq_attr` test is simplified in cases where the value of the attribute being tested is known for all insns matching a particular pattern. This is by far the most common case.

```
(attr_flag name)
```

The value of an `attr_flag` expression is true if the flag specified by *name* is true for the `insn` currently being scheduled.

name is a string specifying one of a fixed set of flags to test. Test the flags `forward` and `backward` to determine the direction of a conditional branch. Test the flags `very_likely`, `likely`, `very_unlikely`, and `unlikely` to determine if a conditional branch is expected to be taken.

If the `very_likely` flag is true, then the `likely` flag is also true. Likewise for the `very_unlikely` and `unlikely` flags.

This example describes a conditional branch delay slot which can be nullified for forward branches that are taken (`annul-true`) or for backward branches which are not taken (`annul-false`).

```
(define_delay (eq_attr "type" "cbranch")
  [(eq_attr "in_branch_delay" "true")
   (and (eq_attr "in_branch_delay" "true")
        (attr_flag "forward"))
   (and (eq_attr "in_branch_delay" "true")
        (attr_flag "backward"))])
```

The `forward` and `backward` flags are false if the current `insn` being scheduled is not a conditional branch.

The `very_likely` and `likely` flags are true if the `insn` being scheduled is not a conditional branch. The `very_unlikely` and `unlikely` flags are false if the `insn` being scheduled is not a conditional branch.

`attr_flag` is only used during delay slot scheduling and has no meaning to other passes of the compiler.

(`attr name`)

The value of another attribute is returned. This is most useful for numeric attributes, as `eq_attr` and `attr_flag` produce more efficient code for non-numeric attributes.

9.18.3 Assigning Attribute Values to Insns

The value assigned to an attribute of an `insn` is primarily determined by which pattern is matched by that `insn` (or which `define_peephole` generated it). Every `define_insn` and `define_peephole` can have an optional last argument to specify the values of attributes for matching `insns`. The value of any attribute not specified in a particular `insn` is set to the default value for that attribute, as specified in its `define_attr`. Extensive use of default values for attributes permits the specification of the values for only one or two attributes in the definition of most `insn` patterns, as seen in the example in the next section.

The optional last argument of `define_insn` and `define_peephole` is a vector of expressions, each of which defines the value for a single attribute. The most general way of assigning an attribute's value is to use a `set` expression whose first operand is an `attr` expression giving the name of the attribute being set. The second operand of the `set` is an attribute expression (see Section 9.18.2 [Expressions], page 176) giving the value of the attribute.

When the attribute value depends on the 'alternative' attribute (i.e., which is the applicable alternative in the constraint of the `insn`), the `set_attr_alternative` expression can be used. It allows the specification of a vector of attribute expressions, one for each alternative.

When the generality of arbitrary attribute expressions is not required, the simpler `set_attr` expression can be used, which allows specifying a string giving either a single attribute value or a list of attribute values, one for each alternative.

The form of each of the above specifications is shown below. In each case, *name* is a string specifying the attribute to be set.

```
(set_attr name value-string)
```

value-string is either a string giving the desired attribute value, or a string containing a comma-separated list giving the values for succeeding alternatives. The number of elements must match the number of alternatives in the constraint of the `insn` pattern.

Note that it may be useful to specify ‘*’ for some alternative, in which case the attribute will assume its default value for `insns` matching that alternative.

```
(set_attr_alternative name [value1 value2 ...])
```

Depending on the alternative of the `insn`, the value will be one of the specified values. This is a shorthand for using a `cond` with tests on the ‘`alternative`’ attribute.

```
(set (attr name) value)
```

The first operand of this `set` must be the special RTL expression `attr`, whose sole operand is a string giving the name of the attribute being set. *value* is the value of the attribute.

The following shows three different ways of representing the same attribute value specification:

```
(set_attr "type" "load,store,arith")
```

```
(set_attr_alternative "type"
  [(const_string "load") (const_string "store")
   (const_string "arith")])
```

```
(set (attr "type")
  (cond [(eq_attr "alternative" "1") (const_string "load")
        (eq_attr "alternative" "2") (const_string "store")]
        (const_string "arith")))
```

The `define_asm_attributes` expression provides a mechanism to specify the attributes assigned to `insns` produced from an `asm` statement. It has the form:

```
(define_asm_attributes [attr-sets])
```

where *attr-sets* is specified the same as for both the `define_insn` and the `define_peephole` expressions.

These values will typically be the “worst case” attribute values. For example, they might indicate that the condition code will be clobbered.

A specification for a `length` attribute is handled specially. The way to compute the length of an `asm` `insn` is to multiply the length specified in the expression `define_asm_attributes` by the number of machine instructions specified in the `asm` statement, determined by counting the number of semicolons and newlines in the string. Therefore, the value of the `length` attribute specified in a `define_asm_attributes` should be the maximum possible length of a single machine instruction.

9.18.4 Example of Attribute Specifications

The judicious use of defaulting is important in the efficient use of `insn` attributes. Typically, `insns` are divided into *types* and an attribute, customarily called `type`, is used to represent this value. This attribute is normally used only to define the default value for other attributes. An example will clarify this usage.

Assume we have a RISC machine with a condition code and in which only full-word operations are performed in registers. Let us assume that we can divide all `insns` into loads, stores, (integer) arithmetic operations, floating point operations, and branches.

Here we will concern ourselves with determining the effect of an `insn` on the condition code and will limit ourselves to the following possible effects: The condition code can be set unpredictably (clobbered), not be changed, be set to agree with the results of the operation, or only changed if the item previously set into the condition code has been modified.

Here is part of a sample ‘`md`’ file for such a machine:

```
(define_attr "type" "load,store,arith,fp,branch" (const_string "arith"))

(define_attr "cc" "clobber,unchanged,set,change0"
  (cond [(eq_attr "type" "load")
         (const_string "change0")
        (eq_attr "type" "store,branch")
         (const_string "unchanged")
        (eq_attr "type" "arith")
         (if_then_else (match_operand:SI 0 "" "")
                      (const_string "set")
                      (const_string "clobber"))]
        (const_string "clobber")))

(define_insn ""
  [(set (match_operand:SI 0 "general_operand" "=r,r,m")
        (match_operand:SI 1 "general_operand" "r,m,r"))]
  ""
  "@
  move %0,%1
  load %0,%1
  store %0,%1"
  [(set_attr "type" "arith,load,store")])
```

Note that we assume in the above example that arithmetic operations performed on quantities smaller than a machine word clobber the condition code since they will set the condition code to a value corresponding to the full-word result.

9.18.5 Computing the Length of an Insn

For many machines, multiple types of branch instructions are provided, each for different length branch displacements. In most cases, the assembler will choose the correct instruction to use. However, when the assembler cannot do so, GCC can when a special attribute, the ‘`length`’ attribute, is defined. This attribute must be defined to have numeric values by specifying a null string in its `define_attr`.

In the case of the ‘`length`’ attribute, two additional forms of arithmetic terms are allowed in test expressions:

(*match_dup n*)

This refers to the address of operand *n* of the current insn, which must be a *label_ref*.

(*pc*)

This refers to the address of the *current* insn. It might have been more consistent with other usage to make this the address of the *next* insn but this would be confusing because the length of the current insn is to be computed.

For normal insns, the length will be determined by value of the ‘*length*’ attribute. In the case of *addr_vec* and *addr_diff_vec* insn patterns, the length is computed as the number of vectors multiplied by the size of each vector.

Lengths are measured in addressable storage units (bytes).

The following macros can be used to refine the length computation:

FIRST_INSN_ADDRESS

When the *length* insn attribute is used, this macro specifies the value to be assigned to the address of the first insn in a function. If not specified, 0 is used.

ADJUST_INSN_LENGTH (*insn*, *length*)

If defined, modifies the length assigned to instruction *insn* as a function of the context in which it is used. *length* is an lvalue that contains the initially computed length of the insn and should be updated with the correct length of the insn.

This macro will normally not be required. A case in which it is required is the ROMP. On this machine, the size of an *addr_vec* insn must be increased by two to compensate for the fact that alignment may be required.

The routine that returns *get_attr_length* (the value of the *length* attribute) can be used by the output routine to determine the form of the branch instruction to be written, as the example below illustrates.

As an example of the specification of variable-length branches, consider the IBM 360. If we adopt the convention that a register will be set to the starting address of a function, we can jump to labels within 4k of the start using a four-byte instruction. Otherwise, we need a six-byte sequence to load the address from memory and then branch to it.

On such a machine, a pattern for a branch instruction might be specified as follows:

```
(define_insn "jump"
  [(set (pc)
        (label_ref (match_operand 0 "" "")))]
  ""
  {
    return (get_attr_length (insn) == 4
           ? "b %l0" : "l r15,=a(%l0); br r15");
  }
  [(set (attr "length")
        (if_then_else (lt (match_dup 0) (const_int 4096))
                      (const_int 4)
                      (const_int 6)))]])
```

9.18.6 Constant Attributes

A special form of `define_attr`, where the expression for the default value is a `const` expression, indicates an attribute that is constant for a given run of the compiler. Constant attributes may be used to specify which variety of processor is used. For example,

```
(define_attr "cpu" "m88100,m88110,m88000"
  (const
    (cond [(symbol_ref "TARGET_88100") (const_string "m88100")
          (symbol_ref "TARGET_88110") (const_string "m88110")]
          (const_string "m88000"))))

(define_attr "memory" "fast,slow"
  (const
    (if_then_else (symbol_ref "TARGET_FAST_MEM")
                  (const_string "fast")
                  (const_string "slow"))))
```

The routine generated for constant attributes has no parameters as it does not depend on any particular insn. RTL expressions used to define the value of a constant attribute may use the `symbol_ref` form, but may not use either the `match_operand` form or `eq_attr` forms involving insn attributes.

9.18.7 Delay Slot Scheduling

The insn attribute mechanism can be used to specify the requirements for delay slots, if any, on a target machine. An instruction is said to require a *delay slot* if some instructions that are physically after the instruction are executed as if they were located before it. Classic examples are branch and call instructions, which often execute the following instruction before the branch or call is performed.

On some machines, conditional branch instructions can optionally *annul* instructions in the delay slot. This means that the instruction will not be executed for certain branch outcomes. Both instructions that annul if the branch is true and instructions that annul if the branch is false are supported.

Delay slot scheduling differs from instruction scheduling in that determining whether an instruction needs a delay slot is dependent only on the type of instruction being generated, not on data flow between the instructions. See the next section for a discussion of data-dependent instruction scheduling.

The requirement of an insn needing one or more delay slots is indicated via the `define_delay` expression. It has the following form:

```
(define_delay test
  [delay-1 annul-true-1 annul-false-1
   delay-2 annul-true-2 annul-false-2
   ...])
```

test is an attribute test that indicates whether this `define_delay` applies to a particular insn. If so, the number of required delay slots is determined by the length of the vector specified as the second argument. An insn placed in delay slot *n* must satisfy attribute test *delay-n*. *annul-true-n* is an attribute test that specifies which insns may be annulled if the branch is true. Similarly, *annul-false-n* specifies which insns in the delay slot may

be annulled if the branch is false. If annulling is not supported for that delay slot, `(nil)` should be coded.

For example, in the common case where branch and call insns require a single delay slot, which may contain any insn other than a branch or call, the following would be placed in the 'md' file:

```
(define_delay (eq_attr "type" "branch,call")
              [(eq_attr "type" "!branch,call") (nil) (nil)])
```

Multiple `define_delay` expressions may be specified. In this case, each such expression specifies different delay slot requirements and there must be no insn for which tests in two `define_delay` expressions are both true.

For example, if we have a machine that requires one delay slot for branches but two for calls, no delay slot can contain a branch or call insn, and any valid insn in the delay slot for the branch can be annulled if the branch is true, we might represent this as follows:

```
(define_delay (eq_attr "type" "branch")
              [(eq_attr "type" "!branch,call")
               (eq_attr "type" "!branch,call")
               (nil)])

(define_delay (eq_attr "type" "call")
              [(eq_attr "type" "!branch,call") (nil) (nil)
               (eq_attr "type" "!branch,call") (nil) (nil)])
```

9.18.8 Specifying processor pipeline description

To achieve better performance, most modern processors (super-pipelined, superscalar RISC, and VLIW processors) have many *functional units* on which several instructions can be executed simultaneously. An instruction starts execution if its issue conditions are satisfied. If not, the instruction is stalled until its conditions are satisfied. Such *interlock (pipeline) delay* causes interruption of the fetching of successor instructions (or demands nop instructions, e.g. for some MIPS processors).

There are two major kinds of interlock delays in modern processors. The first one is a data dependence delay determining *instruction latency time*. The instruction execution is not started until all source data have been evaluated by prior instructions (there are more complex cases when the instruction execution starts even when the data are not available but will be ready in given time after the instruction execution start). Taking the data dependence delays into account is simple. The data dependence (true, output, and anti-dependence) delay between two instructions is given by a constant. In most cases this approach is adequate. The second kind of interlock delays is a reservation delay. The reservation delay means that two instructions under execution will be in need of shared processors resources, i.e. buses, internal registers, and/or functional units, which are reserved for some time. Taking this kind of delay into account is complex especially for modern RISC processors.

The task of exploiting more processor parallelism is solved by an instruction scheduler. For a better solution to this problem, the instruction scheduler has to have an adequate description of the processor parallelism (or *pipeline description*). Currently GCC provides two alternative ways to describe processor parallelism, both described below. The first method is outlined in the next section; it was once the only method provided by GCC,

and thus is used in a number of exiting ports. The second, and preferred method, specifies functional unit reservations for groups of instructions with the aid of *regular expressions*. This is called the *automaton based description*.

The GCC instruction scheduler uses a *pipeline hazard recognizer* to figure out the possibility of the instruction issue by the processor on a given simulated processor cycle. The pipeline hazard recognizer is automatically generated from the processor pipeline description. The pipeline hazard recognizer generated from the automaton based description is more sophisticated and based on a deterministic finite state automaton (DFA) and therefore faster than one generated from the old description. Furthermore, its speed is not dependent on processor complexity. The instruction issue is possible if there is a transition from one automaton state to another one.

You can use any model to describe processor pipeline characteristics or even a mix of them. You could use the old description for some processor submodels and the DFA-based one for the rest processor submodels.

In general, the usage of the automaton based description is more preferable. Its model is more rich. It permits to describe more accurately pipeline characteristics of processors which results in improving code quality (although sometimes only on several percent fractions). It will be also used as an infrastructure to implement sophisticated and practical insn scheduling which will try many instruction sequences to choose the best one.

9.18.8.1 Specifying Function Units

On most RISC machines, there are instructions whose results are not available for a specific number of cycles. Common cases are instructions that load data from memory. On many machines, a pipeline stall will result if the data is referenced too soon after the load instruction.

In addition, many newer microprocessors have multiple function units, usually one for integer and one for floating point, and often will incur pipeline stalls when a result that is needed is not yet ready.

The descriptions in this section allow the specification of how much time must elapse between the execution of an instruction and the time when its result is used. It also allows specification of when the execution of an instruction will delay execution of similar instructions due to function unit conflicts.

For the purposes of the specifications in this section, a machine is divided into *function units*, each of which execute a specific class of instructions in first-in-first-out order. Function units that accept one instruction each cycle and allow a result to be used in the succeeding instruction (usually via forwarding) need not be specified. Classic RISC microprocessors will normally have a single function unit, which we can call ‘memory’. The newer “superscalar” processors will often have function units for floating point operations, usually at least a floating point adder and multiplier.

Each usage of a function units by a class of insns is specified with a `define_function_unit` expression, which looks like this:

```
(define_function_unit name multiplicity simultaneity
  test ready-delay issue-delay
  [conflict-list])
```

name is a string giving the name of the function unit.

multiplicity is an integer specifying the number of identical units in the processor. If more than one unit is specified, they will be scheduled independently. Only truly independent units should be counted; a pipelined unit should be specified as a single unit. (The only common example of a machine that has multiple function units for a single instruction class that are truly independent and not pipelined are the two multiply and two increment units of the CDC 6600.)

simultaneity specifies the maximum number of insns that can be executing in each instance of the function unit simultaneously or zero if the unit is pipelined and has no limit.

All `define_function_unit` definitions referring to function unit *name* must have the same name and values for *multiplicity* and *simultaneity*.

test is an attribute test that selects the insns we are describing in this definition. Note that an insn may use more than one function unit and a function unit may be specified in more than one `define_function_unit`.

ready-delay is an integer that specifies the number of cycles after which the result of the instruction can be used without introducing any stalls.

issue-delay is an integer that specifies the number of cycles after the instruction matching the *test* expression begins using this unit until a subsequent instruction can begin. A cost of *N* indicates an *N-1* cycle delay. A subsequent instruction may also be delayed if an earlier instruction has a longer *ready-delay* value. This blocking effect is computed using the *simultaneity*, *ready-delay*, *issue-delay*, and *conflict-list* terms. For a normal non-pipelined function unit, *simultaneity* is one, the unit is taken to block for the *ready-delay* cycles of the executing insn, and smaller values of *issue-delay* are ignored.

conflict-list is an optional list giving detailed conflict costs for this unit. If specified, it is a list of condition test expressions to be applied to insns chosen to execute in *name* following the particular insn matching *test* that is already executing in *name*. For each insn in the list, *issue-delay* specifies the conflict cost; for insns not in the list, the cost is zero. If not specified, *conflict-list* defaults to all instructions that use the function unit.

Typical uses of this vector are where a floating point function unit can pipeline either single- or double-precision operations, but not both, or where a memory unit can pipeline loads, but not stores, etc.

As an example, consider a classic RISC machine where the result of a load instruction is not available for two cycles (a single “delay” instruction is required) and where only one load instruction can be executed simultaneously. This would be specified as:

```
(define_function_unit "memory" 1 1 (eq_attr "type" "load") 2 0)
```

For the case of a floating point function unit that can pipeline either single or double precision, but not both, the following could be specified:

```
(define_function_unit
  "fp" 1 0 (eq_attr "type" "sp_fp") 4 4 [(eq_attr "type" "dp_fp")])
(define_function_unit
  "fp" 1 0 (eq_attr "type" "dp_fp") 4 4 [(eq_attr "type" "sp_fp")])
```

Note: The scheduler attempts to avoid function unit conflicts and uses all the specifications in the `define_function_unit` expression. It has recently come to our attention that these specifications may not allow modeling of some of the newer “superscalar” processors that have insns using multiple pipelined units. These insns will cause a potential conflict

for the second unit used during their execution and there is no way of representing that conflict. We welcome any examples of how function unit conflicts work in such processors and suggestions for their representation.

9.18.8.2 Describing instruction pipeline characteristics

This section describes constructions of the automaton based processor pipeline description. The order of constructions within the machine description file is not important.

The following optional construction describes names of automata generated and used for the pipeline hazards recognition. Sometimes the generated finite state automaton used by the pipeline hazard recognizer is large. If we use more than one automaton and bind functional units to the automata, the total size of the automata is usually less than the size of the single automaton. If there is no one such construction, only one finite state automaton is generated.

```
(define_automaton automata-names)
```

automata-names is a string giving names of the automata. The names are separated by commas. All the automata should have unique names. The automaton name is used in the constructions `define_cpu_unit` and `define_query_cpu_unit`.

Each processor functional unit used in the description of instruction reservations should be described by the following construction.

```
(define_cpu_unit unit-names [automaton-name])
```

unit-names is a string giving the names of the functional units separated by commas. Don't use name 'nothing', it is reserved for other goals.

automaton-name is a string giving the name of the automaton with which the unit is bound. The automaton should be described in construction `define_automaton`. You should give *automaton-name*, if there is a defined automaton.

The following construction describes CPU functional units analogously to `define_cpu_unit`. If we use automata without their minimization, the reservation of such units can be queried for an automaton state. The instruction scheduler never queries reservation of functional units for given automaton state. So as a rule, you don't need this construction. This construction could be used for future code generation goals (e.g. to generate VLIW insn templates).

```
(define_query_cpu_unit unit-names [automaton-name])
```

unit-names is a string giving names of the functional units separated by commas.

automaton-name is a string giving the name of the automaton with which the unit is bound.

The following construction is the major one to describe pipeline characteristics of an instruction.

```
(define_insn_reservation insn-name default_latency
                        condition regexp)
```

default_latency is a number giving latency time of the instruction. There is an important difference between the old description and the automaton based pipeline description. The latency time is used for all dependencies when we use the old description. In the automaton based pipeline description, the given latency time is only used for true dependencies.

The cost of anti-dependencies is always zero and the cost of output dependencies is the difference between latency times of the producing and consuming insns (if the difference is negative, the cost is considered to be zero). You can always change the default costs for any description by using the target hook `TARGET_SCHED_ADJUST_COST` (see Section 10.17 [Scheduling], page 271).

insn-name is a string giving the internal name of the insn. The internal names are used in constructions `define_bypass` and in the automaton description file generated for debugging. The internal name has nothing in common with the names in `define_insn`. It is a good practice to use insn classes described in the processor manual.

condition defines what RTL insns are described by this construction. You should remember that you will be in trouble if *condition* for two or more different `define_insn_reservation` constructions is TRUE for an insn. In this case what reservation will be used for the insn is not defined. Such cases are not checked during generation of the pipeline hazards recognizer because in general recognizing that two conditions may have the same value is quite difficult (especially if the conditions contain `symbol_ref`). It is also not checked during the pipeline hazard recognizer work because it would slow down the recognizer considerably.

regexp is a string describing the reservation of the cpu's functional units by the instruction. The reservations are described by a regular expression according to the following syntax:

```

regexp = regexp "," oneof
        | oneof

oneof = oneof "|" allof
        | allof

allof = allof "+" repeat
        | repeat

repeat = element "*" number
        | element

element = cpu_function_unit_name
         | reservation_name
         | result_name
         | "nothing"
         | "(" regexp ")"

```

- ‘,’ is used for describing the start of the next cycle in the reservation.
- ‘|’ is used for describing a reservation described by the first regular expression **or** a reservation described by the second regular expression **or** etc.
- ‘+’ is used for describing a reservation described by the first regular expression **and** a reservation described by the second regular expression **and** etc.
- ‘*’ is used for convenience and simply means a sequence in which the regular expression are repeated *number* times with cycle advancing (see ‘,’).
- ‘cpu_function_unit_name’ denotes reservation of the named functional unit.
- ‘reservation_name’ — see description of construction ‘define_reservation’.
- ‘nothing’ denotes no unit reservations.

Sometimes unit reservations for different insns contain common parts. In such case, you can simplify the pipeline description by describing the common part by the following construction

```
(define_reservation reservation-name regexp)
```

reservation-name is a string giving name of *regexp*. Functional unit names and reservation names are in the same name space. So the reservation names should be different from the functional unit names and can not be the reserved name ‘**nothing**’.

The following construction is used to describe exceptions in the latency time for given instruction pair. This is so called bypasses.

```
(define_bypass number out_insn_names in_insn_names
 [guard])
```

number defines when the result generated by the instructions given in string *out_insn_names* will be ready for the instructions given in string *in_insn_names*. The instructions in the string are separated by commas.

guard is an optional string giving the name of a C function which defines an additional guard for the bypass. The function will get the two insns as parameters. If the function returns zero the bypass will be ignored for this case. The additional guard is necessary to recognize complicated bypasses, e.g. when the consumer is only an address of insn ‘**store**’ (not a stored value).

Usually the following three constructions are used to describe VLIW processors (more correctly to describe a placement of small insns into VLIW insn slots). Although they can be used for RISC processors too.

```
(exclusion_set unit-names unit-names)
(presence_set unit-names unit-names)
(absence_set unit-names unit-names)
```

unit-names is a string giving names of functional units separated by commas.

The first construction (‘**exclusion_set**’) means that each functional unit in the first string can not be reserved simultaneously with a unit whose name is in the second string and vice versa. For example, the construction is useful for describing processors (e.g. some SPARC processors) with a fully pipelined floating point functional unit which can execute simultaneously only single floating point insns or only double floating point insns.

The second construction (‘**presence_set**’) means that each functional unit in the first string can not be reserved unless at least one of units whose names are in the second string is reserved. This is an asymmetric relation. For example, it is useful for description that VLIW ‘**slot1**’ is reserved after ‘**slot0**’ reservation.

The third construction (‘**absence_set**’) means that each functional unit in the first string can be reserved only if each unit whose name is in the second string is not reserved. This is an asymmetric relation (actually ‘**exclusion_set**’ is analogous to this one but it is symmetric). For example, it is useful for description that VLIW ‘**slot0**’ can not be reserved after ‘**slot1**’ or ‘**slot2**’ reservation.

All functional units mentioned in a set should belong to the same automaton.

You can control the generator of the pipeline hazard recognizer with the following construction.

```
(automata_option options)
```

options is a string giving options which affect the generated code. Currently there are the following options:

- *no-minimization* makes no minimization of the automaton. This is only worth to do when we are debugging the description and need to look more accurately at reservations of states.
- *time* means printing additional time statistics about generation of automata.
- *v* means a generation of the file describing the result automata. The file has suffix `.dfa` and can be used for the description verification and debugging.
- *w* means a generation of warning instead of error for non-critical errors.
- *ndfa* makes nondeterministic finite state automata. This affects the treatment of operator `|` in the regular expressions. The usual treatment of the operator is to try the first alternative and, if the reservation is not possible, the second alternative. The non-deterministic treatment means trying all alternatives, some of them may be rejected by reservations in the subsequent insns. You can not query functional unit reservations in nondeterministic automaton states.

As an example, consider a superscalar RISC machine which can issue three insns (two integer insns and one floating point insn) on the cycle but can finish only two insns. To describe this, we define the following functional units.

```
(define_cpu_unit "i0_pipeline, i1_pipeline, f_pipeline")
(define_cpu_unit "port0, port1")
```

All simple integer insns can be executed in any integer pipeline and their result is ready in two cycles. The simple integer insns are issued into the first pipeline unless it is reserved, otherwise they are issued into the second pipeline. Integer division and multiplication insns can be executed only in the second integer pipeline and their results are ready correspondingly in 8 and 4 cycles. The integer division is not pipelined, i.e. the subsequent integer division insn can not be issued until the current division insn finished. Floating point insns are fully pipelined and their results are ready in 3 cycles. Where the result of a floating point insn is used by an integer insn, an additional delay of one cycle is incurred. To describe all of this we could specify

```
(define_cpu_unit "div")

(define_insn_reservation "simple" 2 (eq_attr "type" "int")
  "(i0_pipeline | i1_pipeline), (port0 | port1)")

(define_insn_reservation "mult" 4 (eq_attr "type" "mult")
  "i1_pipeline, nothing*2, (port0 | port1)")

(define_insn_reservation "div" 8 (eq_attr "type" "div")
  "i1_pipeline, div*7, div + (port0 | port1)")

(define_insn_reservation "float" 3 (eq_attr "type" "float")
  "f_pipeline, nothing, (port0 | port1)")

(define_bypass 4 "float" "simple,mult,div")
```

To simplify the description we could describe the following reservation

```
(define_reservation "finish" "port0|port1")
```

and use it in all `define_insn_reservation` as in the following construction

```
(define_insn_reservation "simple" 2 (eq_attr "type" "int")
  "(i0_pipeline | i1_pipeline), finish")
```

9.18.8.3 Drawbacks of the old pipeline description

The old instruction level parallelism description and the pipeline hazards recognizer based on it have the following drawbacks in comparison with the DFA-based ones:

- Each functional unit is believed to be reserved at the instruction execution start. This is a very inaccurate model for modern processors.
- An inadequate description of instruction latency times. The latency time is bound with a functional unit reserved by an instruction not with the instruction itself. In other words, the description is oriented to describe at most one unit reservation by each instruction. It also does not permit to describe special bypasses between instruction pairs.
- The implementation of the pipeline hazard recognizer interface has constraints on number of functional units. This is a number of bits in integer on the host machine.
- The interface to the pipeline hazard recognizer is more complex than one to the automaton based pipeline recognizer.
- An unnatural description when you write a unit and a condition which selects instructions using the unit. Writing all unit reservations for an instruction (an instruction class) is more natural.
- The recognition of the interlock delays has a slow implementation. The GCC scheduler supports structures which describe the unit reservations. The more functional units a processor has, the slower its pipeline hazard recognizer will be. Such an implementation would become even slower when we allowed to reserve functional units not only at the instruction execution start. In an automaton based pipeline hazard recognizer, speed is not dependent on processor complexity.

9.19 Conditional Execution

A number of architectures provide for some form of conditional execution, or predication. The hallmark of this feature is the ability to nullify most of the instructions in the instruction set. When the instruction set is large and not entirely symmetric, it can be quite tedious to describe these forms directly in the `.md` file. An alternative is the `define_cond_exec` template.

```
(define_cond_exec
  [predicate-pattern]
  "condition"
  "output-template")
```

predicate-pattern is the condition that must be true for the insn to be executed at runtime and should match a relational operator. One can use `match_operator` to match several relational operators at once. Any `match_operand` operands must have no more than one alternative.

condition is a C expression that must be true for the generated pattern to match.

output-template is a string similar to the `define_insn` output template (see Section 9.5 [Output Template], page 123), except that the ‘*’ and ‘@’ special cases do not apply. This is only useful if the assembly text for the predicate is a simple prefix to the main insn. In order to handle the general case, there is a global variable `current_insn_predicate` that will contain the entire predicate if the current insn is predicated, and will otherwise be NULL.

When `define_cond_exec` is used, an implicit reference to the `predicable` instruction attribute is made. See Section 9.18 [Insn Attributes], page 175. This attribute must be boolean (i.e. have exactly two elements in its *list-of-values*). Further, it must not be used with complex expressions. That is, the default and all uses in the insns must be a simple constant, not dependent on the alternative or anything else.

For each `define_insn` for which the `predicable` attribute is true, a new `define_insn` pattern will be generated that matches a predicated version of the instruction. For example,

```
(define_insn "addsi"
  [(set (match_operand:SI 0 "register_operand" "r")
        (plus:SI (match_operand:SI 1 "register_operand" "r")
                 (match_operand:SI 2 "register_operand" "r")))]
  "test1"
  "add %2,%1,%0")

(define_cond_exec
  [(ne (match_operand:CC 0 "register_operand" "c")
        (const_int 0))]
  "test2"
  "(%0)")
```

generates a new pattern

```
(define_insn ""
  [(cond_exec
    (ne (match_operand:CC 3 "register_operand" "c") (const_int 0))
    (set (match_operand:SI 0 "register_operand" "r")
        (plus:SI (match_operand:SI 1 "register_operand" "r")
                 (match_operand:SI 2 "register_operand" "r"))))]
  "(test2) && (test1)"
  "(%3) add %2,%1,%0")
```

9.20 Constant Definitions

Using literal constants inside instruction patterns reduces legibility and can be a maintenance problem.

To overcome this problem, you may use the `define_constants` expression. It contains a vector of name-value pairs. From that point on, wherever any of the names appears in the MD file, it is as if the corresponding value had been written instead. You may use `define_constants` multiple times; each appearance adds more constants to the table. It is an error to redefine a constant with a different value.

To come back to the a29k load multiple example, instead of

```
(define_insn ""
  [(match_parallel 0 "load_multiple_operation"
    [(set (match_operand:SI 1 "gpc_reg_operand" "=r")
          (match_operand:SI 2 "memory_operand" "m"))
     (use (reg:SI 179))
     (clobber (reg:SI 179))]]]
```

```
""
"loadm 0,0,%1,%2")
```

You could write:

```
(define_constants [
  (R_BP 177)
  (R_FC 178)
  (R_CR 179)
  (R_Q 180)
])

(define_insn ""
  [(match_parallel 0 "load_multiple_operation"
    [(set (match_operand:SI 1 "gpc_reg_operand" "=r")
          (match_operand:SI 2 "memory_operand" "m"))
     (use (reg:SI R_CR))
     (clobber (reg:SI R_CR))]])
  ""
  "loadm 0,0,%1,%2")
```

The constants that are defined with a `define_constant` are also output in the `insn-codes.h` header file as `#defines`.

10 Target Description Macros and Functions

In addition to the file `machine.md`, a machine description includes a C header file conventionally given the name `machine.h` and a C source file named `machine.c`. The header file defines numerous macros that convey the information about the target machine that does not fit into the scheme of the `.md` file. The file `tm.h` should be a link to `machine.h`. The header file `config.h` includes `tm.h` and most compiler source files include `config.h`. The source file defines a variable `targetm`, which is a structure containing pointers to functions and data relating to the target machine. `machine.c` should also contain their definitions, if they are not defined elsewhere in GCC, and other functions called through the macros defined in the `.h` file.

10.1 The Global `targetm` Variable

`struct gcc_target targetm` [Variable]
 The target `.c` file must define the global `targetm` variable which contains pointers to functions and data relating to the target machine. The variable is declared in `target.h`; `target-def.h` defines the macro `TARGET_INITIALIZER` which is used to initialize the variable, and macros for the default initializers for elements of the structure. The `.c` file should override those macros for which the default definition is inappropriate. For example:

```
#include "target.h"
#include "target-def.h"

/* Initialize the GCC target structure.  */

#undef TARGET_COMP_TYPE_ATTRIBUTES
#define TARGET_COMP_TYPE_ATTRIBUTES machine_comp_type_attributes

struct gcc_target targetm = TARGET_INITIALIZER;
```

Where a macro should be defined in the `.c` file in this manner to form part of the `targetm` structure, it is documented below as a “Target Hook” with a prototype. Many macros will change in future from being defined in the `.h` file to being part of the `targetm` structure.

10.2 Controlling the Compilation Driver, ‘gcc’

You can control the compilation driver.

`SWITCH_TAKES_ARG` (*char*)

A C expression which determines whether the option `-char` takes arguments. The value should be the number of arguments that option takes—zero, for many options.

By default, this macro is defined as `DEFAULT_SWITCH_TAKES_ARG`, which handles the standard options properly. You need not define `SWITCH_TAKES_ARG` unless you wish to add additional options which take arguments. Any redefinition should call `DEFAULT_SWITCH_TAKES_ARG` and then check for additional options.

WORD_SWITCH_TAKES_ARG (*name*)

A C expression which determines whether the option ‘*-name*’ takes arguments. The value should be the number of arguments that option takes—zero, for many options. This macro rather than `SWITCH_TAKES_ARG` is used for multi-character option names.

By default, this macro is defined as `DEFAULT_WORD_SWITCH_TAKES_ARG`, which handles the standard options properly. You need not define `WORD_SWITCH_TAKES_ARG` unless you wish to add additional options which take arguments. Any redefinition should call `DEFAULT_WORD_SWITCH_TAKES_ARG` and then check for additional options.

SWITCH_CURTAILS_COMPILATION (*char*)

A C expression which determines whether the option ‘*-char*’ stops compilation before the generation of an executable. The value is boolean, nonzero if the option does stop an executable from being generated, zero otherwise.

By default, this macro is defined as `DEFAULT_SWITCH_CURTAILS_COMPILATION`, which handles the standard options properly. You need not define `SWITCH_CURTAILS_COMPILATION` unless you wish to add additional options which affect the generation of an executable. Any redefinition should call `DEFAULT_SWITCH_CURTAILS_COMPILATION` and then check for additional options.

SWITCHES_NEED_SPACES

A string-valued C expression which enumerates the options for which the linker needs a space between the option and its argument.

If this macro is not defined, the default value is "".

TARGET_OPTION_TRANSLATE_TABLE

If defined, a list of pairs of strings, the first of which is a potential command line target to the ‘gcc’ driver program, and the second of which is a space-separated (tabs and other whitespace are not supported) list of options with which to replace the first option. The target defining this list is responsible for assuring that the results are valid. Replacement options may not be the `--opt` style, they must be the `-opt` style. It is the intention of this macro to provide a mechanism for substitution that affects the multilibs chosen, such as one option that enables many options, some of which select multilibs. Example nonsensical definition, where `-malt-abi`, `-EB`, and `-mspoo` cause different multilibs to be chosen:

```
#define TARGET_OPTION_TRANSLATE_TABLE \
{ "-fast",  "-march=fast-foo -malt-abi -I/usr/fast-foo" }, \
{ "-compat", "-EB -malign=4 -mspoo" }
```

DRIVER_SELF_SPECS

A list of specs for the driver itself. It should be a suitable initializer for an array of strings, with no surrounding braces.

The driver applies these specs to its own command line before choosing the multilib directory or running any subcommands. It applies them in the order given, so each spec can depend on the options added by earlier ones. It is also possible to remove options using ‘`%<option>`’ in the usual way.

This macro can be useful when a port has several interdependent target options. It provides a way of standardizing the command line so that the other specs are easier to write.

Do not define this macro if it does not need to do anything.

CPP_SPEC A C string constant that tells the GCC driver program options to pass to CPP. It can also specify how to translate options you give to GCC into options for GCC to pass to the CPP.

Do not define this macro if it does not need to do anything.

CPLUSPLUS_CPP_SPEC

This macro is just like **CPP_SPEC**, but is used for C++, rather than C. If you do not define this macro, then the value of **CPP_SPEC** (if any) will be used instead.

CC1_SPEC A C string constant that tells the GCC driver program options to pass to **cc1**, **cc1plus**, **f771**, and the other language front ends. It can also specify how to translate options you give to GCC into options for GCC to pass to front ends.

Do not define this macro if it does not need to do anything.

CC1PLUS_SPEC

A C string constant that tells the GCC driver program options to pass to **cc1plus**. It can also specify how to translate options you give to GCC into options for GCC to pass to the **cc1plus**.

Do not define this macro if it does not need to do anything. Note that everything defined in **CC1_SPEC** is already passed to **cc1plus** so there is no need to duplicate the contents of **CC1_SPEC** in **CC1PLUS_SPEC**.

ASM_SPEC A C string constant that tells the GCC driver program options to pass to the assembler. It can also specify how to translate options you give to GCC into options for GCC to pass to the assembler. See the file ‘**sun3.h**’ for an example of this.

Do not define this macro if it does not need to do anything.

ASM_FINAL_SPEC

A C string constant that tells the GCC driver program how to run any programs which cleanup after the normal assembler. Normally, this is not needed. See the file ‘**mips.h**’ for an example of this.

Do not define this macro if it does not need to do anything.

LINK_SPEC

A C string constant that tells the GCC driver program options to pass to the linker. It can also specify how to translate options you give to GCC into options for GCC to pass to the linker.

Do not define this macro if it does not need to do anything.

LIB_SPEC Another C string constant used much like **LINK_SPEC**. The difference between the two is that **LIB_SPEC** is used at the end of the command given to the linker. If this macro is not defined, a default is provided that loads the standard C library from the usual place. See ‘**gcc.c**’.

LIBGCC_SPEC

Another C string constant that tells the GCC driver program how and when to place a reference to `'libgcc.a'` into the linker command line. This constant is placed both before and after the value of `LIB_SPEC`.

If this macro is not defined, the GCC driver provides a default that passes the string `'-lgcc'` to the linker.

STARTFILE_SPEC

Another C string constant used much like `LINK_SPEC`. The difference between the two is that `STARTFILE_SPEC` is used at the very beginning of the command given to the linker.

If this macro is not defined, a default is provided that loads the standard C startup file from the usual place. See `'gcc.c'`.

ENDFILE_SPEC

Another C string constant used much like `LINK_SPEC`. The difference between the two is that `ENDFILE_SPEC` is used at the very end of the command given to the linker.

Do not define this macro if it does not need to do anything.

THREAD_MODEL_SPEC

`GCC -v` will print the thread model GCC was configured to use. However, this doesn't work on platforms that are multilibbed on thread models, such as AIX 4.3. On such platforms, define `THREAD_MODEL_SPEC` such that it evaluates to a string without blanks that names one of the recognized thread models. `%*`, the default value of this macro, will expand to the value of `thread_file` set in `'config.gcc'`.

EXTRA_SPECS

Define this macro to provide additional specifications to put in the `'specs'` file that can be used in various specifications like `CC1_SPEC`.

The definition should be an initializer for an array of structures, containing a string constant, that defines the specification name, and a string constant that provides the specification.

Do not define this macro if it does not need to do anything.

`EXTRA_SPECS` is useful when an architecture contains several related targets, which have various `..._SPECS` which are similar to each other, and the maintainer would like one central place to keep these definitions.

For example, the PowerPC System V.4 targets use `EXTRA_SPECS` to define either `_CALL_SYSV` when the System V calling sequence is used or `_CALL_AIX` when the older AIX-based calling sequence is used.

The `'config/rs6000/rs6000.h'` target file defines:

```
#define EXTRA_SPECS \
  { "cpp_sysv_default", CPP_SYSV_DEFAULT },

#define CPP_SYS_DEFAULT ""
```

The `'config/rs6000/sysv.h'` target file defines:

```

#undef CPP_SPEC
#define CPP_SPEC \
"%{posix: -D_POSIX_SOURCE } \
{%mcall-sysv: -D_CALL_SYSV } {%mcall-aix: -D_CALL_AIX } \
{%!mcall-sysv: {%!mcall-aix: %(cpp_sysv_default) }} \
{%msoft-float: -D_SOFT_FLOAT} {%mcpu=403: -D_SOFT_FLOAT}"

#undef CPP_SYSV_DEFAULT
#define CPP_SYSV_DEFAULT "-D_CALL_SYSV"

```

while the ‘config/rs6000/eabiaix.h’ target file defines CPP_SYSV_DEFAULT as:

```

#undef CPP_SYSV_DEFAULT
#define CPP_SYSV_DEFAULT "-D_CALL_AIX"

```

LINK_LIBGCC_SPECIAL

Define this macro if the driver program should find the library ‘libgcc.a’ itself and should not pass ‘-L’ options to the linker. If you do not define this macro, the driver program will pass the argument ‘-lgcc’ to tell the linker to do the search and will pass ‘-L’ options to it.

LINK_LIBGCC_SPECIAL_1

Define this macro if the driver program should find the library ‘libgcc.a’. If you do not define this macro, the driver program will pass the argument ‘-lgcc’ to tell the linker to do the search. This macro is similar to LINK_LIBGCC_SPECIAL, except that it does not affect ‘-L’ options.

LINK_GCC_C_SEQUENCE_SPEC

The sequence in which libgcc and libc are specified to the linker. By default this is %G %L %G.

LINK_COMMAND_SPEC

A C string constant giving the complete command line need to execute the linker. When you do this, you will need to update your port each time a change is made to the link command line within ‘gcc.c’. Therefore, define this macro only if you need to completely redefine the command line for invoking the linker and there is no other way to accomplish the effect you need. Overriding this macro may be avoidable by overriding LINK_GCC_C_SEQUENCE_SPEC instead.

LINK_ELIMINATE_DUPLICATE_LDIRECTORIES

A nonzero value causes collect2 to remove duplicate ‘-Ldirectory’ search directories from linking commands. Do not give it a nonzero value if removing duplicate search directories changes the linker’s semantics.

MULTILIB_DEFAULTS

Define this macro as a C expression for the initializer of an array of string to tell the driver program which options are defaults for this target and thus do not need to be handled specially when using MULTILIB_OPTIONS.

Do not define this macro if MULTILIB_OPTIONS is not defined in the target makefile fragment or if none of the options listed in MULTILIB_OPTIONS are set by default. See Section 12.1 [Target Fragment], page 331.

RELATIVE_PREFIX_NOT_LINKDIR

Define this macro to tell `gcc` that it should only translate a `'-B'` prefix into a `'-L'` linker option if the prefix indicates an absolute file name.

STANDARD_EXEC_PREFIX

Define this macro as a C string constant if you wish to override the standard choice of `'/usr/local/lib/gcc-lib/'` as the default prefix to try when searching for the executable files of the compiler.

MD_EXEC_PREFIX

If defined, this macro is an additional prefix to try after `STANDARD_EXEC_PREFIX`. `MD_EXEC_PREFIX` is not searched when the `'-b'` option is used, or the compiler is built as a cross compiler. If you define `MD_EXEC_PREFIX`, then be sure to add it to the list of directories used to find the assembler in `'configure.in'`.

STANDARD_STARTFILE_PREFIX

Define this macro as a C string constant if you wish to override the standard choice of `'/usr/local/lib/'` as the default prefix to try when searching for startup files such as `'crt0.o'`.

MD_STARTFILE_PREFIX

If defined, this macro supplies an additional prefix to try after the standard prefixes. `MD_EXEC_PREFIX` is not searched when the `'-b'` option is used, or when the compiler is built as a cross compiler.

MD_STARTFILE_PREFIX_1

If defined, this macro supplies yet another prefix to try after the standard prefixes. It is not searched when the `'-b'` option is used, or when the compiler is built as a cross compiler.

INIT_ENVIRONMENT

Define this macro as a C string constant if you wish to set environment variables for programs called by the driver, such as the assembler and loader. The driver passes the value of this macro to `putenv` to initialize the necessary environment variables.

LOCAL_INCLUDE_DIR

Define this macro as a C string constant if you wish to override the standard choice of `'/usr/local/include'` as the default prefix to try when searching for local header files. `LOCAL_INCLUDE_DIR` comes before `SYSTEM_INCLUDE_DIR` in the search order.

Cross compilers do not search either `'/usr/local/include'` or its replacement.

MODIFY_TARGET_NAME

Define this macro if you wish to define command-line switches that modify the default target name

For each switch, you can include a string to be appended to the first part of the configuration name or a string to be deleted from the configuration name, if present. The definition should be an initializer for an array of structures. Each array element should have three elements: the switch name (a string constant, including the initial dash), one of the enumeration codes `ADD` or `DELETE` to

indicate whether the string should be inserted or deleted, and the string to be inserted or deleted (a string constant).

For example, on a machine where ‘64’ at the end of the configuration name denotes a 64-bit target and you want the ‘-32’ and ‘-64’ switches to select between 32- and 64-bit targets, you would code

```
#define MODIFY_TARGET_NAME \
  { { "-32", DELETE, "64"}, \
    {"-64", ADD, "64"} }
```

SYSTEM_INCLUDE_DIR

Define this macro as a C string constant if you wish to specify a system-specific directory to search for header files before the standard directory. `SYSTEM_INCLUDE_DIR` comes before `STANDARD_INCLUDE_DIR` in the search order.

Cross compilers do not use this macro and do not search the directory specified.

STANDARD_INCLUDE_DIR

Define this macro as a C string constant if you wish to override the standard choice of ‘`/usr/include`’ as the default prefix to try when searching for header files.

Cross compilers do not use this macro and do not search either ‘`/usr/include`’ or its replacement.

STANDARD_INCLUDE_COMPONENT

The “component” corresponding to `STANDARD_INCLUDE_DIR`. See `INCLUDE_DEFAULTS`, below, for the description of components. If you do not define this macro, no component is used.

INCLUDE_DEFAULTS

Define this macro if you wish to override the entire default search path for include files. For a native compiler, the default search path usually consists of `GCC_INCLUDE_DIR`, `LOCAL_INCLUDE_DIR`, `SYSTEM_INCLUDE_DIR`, `GPLUSPLUS_INCLUDE_DIR`, and `STANDARD_INCLUDE_DIR`. In addition, `GPLUSPLUS_INCLUDE_DIR` and `GCC_INCLUDE_DIR` are defined automatically by ‘`Makefile`’, and specify private search areas for GCC. The directory `GPLUSPLUS_INCLUDE_DIR` is used only for C++ programs.

The definition should be an initializer for an array of structures. Each array element should have four elements: the directory name (a string constant), the component name (also a string constant), a flag for C++-only directories, and a flag showing that the includes in the directory don’t need to be wrapped in `extern ‘C’` when compiling C++. Mark the end of the array with a null element.

The component name denotes what GNU package the include file is part of, if any, in all upper-case letters. For example, it might be ‘`GCC`’ or ‘`BINUTILS`’. If the package is part of a vendor-supplied operating system, code the component name as ‘`0`’.

For example, here is the definition used for VAX/VMS:

```
#define INCLUDE_DEFAULTS \
{                                     \
```

```

        { "GNU_GXX_INCLUDE:", "G++", 1, 1},    \
        { "GNU_CC_INCLUDE:", "GCC", 0, 0},    \
        { "SYS$SYSROOT:[SYSLIB.]", 0, 0, 0}, \
        { ".", 0, 0, 0},                    \
        { 0, 0, 0, 0}                       \
    }

```

Here is the order of prefixes tried for exec files:

1. Any prefixes specified by the user with ‘-B’.
2. The environment variable `GCC_EXEC_PREFIX`, if any.
3. The directories specified by the environment variable `COMPILER_PATH`.
4. The macro `STANDARD_EXEC_PREFIX`.
5. ‘`/usr/lib/gcc/`’.
6. The macro `MD_EXEC_PREFIX`, if any.

Here is the order of prefixes tried for startfiles:

1. Any prefixes specified by the user with ‘-B’.
2. The environment variable `GCC_EXEC_PREFIX`, if any.
3. The directories specified by the environment variable `LIBRARY_PATH` (or port-specific name; native only, cross compilers do not use this).
4. The macro `STANDARD_EXEC_PREFIX`.
5. ‘`/usr/lib/gcc/`’.
6. The macro `MD_EXEC_PREFIX`, if any.
7. The macro `MD_STARTFILE_PREFIX`, if any.
8. The macro `STANDARD_STARTFILE_PREFIX`.
9. ‘`/lib/`’.
10. ‘`/usr/lib/`’.

10.3 Run-time Target Specification

Here are run-time target specifications.

`TARGET_CPU_CPP_BUILTINS()`

This function-like macro expands to a block of code that defines built-in preprocessor macros and assertions for the target cpu, using the functions `builtin_define`, `builtin_define_std` and `builtin_assert` defined in ‘`c-common.c`’. When the front end calls this macro it provides a trailing semicolon, and since it has finished command line option processing your code can use those results freely.

`builtin_assert` takes a string in the form you pass to the command-line option ‘-A’, such as `cpu=mips`, and creates the assertion. `builtin_define` takes a string in the form accepted by option ‘-D’ and unconditionally defines the macro. `builtin_define_std` takes a string representing the name of an object-like macro. If it doesn’t lie in the user’s namespace, `builtin_define_std` defines

it unconditionally. Otherwise, it defines a version with two leading underscores, and another version with two leading and trailing underscores, and defines the original only if an ISO standard was not requested on the command line. For example, passing `unix` defines `__unix`, `__unix__` and possibly `unix`; passing `_mips` defines `__mips`, `__mips__` and possibly `_mips`, and passing `_ABI64` defines only `_ABI64`.

You can also test for the C dialect being compiled. The variable `c_language` is set to one of `clk_c`, `clk_cplusplus` or `clk_objective_c`. Note that if we are preprocessing assembler, this variable will be `clk_c` but the function-like macro `preprocessing_asm_p()` will return true, so you might want to check for that first. If you need to check for strict ANSI, the variable `flag_iso` can be used. The function-like macro `preprocessing_trad_p()` can be used to check for traditional preprocessing.

With `TARGET_OS_CPP_BUILTINS` this macro obsoletes the `CPP_PREDEFINES` target macro.

`TARGET_OS_CPP_BUILTINS()`

Similarly to `TARGET_CPU_CPP_BUILTINS` but this macro is optional and is used for the target operating system instead.

With `TARGET_CPU_CPP_BUILTINS` this macro obsoletes the `CPP_PREDEFINES` target macro.

`CPP_PREDEFINES`

Define this to be a string constant containing ‘-D’ options to define the predefined macros that identify this machine and system. These macros will be predefined unless the ‘-ansi’ option (or a ‘-std’ option for strict ISO C conformance) is specified.

In addition, a parallel set of macros are predefined, whose names are made by appending ‘__’ at the beginning and at the end. These ‘__’ macros are permitted by the ISO standard, so they are predefined regardless of whether ‘-ansi’ or a ‘-std’ option is specified.

For example, on the Sun, one can use the following value:

```
"-Dmc68000 -Dsun -Dunix"
```

The result is to define the macros `__mc68000__`, `__sun__` and `__unix__` unconditionally, and the macros `mc68000`, `sun` and `unix` provided ‘-ansi’ is not specified.

```
extern int target_flags;
```

This declaration should be present.

`TARGET_...`

This series of macros is to allow compiler command arguments to enable or disable the use of optional features of the target machine. For example, one machine description serves both the 68000 and the 68020; a command argument tells the compiler whether it should use 68020-only instructions or not. This command argument works by means of a macro `TARGET_68020` that tests a bit in `target_flags`.

Define a macro `TARGET_featurename` for each such option. Its definition should test a bit in `target_flags`. It is recommended that a helper macro `TARGET_MASK_featurename` is defined for each bit-value to test, and used in `TARGET_featurename` and `TARGET_SWITCHES`. For example:

```
#define TARGET_MASK_68020 1
#define TARGET_68020 (target_flags & TARGET_MASK_68020)
```

One place where these macros are used is in the condition-expressions of instruction patterns. Note how `TARGET_68020` appears frequently in the 68000 machine description file, ‘`m68k.md`’. Another place they are used is in the definitions of the other macros in the ‘`machine.h`’ file.

TARGET_SWITCHES

This macro defines names of command options to set and clear bits in `target_flags`. Its definition is an initializer with a subgrouping for each command option.

Each subgrouping contains a string constant, that defines the option name, a number, which contains the bits to set in `target_flags`, and a second string which is the description displayed by ‘`--help`’. If the number is negative then the bits specified by the number are cleared instead of being set. If the description string is present but empty, then no help information will be displayed for that option, but it will not count as an undocumented option. The actual option name is made by appending ‘`-m`’ to the specified name. Non-empty description strings should be marked with `N(...)` for `xgettext`. Please do not mark empty strings because the empty string is reserved by GNU `gettext`. `gettext("")` returns the header entry of the message catalog with meta information, not the empty string.

In addition to the description for ‘`--help`’, more detailed documentation for each option should be added to ‘`invoke.texi`’.

One of the subgroupings should have a null string. The number in this grouping is the default value for `target_flags`. Any target options act starting with that value.

Here is an example which defines ‘`-m68000`’ and ‘`-m68020`’ with opposite meanings, and picks the latter as the default:

```
#define TARGET_SWITCHES \
  { { "68020", TARGET_MASK_68020, "" },          \
    { "68000", -TARGET_MASK_68020, \
      N("Compile for the 68000") }, \
    { "", TARGET_MASK_68020, "" } }
```

TARGET_OPTIONS

This macro is similar to `TARGET_SWITCHES` but defines names of command options that have values. Its definition is an initializer with a subgrouping for each command option.

Each subgrouping contains a string constant, that defines the fixed part of the option name, the address of a variable, and a description string. Non-empty description strings should be marked with `N(...)` for `xgettext`. Please do not mark empty strings because the empty string is reserved by GNU `gettext`.

`gettext("")` returns the header entry of the message catalog with meta information, not the empty string.

The variable, type `char *`, is set to the variable part of the given option if the fixed part matches. The actual option name is made by appending `'-m'` to the specified name. Again, each option should also be documented in `'invoke.texi'`.

Here is an example which defines `'-mshort-data-number'`. If the given option is `'-mshort-data-512'`, the variable `m88k_short_data` will be set to the string `"512"`.

```
extern char *m88k_short_data;
#define TARGET_OPTIONS \
  { { "short-data-", &m88k_short_data, \
      N_("Specify the size of the short data section") } }
```

TARGET_VERSION

This macro is a C statement to print on `stderr` a string describing the particular machine description choice. Every machine description should define `TARGET_VERSION`. For example:

```
#ifdef MOTOROLA
#define TARGET_VERSION \
  fprintf (stderr, " (68k, Motorola syntax)");
#else
#define TARGET_VERSION \
  fprintf (stderr, " (68k, MIT syntax)");
#endif
```

OVERRIDE_OPTIONS

Sometimes certain combinations of command options do not make sense on a particular target machine. You can define a macro `OVERRIDE_OPTIONS` to take account of this. This macro, if defined, is executed once just after all the command options have been parsed.

Don't use this macro to turn on various extra optimizations for `'-O'`. That is what `OPTIMIZATION_OPTIONS` is for.

OPTIMIZATION_OPTIONS (*level*, *size*)

Some machines may desire to change what optimizations are performed for various optimization levels. This macro, if defined, is executed once just after the optimization level is determined and before the remainder of the command options have been parsed. Values set in this macro are used as the default values for the other command line options.

level is the optimization level specified; 2 if `'-O2'` is specified, 1 if `'-O'` is specified, and 0 if neither is specified.

size is nonzero if `'-Os'` is specified and zero otherwise.

You should not use this macro to change options that are not machine-specific. These should uniformly selected by the same optimization level on all supported machines. Use this macro to enable machine-specific optimizations.

Do not examine `write_symbols` in this macro! The debugging options are not supposed to alter the generated code.

CAN_DEBUG_WITHOUT_FP

Define this macro if debugging can be performed even without a frame pointer. If this macro is defined, GCC will turn on the `'-fomit-frame-pointer'` option whenever `'-O'` is specified.

10.4 Defining data structures for per-function information.

If the target needs to store information on a per-function basis, GCC provides a macro and a couple of variables to allow this. Note, just using statics to store the information is a bad idea, since GCC supports nested functions, so you can be halfway through encoding one function when another one comes along.

GCC defines a data structure called `struct function` which contains all of the data specific to an individual function. This structure contains a field called `machine` whose type is `struct machine_function *`, which can be used by targets to point to their own specific data.

If a target needs per-function specific data it should define the type `struct machine_function` and also the macro `INIT_EXPANDERS`. This macro should be used to initialize the function pointer `init_machine_status`. This pointer is explained below.

One typical use of per-function, target specific data is to create an RTX to hold the register containing the function's return address. This RTX can then be used to implement the `__builtin_return_address` function, for level 0.

Note—earlier implementations of GCC used a single data area to hold all of the per-function information. Thus when processing of a nested function began the old per-function data had to be pushed onto a stack, and when the processing was finished, it had to be popped off the stack. GCC used to provide function pointers called `save_machine_status` and `restore_machine_status` to handle the saving and restoring of the target specific information. Since the single data area approach is no longer used, these pointers are no longer supported.

The macro and function pointers are described below.

INIT_EXPANDERS

Macro called to initialize any target specific information. This macro is called once per function, before generation of any RTL has begun. The intention of this macro is to allow the initialization of the function pointers below.

init_machine_status

This is a `void (*)(struct function *)` function pointer. If this pointer is non-NULL it will be called once per function, before function compilation starts, in order to allow the target to perform any target specific initialization of the `struct function` structure. It is intended that this would be used to initialize the `machine` of that structure.

`struct machine_function` structures are expected to be freed by GC. Generally, any memory that they reference must be allocated by using `ggc_alloc`, including the structure itself.

10.5 Storage Layout

Note that the definitions of the macros in this table which are sizes or alignments measured in bits do not need to be constant. They can be C expressions that refer to static variables, such as the `target_flags`. See Section 10.3 [Run-time Target], page 200.

`BITS_BIG_ENDIAN`

Define this macro to have the value 1 if the most significant bit in a byte has the lowest number; otherwise define it to have the value zero. This means that bit-field instructions count from the most significant bit. If the machine has no bit-field instructions, then this must still be defined, but it doesn't matter which value it is defined to. This macro need not be a constant.

This macro does not affect the way structure fields are packed into bytes or words; that is controlled by `BYTES_BIG_ENDIAN`.

`BYTES_BIG_ENDIAN`

Define this macro to have the value 1 if the most significant byte in a word has the lowest number. This macro need not be a constant.

`WORDS_BIG_ENDIAN`

Define this macro to have the value 1 if, in a multiword object, the most significant word has the lowest number. This applies to both memory locations and registers; GCC fundamentally assumes that the order of words in memory is the same as the order in registers. This macro need not be a constant.

`LIBGCC2_WORDS_BIG_ENDIAN`

Define this macro if `WORDS_BIG_ENDIAN` is not constant. This must be a constant value with the same meaning as `WORDS_BIG_ENDIAN`, which will be used only when compiling `'libgcc2.c'`. Typically the value will be set based on preprocessor defines.

`FLOAT_WORDS_BIG_ENDIAN`

Define this macro to have the value 1 if `DFmode`, `XFmode` or `TFmode` floating point numbers are stored in memory with the word containing the sign bit at the lowest address; otherwise define it to have the value 0. This macro need not be a constant.

You need not define this macro if the ordering is the same as for multi-word integers.

`BITS_PER_UNIT`

Define this macro to be the number of bits in an addressable storage unit (byte). If you do not define this macro the default is 8.

`BITS_PER_WORD`

Number of bits in a word. If you do not define this macro, the default is `BITS_PER_UNIT * UNITS_PER_WORD`.

`MAX_BITS_PER_WORD`

Maximum number of bits in a word. If this is undefined, the default is `BITS_PER_WORD`. Otherwise, it is the constant value that is the largest value that `BITS_PER_WORD` can have at run-time.

UNITS_PER_WORD

Number of storage units in a word; normally 4.

MIN_UNITS_PER_WORD

Minimum number of units in a word. If this is undefined, the default is **UNITS_PER_WORD**. Otherwise, it is the constant value that is the smallest value that **UNITS_PER_WORD** can have at run-time.

POINTER_SIZE

Width of a pointer, in bits. You must specify a value no wider than the width of **Pmode**. If it is not equal to the width of **Pmode**, you must define **POINTERS_EXTEND_UNSIGNED**. If you do not specify a value the default is **BITS_PER_WORD**.

POINTERS_EXTEND_UNSIGNED

A C expression whose value is greater than zero if pointers that need to be extended from being **POINTER_SIZE** bits wide to **Pmode** are to be zero-extended and zero if they are to be sign-extended. If the value is less than zero then there must be an "ptr_extend" instruction that extends a pointer from **POINTER_SIZE** to **Pmode**.

You need not define this macro if the **POINTER_SIZE** is equal to the width of **Pmode**.

PROMOTE_MODE (*m*, *unsignedp*, *type*)

A macro to update *m* and *unsignedp* when an object whose type is *type* and which has the specified mode and signedness is to be stored in a register. This macro is only called when *type* is a scalar type.

On most RISC machines, which only have operations that operate on a full register, define this macro to set *m* to **word_mode** if *m* is an integer mode narrower than **BITS_PER_WORD**. In most cases, only integer modes should be widened because wider-precision floating-point operations are usually more expensive than their narrower counterparts.

For most machines, the macro definition does not change *unsignedp*. However, some machines, have instructions that preferentially handle either signed or unsigned quantities of certain modes. For example, on the DEC Alpha, 32-bit loads from memory and 32-bit add instructions sign-extend the result to 64 bits. On such machines, set *unsignedp* according to which kind of extension is more efficient.

Do not define this macro if it would never modify *m*.

PROMOTE_FUNCTION_ARGS

Define this macro if the promotion described by **PROMOTE_MODE** should also be done for outgoing function arguments.

PROMOTE_FUNCTION_RETURN

Define this macro if the promotion described by **PROMOTE_MODE** should also be done for the return value of functions.

If this macro is defined, **FUNCTION_VALUE** must perform the same promotions done by **PROMOTE_MODE**.

PROMOTE_FOR_CALL_ONLY

Define this macro if the promotion described by `PROMOTE_MODE` should *only* be performed for outgoing function arguments or function return values, as specified by `PROMOTE_FUNCTION_ARGS` and `PROMOTE_FUNCTION_RETURN`, respectively.

PARAM_BOUNDARY

Normal alignment required for function parameters on the stack, in bits. All stack parameters receive at least this much alignment regardless of data type. On most machines, this is the same as the size of an integer.

STACK_BOUNDARY

Define this macro to the minimum alignment enforced by hardware for the stack pointer on this machine. The definition is a C expression for the desired alignment (measured in bits). This value is used as a default if `PREFERRED_STACK_BOUNDARY` is not defined. On most machines, this should be the same as `PARAM_BOUNDARY`.

PREFERRED_STACK_BOUNDARY

Define this macro if you wish to preserve a certain alignment for the stack pointer, greater than what the hardware enforces. The definition is a C expression for the desired alignment (measured in bits). This macro must evaluate to a value equal to or larger than `STACK_BOUNDARY`.

FORCE_PREFERRED_STACK_BOUNDARY_IN_MAIN

A C expression that evaluates true if `PREFERRED_STACK_BOUNDARY` is not guaranteed by the runtime and we should emit code to align the stack at the beginning of `main`.

If `PUSH_ROUNDING` is not defined, the stack will always be aligned to the specified boundary. If `PUSH_ROUNDING` is defined and specifies a less strict alignment than `PREFERRED_STACK_BOUNDARY`, the stack may be momentarily unaligned while pushing arguments.

FUNCTION_BOUNDARY

Alignment required for a function entry point, in bits.

BIGGEST_ALIGNMENT

Biggest alignment that any data type can require on this machine, in bits.

MINIMUM_ATOMIC_ALIGNMENT

If defined, the smallest alignment, in bits, that can be given to an object that can be referenced in one operation, without disturbing any nearby object. Normally, this is `BITS_PER_UNIT`, but may be larger on machines that don't have byte or half-word store operations.

BIGGEST_FIELD_ALIGNMENT

Biggest alignment that any structure or union field can require on this machine, in bits. If defined, this overrides `BIGGEST_ALIGNMENT` for structure and union fields only, unless the field alignment has been set by the `__attribute__((aligned (n)))` construct.

ADJUST_FIELD_ALIGN (*field*, *computed*)

An expression for the alignment of a structure field *field* if the alignment computed in the usual way (including applying of `BIGGEST_ALIGNMENT` and

`BIGGEST_FIELD_ALIGNMENT` to the alignment) is *computed*. It overrides alignment only if the field alignment has not been set by the `__attribute__((aligned (n)))` construct.

`MAX_OFFILE_ALIGNMENT`

Biggest alignment supported by the object file format of this machine. Use this macro to limit the alignment which can be specified using the `__attribute__((aligned (n)))` construct. If not defined, the default value is `BIGGEST_ALIGNMENT`.

`DATA_ALIGNMENT (type, basic-align)`

If defined, a C expression to compute the alignment for a variable in the static store. *type* is the data type, and *basic-align* is the alignment that the object would ordinarily have. The value of this macro is used instead of that alignment to align the object.

If this macro is not defined, then *basic-align* is used.

One use of this macro is to increase alignment of medium-size data to make it all fit in fewer cache lines. Another is to cause character arrays to be word-aligned so that `strcpy` calls that copy constants to character arrays can be done inline.

`CONSTANT_ALIGNMENT (constant, basic-align)`

If defined, a C expression to compute the alignment given to a constant that is being placed in memory. *constant* is the constant and *basic-align* is the alignment that the object would ordinarily have. The value of this macro is used instead of that alignment to align the object.

If this macro is not defined, then *basic-align* is used.

The typical use of this macro is to increase alignment for string constants to be word aligned so that `strcpy` calls that copy constants can be done inline.

`LOCAL_ALIGNMENT (type, basic-align)`

If defined, a C expression to compute the alignment for a variable in the local store. *type* is the data type, and *basic-align* is the alignment that the object would ordinarily have. The value of this macro is used instead of that alignment to align the object.

If this macro is not defined, then *basic-align* is used.

One use of this macro is to increase alignment of medium-size data to make it all fit in fewer cache lines.

`EMPTY_FIELD_BOUNDARY`

Alignment in bits to be given to a structure bit-field that follows an empty field such as `int : 0;`.

Note that `PCC_BITFIELD_TYPE_MATTERS` also affects the alignment that results from an empty field.

`STRUCTURE_SIZE_BOUNDARY`

Number of bits which any structure or union's size must be a multiple of. Each structure or union's size is rounded up to a multiple of this.

If you do not define this macro, the default is the same as `BITS_PER_UNIT`.

STRICT_ALIGNMENT

Define this macro to be the value 1 if instructions will fail to work if given data not on the nominal alignment. If instructions will merely go slower in that case, define this macro as 0.

PCC_BITFIELD_TYPE_MATTERS

Define this if you wish to imitate the way many other C compilers handle alignment of bit-fields and the structures that contain them.

The behavior is that the type written for a bit-field (`int`, `short`, or other integer type) imposes an alignment for the entire structure, as if the structure really did contain an ordinary field of that type. In addition, the bit-field is placed within the structure so that it would fit within such a field, not crossing a boundary for it.

Thus, on most machines, a bit-field whose type is written as `int` would not cross a four-byte boundary, and would force four-byte alignment for the whole structure. (The alignment used may not be four bytes; it is controlled by the other alignment parameters.)

If the macro is defined, its definition should be a C expression; a nonzero value for the expression enables this behavior.

Note that if this macro is not defined, or its value is zero, some bit-fields may cross more than one alignment boundary. The compiler can support such references if there are `'insv'`, `'extv'`, and `'extzv'` insns that can directly reference memory.

The other known way of making bit-fields work is to define `STRUCTURE_SIZE_BOUNDARY` as large as `BIGGEST_ALIGNMENT`. Then every structure can be accessed with fullwords.

Unless the machine has bit-field instructions or you define `STRUCTURE_SIZE_BOUNDARY` that way, you must define `PCC_BITFIELD_TYPE_MATTERS` to have a nonzero value.

If your aim is to make GCC use the same conventions for laying out bit-fields as are used by another compiler, here is how to investigate what the other compiler does. Compile and run this program:

```

struct foo1
{
    char x;
    char :0;
    char y;
};

struct foo2
{
    char x;
    int :0;
    char y;
};

```

```

main ()
{
    printf ("Size of foo1 is %d\n",
           sizeof (struct foo1));
    printf ("Size of foo2 is %d\n",
           sizeof (struct foo2));
    exit (0);
}

```

If this prints 2 and 5, then the compiler's behavior is what you would get from `PCC_BITFIELD_TYPE_MATTERS`.

`BITFIELD_NBYTES_LIMITED`

Like `PCC_BITFIELD_TYPE_MATTERS` except that its effect is limited to aligning a bit-field within the structure.

`MEMBER_TYPE_FORCES_BLK` (*field*, *mode*)

Return 1 if a structure or array containing *field* should be accessed using `BLKMODE`.

If *field* is the only field in the structure, *mode* is its mode, otherwise *mode* is `VOIDmode`. *mode* is provided in the case where structures of one field would require the structure's mode to retain the field's mode.

Normally, this is not needed. See the file `'c4x.h'` for an example of how to use this macro to prevent a structure having a floating point field from being accessed in an integer mode.

`ROUND_TYPE_SIZE` (*type*, *computed*, *specified*)

Define this macro as an expression for the overall size of a type (given by *type* as a tree node) when the size computed in the usual way is *computed* and the alignment is *specified*.

The default is to round *computed* up to a multiple of *specified*.

`ROUND_TYPE_SIZE_UNIT` (*type*, *computed*, *specified*)

Similar to `ROUND_TYPE_SIZE`, but sizes and alignments are specified in units (bytes). If you define `ROUND_TYPE_SIZE`, you must also define this macro and they must be defined consistently with each other.

`ROUND_TYPE_ALIGN` (*type*, *computed*, *specified*)

Define this macro as an expression for the alignment of a type (given by *type* as a tree node) if the alignment computed in the usual way is *computed* and the alignment explicitly specified was *specified*.

The default is to use *specified* if it is larger; otherwise, use the smaller of *computed* and `BIGGEST_ALIGNMENT`

`MAX_FIXED_MODE_SIZE`

An integer expression for the size in bits of the largest integer machine mode that should actually be used. All integer machine modes of this size or smaller can be used for structures and unions with the appropriate sizes. If this macro is undefined, `GET_MODE_BITSIZE` (`DI mode`) is assumed.

VECTOR_MODE_SUPPORTED_P(*mode*)

Define this macro to be nonzero if the port is prepared to handle insns involving vector mode *mode*. At the very least, it must have move patterns for this mode.

STACK_SAVEAREA_MODE(*save_level*)

If defined, an expression of type `enum machine_mode` that specifies the mode of the save area operand of a `save_stack_level` named pattern (see Section 9.8 [Standard Names], page 144). *save_level* is one of `SAVE_BLOCK`, `SAVE_FUNCTION`, or `SAVE_NONLOCAL` and selects which of the three named patterns is having its mode specified.

You need not define this macro if it always returns `Pmode`. You would most commonly define this macro if the `save_stack_level` patterns need to support both a 32- and a 64-bit mode.

STACK_SIZE_MODE

If defined, an expression of type `enum machine_mode` that specifies the mode of the size increment operand of an `allocate_stack` named pattern (see Section 9.8 [Standard Names], page 144).

You need not define this macro if it always returns `word_mode`. You would most commonly define this macro if the `allocate_stack` pattern needs to support both a 32- and a 64-bit mode.

TARGET_FLOAT_FORMAT

A code distinguishing the floating point format of the target machine. There are five defined values:

IEEE_FLOAT_FORMAT

This code indicates IEEE floating point. It is the default; there is no need to define this macro when the format is IEEE.

VAX_FLOAT_FORMAT

This code indicates the “F float” (for `float`) and “D float” or “G float” formats (for `double`) used on the VAX and PDP-11.

IBM_FLOAT_FORMAT

This code indicates the format used on the IBM System/370.

C4X_FLOAT_FORMAT

This code indicates the format used on the TMS320C3x/C4x.

UNKNOWN_FLOAT_FORMAT

This code indicates any other format.

If any other formats are actually in use on supported machines, new codes should be defined for them.

The ordering of the component words of floating point values stored in memory is controlled by `FLOAT_WORDS_BIG_ENDIAN`.

MODE_HAS_NANS(*mode*)

When defined, this macro should be true if *mode* has a NaN representation. The compiler assumes that NaNs are not equal to anything (including themselves)

and that addition, subtraction, multiplication and division all return NaNs when one operand is NaN.

By default, this macro is true if *mode* is a floating-point mode and the target floating-point format is IEEE.

MODE_HAS_INFINITYS (*mode*)

This macro should be true if *mode* can represent infinity. At present, the compiler uses this macro to decide whether ‘*x* - *x*’ is always defined. By default, the macro is true when *mode* is a floating-point mode and the target format is IEEE.

MODE_HAS_SIGNED_ZEROS (*mode*)

True if *mode* distinguishes between positive and negative zero. The rules are expected to follow the IEEE standard:

- ‘*x* + *x*’ has the same sign as ‘*x*’.
- If the sum of two values with opposite sign is zero, the result is positive for all rounding modes except towards $-\infty$, for which it is negative.
- The sign of a product or quotient is negative when exactly one of the operands is negative.

The default definition is true if *mode* is a floating-point mode and the target format is IEEE.

MODE_HAS_SIGN_DEPENDENT_ROUNDING (*mode*)

If defined, this macro should be true for *mode* if it has at least one rounding mode in which ‘*x*’ and ‘-*x*’ can be rounded to numbers of different magnitude. Two such modes are towards $-\infty$ and towards $+\infty$.

The default definition of this macro is true if *mode* is a floating-point mode and the target format is IEEE.

ROUND_TOWARDS_ZERO

If defined, this macro should be true if the prevailing rounding mode is towards zero. A true value has the following effects:

- `MODE_HAS_SIGN_DEPENDENT_ROUNDING` will be false for all modes.
- ‘`libgcc.a`’s floating-point emulator will round towards zero rather than towards nearest.
- The compiler’s floating-point emulator will round towards zero after doing arithmetic, and when converting from the internal float format to the target format.

The macro does not affect the parsing of string literals. When the primary rounding mode is towards zero, library functions like `strtod` might still round towards nearest, and the compiler’s parser should behave like the target’s `strtod` where possible.

Not defining this macro is equivalent to returning zero.

LARGEST_EXPONENT_IS_NORMAL (*size*)

This macro should return true if floats with *size* bits do not have a NaN or infinity representation, but use the largest exponent for normal numbers instead.

Defining this macro to true for *size* causes `MODE_HAS_NANS` and `MODE_HAS_INFINITIES` to be false for *size*-bit modes. It also affects the way ‘`libgcc.a`’ and ‘`real.c`’ emulate floating-point arithmetic.

The default definition of this macro returns false for all sizes.

`bool TARGET_MS_BITFIELD_LAYOUT_P (tree record_type)` [Target Hook]

This target hook returns `true` if bit-fields in the given *record_type* are to be laid out following the rules of Microsoft Visual C/C++, namely: (i) a bit-field won’t share the same storage unit with the previous bit-field if their underlying types have different sizes, and the bit-field will be aligned to the highest alignment of the underlying types of itself and of the previous bit-field; (ii) a zero-sized bit-field will affect the alignment of the whole enclosing structure, even if it is unnamed; except that (iii) a zero-sized bit-field will be disregarded unless it follows another bit-field of nonzero size. If this hook returns `true`, other macros that control bit-field layout are ignored.

When a bit-field is inserted into a packed record, the whole size of the underlying type is used by one or more same-size adjacent bit-fields (that is, if its `long:3`, 32 bits is used in the record, and any additional adjacent long bit-fields are packed into the same chunk of 32 bits. However, if the size changes, a new field of that size is allocated). In an unpacked record, this is the same as using alignment, but not equivalent when packing.

If both MS bit-fields and ‘`__attribute__((packed))`’ are used, the latter will take precedence. If ‘`__attribute__((packed))`’ is used on a single field when MS bit-fields are in use, it will take precedence for that field, but the alignment of the rest of the structure may affect its placement.

10.6 Layout of Source Language Data Types

These macros define the sizes and other characteristics of the standard basic data types used in programs being compiled. Unlike the macros in the previous section, these apply to specific features of C and related languages, rather than to fundamental aspects of storage layout.

INT_TYPE_SIZE

A C expression for the size in bits of the type `int` on the target machine. If you don’t define this, the default is one word.

SHORT_TYPE_SIZE

A C expression for the size in bits of the type `short` on the target machine. If you don’t define this, the default is half a word. (If this would be less than one storage unit, it is rounded up to one unit.)

LONG_TYPE_SIZE

A C expression for the size in bits of the type `long` on the target machine. If you don’t define this, the default is one word.

ADA_LONG_TYPE_SIZE

On some machines, the size used for the Ada equivalent of the type `long` by a native Ada compiler differs from that used by C. In that situation, define this

macro to be a C expression to be used for the size of that type. If you don't define this, the default is the value of `LONG_TYPE_SIZE`.

`MAX_LONG_TYPE_SIZE`

Maximum number for the size in bits of the type `long` on the target machine. If this is undefined, the default is `LONG_TYPE_SIZE`. Otherwise, it is the constant value that is the largest value that `LONG_TYPE_SIZE` can have at run-time. This is used in `cpp`.

`LONG_LONG_TYPE_SIZE`

A C expression for the size in bits of the type `long long` on the target machine. If you don't define this, the default is two words. If you want to support GNU Ada on your machine, the value of this macro must be at least 64.

`CHAR_TYPE_SIZE`

A C expression for the size in bits of the type `char` on the target machine. If you don't define this, the default is `BITS_PER_UNIT`.

`BOOL_TYPE_SIZE`

A C expression for the size in bits of the C++ type `bool` and C99 type `_Bool` on the target machine. If you don't define this, and you probably shouldn't, the default is `CHAR_TYPE_SIZE`.

`FLOAT_TYPE_SIZE`

A C expression for the size in bits of the type `float` on the target machine. If you don't define this, the default is one word.

`DOUBLE_TYPE_SIZE`

A C expression for the size in bits of the type `double` on the target machine. If you don't define this, the default is two words.

`LONG_DOUBLE_TYPE_SIZE`

A C expression for the size in bits of the type `long double` on the target machine. If you don't define this, the default is two words.

Maximum number for the size in bits of the type `long double` on the target machine. If this is undefined, the default is `LONG_DOUBLE_TYPE_SIZE`. Otherwise, it is the constant value that is the largest value that `LONG_DOUBLE_TYPE_SIZE` can have at run-time. This is used in `cpp`.

`TARGET_FLT_EVAL_METHOD`

A C expression for the value for `FLT_EVAL_METHOD` in `'float.h'`, assuming, if applicable, that the floating-point control word is in its default state. If you do not define this macro the value of `FLT_EVAL_METHOD` will be zero.

`WIDEST_HARDWARE_FP_SIZE`

A C expression for the size in bits of the widest floating-point format supported by the hardware. If you define this macro, you must specify a value less than or equal to the value of `LONG_DOUBLE_TYPE_SIZE`. If you do not define this macro, the value of `LONG_DOUBLE_TYPE_SIZE` is the default.

DEFAULT_SIGNED_CHAR

An expression whose value is 1 or 0, according to whether the type `char` should be signed or unsigned by default. The user can always override this default with the options `'-fsigned-char'` and `'-funsigned-char'`.

DEFAULT_SHORT_ENUMS

A C expression to determine whether to give an `enum` type only as many bytes as it takes to represent the range of possible values of that type. A nonzero value means to do that; a zero value means all `enum` types should be allocated like `int`.

If you don't define the macro, the default is 0.

SIZE_TYPE

A C expression for a string describing the name of the data type to use for size values. The typedef name `size_t` is defined using the contents of the string.

The string can contain more than one keyword. If so, separate them with spaces, and write first any length keyword, then `unsigned` if appropriate, and finally `int`. The string must exactly match one of the data type names defined in the function `init_decl_processing` in the file `'c-decl.c'`. You may not omit `int` or change the order—that would cause the compiler to crash on startup.

If you don't define this macro, the default is `"long unsigned int"`.

PTRDIFF_TYPE

A C expression for a string describing the name of the data type to use for the result of subtracting two pointers. The typedef name `ptrdiff_t` is defined using the contents of the string. See `SIZE_TYPE` above for more information.

If you don't define this macro, the default is `"long int"`.

WCHAR_TYPE

A C expression for a string describing the name of the data type to use for wide characters. The typedef name `wchar_t` is defined using the contents of the string. See `SIZE_TYPE` above for more information.

If you don't define this macro, the default is `"int"`.

WCHAR_TYPE_SIZE

A C expression for the size in bits of the data type for wide characters. This is used in `cpp`, which cannot make use of `WCHAR_TYPE`.

MAX_WCHAR_TYPE_SIZE

Maximum number for the size in bits of the data type for wide characters. If this is undefined, the default is `WCHAR_TYPE_SIZE`. Otherwise, it is the constant value that is the largest value that `WCHAR_TYPE_SIZE` can have at run-time. This is used in `cpp`.

GCOV_TYPE_SIZE

A C expression for the size in bits of the type used for gcov counters on the target machine. If you don't define this, the default is one `LONG_TYPE_SIZE` in case it is greater or equal to 64-bit and `LONG_LONG_TYPE_SIZE` otherwise. You may want to re-define the type to ensure atomicity for counters in multithreaded programs.

WINT_TYPE

A C expression for a string describing the name of the data type to use for wide characters passed to `printf` and returned from `getwc`. The typedef name `wint_t` is defined using the contents of the string. See **SIZE_TYPE** above for more information.

If you don't define this macro, the default is "unsigned int".

INTMAX_TYPE

A C expression for a string describing the name of the data type that can represent any value of any standard or extended signed integer type. The typedef name `intmax_t` is defined using the contents of the string. See **SIZE_TYPE** above for more information.

If you don't define this macro, the default is the first of "int", "long int", or "long long int" that has as much precision as long long int.

UINTMAX_TYPE

A C expression for a string describing the name of the data type that can represent any value of any standard or extended unsigned integer type. The typedef name `uintmax_t` is defined using the contents of the string. See **SIZE_TYPE** above for more information.

If you don't define this macro, the default is the first of "unsigned int", "long unsigned int", or "long long unsigned int" that has as much precision as long long unsigned int.

TARGET_PTRMEMFUNC_VBIT_LOCATION

The C++ compiler represents a pointer-to-member-function with a struct that looks like:

```
struct {
  union {
    void (*fn)();
    ptrdiff_t vtable_index;
  };
  ptrdiff_t delta;
};
```

The C++ compiler must use one bit to indicate whether the function that will be called through a pointer-to-member-function is virtual. Normally, we assume that the low-order bit of a function pointer must always be zero. Then, by ensuring that the `vtable_index` is odd, we can distinguish which variant of the union is in use. But, on some platforms function pointers can be odd, and so this doesn't work. In that case, we use the low-order bit of the `delta` field, and shift the remainder of the `delta` field to the left.

GCC will automatically make the right selection about where to store this bit using the `FUNCTION_BOUNDARY` setting for your platform. However, some platforms such as ARM/Thumb have `FUNCTION_BOUNDARY` set such that functions always start at even addresses, but the lowest bit of pointers to functions indicate whether the function at that address is in ARM or Thumb mode. If this is the case of your architecture, you should define this macro to `ptrmemfunc_vbit_in_delta`.

In general, you should not have to define this macro. On architectures in which function addresses are always even, according to `FUNCTION_BOUNDARY`, GCC will automatically define this macro to `ptrdiff_t`.

`TARGET_VTABLE_USES_DESCRIPTOR`

Normally, the C++ compiler uses function pointers in vtables. This macro allows the target to change to use “function descriptors” instead. Function descriptors are found on targets for whom a function pointer is actually a small data structure. Normally the data structure consists of the actual code address plus a data pointer to which the function’s data is relative.

If vtables are used, the value of this macro should be the number of words that the function descriptor occupies.

`TARGET_VTABLE_ENTRY_ALIGN`

By default, the vtable entries are void pointers, so the alignment is the same as pointer alignment. The value of this macro specifies the alignment of the vtable entry in bits. It should be defined only when special alignment is necessary. */

`TARGET_VTABLE_DATA_ENTRY_DISTANCE`

There are a few non-descriptor entries in the vtable at offsets below zero. If these entries must be padded (say, to preserve the alignment specified by `TARGET_VTABLE_ENTRY_ALIGN`), set this to the number of words in each data entry.

10.7 Target Character Escape Sequences

By default, GCC assumes that the C character escape sequences take on their ASCII values for the target. If this is not correct, you must explicitly define all of the macros below.

`TARGET_BELL`

A C constant expression for the integer value for escape sequence `'\a'`.

`TARGET_ESC`

A C constant expression for the integer value of the target escape character. As an extension, GCC evaluates the escape sequences `'\e'` and `'\E'` to this.

`TARGET_BS`

`TARGET_TAB`

`TARGET_NEWLINE`

C constant expressions for the integer values for escape sequences `'\b'`, `'\t'` and `'\n'`.

`TARGET_VT`

`TARGET_FF`

`TARGET_CR`

C constant expressions for the integer values for escape sequences `'\v'`, `'\f'` and `'\r'`.

10.8 Register Usage

This section explains how to describe what registers the target machine has, and how (in general) they can be used.

The description of which registers a specific instruction can use is done with register classes; see Section 10.9 [Register Classes], page 223. For information on using registers to access a stack frame, see Section 10.10.4 [Frame Registers], page 236. For passing values in registers, see Section 10.10.7 [Register Arguments], page 242. For returning values in registers, see Section 10.10.8 [Scalar Return], page 245.

10.8.1 Basic Characteristics of Registers

Registers have various characteristics.

FIRST_PSEUDO_REGISTER

Number of hardware registers known to the compiler. They receive numbers 0 through `FIRST_PSEUDO_REGISTER-1`; thus, the first pseudo register's number really is assigned the number `FIRST_PSEUDO_REGISTER`.

FIXED_REGISTERS

An initializer that says which registers are used for fixed purposes all throughout the compiled code and are therefore not available for general allocation. These would include the stack pointer, the frame pointer (except on machines where that can be used as a general register when no frame pointer is needed), the program counter on machines where that is considered one of the addressable registers, and any other numbered register with a standard use.

This information is expressed as a sequence of numbers, separated by commas and surrounded by braces. The n th number is 1 if register n is fixed, 0 otherwise.

The table initialized from this macro, and the table initialized by the following one, may be overridden at run time either automatically, by the actions of the macro `CONDITIONAL_REGISTER_USAGE`, or by the user with the command options `'-ffixed-reg'`, `'-fcall-used-reg'` and `'-fcall-saved-reg'`.

CALL_USED_REGISTERS

Like `FIXED_REGISTERS` but has 1 for each register that is clobbered (in general) by function calls as well as for fixed registers. This macro therefore identifies the registers that are not available for general allocation of values that must live across function calls.

If a register has 0 in `CALL_USED_REGISTERS`, the compiler automatically saves it on function entry and restores it on function exit, if the register is used within the function.

CALL_REALLY_USED_REGISTERS

Like `CALL_USED_REGISTERS` except this macro doesn't require that the entire set of `FIXED_REGISTERS` be included. (`CALL_USED_REGISTERS` must be a superset of `FIXED_REGISTERS`). This macro is optional. If not specified, it defaults to the value of `CALL_USED_REGISTERS`.

HARD_REGNO_CALL_PART_CLOBBERED (*regno*, *mode*)

A C expression that is nonzero if it is not permissible to store a value of mode *mode* in hard register number *regno* across a call without some part of it being clobbered. For most machines this macro need not be defined. It is only required for machines that do not preserve the entire contents of a register across a call.

CONDITIONAL_REGISTER_USAGE

Zero or more C statements that may conditionally modify five variables *fixed_regs*, *call_used_regs*, *global_regs*, *reg_names*, and *reg_class_contents*, to take into account any dependence of these register sets on target flags. The first three of these are of type `char []` (interpreted as Boolean vectors). *global_regs* is a `const char *[]`, and *reg_class_contents* is a `HARD_REG_SET`. Before the macro is called, *fixed_regs*, *call_used_regs*, *reg_class_contents*, and *reg_names* have been initialized from `FIXED_REGISTERS`, `CALL_USED_REGISTERS`, `REG_CLASS_CONTENTS`, and `REGISTER_NAMES`, respectively. *global_regs* has been cleared, and any ‘-ffixed-reg’, ‘-fcall-used-reg’ and ‘-fcall-saved-reg’ command options have been applied.

You need not define this macro if it has no work to do.

If the usage of an entire class of registers depends on the target flags, you may indicate this to GCC by using this macro to modify *fixed_regs* and *call_used_regs* to 1 for each of the registers in the classes which should not be used by GCC. Also define the macro `REG_CLASS_FROM_LETTER` to return `NO_REGS` if it is called with a letter for a class that shouldn't be used.

(However, if this class is not included in `GENERAL_REGS` and all of the insn patterns whose constraints permit this class are controlled by target switches, then GCC will automatically avoid using these registers when the target switches are opposed to them.)

NON_SAVING_SETJMP

If this macro is defined and has a nonzero value, it means that `setjmp` and related functions fail to save the registers, or that `longjmp` fails to restore them. To compensate, the compiler avoids putting variables in registers in functions that use `setjmp`.

INCOMING_REGNO (*out*)

Define this macro if the target machine has register windows. This C expression returns the register number as seen by the called function corresponding to the register number *out* as seen by the calling function. Return *out* if register number *out* is not an outbound register.

OUTGOING_REGNO (*in*)

Define this macro if the target machine has register windows. This C expression returns the register number as seen by the calling function corresponding to the register number *in* as seen by the called function. Return *in* if register number *in* is not an inbound register.

LOCAL_REGNO (*regno*)

Define this macro if the target machine has register windows. This C expression returns true if the register is call-saved but is in the register window. Unlike

most call-saved registers, such registers need not be explicitly restored on function exit or during non-local gotos.

10.8.2 Order of Allocation of Registers

Registers are allocated in order.

REG_ALLOC_ORDER

If defined, an initializer for a vector of integers, containing the numbers of hard registers in the order in which GCC should prefer to use them (from most preferred to least).

If this macro is not defined, registers are used lowest numbered first (all else being equal).

One use of this macro is on machines where the highest numbered registers must always be saved and the save-multiple-registers instruction supports only sequences of consecutive registers. On such machines, define `REG_ALLOC_ORDER` to be an initializer that lists the highest numbered allocable register first.

ORDER_REGS_FOR_LOCAL_ALLOC

A C statement (sans semicolon) to choose the order in which to allocate hard registers for pseudo-registers local to a basic block.

Store the desired register order in the array `reg_alloc_order`. Element 0 should be the register to allocate first; element 1, the next register; and so on.

The macro body should not assume anything about the contents of `reg_alloc_order` before execution of the macro.

On most machines, it is not necessary to define this macro.

10.8.3 How Values Fit in Registers

This section discusses the macros that describe which kinds of values (specifically, which machine modes) each register can hold, and how many consecutive registers are needed for a given mode.

HARD_REGNO_NREGS (*regno*, *mode*)

A C expression for the number of consecutive hard registers, starting at register number *regno*, required to hold a value of mode *mode*.

On a machine where all registers are exactly one word, a suitable definition of this macro is

```
#define HARD_REGNO_NREGS(REGNO, MODE)      \
  ((GET_MODE_SIZE (MODE) + UNITS_PER_WORD - 1) \
   / UNITS_PER_WORD)
```

HARD_REGNO_MODE_OK (*regno*, *mode*)

A C expression that is nonzero if it is permissible to store a value of mode *mode* in hard register number *regno* (or in several registers starting with that one). For a machine where all registers are equivalent, a suitable definition is

```
#define HARD_REGNO_MODE_OK(REGNO, MODE) 1
```

You need not include code to check for the numbers of fixed registers, because the allocation mechanism considers them to be always occupied.

On some machines, double-precision values must be kept in even/odd register pairs. You can implement that by defining this macro to reject odd register numbers for such modes.

The minimum requirement for a mode to be OK in a register is that the ‘*movmode*’ instruction pattern support moves between the register and other hard register in the same class and that moving a value into the register and back out not alter it.

Since the same instruction used to move *word_mode* will work for all narrower integer modes, it is not necessary on any machine for *HARD_REGNO_MODE_OK* to distinguish between these modes, provided you define patterns ‘*movhi*’, etc., to take advantage of this. This is useful because of the interaction between *HARD_REGNO_MODE_OK* and *MODES_TIEABLE_P*; it is very desirable for all integer modes to be tieable.

Many machines have special registers for floating point arithmetic. Often people assume that floating point machine modes are allowed only in floating point registers. This is not true. Any registers that can hold integers can safely *hold* a floating point machine mode, whether or not floating arithmetic can be done on it in those registers. Integer move instructions can be used to move the values.

On some machines, though, the converse is true: fixed-point machine modes may not go in floating registers. This is true if the floating registers normalize any value stored in them, because storing a non-floating value there would garble it. In this case, *HARD_REGNO_MODE_OK* should reject fixed-point machine modes in floating registers. But if the floating registers do not automatically normalize, if you can store any bit pattern in one and retrieve it unchanged without a trap, then any machine mode may go in a floating register, so you can define this macro to say so.

The primary significance of special floating registers is rather that they are the registers acceptable in floating point arithmetic instructions. However, this is of no concern to *HARD_REGNO_MODE_OK*. You handle it by writing the proper constraints for those instructions.

On some machines, the floating registers are especially slow to access, so that it is better to store a value in a stack frame than in such a register if floating point arithmetic is not being done. As long as the floating registers are not in class *GENERAL_REGS*, they will not be used unless some pattern’s constraint asks for one.

MODES_TIEABLE_P (*mode1*, *mode2*)

A C expression that is nonzero if a value of mode *mode1* is accessible in mode *mode2* without copying.

If *HARD_REGNO_MODE_OK* (*r*, *mode1*) and *HARD_REGNO_MODE_OK* (*r*, *mode2*) are always the same for any *r*, then *MODES_TIEABLE_P* (*mode1*, *mode2*) should be nonzero. If they differ for any *r*, you should define this macro to return

zero unless some other mechanism ensures the accessibility of the value in a narrower mode.

You should define this macro to return nonzero in as many cases as possible since doing so will allow GCC to perform better register allocation.

AVOID_CCmode_COPIES

Define this macro if the compiler should avoid copies to/from `CCmode` registers. You should only define this macro if support for copying to/from `CCmode` is incomplete.

10.8.4 Handling Leaf Functions

On some machines, a leaf function (i.e., one which makes no calls) can run more efficiently if it does not make its own register window. Often this means it is required to receive its arguments in the registers where they are passed by the caller, instead of the registers where they would normally arrive.

The special treatment for leaf functions generally applies only when other conditions are met; for example, often they may use only those registers for its own variables and temporaries. We use the term “leaf function” to mean a function that is suitable for this special handling, so that functions with no calls are not necessarily “leaf functions”.

GCC assigns register numbers before it knows whether the function is suitable for leaf function treatment. So it needs to renumber the registers in order to output a leaf function. The following macros accomplish this.

LEAF_REGISTERS

Name of a char vector, indexed by hard register number, which contains 1 for a register that is allowable in a candidate for leaf function treatment.

If leaf function treatment involves renumbering the registers, then the registers marked here should be the ones before renumbering—those that GCC would ordinarily allocate. The registers which will actually be used in the assembler code, after renumbering, should not be marked with 1 in this vector.

Define this macro only if the target machine offers a way to optimize the treatment of leaf functions.

LEAF_REG_REMAP (*regno*)

A C expression whose value is the register number to which *regno* should be renumbered, when a function is treated as a leaf function.

If *regno* is a register number which should not appear in a leaf function before renumbering, then the expression should yield -1 , which will cause the compiler to abort.

Define this macro only if the target machine offers a way to optimize the treatment of leaf functions, and registers need to be renumbered to do this.

`TARGET_ASM_FUNCTION_PROLOGUE` and `TARGET_ASM_FUNCTION_EPILOGUE` must usually treat leaf functions specially. They can test the C variable `current_function_is_leaf` which is nonzero for leaf functions. `current_function_is_leaf` is set prior to local register allocation and is valid for the remaining compiler passes. They can also test the C variable

`current_function_uses_only_leaf_regs` which is nonzero for leaf functions which only use leaf registers. `current_function_uses_only_leaf_regs` is valid after reload and is only useful if `LEAF_REGISTERS` is defined.

10.8.5 Registers That Form a Stack

There are special features to handle computers where some of the “registers” form a stack, as in the 80387 coprocessor for the 80386. Stack registers are normally written by pushing onto the stack, and are numbered relative to the top of the stack.

Currently, GCC can only handle one group of stack-like registers, and they must be consecutively numbered.

`STACK_REGS`

Define this if the machine has any stack-like registers.

`FIRST_STACK_REG`

The number of the first stack-like register. This one is the top of the stack.

`LAST_STACK_REG`

The number of the last stack-like register. This one is the bottom of the stack.

10.9 Register Classes

On many machines, the numbered registers are not all equivalent. For example, certain registers may not be allowed for indexed addressing; certain registers may not be allowed in some instructions. These machine restrictions are described to the compiler using *register classes*.

You define a number of register classes, giving each one a name and saying which of the registers belong to it. Then you can specify register classes that are allowed as operands to particular instruction patterns.

In general, each register will belong to several classes. In fact, one class must be named `ALL_REGS` and contain all the registers. Another class must be named `NO_REGS` and contain no registers. Often the union of two classes will be another class; however, this is not required.

One of the classes must be named `GENERAL_REGS`. There is nothing terribly special about the name, but the operand constraint letters ‘r’ and ‘g’ specify this class. If `GENERAL_REGS` is the same as `ALL_REGS`, just define it as a macro which expands to `ALL_REGS`.

Order the classes so that if class *x* is contained in class *y* then *x* has a lower class number than *y*.

The way classes other than `GENERAL_REGS` are specified in operand constraints is through machine-dependent operand constraint letters. You can define such letters to correspond to various classes, then use them in operand constraints.

You should define a class for the union of two classes whenever some instruction allows both classes. For example, if an instruction allows either a floating point (coprocessor) register or a general register for a certain operand, you should define a class `FLOAT_OR_GENERAL_REGS` which includes both of them. Otherwise you will get suboptimal code.

You must also specify certain redundant information about the register classes: for each class, which classes contain it and which ones are contained in it; for each pair of classes, the largest class contained in their union.

When a value occupying several consecutive registers is expected in a certain class, all the registers used must belong to that class. Therefore, register classes cannot be used to enforce a requirement for a register pair to start with an even-numbered register. The way to specify this requirement is with `HARD_REGNO_MODE_OK`.

Register classes used for input-operands of bitwise-and or shift instructions have a special requirement: each such class must have, for each fixed-point machine mode, a subclass whose registers can transfer that mode to or from memory. For example, on some machines, the operations for single-byte values (`QImode`) are limited to certain registers. When this is so, each register class that is used in a bitwise-and or shift instruction must have a subclass consisting of registers from which single-byte values can be loaded or stored. This is so that `PREFERRED_RELOAD_CLASS` can always have a possible value to return.

`enum reg_class`

An enumerational type that must be defined with all the register class names as enumerational values. `NO_REGS` must be first. `ALL_REGS` must be the last register class, followed by one more enumerational value, `LIM_REG_CLASSES`, which is not a register class but rather tells how many classes there are.

Each register class has a number, which is the value of casting the class name to type `int`. The number serves as an index in many of the tables described below.

`N_REG_CLASSES`

The number of distinct register classes, defined as follows:

```
#define N_REG_CLASSES (int) LIM_REG_CLASSES
```

`REG_CLASS_NAMES`

An initializer containing the names of the register classes as C string constants. These names are used in writing some of the debugging dumps.

`REG_CLASS_CONTENTS`

An initializer containing the contents of the register classes, as integers which are bit masks. The n th integer specifies the contents of class n . The way the integer $mask$ is interpreted is that register r is in the class if $mask \& (1 \ll r)$ is 1.

When the machine has more than 32 registers, an integer does not suffice. Then the integers are replaced by sub-initializers, braced groupings containing several integers. Each sub-initializer must be suitable as an initializer for the type `HARD_REG_SET` which is defined in ‘`hard-reg-set.h`’. In this situation, the first integer in each sub-initializer corresponds to registers 0 through 31, the second integer to registers 32 through 63, and so on.

`REGNO_REG_CLASS (regno)`

A C expression whose value is a register class containing hard register $regno$. In general there is more than one such class; choose a class which is *minimal*, meaning that no smaller class also contains the register.

BASE_REG_CLASS

A macro whose definition is the name of the class to which a valid base register must belong. A base register is one used in an address which is the register value plus a displacement.

MODE_BASE_REG_CLASS (*mode*)

This is a variation of the **BASE_REG_CLASS** macro which allows the selection of a base register in a mode dependent manner. If *mode* is **VOIDmode** then it should return the same value as **BASE_REG_CLASS**.

INDEX_REG_CLASS

A macro whose definition is the name of the class to which a valid index register must belong. An index register is one used in an address where its value is either multiplied by a scale factor or added to another register (as well as added to a displacement).

REG_CLASS_FROM_LETTER (*char*)

A C expression which defines the machine-dependent operand constraint letters for register classes. If *char* is such a letter, the value should be the register class corresponding to it. Otherwise, the value should be **NO_REGS**. The register letter ‘r’, corresponding to class **GENERAL_REGS**, will not be passed to this macro; you do not need to handle it.

REGNO_OK_FOR_BASE_P (*num*)

A C expression which is nonzero if register number *num* is suitable for use as a base register in operand addresses. It may be either a suitable hard register or a pseudo register that has been allocated such a hard register.

REGNO_MODE_OK_FOR_BASE_P (*num*, *mode*)

A C expression that is just like **REGNO_OK_FOR_BASE_P**, except that that expression may examine the mode of the memory reference in *mode*. You should define this macro if the mode of the memory reference affects whether a register may be used as a base register. If you define this macro, the compiler will use it instead of **REGNO_OK_FOR_BASE_P**.

REGNO_OK_FOR_INDEX_P (*num*)

A C expression which is nonzero if register number *num* is suitable for use as an index register in operand addresses. It may be either a suitable hard register or a pseudo register that has been allocated such a hard register.

The difference between an index register and a base register is that the index register may be scaled. If an address involves the sum of two registers, neither one of them scaled, then either one may be labeled the “base” and the other the “index”; but whichever labeling is used must fit the machine’s constraints of which registers may serve in each capacity. The compiler will try both labelings, looking for one that is valid, and will reload one or both registers only if neither labeling works.

PREFERRED_RELOAD_CLASS (*x*, *class*)

A C expression that places additional restrictions on the register class to use when it is necessary to copy value *x* into a register in class *class*. The value

is a register class; perhaps *class*, or perhaps another, smaller class. On many machines, the following definition is safe:

```
#define PREFERRED_RELOAD_CLASS(X,CLASS) CLASS
```

Sometimes returning a more restrictive class makes better code. For example, on the 68000, when *x* is an integer constant that is in range for a ‘moveq’ instruction, the value of this macro is always `DATA_REGS` as long as *class* includes the data registers. Requiring a data register guarantees that a ‘moveq’ will be used.

If *x* is a `const_double`, by returning `NO_REGS` you can force *x* into a memory constant. This is useful on certain machines where immediate floating values cannot be loaded into certain kinds of registers.

`PREFERRED_OUTPUT_RELOAD_CLASS (x, class)`

Like `PREFERRED_RELOAD_CLASS`, but for output reloads instead of input reloads. If you don’t define this macro, the default is to use *class*, unchanged.

`LIMIT_RELOAD_CLASS (mode, class)`

A C expression that places additional restrictions on the register class to use when it is necessary to be able to hold a value of mode *mode* in a reload register for which class *class* would ordinarily be used.

Unlike `PREFERRED_RELOAD_CLASS`, this macro should be used when there are certain modes that simply can’t go in certain reload classes.

The value is a register class; perhaps *class*, or perhaps another, smaller class.

Don’t define this macro unless the target machine has limitations which require the macro to do something nontrivial.

`SECONDARY_RELOAD_CLASS (class, mode, x)`

`SECONDARY_INPUT_RELOAD_CLASS (class, mode, x)`

`SECONDARY_OUTPUT_RELOAD_CLASS (class, mode, x)`

Many machines have some registers that cannot be copied directly to or from memory or even from other types of registers. An example is the ‘MQ’ register, which on most machines, can only be copied to or from general registers, but not memory. Some machines allow copying all registers to and from memory, but require a scratch register for stores to some memory locations (e.g., those with symbolic address on the RT, and those with certain symbolic address on the SPARC when compiling PIC). In some cases, both an intermediate and a scratch register are required.

You should define these macros to indicate to the reload phase that it may need to allocate at least one register for a reload in addition to the register to contain the data. Specifically, if copying *x* to a register *class* in *mode* requires an intermediate register, you should define `SECONDARY_INPUT_RELOAD_CLASS` to return the largest register class all of whose registers can be used as intermediate registers or scratch registers.

If copying a register *class* in *mode* to *x* requires an intermediate or scratch register, `SECONDARY_OUTPUT_RELOAD_CLASS` should be defined to return the largest register class required. If the requirements for input and output reloads are the

same, the macro `SECONDARY_RELOAD_CLASS` should be used instead of defining both macros identically.

The values returned by these macros are often `GENERAL_REGS`. Return `NO_REGS` if no spare register is needed; i.e., if x can be directly copied to or from a register of *class* in *mode* without requiring a scratch register. Do not define this macro if it would always return `NO_REGS`.

If a scratch register is required (either with or without an intermediate register), you should define patterns for `'reload_inm'` or `'reload_outm'`, as required (see Section 9.8 [Standard Names], page 144. These patterns, which will normally be implemented with a `define_expand`, should be similar to the `'movm'` patterns, except that operand 2 is the scratch register.

Define constraints for the reload register and scratch register that contain a single register class. If the original reload register (whose class is *class*) can meet the constraint given in the pattern, the value returned by these macros is used for the class of the scratch register. Otherwise, two additional reload registers are required. Their classes are obtained from the constraints in the `insn` pattern.

x might be a pseudo-register or a `subreg` of a pseudo-register, which could either be in a hard register or in memory. Use `true_regnum` to find out; it will return -1 if the pseudo is in memory and the hard register number if it is in a register.

These macros should not be used in the case where a particular class of registers can only be copied to memory and not to another class of registers. In that case, secondary reload registers are not needed and would not be helpful. Instead, a stack location must be used to perform the copy and the `movm` pattern should use memory as an intermediate storage. This case often occurs between floating-point and general registers.

`SECONDARY_MEMORY_NEEDED` (*class1*, *class2*, *m*)

Certain machines have the property that some registers cannot be copied to some other registers without using memory. Define this macro on those machines to be a C expression that is nonzero if objects of mode *m* in registers of *class1* can only be copied to registers of class *class2* by storing a register of *class1* into memory and loading that memory location into a register of *class2*.

Do not define this macro if its value would always be zero.

`SECONDARY_MEMORY_NEEDED_RTX` (*mode*)

Normally when `SECONDARY_MEMORY_NEEDED` is defined, the compiler allocates a stack slot for a memory location needed for register copies. If this macro is defined, the compiler instead uses the memory location defined by this macro.

Do not define this macro if you do not define `SECONDARY_MEMORY_NEEDED`.

`SECONDARY_MEMORY_NEEDED_MODE` (*mode*)

When the compiler needs a secondary memory location to copy between two registers of mode *mode*, it normally allocates sufficient memory to hold a quantity of `BITS_PER_WORD` bits and performs the store and load operations in a mode that many bits wide and whose class is the same as that of *mode*.

This is right thing to do on most machines because it ensures that all bits of the register are copied and prevents accesses to the registers in a narrower mode, which some machines prohibit for floating-point registers.

However, this default behavior is not correct on some machines, such as the DEC Alpha, that store short integers in floating-point registers differently than in integer registers. On those machines, the default widening will not work correctly and you must define this macro to suppress that widening in some cases. See the file ‘`alpha.h`’ for details.

Do not define this macro if you do not define `SECONDARY_MEMORY_NEEDED` or if widening *mode* to a mode that is `BITS_PER_WORD` bits wide is correct for your machine.

`SMALL_REGISTER_CLASSES`

On some machines, it is risky to let hard registers live across arbitrary insns. Typically, these machines have instructions that require values to be in specific registers (like an accumulator), and reload will fail if the required hard register is used for another purpose across such an insn.

Define `SMALL_REGISTER_CLASSES` to be an expression with a nonzero value on these machines. When this macro has a nonzero value, the compiler will try to minimize the lifetime of hard registers.

It is always safe to define this macro with a nonzero value, but if you unnecessarily define it, you will reduce the amount of optimizations that can be performed in some cases. If you do not define this macro with a nonzero value when it is required, the compiler will run out of spill registers and print a fatal error message. For most machines, you should not define this macro at all.

`CLASS_LIKELY_SPILLED_P` (*class*)

A C expression whose value is nonzero if pseudos that have been assigned to registers of class *class* would likely be spilled because registers of *class* are needed for spill registers.

The default value of this macro returns 1 if *class* has exactly one register and zero otherwise. On most machines, this default should be used. Only define this macro to some other expression if pseudos allocated by ‘`local-alloc.c`’ end up in memory because their hard registers were needed for spill registers. If this macro returns nonzero for those classes, those pseudos will only be allocated by ‘`global.c`’, which knows how to reallocate the pseudo to another register. If there would not be another register available for reallocation, you should not change the definition of this macro since the only effect of such a definition would be to slow down register allocation.

`CLASS_MAX_NREGS` (*class*, *mode*)

A C expression for the maximum number of consecutive registers of class *class* needed to hold a value of mode *mode*.

This is closely related to the macro `HARD_REGNO_NREGS`. In fact, the value of the macro `CLASS_MAX_NREGS` (*class*, *mode*) should be the maximum value of `HARD_REGNO_NREGS` (*regno*, *mode*) for all *regno* values in the class *class*.

This macro helps control the handling of multiple-word values in the reload pass.

CANNOT_CHANGE_MODE_CLASS(*from*, *to*, *class*)

If defined, a C expression that returns nonzero for a *class* for which a change from mode *from* to mode *to* is invalid.

For the example, loading 32-bit integer or floating-point objects into floating-point registers on the Alpha extends them to 64 bits. Therefore loading a 64-bit object and then storing it as a 32-bit object does not store the low-order 32 bits, as would be the case for a normal register. Therefore, ‘alpha.h’ defines CANNOT_CHANGE_MODE_CLASS as below:

```
#define CANNOT_CHANGE_MODE_CLASS(FROM, TO, CLASS) \
  (GET_MODE_SIZE (FROM) != GET_MODE_SIZE (TO) \
   ? reg_classes_intersect_p (FLOAT_REGS, (CLASS)) : 0)
```

Three other special macros describe which operands fit which constraint letters.

CONST_OK_FOR_LETTER_P(*value*, *c*)

A C expression that defines the machine-dependent operand constraint letters (‘I’, ‘J’, ‘K’, . . . ‘P’) that specify particular ranges of integer values. If *c* is one of those letters, the expression should check that *value*, an integer, is in the appropriate range and return 1 if so, 0 otherwise. If *c* is not one of those letters, the value should be 0 regardless of *value*.

CONST_DOUBLE_OK_FOR_LETTER_P(*value*, *c*)

A C expression that defines the machine-dependent operand constraint letters that specify particular ranges of `const_double` values (‘G’ or ‘H’).

If *c* is one of those letters, the expression should check that *value*, an RTX of code `const_double`, is in the appropriate range and return 1 if so, 0 otherwise. If *c* is not one of those letters, the value should be 0 regardless of *value*.

`const_double` is used for all floating-point constants and for DImode fixed-point constants. A given letter can accept either or both kinds of values. It can use GET_MODE to distinguish between these kinds.

EXTRA_CONSTRAINT(*value*, *c*)

A C expression that defines the optional machine-dependent constraint letters that can be used to segregate specific types of operands, usually memory references, for the target machine. Any letter that is not elsewhere defined and not matched by REG_CLASS_FROM_LETTER may be used. Normally this macro will not be defined.

If it is required for a particular target machine, it should return 1 if *value* corresponds to the operand type represented by the constraint letter *c*. If *c* is not defined as an extra constraint, the value returned should be 0 regardless of *value*.

For example, on the ROMP, load instructions cannot have their output in r0 if the memory reference contains a symbolic address. Constraint letter ‘Q’ is defined as representing a memory address that does *not* contain a symbolic address. An alternative is specified with a ‘Q’ constraint on the input and ‘r’ on the output. The next alternative specifies ‘m’ on the input and a register class that does not include r0 on the output.

EXTRA_MEMORY_CONSTRAINT (*c*)

A C expression that defines the optional machine-dependent constraint letters, amongst those accepted by `EXTRA_CONSTRAINT`, that should be treated like memory constraints by the reload pass.

It should return 1 if the operand type represented by the constraint letter *c* comprises a subset of all memory references including all those whose address is simply a base register. This allows the reload pass to reload an operand, if it does not directly correspond to the operand type of *c*, by copying its address into a base register.

For example, on the S/390, some instructions do not accept arbitrary memory references, but only those that do not make use of an index register. The constraint letter ‘Q’ is defined via `EXTRA_CONSTRAINT` as representing a memory address of this type. If the letter ‘Q’ is marked as `EXTRA_MEMORY_CONSTRAINT`, a ‘Q’ constraint can handle any memory operand, because the reload pass knows it can be reloaded by copying the memory address into a base register if required. This is analogous to the way a ‘o’ constraint can handle any memory operand.

EXTRA_ADDRESS_CONSTRAINT (*c*)

A C expression that defines the optional machine-dependent constraint letters, amongst those accepted by `EXTRA_CONSTRAINT`, that should be treated like address constraints by the reload pass.

It should return 1 if the operand type represented by the constraint letter *c* comprises a subset of all memory addresses including all those that consist of just a base register. This allows the reload pass to reload an operand, if it does not directly correspond to the operand type of *c*, by copying it into a base register.

Any constraint marked as `EXTRA_ADDRESS_CONSTRAINT` can only be used with the `address_operand` predicate. It is treated analogously to the ‘p’ constraint.

10.10 Stack Layout and Calling Conventions

This describes the stack layout and calling conventions.

10.10.1 Basic Stack Layout

Here is the basic stack layout.

STACK_GROWS_DOWNWARD

Define this macro if pushing a word onto the stack moves the stack pointer to a smaller address.

When we say, “define this macro if . . .,” it means that the compiler checks this macro only with `#ifdef` so the precise definition used does not matter.

STACK_PUSH_CODE

This macro defines the operation used when something is pushed on the stack. In RTL, a push operation will be `(set (mem (STACK_PUSH_CODE (reg sp))) ...)`

The choices are `PRE_DEC`, `POST_DEC`, `PRE_INC`, and `POST_INC`. Which of these is correct depends on the stack direction and on whether the stack pointer points to the last item on the stack or whether it points to the space for the next item on the stack.

The default is `PRE_DEC` when `STACK_GROWS_DOWNWARD` is defined, which is almost always right, and `PRE_INC` otherwise, which is often wrong.

`FRAME_GROWS_DOWNWARD`

Define this macro if the addresses of local variable slots are at negative offsets from the frame pointer.

`ARGS_GROW_DOWNWARD`

Define this macro if successive arguments to a function occupy decreasing addresses on the stack.

`STARTING_FRAME_OFFSET`

Offset from the frame pointer to the first local variable slot to be allocated.

If `FRAME_GROWS_DOWNWARD`, find the next slot's offset by subtracting the first slot's length from `STARTING_FRAME_OFFSET`. Otherwise, it is found by adding the length of the first slot to the value `STARTING_FRAME_OFFSET`.

`STACK_POINTER_OFFSET`

Offset from the stack pointer register to the first location at which outgoing arguments are placed. If not specified, the default value of zero is used. This is the proper value for most machines.

If `ARGS_GROW_DOWNWARD`, this is the offset to the location above the first location at which outgoing arguments are placed.

`FIRST_PARM_OFFSET` (*fundecl*)

Offset from the argument pointer register to the first argument's address. On some machines it may depend on the data type of the function.

If `ARGS_GROW_DOWNWARD`, this is the offset to the location above the first argument's address.

`STACK_DYNAMIC_OFFSET` (*fundecl*)

Offset from the stack pointer register to an item dynamically allocated on the stack, e.g., by `alloca`.

The default value for this macro is `STACK_POINTER_OFFSET` plus the length of the outgoing arguments. The default is correct for most machines. See `'function.c'` for details.

`DYNAMIC_CHAIN_ADDRESS` (*frameaddr*)

A C expression whose value is RTL representing the address in a stack frame where the pointer to the caller's frame is stored. Assume that *frameaddr* is an RTL expression for the address of the stack frame itself.

If you don't define this macro, the default is to return the value of *frameaddr*—that is, the stack frame address is also the address of the stack word that points to the previous frame.

SETUP_FRAME_ADDRESSES

If defined, a C expression that produces the machine-specific code to setup the stack so that arbitrary frames can be accessed. For example, on the SPARC, we must flush all of the register windows to the stack before we can access arbitrary stack frames. You will seldom need to define this macro.

BUILTIN_SETJMP_FRAME_VALUE

If defined, a C expression that contains an rtx that is used to store the address of the current frame into the built in `setjmp` buffer. The default value, `virtual_stack_vars_rtx`, is correct for most machines. One reason you may need to define this macro is if `hard_frame_pointer_rtx` is the appropriate value on your machine.

RETURN_ADDR_RTX (*count*, *frameaddr*)

A C expression whose value is RTL representing the value of the return address for the frame *count* steps up from the current frame, after the prologue. *frameaddr* is the frame pointer of the *count* frame, or the frame pointer of the *count* - 1 frame if `RETURN_ADDR_IN_PREVIOUS_FRAME` is defined.

The value of the expression must always be the correct address when *count* is zero, but may be `NULL_RTX` if there is not way to determine the return address of other frames.

RETURN_ADDR_IN_PREVIOUS_FRAME

Define this if the return address of a particular stack frame is accessed from the frame pointer of the previous stack frame.

INCOMING_RETURN_ADDR_RTX

A C expression whose value is RTL representing the location of the incoming return address at the beginning of any function, before the prologue. This RTL is either a `REG`, indicating that the return value is saved in 'REG', or a `MEM` representing a location in the stack.

You only need to define this macro if you want to support call frame debugging information like that provided by DWARF 2.

If this RTL is a `REG`, you should also define `DWARF_FRAME_RETURN_COLUMN` to `DWARF_FRAME_REGNUM (REGNO)`.

INCOMING_FRAME_SP_OFFSET

A C expression whose value is an integer giving the offset, in bytes, from the value of the stack pointer register to the top of the stack frame at the beginning of any function, before the prologue. The top of the frame is defined to be the value of the stack pointer in the previous frame, just before the call instruction.

You only need to define this macro if you want to support call frame debugging information like that provided by DWARF 2.

ARG_POINTER_CFA_OFFSET (*fundecl*)

A C expression whose value is an integer giving the offset, in bytes, from the argument pointer to the canonical frame address (*cfa*). The final value should coincide with that calculated by `INCOMING_FRAME_SP_OFFSET`. Which is unfortunately not usable during virtual register instantiation.

The default value for this macro is `FIRST_PARM_OFFSET` (`fundec1`), which is correct for most machines; in general, the arguments are found immediately before the stack frame. Note that this is not the case on some targets that save registers into the caller's frame, such as SPARC and rs6000, and so such targets need to define this macro.

You only need to define this macro if the default is incorrect, and you want to support call frame debugging information like that provided by DWARF 2.

`SMALL_STACK`

Define this macro if the stack size for the target is very small. This has the effect of disabling gcc's built-in `'alloca'`, though `'__builtin_alloca'` is not affected.

10.10.2 Exception Handling Support

`EH_RETURN_DATA_REGNO` (*N*)

A C expression whose value is the *N*th register number used for data by exception handlers, or `INVALID_REGNUM` if fewer than *N* registers are usable.

The exception handling library routines communicate with the exception handlers via a set of agreed upon registers. Ideally these registers should be call-clobbered; it is possible to use call-saved registers, but may negatively impact code size. The target must support at least 2 data registers, but should define 4 if there are enough free registers.

You must define this macro if you want to support call frame exception handling like that provided by DWARF 2.

`EH_RETURN_STACKADJ_RTX`

A C expression whose value is RTL representing a location in which to store a stack adjustment to be applied before function return. This is used to unwind the stack to an exception handler's call frame. It will be assigned zero on code paths that return normally.

Typically this is a call-clobbered hard register that is otherwise untouched by the epilogue, but could also be a stack slot.

Do not define this macro if the stack pointer is saved and restored by the regular prolog and epilog code in the call frame itself; in this case, the exception handling library routines will update the stack location to be restored in place. Otherwise, you must define this macro if you want to support call frame exception handling like that provided by DWARF 2.

`EH_RETURN_HANDLER_RTX`

A C expression whose value is RTL representing a location in which to store the address of an exception handler to which we should return. It will not be assigned on code paths that return normally.

Typically this is the location in the call frame at which the normal return address is stored. For targets that return by popping an address off the stack, this might be a memory address just below the *target* call frame rather than inside the current call frame. If defined, `EH_RETURN_STACKADJ_RTX` will have

already been assigned, so it may be used to calculate the location of the target call frame.

Some targets have more complex requirements than storing to an address calculable during initial code generation. In that case the `eh_return` instruction pattern should be used instead.

If you want to support call frame exception handling, you must define either this macro or the `eh_return` instruction pattern.

`ASM_PREFERRED_EH_DATA_FORMAT(code, global)`

This macro chooses the encoding of pointers embedded in the exception handling sections. If at all possible, this should be defined such that the exception handling section will not require dynamic relocations, and so may be read-only. *code* is 0 for data, 1 for code labels, 2 for function pointers. *global* is true if the symbol may be affected by dynamic relocations. The macro should return a combination of the `DW_EH_PE_*` defines as found in `'dwarf2.h'`.

If this macro is not defined, pointers will not be encoded but represented directly.

`ASM_MAYBE_OUTPUT_ENCODED_ADDR_RTX(file, encoding, size, addr, done)`

This macro allows the target to emit whatever special magic is required to represent the encoding chosen by `ASM_PREFERRED_EH_DATA_FORMAT`. Generic code takes care of pc-relative and indirect encodings; this must be defined if the target uses text-relative or data-relative encodings.

This is a C statement that branches to *done* if the format was handled. *encoding* is the format chosen, *size* is the number of bytes that the format occupies, *addr* is the `SYMBOL_REF` to be emitted.

`MD_FALLBACK_FRAME_STATE_FOR(context, fs, success)`

This macro allows the target to add cpu and operating system specific code to the call-frame unwinder for use when there is no unwind data available. The most common reason to implement this macro is to unwind through signal frames.

This macro is called from `uw_frame_state_for` in `'unwind-dw2.c'` and `'unwind-ia64.c'`. *context* is an `_Unwind_Context`; *fs* is an `_Unwind_FrameState`. Examine `context->ra` for the address of the code being executed and `context->cfa` for the stack pointer value. If the frame can be decoded, the register save addresses should be updated in *fs* and the macro should branch to *success*. If the frame cannot be decoded, the macro should do nothing.

`MD_HANDLE_UNWABI(context, fs)`

This macro allows the target to add operating system specific code to the call-frame unwinder to handle the IA-64 `.unwabi` unwinding directive, usually used for signal or interrupt frames.

This macro is called from `uw_update_context` in `'unwind-ia64.c'`. *context* is an `_Unwind_Context`; *fs* is an `_Unwind_FrameState`. Examine `fs->unwabi` for the abi and context in the `.unwabi` directive. If the `.unwabi` directive can be handled, the register save addresses should be updated in *fs*.

10.10.3 Specifying How Stack Checking is Done

GCC will check that stack references are within the boundaries of the stack, if the `-fstack-check` is specified, in one of three ways:

1. If the value of the `STACK_CHECK_BUILTIN` macro is nonzero, GCC will assume that you have arranged for stack checking to be done at appropriate places in the configuration files, e.g., in `TARGET_ASM_FUNCTION_PROLOGUE`. GCC will do not other special processing.
2. If `STACK_CHECK_BUILTIN` is zero and you defined a named pattern called `check_stack` in your `'md'` file, GCC will call that pattern with one argument which is the address to compare the stack value against. You must arrange for this pattern to report an error if the stack pointer is out of range.
3. If neither of the above are true, GCC will generate code to periodically “probe” the stack pointer using the values of the macros defined below.

Normally, you will use the default values of these macros, so GCC will use the third approach.

`STACK_CHECK_BUILTIN`

A nonzero value if stack checking is done by the configuration files in a machine-dependent manner. You should define this macro if stack checking is require by the ABI of your machine or if you would like to have to stack checking in some more efficient way than GCC’s portable approach. The default value of this macro is zero.

`STACK_CHECK_PROBE_INTERVAL`

An integer representing the interval at which GCC must generate stack probe instructions. You will normally define this macro to be no larger than the size of the “guard pages” at the end of a stack area. The default value of 4096 is suitable for most systems.

`STACK_CHECK_PROBE_LOAD`

A integer which is nonzero if GCC should perform the stack probe as a load instruction and zero if GCC should use a store instruction. The default is zero, which is the most efficient choice on most systems.

`STACK_CHECK_PROTECT`

The number of bytes of stack needed to recover from a stack overflow, for languages where such a recovery is supported. The default value of 75 words should be adequate for most machines.

`STACK_CHECK_MAX_FRAME_SIZE`

The maximum size of a stack frame, in bytes. GCC will generate probe instructions in non-leaf functions to ensure at least this many bytes of stack are available. If a stack frame is larger than this size, stack checking will not be reliable and GCC will issue a warning. The default is chosen so that GCC only generates one instruction on most systems. You should normally not change the default value of this macro.

STACK_CHECK_FIXED_FRAME_SIZE

GCC uses this value to generate the above warning message. It represents the amount of fixed frame used by a function, not including space for any callee-saved registers, temporaries and user variables. You need only specify an upper bound for this amount and will normally use the default of four words.

STACK_CHECK_MAX_VAR_SIZE

The maximum size, in bytes, of an object that GCC will place in the fixed area of the stack frame when the user specifies `-fstack-check`. GCC computed the default from the values of the above macros and you will normally not need to override that default.

10.10.4 Registers That Address the Stack Frame

This discusses registers that address the stack frame.

STACK_POINTER_REGNUM

The register number of the stack pointer register, which must also be a fixed register according to `FIXED_REGISTERS`. On most machines, the hardware determines which register this is.

FRAME_POINTER_REGNUM

The register number of the frame pointer register, which is used to access automatic variables in the stack frame. On some machines, the hardware determines which register this is. On other machines, you can choose any register you wish for this purpose.

HARD_FRAME_POINTER_REGNUM

On some machines the offset between the frame pointer and starting offset of the automatic variables is not known until after register allocation has been done (for example, because the saved registers are between these two locations). On those machines, define `FRAME_POINTER_REGNUM` the number of a special, fixed register to be used internally until the offset is known, and define `HARD_FRAME_POINTER_REGNUM` to be the actual hard register number used for the frame pointer.

You should define this macro only in the very rare circumstances when it is not possible to calculate the offset between the frame pointer and the automatic variables until after register allocation has been completed. When this macro is defined, you must also indicate in your definition of `ELIMINABLE_REGS` how to eliminate `FRAME_POINTER_REGNUM` into either `HARD_FRAME_POINTER_REGNUM` or `STACK_POINTER_REGNUM`.

Do not define this macro if it would be the same as `FRAME_POINTER_REGNUM`.

ARG_POINTER_REGNUM

The register number of the arg pointer register, which is used to access the function's argument list. On some machines, this is the same as the frame pointer register. On some machines, the hardware determines which register this is. On other machines, you can choose any register you wish for this purpose. If this is not the same register as the frame pointer register, then you

must mark it as a fixed register according to `FIXED_REGISTERS`, or arrange to be able to eliminate it (see Section 10.10.5 [Elimination], page 238).

`RETURN_ADDRESS_POINTER_REGNUM`

The register number of the return address pointer register, which is used to access the current function's return address from the stack. On some machines, the return address is not at a fixed offset from the frame pointer or stack pointer or argument pointer. This register can be defined to point to the return address on the stack, and then be converted by `ELIMINABLE_REGS` into either the frame pointer or stack pointer.

Do not define this macro unless there is no other way to get the return address from the stack.

`STATIC_CHAIN_REGNUM`

`STATIC_CHAIN_INCOMING_REGNUM`

Register numbers used for passing a function's static chain pointer. If register windows are used, the register number as seen by the called function is `STATIC_CHAIN_INCOMING_REGNUM`, while the register number as seen by the calling function is `STATIC_CHAIN_REGNUM`. If these registers are the same, `STATIC_CHAIN_INCOMING_REGNUM` need not be defined.

The static chain register need not be a fixed register.

If the static chain is passed in memory, these macros should not be defined; instead, the next two macros should be defined.

`STATIC_CHAIN`

`STATIC_CHAIN_INCOMING`

If the static chain is passed in memory, these macros provide `rtx` giving `mem` expressions that denote where they are stored. `STATIC_CHAIN` and `STATIC_CHAIN_INCOMING` give the locations as seen by the calling and called functions, respectively. Often the former will be at an offset from the stack pointer and the latter at an offset from the frame pointer.

The variables `stack_pointer_rtx`, `frame_pointer_rtx`, and `arg_pointer_rtx` will have been initialized prior to the use of these macros and should be used to refer to those items.

If the static chain is passed in a register, the two previous macros should be defined instead.

`DWARF_FRAME_REGISTERS`

This macro specifies the maximum number of hard registers that can be saved in a call frame. This is used to size data structures used in DWARF2 exception handling.

Prior to GCC 3.0, this macro was needed in order to establish a stable exception handling ABI in the face of adding new hard registers for ISA extensions. In GCC 3.0 and later, the EH ABI is insulated from changes in the number of hard registers. Nevertheless, this macro can still be used to reduce the runtime memory requirements of the exception handling routines, which can be substantial if the ISA contains a lot of registers that are not call-saved.

If this macro is not defined, it defaults to `FIRST_PSEUDO_REGISTER`.

PRE_GCC3_DWARF_FRAME_REGISTERS

This macro is similar to `DWARF_FRAME_REGISTERS`, but is provided for backward compatibility in pre GCC 3.0 compiled code.

If this macro is not defined, it defaults to `DWARF_FRAME_REGISTERS`.

10.10.5 Eliminating Frame Pointer and Arg Pointer

This is about eliminating the frame pointer and arg pointer.

FRAME_POINTER_REQUIRED

A C expression which is nonzero if a function must have and use a frame pointer. This expression is evaluated in the reload pass. If its value is nonzero the function will have a frame pointer.

The expression can in principle examine the current function and decide according to the facts, but on most machines the constant 0 or the constant 1 suffices. Use 0 when the machine allows code to be generated with no frame pointer, and doing so saves some time or space. Use 1 when there is no possible advantage to avoiding a frame pointer.

In certain cases, the compiler does not know how to produce valid code without a frame pointer. The compiler recognizes those cases and automatically gives the function a frame pointer regardless of what `FRAME_POINTER_REQUIRED` says. You don't need to worry about them.

In a function that does not require a frame pointer, the frame pointer register can be allocated for ordinary usage, unless you mark it as a fixed register. See `FIXED_REGISTERS` for more information.

INITIAL_FRAME_POINTER_OFFSET (*depth-var*)

A C statement to store in the variable *depth-var* the difference between the frame pointer and the stack pointer values immediately after the function prologue. The value would be computed from information such as the result of `get_frame_size()` and the tables of registers `regs_ever_live` and `call_used_regs`.

If `ELIMINABLE_REGS` is defined, this macro will be not be used and need not be defined. Otherwise, it must be defined even if `FRAME_POINTER_REQUIRED` is defined to always be true; in that case, you may set *depth-var* to anything.

ELIMINABLE_REGS

If defined, this macro specifies a table of register pairs used to eliminate unneeded registers that point into the stack frame. If it is not defined, the only elimination attempted by the compiler is to replace references to the frame pointer with references to the stack pointer.

The definition of this macro is a list of structure initializations, each of which specifies an original and replacement register.

On some machines, the position of the argument pointer is not known until the compilation is completed. In such a case, a separate hard register must be used for the argument pointer. This register can be eliminated by replacing it with

either the frame pointer or the argument pointer, depending on whether or not the frame pointer has been eliminated.

In this case, you might specify:

```
#define ELIMINABLE_REGS \
  {{ARG_POINTER_REGNUM, STACK_POINTER_REGNUM}, \
   {ARG_POINTER_REGNUM, FRAME_POINTER_REGNUM}, \
   {FRAME_POINTER_REGNUM, STACK_POINTER_REGNUM}}
```

Note that the elimination of the argument pointer with the stack pointer is specified first since that is the preferred elimination.

CAN_ELIMINATE (*from-reg*, *to-reg*)

A C expression that returns nonzero if the compiler is allowed to try to replace register number *from-reg* with register number *to-reg*. This macro need only be defined if `ELIMINABLE_REGS` is defined, and will usually be the constant 1, since most of the cases preventing register elimination are things that the compiler already knows about.

INITIAL_ELIMINATION_OFFSET (*from-reg*, *to-reg*, *offset-var*)

This macro is similar to `INITIAL_FRAME_POINTER_OFFSET`. It specifies the initial difference between the specified pair of registers. This macro must be defined if `ELIMINABLE_REGS` is defined.

10.10.6 Passing Function Arguments on the Stack

The macros in this section control how arguments are passed on the stack. See the following section for other macros that control passing certain arguments in registers.

PROMOTE_PROTOTYPES

A C expression whose value is nonzero if an argument declared in a prototype as an integral type smaller than `int` should actually be passed as an `int`. In addition to avoiding errors in certain cases of mismatch, it also makes for better code on certain machines. If the macro is not defined in target header files, it defaults to 0.

PUSH_ARGS

A C expression. If nonzero, push insns will be used to pass outgoing arguments. If the target machine does not have a push instruction, set it to zero. That directs GCC to use an alternate strategy: to allocate the entire argument block and then store the arguments into it. When `PUSH_ARGS` is nonzero, `PUSH_ROUNDING` must be defined too.

PUSH_ROUNDING (*npushed*)

A C expression that is the number of bytes actually pushed onto the stack when an instruction attempts to push *npushed* bytes.

On some machines, the definition

```
#define PUSH_ROUNDING(BYTES) (BYTES)
```

will suffice. But on other machines, instructions that appear to push one byte actually push two bytes in an attempt to maintain alignment. Then the definition should be

```
#define PUSH_ROUNDING(BYTES) (((BYTES) + 1) & ~1)
```

ACCUMULATE_OUTGOING_ARGS

A C expression. If nonzero, the maximum amount of space required for outgoing arguments will be computed and placed into the variable `current_function_outgoing_args_size`. No space will be pushed onto the stack for each call; instead, the function prologue should increase the stack frame size by this amount. Setting both `PUSH_ARGS` and `ACCUMULATE_OUTGOING_ARGS` is not proper.

REG_PARM_STACK_SPACE (*fndecl*)

Define this macro if functions should assume that stack space has been allocated for arguments even when their values are passed in registers.

The value of this macro is the size, in bytes, of the area reserved for arguments passed in registers for the function represented by *fndecl*, which can be zero if GCC is calling a library function.

This space can be allocated by the caller, or be a part of the machine-dependent stack frame: `OUTGOING_REG_PARM_STACK_SPACE` says which.

MAYBE_REG_PARM_STACK_SPACE

FINAL_REG_PARM_STACK_SPACE (*const_size*, *var_size*)

Define these macros in addition to the one above if functions might allocate stack space for arguments even when their values are passed in registers. These should be used when the stack space allocated for arguments in registers is not a simple constant independent of the function declaration.

The value of the first macro is the size, in bytes, of the area that we should initially assume would be reserved for arguments passed in registers.

The value of the second macro is the actual size, in bytes, of the area that will be reserved for arguments passed in registers. This takes two arguments: an integer representing the number of bytes of fixed sized arguments on the stack, and a tree representing the number of bytes of variable sized arguments on the stack.

When these macros are defined, `REG_PARM_STACK_SPACE` will only be called for libcall functions, the current function, or for a function being called when it is known that such stack space must be allocated. In each case this value can be easily computed.

When deciding whether a called function needs such stack space, and how much space to reserve, GCC uses these two macros instead of `REG_PARM_STACK_SPACE`.

OUTGOING_REG_PARM_STACK_SPACE

Define this if it is the responsibility of the caller to allocate the area reserved for arguments passed in registers.

If `ACCUMULATE_OUTGOING_ARGS` is defined, this macro controls whether the space for these arguments counts in the value of `current_function_outgoing_args_size`.

STACK_PARAMS_IN_REG_PARM_AREA

Define this macro if `REG_PARM_STACK_SPACE` is defined, but the stack parameters don't skip the area specified by it.

Normally, when a parameter is not passed in registers, it is placed on the stack beyond the `REG_PARM_STACK_SPACE` area. Defining this macro suppresses this behavior and causes the parameter to be passed on the stack in its natural location.

`RETURN_POPS_ARGS` (*fundecl*, *funtype*, *stack-size*)

A C expression that should indicate the number of bytes of its own arguments that a function pops on returning, or 0 if the function pops no arguments and the caller must therefore pop them all after the function returns.

fundecl is a C variable whose value is a tree node that describes the function in question. Normally it is a node of type `FUNCTION_DECL` that describes the declaration of the function. From this you can obtain the `DECL_ATTRIBUTES` of the function.

funtype is a C variable whose value is a tree node that describes the function in question. Normally it is a node of type `FUNCTION_TYPE` that describes the data type of the function. From this it is possible to obtain the data types of the value and arguments (if known).

When a call to a library function is being considered, *fundecl* will contain an identifier node for the library function. Thus, if you need to distinguish among various library functions, you can do so by their names. Note that “library function” in this context means a function used to perform arithmetic, whose name is known specially in the compiler and was not mentioned in the C code being compiled.

stack-size is the number of bytes of arguments passed on the stack. If a variable number of bytes is passed, it is zero, and argument popping will always be the responsibility of the calling function.

On the VAX, all functions always pop their arguments, so the definition of this macro is *stack-size*. On the 68000, using the standard calling convention, no functions pop their arguments, so the value of the macro is always 0 in this case. But an alternative calling convention is available in which functions that take a fixed number of arguments pop them but other functions (such as `printf`) pop nothing (the caller pops all). When this convention is in use, *funtype* is examined to determine whether a function takes a fixed number of arguments.

`CALL_POPS_ARGS` (*cum*)

A C expression that should indicate the number of bytes a call sequence pops off the stack. It is added to the value of `RETURN_POPS_ARGS` when compiling a function call.

cum is the variable in which all arguments to the called function have been accumulated.

On certain architectures, such as the SH5, a call trampoline is used that pops certain registers off the stack, depending on the arguments that have been passed to the function. Since this is a property of the call site, not of the called function, `RETURN_POPS_ARGS` is not appropriate.

10.10.7 Passing Arguments in Registers

This section describes the macros which let you control how various types of arguments are passed in registers or how they are arranged in the stack.

FUNCTION_ARG (*cum*, *mode*, *type*, *named*)

A C expression that controls whether a function argument is passed in a register, and which register.

The arguments are *cum*, which summarizes all the previous arguments; *mode*, the machine mode of the argument; *type*, the data type of the argument as a tree node or 0 if that is not known (which happens for C support library functions); and *named*, which is 1 for an ordinary argument and 0 for nameless arguments that correspond to ‘...’ in the called function’s prototype. *type* can be an incomplete type if a syntax error has previously occurred.

The value of the expression is usually either a **reg** RTX for the hard register in which to pass the argument, or zero to pass the argument on the stack.

For machines like the VAX and 68000, where normally all arguments are pushed, zero suffices as a definition.

The value of the expression can also be a **parallel** RTX. This is used when an argument is passed in multiple locations. The mode of the **parallel** should be the mode of the entire argument. The **parallel** holds any number of **expr_list** pairs; each one describes where part of the argument is passed. In each **expr_list** the first operand must be a **reg** RTX for the hard register in which to pass this part of the argument, and the mode of the register RTX indicates how large this part of the argument is. The second operand of the **expr_list** is a **const_int** which gives the offset in bytes into the entire argument of where this part starts. As a special exception the first **expr_list** in the **parallel** RTX may have a first operand of zero. This indicates that the entire argument is also stored on the stack.

The last time this macro is called, it is called with `MODE == VOIDmode`, and its result is passed to the **call** or **call_value** pattern as operands 2 and 3 respectively.

The usual way to make the ISO library ‘**stdarg.h**’ work on a machine where some arguments are usually passed in registers, is to cause nameless arguments to be passed on the stack instead. This is done by making **FUNCTION_ARG** return 0 whenever *named* is 0.

You may use the macro **MUST_PASS_IN_STACK** (*mode*, *type*) in the definition of this macro to determine if this argument is of a type that must be passed in the stack. If **REG_PARM_STACK_SPACE** is not defined and **FUNCTION_ARG** returns nonzero for such an argument, the compiler will abort. If **REG_PARM_STACK_SPACE** is defined, the argument will be computed in the stack and then loaded into a register.

MUST_PASS_IN_STACK (*mode*, *type*)

Define as a C expression that evaluates to nonzero if we do not know how to pass **TYPE** solely in registers. The file ‘**expr.h**’ defines a definition that is usually appropriate, refer to ‘**expr.h**’ for additional documentation.

FUNCTION_INCOMING_ARG (*cum, mode, type, named*)

Define this macro if the target machine has “register windows”, so that the register in which a function sees an arguments is not necessarily the same as the one in which the caller passed the argument.

For such machines, **FUNCTION_ARG** computes the register in which the caller passes the value, and **FUNCTION_INCOMING_ARG** should be defined in a similar fashion to tell the function being called where the arguments will arrive.

If **FUNCTION_INCOMING_ARG** is not defined, **FUNCTION_ARG** serves both purposes.

FUNCTION_ARG_PARTIAL_NREGS (*cum, mode, type, named*)

A C expression for the number of words, at the beginning of an argument, that must be put in registers. The value must be zero for arguments that are passed entirely in registers or that are entirely pushed on the stack.

On some machines, certain arguments must be passed partially in registers and partially in memory. On these machines, typically the first *n* words of arguments are passed in registers, and the rest on the stack. If a multi-word argument (a **double** or a structure) crosses that boundary, its first few words must be passed in registers and the rest must be pushed. This macro tells the compiler when this occurs, and how many of the words should go in registers.

FUNCTION_ARG for these arguments should return the first register to be used by the caller for this argument; likewise **FUNCTION_INCOMING_ARG**, for the called function.

FUNCTION_ARG_PASS_BY_REFERENCE (*cum, mode, type, named*)

A C expression that indicates when an argument must be passed by reference. If nonzero for an argument, a copy of that argument is made in memory and a pointer to the argument is passed instead of the argument itself. The pointer is passed in whatever way is appropriate for passing a pointer to that type.

On machines where **REG_PARM_STACK_SPACE** is not defined, a suitable definition of this macro might be

```
#define FUNCTION_ARG_PASS_BY_REFERENCE\
(CUM, MODE, TYPE, NAMED) \
    MUST_PASS_IN_STACK (MODE, TYPE)
```

FUNCTION_ARG_CALLEE_COPIES (*cum, mode, type, named*)

If defined, a C expression that indicates when it is the called function’s responsibility to make a copy of arguments passed by invisible reference. Normally, the caller makes a copy and passes the address of the copy to the routine being called. When **FUNCTION_ARG_CALLEE_COPIES** is defined and is nonzero, the caller does not make a copy. Instead, it passes a pointer to the “live” value. The called function must not modify this value. If it can be determined that the value won’t be modified, it need not make a copy; otherwise a copy must be made.

CUMULATIVE_ARGS

A C type for declaring a variable that is used as the first argument of **FUNCTION_ARG** and other related values. For some target machines, the type **int** suffices and can hold the number of bytes of argument so far.

There is no need to record in `CUMULATIVE_ARGS` anything about the arguments that have been passed on the stack. The compiler has other variables to keep track of that. For target machines on which all arguments are passed on the stack, there is no need to store anything in `CUMULATIVE_ARGS`; however, the data structure must exist and should not be empty, so use `int`.

`INIT_CUMULATIVE_ARGS` (*cum*, *fntype*, *libname*, *indirect*)

A C statement (sans semicolon) for initializing the variable *cum* for the state at the beginning of the argument list. The variable has type `CUMULATIVE_ARGS`. The value of *fntype* is the tree node for the data type of the function which will receive the args, or 0 if the args are to a compiler support library function. The value of *indirect* is nonzero when processing an indirect call, for example a call through a function pointer. The value of *indirect* is zero for a call to an explicitly named function, a library function call, or when `INIT_CUMULATIVE_ARGS` is used to find arguments for the function being compiled.

When processing a call to a compiler support library function, *libname* identifies which one. It is a `symbol_ref` rtx which contains the name of the function, as a string. *libname* is 0 when an ordinary C function call is being processed. Thus, each time this macro is called, either *libname* or *fntype* is nonzero, but never both of them at once.

`INIT_CUMULATIVE_LIBCALL_ARGS` (*cum*, *mode*, *libname*)

Like `INIT_CUMULATIVE_ARGS` but only used for outgoing libcalls, it gets a `MODE` argument instead of *fntype*, that would be `NULL`. *indirect* would always be zero, too. If this macro is not defined, `INIT_CUMULATIVE_ARGS` (*cum*, `NULL_RTX`, *libname*, 0) is used instead.

`INIT_CUMULATIVE_INCOMING_ARGS` (*cum*, *fntype*, *libname*)

Like `INIT_CUMULATIVE_ARGS` but overrides it for the purposes of finding the arguments for the function being compiled. If this macro is undefined, `INIT_CUMULATIVE_ARGS` is used instead.

The value passed for *libname* is always 0, since library routines with special calling conventions are never compiled with GCC. The argument *libname* exists for symmetry with `INIT_CUMULATIVE_ARGS`.

`FUNCTION_ARG_ADVANCE` (*cum*, *mode*, *type*, *named*)

A C statement (sans semicolon) to update the summarizer variable *cum* to advance past an argument in the argument list. The values *mode*, *type* and *named* describe that argument. Once this is done, the variable *cum* is suitable for analyzing the *following* argument with `FUNCTION_ARG`, etc.

This macro need not do anything if the argument in question was passed on the stack. The compiler knows how to track the amount of stack space used for arguments without any special help.

`FUNCTION_ARG_PADDING` (*mode*, *type*)

If defined, a C expression which determines whether, and in which direction, to pad out an argument with extra space. The value should be of type `enum direction`: either `upward` to pad above the argument, `downward` to pad below, or `none` to inhibit padding.

The *amount* of padding is always just enough to reach the next multiple of `FUNCTION_ARG_BOUNDARY`; this macro does not control it.

This macro has a default definition which is right for most systems. For little-endian machines, the default is to pad upward. For big-endian machines, the default is to pad downward for an argument of constant size shorter than an `int`, and upward otherwise.

`PAD_VARARGS_DOWN`

If defined, a C expression which determines whether the default implementation of `va_arg` will attempt to pad down before reading the next argument, if that argument is smaller than its aligned space as controlled by `PARAM_BOUNDARY`. If this macro is not defined, all such arguments are padded down if `BYTES_BIG_ENDIAN` is true.

`FUNCTION_ARG_BOUNDARY` (*mode*, *type*)

If defined, a C expression that gives the alignment boundary, in bits, of an argument with the specified mode and type. If it is not defined, `PARAM_BOUNDARY` is used for all arguments.

`FUNCTION_ARG_REGNO_P` (*regno*)

A C expression that is nonzero if *regno* is the number of a hard register in which function arguments are sometimes passed. This does *not* include implicit arguments such as the static chain and the structure-value address. On many machines, no registers can be used for this purpose since all function arguments are pushed on the stack.

`LOAD_ARGS_REVERSED`

If defined, the order in which arguments are loaded into their respective argument registers is reversed so that the last argument is loaded first. This macro only affects arguments passed in registers.

10.10.8 How Scalar Function Values Are Returned

This section discusses the macros that control returning scalars as values—values that can fit in registers.

`FUNCTION_VALUE` (*valtype*, *func*)

A C expression to create an RTX representing the place where a function returns a value of data type *valtype*. *valtype* is a tree node representing a data type. Write `TYPE_MODE (valtype)` to get the machine mode used to represent that type. On many machines, only the mode is relevant. (Actually, on most machines, scalar values are returned in the same place regardless of mode).

The value of the expression is usually a `reg` RTX for the hard register where the return value is stored. The value can also be a `parallel` RTX, if the return value is in multiple places. See `FUNCTION_ARG` for an explanation of the `parallel` form.

If `PROMOTE_FUNCTION_RETURN` is defined, you must apply the same promotion rules specified in `PROMOTE_MODE` if *valtype* is a scalar type.

If the precise function being called is known, *func* is a tree node (`FUNCTION_DECL`) for it; otherwise, *func* is a null pointer. This makes it possible to use a different value-returning convention for specific functions when all their calls are known.

`FUNCTION_VALUE` is not used for return vales with aggregate data types, because these are returned in another way. See `STRUCT_VALUE_REGNUM` and related macros, below.

`FUNCTION_OUTGOING_VALUE` (*valtype*, *func*)

Define this macro if the target machine has “register windows” so that the register in which a function returns its value is not the same as the one in which the caller sees the value.

For such machines, `FUNCTION_VALUE` computes the register in which the caller will see the value. `FUNCTION_OUTGOING_VALUE` should be defined in a similar fashion to tell the function where to put the value.

If `FUNCTION_OUTGOING_VALUE` is not defined, `FUNCTION_VALUE` serves both purposes.

`FUNCTION_OUTGOING_VALUE` is not used for return vales with aggregate data types, because these are returned in another way. See `STRUCT_VALUE_REGNUM` and related macros, below.

`LIBCALL_VALUE` (*mode*)

A C expression to create an RTX representing the place where a library function returns a value of mode *mode*. If the precise function being called is known, *func* is a tree node (`FUNCTION_DECL`) for it; otherwise, *func* is a null pointer. This makes it possible to use a different value-returning convention for specific functions when all their calls are known.

Note that “library function” in this context means a compiler support routine, used to perform arithmetic, whose name is known specially by the compiler and was not mentioned in the C code being compiled.

The definition of `LIBRARY_VALUE` need not be concerned aggregate data types, because none of the library functions returns such types.

`FUNCTION_VALUE_REGNO_P` (*regno*)

A C expression that is nonzero if *regno* is the number of a hard register in which the values of called function may come back.

A register whose use for returning values is limited to serving as the second of a pair (for a value of type `double`, say) need not be recognized by this macro. So for most machines, this definition suffices:

```
#define FUNCTION_VALUE_REGNO_P(N) ((N) == 0)
```

If the machine has register windows, so that the caller and the called function use different registers for the return value, this macro should recognize only the caller’s register numbers.

`APPLY_RESULT_SIZE`

Define this macro if ‘`untyped_call`’ and ‘`untyped_return`’ need more space than is implied by `FUNCTION_VALUE_REGNO_P` for saving and restoring an arbitrary return value.

10.10.9 How Large Values Are Returned

When a function value's mode is `BLKmode` (and in some other cases), the value is not returned according to `FUNCTION_VALUE` (see Section 10.10.8 [Scalar Return], page 245). Instead, the caller passes the address of a block of memory in which the value should be stored. This address is called the *structure value address*.

This section describes how to control returning structure values in memory.

`RETURN_IN_MEMORY` (*type*)

A C expression which can inhibit the returning of certain function values in registers, based on the type of value. A nonzero value says to return the function value in memory, just as large structures are always returned. Here *type* will be a C expression of type `tree`, representing the data type of the value.

Note that values of mode `BLKmode` must be explicitly handled by this macro. Also, the option `'-fpcc-struct-return'` takes effect regardless of this macro. On most systems, it is possible to leave the macro undefined; this causes a default definition to be used, whose value is the constant 1 for `BLKmode` values, and 0 otherwise.

Do not use this macro to indicate that structures and unions should always be returned in memory. You should instead use `DEFAULT_PCC_STRUCT_RETURN` to indicate this.

`DEFAULT_PCC_STRUCT_RETURN`

Define this macro to be 1 if all structure and union return values must be in memory. Since this results in slower code, this should be defined only if needed for compatibility with other compilers or with an ABI. If you define this macro to be 0, then the conventions used for structure and union return values are decided by the `RETURN_IN_MEMORY` macro.

If not defined, this defaults to the value 1.

`STRUCT_VALUE_REGNUM`

If the structure value address is passed in a register, then `STRUCT_VALUE_REGNUM` should be the number of that register.

`STRUCT_VALUE`

If the structure value address is not passed in a register, define `STRUCT_VALUE` as an expression returning an RTX for the place where the address is passed. If it returns 0, the address is passed as an "invisible" first argument.

`STRUCT_VALUE_INCOMING_REGNUM`

On some architectures the place where the structure value address is found by the called function is not the same place that the caller put it. This can be due to register windows, or it could be because the function prologue moves it to a different place.

If the incoming location of the structure value address is in a register, define this macro as the register number.

`STRUCT_VALUE_INCOMING`

If the incoming location is not a register, then you should define `STRUCT_VALUE_INCOMING` as an expression for an RTX for where the called function should find

the value. If it should find the value on the stack, define this to create a `mem` which refers to the frame pointer. A definition of 0 means that the address is passed as an “invisible” first argument.

PCC_STATIC_STRUCT_RETURN

Define this macro if the usual system convention on the target machine for returning structures and unions is for the called function to return the address of a static variable containing the value.

Do not define this if the usual system convention is for the caller to pass an address to the subroutine.

This macro has effect in ‘`-fpcc-struct-return`’ mode, but it does nothing when you use ‘`-freg-struct-return`’ mode.

10.10.10 Caller-Saves Register Allocation

If you enable it, GCC can save registers around function calls. This makes it possible to use call-clobbered registers to hold variables that must live across calls.

DEFAULT_CALLER_SAVES

Define this macro if function calls on the target machine do not preserve any registers; in other words, if `CALL_USED_REGISTERS` has 1 for all registers. When defined, this macro enables ‘`-fcaller-saves`’ by default for all optimization levels. It has no effect for optimization levels 2 and higher, where ‘`-fcaller-saves`’ is the default.

CALLER_SAVE_PROFITABLE (*refs*, *calls*)

A C expression to determine whether it is worthwhile to consider placing a pseudo-register in a call-clobbered hard register and saving and restoring it around each function call. The expression should be 1 when this is worth doing, and 0 otherwise.

If you don’t define this macro, a default is used which is good on most machines: $4 * \text{calls} < \text{refs}$.

HARD_REGNO_CALLER_SAVE_MODE (*regno*, *nregs*)

A C expression specifying which mode is required for saving *nregs* of a pseudo-register in call-clobbered hard register *regno*. If *regno* is unsuitable for caller save, `VOIDmode` should be returned. For most machines this macro need not be defined since GCC will select the smallest suitable mode.

10.10.11 Function Entry and Exit

This section describes the macros that output function entry (*prologue*) and exit (*epilogue*) code.

```
void TARGET_ASM_FUNCTION_PROLOGUE (FILE *file, [Target Hook]
    HOST_WIDE_INT size)
```

If defined, a function that outputs the assembler code for entry to a function. The prologue is responsible for setting up the stack frame, initializing the frame pointer

register, saving registers that must be saved, and allocating *size* additional bytes of storage for the local variables. *size* is an integer. *file* is a stdio stream to which the assembler code should be output.

The label for the beginning of the function need not be output by this macro. That has already been done when the macro is run.

To determine which registers to save, the macro can refer to the array `regs_ever_live`: element *r* is nonzero if hard register *r* is used anywhere within the function. This implies the function prologue should save register *r*, provided it is not one of the call-used registers. (`TARGET_ASM_FUNCTION_EPILOGUE` must likewise use `regs_ever_live`.)

On machines that have “register windows”, the function entry code does not save on the stack the registers that are in the windows, even if they are supposed to be preserved by function calls; instead it takes appropriate steps to “push” the register stack, if any non-call-used registers are used in the function.

On machines where functions may or may not have frame-pointers, the function entry code must vary accordingly; it must set up the frame pointer if one is wanted, and not otherwise. To determine whether a frame pointer is in wanted, the macro can refer to the variable `frame_pointer_needed`. The variable’s value will be 1 at run time in a function that needs a frame pointer. See Section 10.10.5 [Elimination], page 238.

The function entry code is responsible for allocating any stack space required for the function. This stack space consists of the regions listed below. In most cases, these regions are allocated in the order listed, with the last listed region closest to the top of the stack (the lowest address if `STACK_GROWS_DOWNWARD` is defined, and the highest address if it is not defined). You can use a different order for a machine if doing so is more convenient or required for compatibility reasons. Except in cases where required by standard or by a debugger, there is no reason why the stack layout used by GCC need agree with that used by other compilers for a machine.

`void TARGET_ASM_FUNCTION_END_PROLOGUE (FILE *file)` [Target Hook]

If defined, a function that outputs assembler code at the end of a prologue. This should be used when the function prologue is being emitted as RTL, and you have some extra assembler that needs to be emitted. See [prologue instruction pattern], page 157.

`void TARGET_ASM_FUNCTION_BEGIN_EPILOGUE (FILE *file)` [Target Hook]

If defined, a function that outputs assembler code at the start of an epilogue. This should be used when the function epilogue is being emitted as RTL, and you have some extra assembler that needs to be emitted. See [epilogue instruction pattern], page 157.

`void TARGET_ASM_FUNCTION_EPILOGUE (FILE *file, HOST_WIDE_INT size)` [Target Hook]

If defined, a function that outputs the assembler code for exit from a function. The epilogue is responsible for restoring the saved registers and stack pointer to their values when the function was called, and returning control to the caller. This macro takes the same arguments as the macro `TARGET_ASM_FUNCTION_PROLOGUE`, and the

registers to restore are determined from `regs_ever_live` and `CALL_USED_REGISTERS` in the same way.

On some machines, there is a single instruction that does all the work of returning from the function. On these machines, give that instruction the name `'return'` and do not define the macro `TARGET_ASM_FUNCTION_EPILOGUE` at all.

Do not define a pattern named `'return'` if you want the `TARGET_ASM_FUNCTION_EPILOGUE` to be used. If you want the target switches to control whether return instructions or epilogues are used, define a `'return'` pattern with a validity condition that tests the target switches appropriately. If the `'return'` pattern's validity condition is false, epilogues will be used.

On machines where functions may or may not have frame-pointers, the function exit code must vary accordingly. Sometimes the code for these two cases is completely different. To determine whether a frame pointer is wanted, the macro can refer to the variable `frame_pointer_needed`. The variable's value will be 1 when compiling a function that needs a frame pointer.

Normally, `TARGET_ASM_FUNCTION_PROLOGUE` and `TARGET_ASM_FUNCTION_EPILOGUE` must treat leaf functions specially. The C variable `current_function_is_leaf` is nonzero for such a function. See Section 10.8.4 [Leaf Functions], page 222.

On some machines, some functions pop their arguments on exit while others leave that for the caller to do. For example, the 68020 when given `'-mrttd'` pops arguments in functions that take a fixed number of arguments.

Your definition of the macro `RETURN_POPS_ARGS` decides which functions pop their own arguments. `TARGET_ASM_FUNCTION_EPILOGUE` needs to know what was decided. The variable that is called `current_function_pops_args` is the number of bytes of its arguments that a function should pop. See Section 10.10.8 [Scalar Return], page 245.

- A region of `current_function_pretend_args_size` bytes of uninitialized space just underneath the first argument arriving on the stack. (This may not be at the very start of the allocated stack region if the calling sequence has pushed anything else since pushing the stack arguments. But usually, on such machines, nothing else has been pushed yet, because the function prologue itself does all the pushing.) This region is used on machines where an argument may be passed partly in registers and partly in memory, and, in some cases to support the features in `<stdarg.h>`.
- An area of memory used to save certain registers used by the function. The size of this area, which may also include space for such things as the return address and pointers to previous stack frames, is machine-specific and usually depends on which registers have been used in the function. Machines with register windows often do not require a save area.
- A region of at least `size` bytes, possibly rounded up to an allocation boundary, to contain the local variables of the function. On some machines, this region and the save area may occur in the opposite order, with the save area closer to the top of the stack.
- Optionally, when `ACCUMULATE_OUTGOING_ARGS` is defined, a region of `current_function_outgoing_args_size` bytes to be used for outgoing

argument lists of the function. See Section 10.10.6 [Stack Arguments], page 239.

Normally, it is necessary for the macros `TARGET_ASM_FUNCTION_PROLOGUE` and `TARGET_ASM_FUNCTION_EPILOGUE` to treat leaf functions specially. The C variable `current_function_is_leaf` is nonzero for such a function.

`EXIT_IGNORE_STACK`

Define this macro as a C expression that is nonzero if the return instruction or the function epilogue ignores the value of the stack pointer; in other words, if it is safe to delete an instruction to adjust the stack pointer before a return from the function.

Note that this macro's value is relevant only for functions for which frame pointers are maintained. It is never safe to delete a final stack adjustment in a function that has no frame pointer, and the compiler knows this regardless of `EXIT_IGNORE_STACK`.

`EPILOGUE_USES (regno)`

Define this macro as a C expression that is nonzero for registers that are used by the epilogue or the 'return' pattern. The stack and frame pointer registers are already be assumed to be used as needed.

`EH_USES (regno)`

Define this macro as a C expression that is nonzero for registers that are used by the exception handling mechanism, and so should be considered live on entry to an exception edge.

`DELAY_SLOTS_FOR_EPILOGUE`

Define this macro if the function epilogue contains delay slots to which instructions from the rest of the function can be "moved". The definition should be a C expression whose value is an integer representing the number of delay slots there.

`ELIGIBLE_FOR_EPILOGUE_DELAY (insn, n)`

A C expression that returns 1 if `insn` can be placed in delay slot number `n` of the epilogue.

The argument `n` is an integer which identifies the delay slot now being considered (since different slots may have different rules of eligibility). It is never negative and is always less than the number of epilogue delay slots (what `DELAY_SLOTS_FOR_EPILOGUE` returns). If you reject a particular `insn` for a given delay slot, in principle, it may be reconsidered for a subsequent delay slot. Also, other `insns` may (at least in principle) be considered for the so far unfilled delay slot.

The `insns` accepted to fill the epilogue delay slots are put in an RTL list made with `insn_list` objects, stored in the variable `current_function_epilogue_delay_list`. The `insn` for the first delay slot comes first in the list. Your definition of the macro `TARGET_ASM_FUNCTION_EPILOGUE` should fill the delay slots by outputting the `insns` in this list, usually by calling `final_scan_insn`.

You need not define this macro if you did not define `DELAY_SLOTS_FOR_EPILOGUE`.

```
void TARGET_ASM_OUTPUT_MI_THUNK (FILE *file, tree [Target Hook]
    thunk_fndecl, HOST_WIDE_INT delta, tree function)
```

A function that outputs the assembler code for a thunk function, used to implement C++ virtual function calls with multiple inheritance. The thunk acts as a wrapper around a virtual function, adjusting the implicit object parameter before handing control off to the real function.

First, emit code to add the integer *delta* to the location that contains the incoming first argument. Assume that this argument contains a pointer, and is the one used to pass the `this` pointer in C++. This is the incoming argument *before* the function prologue, e.g. `%o0` on a sparc. The addition must preserve the values of all other incoming arguments.

After the addition, emit code to jump to *function*, which is a `FUNCTION_DECL`. This is a direct pure jump, not a call, and does not touch the return address. Hence returning from *FUNCTION* will return to whoever called the current `thunk`.

The effect must be as if *function* had been called directly with the adjusted first argument. This macro is responsible for emitting all of the code for a thunk function; `TARGET_ASM_FUNCTION_PROLOGUE` and `TARGET_ASM_FUNCTION_EPILOGUE` are not invoked.

The *thunk_fndecl* is redundant. (*delta* and *function* have already been extracted from it.) It might possibly be useful on some targets, but probably not.

If you do not define this macro, the target-independent code in the C++ front end will generate a less efficient heavyweight thunk that calls *function* instead of jumping to it. The generic approach does not support varargs.

```
void TARGET_ASM_OUTPUT_MI_VCALL_THUNK (FILE *file, tree [Target Hook]
    thunk_fndecl, HOST_WIDE_INT delta, int vcall_offset, tree function)
```

A function like `TARGET_ASM_OUTPUT_MI_THUNK`, except that if *vcall_offset* is nonzero, an additional adjustment should be made after adding *delta*. In particular, if *p* is the adjusted pointer, the following adjustment should be made:

```
p += (*((ptrdiff_t **)p))[vcall_offset/sizeof(ptrdiff_t)]
```

If this function is defined, it will always be used in place of `TARGET_ASM_OUTPUT_MI_THUNK`.

10.10.12 Generating Code for Profiling

These macros will help you generate code for profiling.

```
FUNCTION_PROFILER (file, labelno)
```

A C statement or compound statement to output to *file* some assembler code to call the profiling subroutine `mcount`.

The details of how `mcount` expects to be called are determined by your operating system environment, not by GCC. To figure them out, compile a small program for profiling using the system's installed C compiler and look at the assembler code that results.

Older implementations of `mcount` expect the address of a counter variable to be loaded into some register. The name of this variable is ‘LP’ followed by the number *labelno*, so you would generate the name using ‘LP%d’ in a `fprintf`.

PROFILE_HOOK

A C statement or compound statement to output to *file* some assembly code to call the profiling subroutine `mcount` even the target does not support profiling.

NO_PROFILE_COUNTERS

Define this macro if the `mcount` subroutine on your system does not need a counter variable allocated for each function. This is true for almost all modern implementations. If you define this macro, you must not use the *labelno* argument to `FUNCTION_PROFILER`.

PROFILE_BEFORE_PROLOGUE

Define this macro if the code for function profiling should come before the function prologue. Normally, the profiling code comes after.

10.10.13 Permitting tail calls

FUNCTION_OK_FOR_SIBCALL (*decl*)

A C expression that evaluates to true if it is ok to perform a sibling call to *decl* from the current function.

It is not uncommon for limitations of calling conventions to prevent tail calls to functions outside the current unit of translation, or during PIC compilation. Use this macro to enforce these restrictions, as the `sibcall` md pattern can not fail, or fall over to a “normal” call.

10.11 Implementing the Varargs Macros

GCC comes with an implementation of `<varargs.h>` and `<stdarg.h>` that work without change on machines that pass arguments on the stack. Other machines require their own implementations of `varargs`, and the two machine independent header files must have conditionals to include it.

ISO `<stdarg.h>` differs from traditional `<varargs.h>` mainly in the calling convention for `va_start`. The traditional implementation takes just one argument, which is the variable in which to store the argument pointer. The ISO implementation of `va_start` takes an additional second argument. The user is supposed to write the last named argument of the function here.

However, `va_start` should not use this argument. The way to find the end of the named arguments is with the built-in functions described below.

__builtin_saveregs ()

Use this built-in function to save the argument registers in memory so that the `varargs` mechanism can access them. Both ISO and traditional versions of `va_start` must use `__builtin_saveregs`, unless you use `SETUP_INCOMING_VARARGS` (see below) instead.

On some machines, `__builtin_saveregs` is open-coded under the control of the macro `EXPAND_BUILTIN_SAVEREGS`. On other machines, it calls a routine written in assembler language, found in `'libgcc2.c'`.

Code generated for the call to `__builtin_saveregs` appears at the beginning of the function, as opposed to where the call to `__builtin_saveregs` is written, regardless of what the code is. This is because the registers must be saved before the function starts to use them for its own purposes.

`__builtin_args_info (category)`

Use this built-in function to find the first anonymous arguments in registers.

In general, a machine may have several categories of registers used for arguments, each for a particular category of data types. (For example, on some machines, floating-point registers are used for floating-point arguments while other arguments are passed in the general registers.) To make non-varargs functions use the proper calling convention, you have defined the `CUMULATIVE_ARGS` data type to record how many registers in each category have been used so far

`__builtin_args_info` accesses the same data structure of type `CUMULATIVE_ARGS` after the ordinary argument layout is finished with it, with *category* specifying which word to access. Thus, the value indicates the first unused register in a given category.

Normally, you would use `__builtin_args_info` in the implementation of `va_start`, accessing each category just once and storing the value in the `va_list` object. This is because `va_list` will have to update the values, and there is no way to alter the values accessed by `__builtin_args_info`.

`__builtin_next_arg (lastarg)`

This is the equivalent of `__builtin_args_info`, for stack arguments. It returns the address of the first anonymous stack argument, as type `void *`. If `ARGS_GROW_DOWNWARD`, it returns the address of the location above the first anonymous stack argument. Use it in `va_start` to initialize the pointer for fetching arguments from the stack. Also use it in `va_start` to verify that the second parameter *lastarg* is the last named argument of the current function.

`__builtin_classify_type (object)`

Since each machine has its own conventions for which data types are passed in which kind of register, your implementation of `va_arg` has to embody these conventions. The easiest way to categorize the specified data type is to use `__builtin_classify_type` together with `sizeof` and `__alignof__`.

`__builtin_classify_type` ignores the value of *object*, considering only its data type. It returns an integer describing what kind of type that is—integer, floating, pointer, structure, and so on.

The file `'typeclass.h'` defines an enumeration that you can use to interpret the values of `__builtin_classify_type`.

These machine description macros help implement varargs:

EXPAND_BUILTIN_SAVEREGS ()

If defined, is a C expression that produces the machine-specific code for a call to `__builtin_saveregs`. This code will be moved to the very beginning of the function, before any parameter access are made. The return value of this function should be an RTX that contains the value to use as the return of `__builtin_saveregs`.

SETUP_INCOMING_VARARGS (*args_so_far*, *mode*, *type*, *pretend_args_size*, *second_time*)

This macro offers an alternative to using `__builtin_saveregs` and defining the macro `EXPAND_BUILTIN_SAVEREGS`. Use it to store the anonymous register arguments into the stack so that all the arguments appear to have been passed consecutively on the stack. Once this is done, you can use the standard implementation of `varargs` that works for machines that pass all their arguments on the stack.

The argument *args_so_far* is the `CUMULATIVE_ARGS` data structure, containing the values that are obtained after processing the named arguments. The arguments *mode* and *type* describe the last named argument—its machine mode and its data type as a tree node.

The macro implementation should do two things: first, push onto the stack all the argument registers *not* used for the named arguments, and second, store the size of the data thus pushed into the `int`-valued variable whose name is supplied as the argument *pretend_args_size*. The value that you store here will serve as additional offset for setting up the stack frame.

Because you must generate code to push the anonymous arguments at compile time without knowing their data types, `SETUP_INCOMING_VARARGS` is only useful on machines that have just a single category of argument register and use it uniformly for all data types.

If the argument *second_time* is nonzero, it means that the arguments of the function are being analyzed for the second time. This happens for an inline function, which is not actually compiled until the end of the source file. The macro `SETUP_INCOMING_VARARGS` should not generate any instructions in this case.

STRICT_ARGUMENT_NAMING

Define this macro to be a nonzero value if the location where a function argument is passed depends on whether or not it is a named argument.

This macro controls how the *named* argument to `FUNCTION_ARG` is set for `varargs` and `stdarg` functions. If this macro returns a nonzero value, the *named* argument is always true for named arguments, and false for unnamed arguments. If it returns a value of zero, but `SETUP_INCOMING_VARARGS` is defined, then all arguments are treated as named. Otherwise, all named arguments except the last are treated as named.

You need not define this macro if it always returns zero.

PRETEND_OUTGOING_VARARGS_NAMED

If you need to conditionally change ABIs so that one works with `SETUP_INCOMING_VARARGS`, but the other works like neither `SETUP_INCOMING_VARARGS`

nor `STRICT_ARGUMENT_NAMING` was defined, then define this macro to return nonzero if `SETUP_INCOMING_VARARGS` is used, zero otherwise. Otherwise, you should not define this macro.

10.12 Trampolines for Nested Functions

A *trampoline* is a small piece of code that is created at run time when the address of a nested function is taken. It normally resides on the stack, in the stack frame of the containing function. These macros tell GCC how to generate code to allocate and initialize a trampoline.

The instructions in the trampoline must do two things: load a constant address into the static chain register, and jump to the real address of the nested function. On CISC machines such as the m68k, this requires two instructions, a move immediate and a jump. Then the two addresses exist in the trampoline as word-long immediate operands. On RISC machines, it is often necessary to load each address into a register in two parts. Then pieces of each address form separate immediate operands.

The code generated to initialize the trampoline must store the variable parts—the static chain value and the function address—into the immediate operands of the instructions. On a CISC machine, this is simply a matter of copying each address to a memory reference at the proper offset from the start of the trampoline. On a RISC machine, it may be necessary to take out pieces of the address and store them separately.

TRAMPOLINE_TEMPLATE (*file*)

A C statement to output, on the stream *file*, assembler code for a block of data that contains the constant parts of a trampoline. This code should not include a label—the label is taken care of automatically.

If you do not define this macro, it means no template is needed for the target. Do not define this macro on systems where the block move code to copy the trampoline into place would be larger than the code to generate it on the spot.

TRAMPOLINE_SECTION

The name of a subroutine to switch to the section in which the trampoline template is to be placed (see Section 10.18 [Sections], page 275). The default is a value of `'readonly_data_section'`, which places the trampoline in the section containing read-only data.

TRAMPOLINE_SIZE

A C expression for the size in bytes of the trampoline, as an integer.

TRAMPOLINE_ALIGNMENT

Alignment required for trampolines, in bits.

If you don't define this macro, the value of `BIGGEST_ALIGNMENT` is used for aligning trampolines.

INITIALIZE_TRAMPOLINE (*addr*, *fnaddr*, *static_chain*)

A C statement to initialize the variable parts of a trampoline. *addr* is an RTX for the address of the trampoline; *fnaddr* is an RTX for the address of the nested function; *static_chain* is an RTX for the static chain value that should be passed to the function when it is called.

TRAMPOLINE_ADJUST_ADDRESS (*addr*)

A C statement that should perform any machine-specific adjustment in the address of the trampoline. Its argument contains the address that was passed to `INITIALIZE_TRAMPOLINE`. In case the address to be used for a function call should be different from the address in which the template was stored, the different address should be assigned to *addr*. If this macro is not defined, *addr* will be used for function calls.

ALLOCATE_TRAMPOLINE (*fp*)

A C expression to allocate run-time space for a trampoline. The expression value should be an RTX representing a memory reference to the space for the trampoline.

If this macro is not defined, by default the trampoline is allocated as a stack slot. This default is right for most machines. The exceptions are machines where it is impossible to execute instructions in the stack area. On such machines, you may have to implement a separate stack, using this macro in conjunction with `TARGET_ASM_FUNCTION_PROLOGUE` and `TARGET_ASM_FUNCTION_EPILOGUE`.

fp points to a data structure, a `struct function`, which describes the compilation status of the immediate containing function of the function which the trampoline is for. Normally (when `ALLOCATE_TRAMPOLINE` is not defined), the stack slot for the trampoline is in the stack frame of this containing function. Other allocation strategies probably must do something analogous with this information.

Implementing trampolines is difficult on many machines because they have separate instruction and data caches. Writing into a stack location fails to clear the memory in the instruction cache, so when the program jumps to that location, it executes the old contents.

Here are two possible solutions. One is to clear the relevant parts of the instruction cache whenever a trampoline is set up. The other is to make all trampolines identical, by having them jump to a standard subroutine. The former technique makes trampoline execution faster; the latter makes initialization faster.

To clear the instruction cache when a trampoline is initialized, define the following macros which describe the shape of the cache.

INSN_CACHE_SIZE

The total size in bytes of the cache.

INSN_CACHE_LINE_WIDTH

The length in bytes of each cache line. The cache is divided into cache lines which are disjoint slots, each holding a contiguous chunk of data fetched from memory. Each time data is brought into the cache, an entire line is read at once. The data loaded into a cache line is always aligned on a boundary equal to the line size.

INSN_CACHE_DEPTH

The number of alternative cache lines that can hold any particular memory location.

Alternatively, if the machine has system calls or instructions to clear the instruction cache directly, you can define the following macro.

CLEAR_INSN_CACHE (*beg*, *end*)

If defined, expands to a C expression clearing the *instruction cache* in the specified interval. If it is not defined, and the macro `INSN_CACHE_SIZE` is defined, some generic code is generated to clear the cache. The definition of this macro would typically be a series of `asm` statements. Both *beg* and *end* are both pointer expressions.

To use a standard subroutine, define the following macro. In addition, you must make sure that the instructions in a trampoline fill an entire cache line with identical instructions, or else ensure that the beginning of the trampoline code is always aligned at the same point in its cache line. Look in `'m68k.h'` as a guide.

TRANSFER_FROM_TRAMPOLINE

Define this macro if trampolines need a special subroutine to do their work. The macro should expand to a series of `asm` statements which will be compiled with GCC. They go in a library function named `__transfer_from_trampoline`.

If you need to avoid executing the ordinary prologue code of a compiled C function when you jump to the subroutine, you can do so by placing a special label of your own in the assembler code. Use one `asm` statement to generate an assembler label, and another to make the label global. Then trampolines can use that label to jump directly to your special assembler code.

10.13 Implicit Calls to Library Routines

Here is an explanation of implicit calls to library routines.

MULSI3_LIBCALL

A C string constant giving the name of the function to call for multiplication of one signed full-word by another. If you do not define this macro, the default name is used, which is `__mulsi3`, a function defined in `'libgcc.a'`.

DIVSI3_LIBCALL

A C string constant giving the name of the function to call for division of one signed full-word by another. If you do not define this macro, the default name is used, which is `__divsi3`, a function defined in `'libgcc.a'`.

UDIVSI3_LIBCALL

A C string constant giving the name of the function to call for division of one unsigned full-word by another. If you do not define this macro, the default name is used, which is `__udivsi3`, a function defined in `'libgcc.a'`.

MODSI3_LIBCALL

A C string constant giving the name of the function to call for the remainder in division of one signed full-word by another. If you do not define this macro, the default name is used, which is `__modsi3`, a function defined in `'libgcc.a'`.

UMODSI3_LIBCALL

A C string constant giving the name of the function to call for the remainder in division of one unsigned full-word by another. If you do not define this macro, the default name is used, which is `__umodsi3`, a function defined in `'libgcc.a'`.

MULDI3_LIBCALL

A C string constant giving the name of the function to call for multiplication of one signed double-word by another. If you do not define this macro, the default name is used, which is `__muldi3`, a function defined in `'libgcc.a'`.

DIVDI3_LIBCALL

A C string constant giving the name of the function to call for division of one signed double-word by another. If you do not define this macro, the default name is used, which is `__divdi3`, a function defined in `'libgcc.a'`.

UDIVDI3_LIBCALL

A C string constant giving the name of the function to call for division of one unsigned full-word by another. If you do not define this macro, the default name is used, which is `__udivdi3`, a function defined in `'libgcc.a'`.

MODDI3_LIBCALL

A C string constant giving the name of the function to call for the remainder in division of one signed double-word by another. If you do not define this macro, the default name is used, which is `__moddi3`, a function defined in `'libgcc.a'`.

UMODDI3_LIBCALL

A C string constant giving the name of the function to call for the remainder in division of one unsigned full-word by another. If you do not define this macro, the default name is used, which is `__umoddi3`, a function defined in `'libgcc.a'`.

DECLARE_LIBRARY_RENAMES

This macro, if defined, should expand to a piece of C code that will get expanded when compiling functions for `libgcc.a`. It can be used to provide alternate names for `gcc`'s internal library functions if there are ABI-mandated names that the compiler should provide.

INIT_TARGET_OPTABS

Define this macro as a C statement that declares additional library routines renames existing ones. `init_optabs` calls this macro after initializing all the normal library routines.

FLOAT_LIB_COMPARE_RETURNS_BOOL

Define this macro as a C statement that returns nonzero if a call to the floating point comparison library function will return a boolean value that indicates the result of the comparison. It should return zero if one of `gcc`'s own `libgcc` functions is called.

Most ports don't need to define this macro.

TARGET_EDOM

The value of `EDOM` on the target machine, as a C integer constant expression. If you don't define this macro, `GCC` does not attempt to deposit the value of `EDOM` into `errno` directly. Look in `'/usr/include/errno.h'` to find the value of `EDOM` on your system.

If you do not define `TARGET_EDOM`, then compiled code reports domain errors by calling the library function and letting it report the error. If mathematical functions on your system use `matherr` when there is an error, then you should leave `TARGET_EDOM` undefined so that `matherr` is used normally.

GEN_ERRNO_RTX

Define this macro as a C expression to create an rtl expression that refers to the global “variable” `errno`. (On certain systems, `errno` may not actually be a variable.) If you don’t define this macro, a reasonable default is used.

TARGET_MEM_FUNCTIONS

Define this macro if GCC should generate calls to the ISO C (and System V) library functions `memcpy`, `memmove` and `memset` rather than the BSD functions `bcopy` and `bzero`.

LIBGCC_NEEDS_DOUBLE

Define this macro if `float` arguments cannot be passed to library routines (so they must be converted to `double`). This macro affects both how library calls are generated and how the library routines in ‘`libgcc.a`’ accept their arguments. It is useful on machines where floating and fixed point arguments are passed differently, such as the i860.

NEXT_OBJC_RUNTIME

Define this macro to generate code for Objective-C message sending using the calling convention of the NeXT system. This calling convention involves passing the object, the selector and the method arguments all at once to the method-lookup library function.

The default calling convention passes just the object and the selector to the lookup function, which returns a pointer to the method.

10.14 Addressing Modes

This is about addressing modes.

HAVE_PRE_INCREMENT**HAVE_PRE_DECREMENT****HAVE_POST_INCREMENT****HAVE_POST_DECREMENT**

A C expression that is nonzero if the machine supports pre-increment, pre-decrement, post-increment, or post-decrement addressing respectively.

HAVE_PRE_MODIFY_DISP**HAVE_POST_MODIFY_DISP**

A C expression that is nonzero if the machine supports pre- or post-address side-effect generation involving constants other than the size of the memory operand.

HAVE_PRE_MODIFY_REG**HAVE_POST_MODIFY_REG**

A C expression that is nonzero if the machine supports pre- or post-address side-effect generation involving a register displacement.

CONSTANT_ADDRESS_P (x)

A C expression that is 1 if the RTX `x` is a constant which is a valid address. On most machines, this can be defined as `CONSTANT_P (x)`, but a few machines are more restrictive in which constant addresses are supported.

`CONSTANT_P` accepts integer-values expressions whose values are not explicitly known, such as `symbol_ref`, `label_ref`, and `high` expressions and `const` arithmetic expressions, in addition to `const_int` and `const_double` expressions.

`MAX_REGS_PER_ADDRESS`

A number, the maximum number of registers that can appear in a valid memory address. Note that it is up to you to specify a value equal to the maximum number that `GO_IF_LEGITIMATE_ADDRESS` would ever accept.

`GO_IF_LEGITIMATE_ADDRESS (mode, x, label)`

A C compound statement with a conditional `goto label`; executed if `x` (an RTX) is a legitimate memory address on the target machine for a memory operand of mode `mode`.

It usually pays to define several simpler macros to serve as subroutines for this one. Otherwise it may be too complicated to understand.

This macro must exist in two variants: a strict variant and a non-strict one. The strict variant is used in the reload pass. It must be defined so that any pseudo-register that has not been allocated a hard register is considered a memory reference. In contexts where some kind of register is required, a pseudo-register with no hard register must be rejected.

The non-strict variant is used in other passes. It must be defined to accept all pseudo-registers in every context where some kind of register is required.

Compiler source files that want to use the strict variant of this macro define the macro `REG_OK_STRICT`. You should use an `#ifdef REG_OK_STRICT` conditional to define the strict variant in that case and the non-strict variant otherwise.

Subroutines to check for acceptable registers for various purposes (one for base registers, one for index registers, and so on) are typically among the subroutines used to define `GO_IF_LEGITIMATE_ADDRESS`. Then only these subroutine macros need have two variants; the higher levels of macros may be the same whether strict or not.

Normally, constant addresses which are the sum of a `symbol_ref` and an integer are stored inside a `const` RTX to mark them as constant. Therefore, there is no need to recognize such sums specifically as legitimate addresses. Normally you would simply recognize any `const` as legitimate.

Usually `PRINT_OPERAND_ADDRESS` is not prepared to handle constant sums that are not marked with `const`. It assumes that a naked `plus` indicates indexing. If so, then you *must* reject such naked constant sums as illegitimate addresses, so that none of them will be given to `PRINT_OPERAND_ADDRESS`.

On some machines, whether a symbolic address is legitimate depends on the section that the address refers to. On these machines, define the target hook `TARGET_ENCODE_SECTION_INFO` to store the information into the `symbol_ref`, and then check for it here. When you see a `const`, you will have to look inside it to find the `symbol_ref` in order to determine the section. See Section 10.20 [Assembler Format], page 279.

The best way to modify the name string is by adding text to the beginning, with suitable punctuation to prevent any ambiguity. Allocate the new name in

`saveable_obstack`. You will have to modify `ASM_OUTPUT_LABELREF` to remove and decode the added text and output the name accordingly, and define `TARGET_STRIP_NAME_ENCODING` to access the original name string.

You can check the information stored here into the `symbol_ref` in the definitions of the macros `GO_IF_LEGITIMATE_ADDRESS` and `PRINT_OPERAND_ADDRESS`.

`REG_OK_FOR_BASE_P (x)`

A C expression that is nonzero if `x` (assumed to be a `reg RTX`) is valid for use as a base register. For hard registers, it should always accept those which the hardware permits and reject the others. Whether the macro accepts or rejects pseudo registers must be controlled by `REG_OK_STRICT` as described above. This usually requires two variant definitions, of which `REG_OK_STRICT` controls the one actually used.

`REG_MODE_OK_FOR_BASE_P (x, mode)`

A C expression that is just like `REG_OK_FOR_BASE_P`, except that that expression may examine the mode of the memory reference in `mode`. You should define this macro if the mode of the memory reference affects whether a register may be used as a base register. If you define this macro, the compiler will use it instead of `REG_OK_FOR_BASE_P`.

`REG_OK_FOR_INDEX_P (x)`

A C expression that is nonzero if `x` (assumed to be a `reg RTX`) is valid for use as an index register.

The difference between an index register and a base register is that the index register may be scaled. If an address involves the sum of two registers, neither one of them scaled, then either one may be labeled the “base” and the other the “index”; but whichever labeling is used must fit the machine’s constraints of which registers may serve in each capacity. The compiler will try both labelings, looking for one that is valid, and will reload one or both registers only if neither labeling works.

`FIND_BASE_TERM (x)`

A C expression to determine the base term of address `x`. This macro is used in only one place: ‘`find_base_term`’ in `alias.c`.

It is always safe for this macro to not be defined. It exists so that alias analysis can understand machine-dependent addresses.

The typical use of this macro is to handle addresses containing a `label_ref` or `symbol_ref` within an `UNSPEC`.

`LEGITIMIZE_ADDRESS (x, oldx, mode, win)`

A C compound statement that attempts to replace `x` with a valid memory address for an operand of mode `mode`. `win` will be a C statement label elsewhere in the code; the macro definition may use

```
GO_IF_LEGITIMATE_ADDRESS (mode, x, win);
```

to avoid further processing if the address has become legitimate.

`x` will always be the result of a call to `break_out_memory_refs`, and `oldx` will be the operand that was given to that function to produce `x`.

The code generated by this macro should not alter the substructure of *x*. If it transforms *x* into a more legitimate form, it should assign *x* (which will always be a C variable) a new value.

It is not necessary for this macro to come up with a legitimate address. The compiler has standard ways of doing so in all cases. In fact, it is safe for this macro to do nothing. But often a machine-dependent strategy can generate better code.

LEGITIMIZE_RELOAD_ADDRESS (*x*, *mode*, *opnum*, *type*, *ind_levels*, *win*)

A C compound statement that attempts to replace *x*, which is an address that needs reloading, with a valid memory address for an operand of mode *mode*. *win* will be a C statement label elsewhere in the code. It is not necessary to define this macro, but it might be useful for performance reasons.

For example, on the i386, it is sometimes possible to use a single reload register instead of two by reloading a sum of two pseudo registers into a register. On the other hand, for number of RISC processors offsets are limited so that often an intermediate address needs to be generated in order to address a stack slot. By defining **LEGITIMIZE_RELOAD_ADDRESS** appropriately, the intermediate addresses generated for adjacent some stack slots can be made identical, and thus be shared.

Note: This macro should be used with caution. It is necessary to know something of how reload works in order to effectively use this, and it is quite easy to produce macros that build in too much knowledge of reload internals.

Note: This macro must be able to reload an address created by a previous invocation of this macro. If it fails to handle such addresses then the compiler may generate incorrect code or abort.

The macro definition should use **push_reload** to indicate parts that need reloading; *opnum*, *type* and *ind_levels* are usually suitable to be passed unaltered to **push_reload**.

The code generated by this macro must not alter the substructure of *x*. If it transforms *x* into a more legitimate form, it should assign *x* (which will always be a C variable) a new value. This also applies to parts that you change indirectly by calling **push_reload**.

The macro definition may use **strict_memory_address_p** to test if the address has become legitimate.

If you want to change only a part of *x*, one standard way of doing this is to use **copy_rtx**. Note, however, that it unshares only a single level of rtl. Thus, if the part to be changed is not at the top level, you'll need to replace first the top level. It is not necessary for this macro to come up with a legitimate address; but often a machine-dependent strategy can generate better code.

GO_IF_MODE_DEPENDENT_ADDRESS (*addr*, *label*)

A C statement or compound statement with a conditional **goto label**; executed if memory address *x* (an RTX) can have different meanings depending on the machine mode of the memory reference it is used for or if the address is valid for some modes but not others.

Autoincrement and autodecrement addresses typically have mode-dependent effects because the amount of the increment or decrement is the size of the operand being addressed. Some machines have other mode-dependent addresses. Many RISC machines have no mode-dependent addresses.

You may assume that *addr* is a valid address for the machine.

LEGITIMATE_CONSTANT_P (*x*)

A C expression that is nonzero if *x* is a legitimate constant for an immediate operand on the target machine. You can assume that *x* satisfies `CONSTANT_P`, so you need not check this. In fact, ‘1’ is a suitable definition for this macro on machines where anything `CONSTANT_P` is valid.

10.15 Condition Code Status

This describes the condition code status.

The file ‘`conditions.h`’ defines a variable `cc_status` to describe how the condition code was computed (in case the interpretation of the condition code depends on the instruction that it was set by). This variable contains the RTL expressions on which the condition code is currently based, and several standard flags.

Sometimes additional machine-specific flags must be defined in the machine description header file. It can also add additional machine-specific information by defining `CC_STATUS_MDEP`.

CC_STATUS_MDEP

C code for a data type which is used for declaring the `mdep` component of `cc_status`. It defaults to `int`.

This macro is not used on machines that do not use `cc0`.

CC_STATUS_MDEP_INIT

A C expression to initialize the `mdep` field to “empty”. The default definition does nothing, since most machines don’t use the field anyway. If you want to use the field, you should probably define this macro to initialize it.

This macro is not used on machines that do not use `cc0`.

NOTICE_UPDATE_CC (*exp*, *insn*)

A C compound statement to set the components of `cc_status` appropriately for an `insn` whose body is *exp*. It is this macro’s responsibility to recognize insns that set the condition code as a byproduct of other activity as well as those that explicitly set (`cc0`).

This macro is not used on machines that do not use `cc0`.

If there are insns that do not set the condition code but do alter other machine registers, this macro must check to see whether they invalidate the expressions that the condition code is recorded as reflecting. For example, on the 68000, insns that store in address registers do not set the condition code, which means that usually `NOTICE_UPDATE_CC` can leave `cc_status` unaltered for such insns. But suppose that the previous `insn` set the condition code based on location

'a4@102)' and the current insn stores a new value in 'a4'. Although the condition code is not changed by this, it will no longer be true that it reflects the contents of 'a4@102'. Therefore, NOTICE_UPDATE_CC must alter `cc_status` in this case to say that nothing is known about the condition code value.

The definition of NOTICE_UPDATE_CC must be prepared to deal with the results of peephole optimization: insns whose patterns are `parallel` RTXs containing various `reg`, `mem` or constants which are just the operands. The RTL structure of these insns is not sufficient to indicate what the insns actually do. What NOTICE_UPDATE_CC should do when it sees one is just to run `CC_STATUS_INIT`.

A possible definition of NOTICE_UPDATE_CC is to call a function that looks at an attribute (see Section 9.18 [Insn Attributes], page 175) named, for example, 'cc'. This avoids having detailed information about patterns in two places, the 'md' file and in NOTICE_UPDATE_CC.

EXTRA_CC_MODES

Condition codes are represented in registers by machine modes of class `MODE_CC`. By default, there is just one mode, `CCmode`, with this class. If you need more such modes, create a file named '`machine-modes.def`' in your '`config/machine`' directory (see Section 5.3.9 [Anatomy of a Target Back End], page 23), containing a list of these modes. Each entry in the list should be a call to the macro `CC`. This macro takes one argument, which is the name of the mode: it should begin with 'CC'. Do not put quotation marks around the name, or include the trailing 'mode'; these are automatically added. There should not be anything else in the file except comments.

A sample '`machine-modes.def`' file might look like this:

```
CC (CC_NOOV) /* Comparison only valid if there was no overflow. */
CC (CCFP) /* Floating point comparison that cannot trap. */
CC (CCFPE) /* Floating point comparison that may trap. */
```

When you create this file, the macro `EXTRA_CC_MODES` is automatically defined by `configure`, with value '1'.

SELECT_CC_MODE (*op*, *x*, *y*)

Returns a mode from class `MODE_CC` to be used when comparison operation code *op* is applied to rtx *x* and *y*. For example, on the SPARC, `SELECT_CC_MODE` is defined as (see Section 9.11 [Jump Patterns], page 160 for a description of the reason for this definition)

```
#define SELECT_CC_MODE(OP,X,Y) \
  (GET_MODE_CLASS (GET_MODE (X)) == MODE_FLOAT \
   ? ((OP == EQ || OP == NE) ? CCFPmode : CCFPEmode) \
   : ((GET_CODE (X) == PLUS || GET_CODE (X) == MINUS \
      || GET_CODE (X) == NEG) \
      ? CC_NOOVmode : CCmode))
```

You need not define this macro if `EXTRA_CC_MODES` is not defined.

CANONICALIZE_COMPARISON (*code*, *op0*, *op1*)

On some machines not all possible comparisons are defined, but you can convert an invalid comparison into a valid one. For example, the Alpha does not have a `GT` comparison, but you can use an `LT` comparison instead and swap the order of the operands.

On such machines, define this macro to be a C statement to do any required conversions. *code* is the initial comparison code and *op0* and *op1* are the left and right operands of the comparison, respectively. You should modify *code*, *op0*, and *op1* as required.

GCC will not assume that the comparison resulting from this macro is valid but will see if the resulting insn matches a pattern in the ‘md’ file.

You need not define this macro if it would never change the comparison code or operands.

REVERSIBLE_CC_MODE (*mode*)

A C expression whose value is one if it is always safe to reverse a comparison whose mode is *mode*. If SELECT_CC_MODE can ever return *mode* for a floating-point inequality comparison, then REVERSIBLE_CC_MODE (*mode*) must be zero. You need not define this macro if it would always returns zero or if the floating-point format is anything other than IEEE_FLOAT_FORMAT. For example, here is the definition used on the SPARC, where floating-point inequality comparisons are always given CCFPEmode:

```
#define REVERSIBLE_CC_MODE(MODE) ((MODE) != CCFPEmode)
```

A C expression whose value is reversed condition code of the *code* for comparison done in CC_MODE *mode*. The macro is used only in case REVERSIBLE_CC_MODE (*mode*) is nonzero. Define this macro in case machine has some non-standard way how to reverse certain conditionals. For instance in case all floating point conditions are non-trapping, compiler may freely convert unordered compares to ordered one. Then definition may look like:

```
#define REVERSE_CONDITION(CODE, MODE) \
((MODE) != CCFPmode ? reverse_condition (CODE) \
: reverse_condition_maybe_unordered (CODE))
```

REVERSE_CONDEXEC_PREDICATES_P (*code1*, *code2*)

A C expression that returns true if the conditional execution predicate *code1* is the inverse of *code2* and vice versa. Define this to return 0 if the target has conditional execution predicates that cannot be reversed safely. If no expansion is specified, this macro is defined as follows:

```
#define REVERSE_CONDEXEC_PREDICATES_P (x, y) \
((x) == reverse_condition (y))
```

bool TARGET_FIXED_CONDITION_CODE_REGS (unsigned int *, [Target Hook] unsigned int *)

On targets which do not use (cc0), and which use a hard register rather than a pseudo-register to hold condition codes, the regular CSE passes are often not able to identify cases in which the hard register is set to a common value. Use this hook to enable a small pass which optimizes such cases. This hook should return true to enable this pass, and it should set the integers to which its arguments point to the hard register numbers used for condition codes. When there is only one such register, as is true on most systems, the integer pointed to by the second argument should be set to INVALID_REGNUM.

The default version of this hook returns false.

```
enum machine_mode TARGET_CC_MODES_COMPATIBLE (enum machine_mode, enum machine_mode) [Target Hook]
```

On targets which use multiple condition code modes in class `MODE_CC`, it is sometimes the case that a comparison can be validly done in more than one mode. On such a system, define this target hook to take two mode arguments and to return a mode in which both comparisons may be validly done. If there is no such mode, return `VOIDmode`.

The default version of this hook checks whether the modes are the same. If they are, it returns that mode. If they are different, it returns `VOIDmode`.

10.16 Describing Relative Costs of Operations

These macros let you describe the relative speed of various operations on the target machine.

```
CONST_COSTS (x, code, outer_code)
```

A part of a C `switch` statement that describes the relative costs of constant RTL expressions. It must contain `case` labels for expression codes `const_int`, `const`, `symbol_ref`, `label_ref` and `const_double`. Each case must ultimately reach a `return` statement to return the relative cost of the use of that kind of constant value in an expression. The cost may depend on the precise value of the constant, which is available for examination in `x`, and the rtx code of the expression in which it is contained, found in `outer_code`.

`code` is the expression code—redundant, since it can be obtained with `GET_CODE (x)`.

```
RTX_COSTS (x, code, outer_code)
```

Like `CONST_COSTS` but applies to nonconstant RTL expressions. This can be used, for example, to indicate how costly a multiply instruction is. In writing this macro, you can use the construct `COSTS_N_INSNS (n)` to specify a cost equal to `n` fast instructions. `outer_code` is the code of the expression in which `x` is contained.

This macro is optional; do not define it if the default cost assumptions are adequate for the target machine.

```
DEFAULT_RTX_COSTS (x, code, outer_code)
```

This macro, if defined, is called for any case not handled by the `RTX_COSTS` or `CONST_COSTS` macros. This eliminates the need to put case labels into the macro, but the code, or any functions it calls, must assume that the RTL in `x` could be of any type that has not already been handled. The arguments are the same as for `RTX_COSTS`, and the macro should execute a return statement giving the cost of any RTL expressions that it can handle. The default cost calculation is used for any RTL for which this macro does not return a value.

This macro is optional; do not define it if the default cost assumptions are adequate for the target machine.

ADDRESS_COST (*address*)

An expression giving the cost of an addressing mode that contains *address*. If not defined, the cost is computed from the *address* expression and the `CONST_COSTS` values.

For most CISC machines, the default cost is a good approximation of the true cost of the addressing mode. However, on RISC machines, all instructions normally have the same length and execution time. Hence all addresses will have equal costs.

In cases where more than one form of an address is known, the form with the lowest cost will be used. If multiple forms have the same, lowest, cost, the one that is the most complex will be used.

For example, suppose an address that is equal to the sum of a register and a constant is used twice in the same basic block. When this macro is not defined, the address will be computed in a register and memory references will be indirect through that register. On machines where the cost of the addressing mode containing the sum is no higher than that of a simple indirect reference, this will produce an additional instruction and possibly require an additional register. Proper specification of this macro eliminates this overhead for such machines.

Similar use of this macro is made in strength reduction of loops.

address need not be valid as an address. In such a case, the cost is not relevant and can be any value; invalid addresses need not be assigned a different cost.

On machines where an address involving more than one register is as cheap as an address computation involving only one register, defining `ADDRESS_COST` to reflect this can cause two registers to be live over a region of code where only one would have been if `ADDRESS_COST` were not defined in that manner. This effect should be considered in the definition of this macro. Equivalent costs should probably only be given to addresses with different numbers of registers on machines with lots of registers.

This macro will normally either not be defined or be defined as a constant.

REGISTER_MOVE_COST (*mode, from, to*)

A C expression for the cost of moving data of mode *mode* from a register in class *from* to one in class *to*. The classes are expressed using the enumeration values such as `GENERAL_REGS`. A value of 2 is the default; other values are interpreted relative to that.

It is not required that the cost always equal 2 when *from* is the same as *to*; on some machines it is expensive to move between registers if they are not general registers.

If `reload` sees an insn consisting of a single `set` between two hard registers, and if `REGISTER_MOVE_COST` applied to their classes returns a value of 2, `reload` does not check to ensure that the constraints of the insn are met. Setting a cost of other than 2 will allow `reload` to verify that the constraints are met. You should do this if the `'movm'` pattern's constraints do not allow such copying.

MEMORY_MOVE_COST (*mode*, *class*, *in*)

A C expression for the cost of moving data of mode *mode* between a register of class *class* and memory; *in* is zero if the value is to be written to memory, nonzero if it is to be read in. This cost is relative to those in **REGISTER_MOVE_COST**. If moving between registers and memory is more expensive than between two registers, you should define this macro to express the relative cost.

If you do not define this macro, GCC uses a default cost of 4 plus the cost of copying via a secondary reload register, if one is needed. If your machine requires a secondary reload register to copy between memory and a register of *class* but the reload mechanism is more complex than copying via an intermediate, define this macro to reflect the actual cost of the move.

GCC defines the function `memory_move_secondary_cost` if secondary reloads are needed. It computes the costs due to copying via a secondary register. If your machine copies from memory using a secondary register in the conventional way but the default base value of 4 is not correct for your machine, define this macro to add some other value to the result of that function. The arguments to that function are the same as to this macro.

BRANCH_COST

A C expression for the cost of a branch instruction. A value of 1 is the default; other values are interpreted relative to that.

Here are additional macros which do not specify precise relative costs, but only that certain actions are more expensive than GCC would ordinarily expect.

SLOW_BYTE_ACCESS

Define this macro as a C expression which is nonzero if accessing less than a word of memory (i.e. a `char` or a `short`) is no faster than accessing a word of memory, i.e., if such access require more than one instruction or if there is no difference in cost between byte and (aligned) word loads.

When this macro is not defined, the compiler will access a field by finding the smallest containing object; when it is defined, a fullword load will be used if alignment permits. Unless bytes accesses are faster than word accesses, using word accesses is preferable since it may eliminate subsequent memory access if subsequent accesses occur to other fields in the same word of the structure, but to different bytes.

SLOW_UNALIGNED_ACCESS (*mode*, *alignment*)

Define this macro to be the value 1 if memory accesses described by the *mode* and *alignment* parameters have a cost many times greater than aligned accesses, for example if they are emulated in a trap handler.

When this macro is nonzero, the compiler will act as if **STRICT_ALIGNMENT** were nonzero when generating code for block moves. This can cause significantly more instructions to be produced. Therefore, do not set this macro nonzero if unaligned accesses only add a cycle or two to the time for a memory access.

If the value of this macro is always zero, it need not be defined. If this macro is defined, it should produce a nonzero value when **STRICT_ALIGNMENT** is nonzero.

DONT_REDUCE_ADDR

Define this macro to inhibit strength reduction of memory addresses. (On some machines, such strength reduction seems to do harm rather than good.)

MOVE_RATIO

The threshold of number of scalar memory-to-memory move insns, *below* which a sequence of insns should be generated instead of a string move insn or a library call. Increasing the value will always make code faster, but eventually incurs high cost in increased code size.

Note that on machines where the corresponding move insn is a `define_expand` that emits a sequence of insns, this macro counts the number of such sequences.

If you don't define this, a reasonable default is used.

MOVE_BY_PIECES_P (*size, alignment*)

A C expression used to determine whether `move_by_pieces` will be used to copy a chunk of memory, or whether some other block move mechanism will be used. Defaults to 1 if `move_by_pieces_ninsns` returns less than `MOVE_RATIO`.

MOVE_MAX_PIECES

A C expression used by `move_by_pieces` to determine the largest unit a load or store used to copy memory is. Defaults to `MOVE_MAX`.

CLEAR_RATIO

The threshold of number of scalar move insns, *below* which a sequence of insns should be generated to clear memory instead of a string clear insn or a library call. Increasing the value will always make code faster, but eventually incurs high cost in increased code size.

If you don't define this, a reasonable default is used.

CLEAR_BY_PIECES_P (*size, alignment*)

A C expression used to determine whether `clear_by_pieces` will be used to clear a chunk of memory, or whether some other block clear mechanism will be used. Defaults to 1 if `move_by_pieces_ninsns` returns less than `CLEAR_RATIO`.

USE_LOAD_POST_INCREMENT (*mode*)

A C expression used to determine whether a load postincrement is a good thing to use for a given mode. Defaults to the value of `HAVE_POST_INCREMENT`.

USE_LOAD_POST_DECREMENT (*mode*)

A C expression used to determine whether a load postdecrement is a good thing to use for a given mode. Defaults to the value of `HAVE_POST_DECREMENT`.

USE_LOAD_PRE_INCREMENT (*mode*)

A C expression used to determine whether a load preincrement is a good thing to use for a given mode. Defaults to the value of `HAVE_PRE_INCREMENT`.

USE_LOAD_PRE_DECREMENT (*mode*)

A C expression used to determine whether a load predecrement is a good thing to use for a given mode. Defaults to the value of `HAVE_PRE_DECREMENT`.

USE_STORE_POST_INCREMENT (*mode*)

A C expression used to determine whether a store postincrement is a good thing to use for a given mode. Defaults to the value of `HAVE_POST_INCREMENT`.

USE_STORE_POST_DECREMENT (*mode*)

A C expression used to determine whether a store postdecrement is a good thing to use for a given mode. Defaults to the value of `HAVE_POST_DECREMENT`.

USE_STORE_PRE_INCREMENT (*mode*)

This macro is used to determine whether a store preincrement is a good thing to use for a given mode. Defaults to the value of `HAVE_PRE_INCREMENT`.

USE_STORE_PRE_DECREMENT (*mode*)

This macro is used to determine whether a store predecrement is a good thing to use for a given mode. Defaults to the value of `HAVE_PRE_DECREMENT`.

NO_FUNCTION_CSE

Define this macro if it is as good or better to call a constant function address than to call an address kept in a register.

NO_RECURSIVE_FUNCTION_CSE

Define this macro if it is as good or better for a function to call itself with an explicit address than to call an address kept in a register.

10.17 Adjusting the Instruction Scheduler

The instruction scheduler may need a fair amount of machine-specific adjustment in order to produce good code. GCC provides several target hooks for this purpose. It is usually enough to define just a few of them: try the first ones in this list first.

int TARGET_SCHED_ISSUE_RATE (**void**) [Target Hook]

This hook returns the maximum number of instructions that can ever issue at the same time on the target machine. The default is one. Although the insn scheduler can define itself the possibility of issue an insn on the same cycle, the value can serve as an additional constraint to issue insns on the same simulated processor cycle (see hooks ‘TARGET_SCHED_REORDER’ and ‘TARGET_SCHED_REORDER2’). This value must be constant over the entire compilation. If you need it to vary depending on what the instructions are, you must use ‘TARGET_SCHED_VARIABLE_ISSUE’.

For the automaton based pipeline interface, you could define this hook to return the value of the macro `MAX_DFA_ISSUE_RATE`.

int TARGET_SCHED_VARIABLE_ISSUE (**FILE** **file*, **int** *verbose*, [Target Hook]
rtx *insn*, **int** *more*)

This hook is executed by the scheduler after it has scheduled an insn from the ready list. It should return the number of insns which can still be issued in the current cycle. Normally this is ‘*more* - 1’. You should define this hook if some insns take more machine resources than others, so that fewer insns can follow them in the same cycle. *file* is either a null pointer, or a stdio stream to write any debug output to. *verbose* is the verbose level provided by ‘`-fsched-verbose-n`’. *insn* is the instruction that was scheduled.

`int TARGET_SCHED_ADJUST_COST (rtx insn, rtx link, rtx dep_insn, int cost)` [Target Hook]

This function corrects the value of *cost* based on the relationship between *insn* and *dep_insn* through the dependence *link*. It should return the new value. The default is to make no adjustment to *cost*. This can be used for example to specify to the scheduler using the traditional pipeline description that an output- or anti-dependence does not incur the same cost as a data-dependence. If the scheduler using the automaton based pipeline description, the cost of anti-dependence is zero and the cost of output-dependence is maximum of one and the difference of latency times of the first and the second insns. If these values are not acceptable, you could use the hook to modify them too. See also see Section 9.18.8.2 [Automaton pipeline description], page 186.

`int TARGET_SCHED_ADJUST_PRIORITY (rtx insn, int priority)` [Target Hook]

This hook adjusts the integer scheduling priority *priority* of *insn*. It should return the new priority. Reduce the priority to execute *insn* earlier, increase the priority to execute *insn* later. Do not define this hook if you do not need to adjust the scheduling priorities of insns.

`int TARGET_SCHED_REORDER (FILE *file, int verbose, rtx *ready, int *n_ready, int clock)` [Target Hook]

This hook is executed by the scheduler after it has scheduled the ready list, to allow the machine description to reorder it (for example to combine two small instructions together on ‘VLIW’ machines). *file* is either a null pointer, or a stdio stream to write any debug output to. *verbose* is the verbose level provided by ‘-fsched-verbose-*n*’. *ready* is a pointer to the ready list of instructions that are ready to be scheduled. *n_ready* is a pointer to the number of elements in the ready list. The scheduler reads the ready list in reverse order, starting with *ready*[**n_ready*-1] and going to *ready*[0]. *clock* is the timer tick of the scheduler. You may modify the ready list and the number of ready insns. The return value is the number of insns that can issue this cycle; normally this is just *issue_rate*. See also ‘TARGET_SCHED_REORDER2’.

`int TARGET_SCHED_REORDER2 (FILE *file, int verbose, rtx *ready, int *n_ready, int clock)` [Target Hook]

Like ‘TARGET_SCHED_REORDER’, but called at a different time. That function is called whenever the scheduler starts a new cycle. This one is called once per iteration over a cycle, immediately after ‘TARGET_SCHED_VARIABLE_ISSUE’; it can reorder the ready list and return the number of insns to be scheduled in the same cycle. Defining this hook can be useful if there are frequent situations where scheduling one insn causes other insns to become ready in the same cycle. These other insns can then be taken into account properly.

`void TARGET_SCHED_INIT (FILE *file, int verbose, int max_ready)` [Target Hook]

This hook is executed by the scheduler at the beginning of each block of instructions that are to be scheduled. *file* is either a null pointer, or a stdio stream to write any debug output to. *verbose* is the verbose level provided by ‘-fsched-verbose-*n*’. *max_ready* is the maximum number of insns in the current scheduling region that can be live at the same time. This can be used to allocate scratch space if it is needed, e.g. by ‘TARGET_SCHED_REORDER’.

void TARGET_SCHED_FINISH (*FILE *file*, *int verbose*) [Target Hook]

This hook is executed by the scheduler at the end of each block of instructions that are to be scheduled. It can be used to perform cleanup of any actions done by the other scheduling hooks. *file* is either a null pointer, or a stdio stream to write any debug output to. *verbose* is the verbose level provided by ‘-fsched-verbose-*n*’.

int TARGET_SCHED_USE_DFA_PIPELINE_INTERFACE (*void*) [Target Hook]

This hook is called many times during insn scheduling. If the hook returns nonzero, the automaton based pipeline description is used for insn scheduling. Otherwise the traditional pipeline description is used. The default is usage of the traditional pipeline description.

You should also remember that to simplify the insn scheduler sources an empty traditional pipeline description interface is generated even if there is no a traditional pipeline description in the ‘.md’ file. The same is true for the automaton based pipeline description. That means that you should be accurate in defining the hook.

int TARGET_SCHED_DFA_PRE_CYCLE_INSN (*void*) [Target Hook]

The hook returns an RTL insn. The automaton state used in the pipeline hazard recognizer is changed as if the insn were scheduled when the new simulated processor cycle starts. Usage of the hook may simplify the automaton pipeline description for some VLIW processors. If the hook is defined, it is used only for the automaton based pipeline description. The default is not to change the state when the new simulated processor cycle starts.

void TARGET_SCHED_INIT_DFA_PRE_CYCLE_INSN (*void*) [Target Hook]

The hook can be used to initialize data used by the previous hook.

int TARGET_SCHED_DFA_POST_CYCLE_INSN (*void*) [Target Hook]

The hook is analogous to ‘TARGET_SCHED_DFA_PRE_CYCLE_INSN’ but used to changed the state as if the insn were scheduled when the new simulated processor cycle finishes.

void TARGET_SCHED_INIT_DFA_POST_CYCLE_INSN (*void*) [Target Hook]

The hook is analogous to ‘TARGET_SCHED_INIT_DFA_PRE_CYCLE_INSN’ but used to initialize data used by the previous hook.

int TARGET_SCHED_FIRST_CYCLE_MULTIPASS_DFA_LOOKAHEAD (*void*) [Target Hook]

This hook controls better choosing an insn from the ready insn queue for the DFA-based insn scheduler. Usually the scheduler chooses the first insn from the queue. If the hook returns a positive value, an additional scheduler code tries all permutations of ‘TARGET_SCHED_FIRST_CYCLE_MULTIPASS_DFA_LOOKAHEAD ()’ subsequent ready insns to choose an insn whose issue will result in maximal number of issued insns on the same cycle. For the VLIW processor, the code could actually solve the problem of packing simple insns into the VLIW insn. Of course, if the rules of VLIW packing are described in the automaton.

This code also could be used for superscalar RISC processors. Let us consider a superscalar RISC processor with 3 pipelines. Some insns can be executed in pipelines *A* or *B*, some insns can be executed only in pipelines *B* or *C*, and one insn can be

executed in pipeline *B*. The processor may issue the 1st insn into *A* and the 2nd one into *B*. In this case, the 3rd insn will wait for freeing *B* until the next cycle. If the scheduler issues the 3rd insn the first, the processor could issue all 3 insns per cycle.

Actually this code demonstrates advantages of the automaton based pipeline hazard recognizer. We try quickly and easy many insn schedules to choose the best one.

The default is no multipass scheduling.

`void TARGET_SCHED_INIT_DFA_BUBBLES (void)` [Target Hook]

The DFA-based scheduler could take the insertion of nop operations for better insn scheduling into account. It can be done only if the multi-pass insn scheduling works (see hook ‘TARGET_SCHED_FIRST_CYCLE_MULTIPASS_DFA_LOOKAHEAD’).

Let us consider a VLIW processor insn with 3 slots. Each insn can be placed only in one of the three slots. We have 3 ready insns *A*, *B*, and *C*. *A* and *C* can be placed only in the 1st slot, *B* can be placed only in the 3rd slot. We described the automaton which does not permit empty slot gaps between insns (usually such description is simpler). Without this code the scheduler would place each insn in 3 separate VLIW insns. If the scheduler places a nop insn into the 2nd slot, it could place the 3 insns into 2 VLIW insns. What is the nop insn is returned by hook ‘TARGET_SCHED_DFA_BUBBLE’. Hook ‘TARGET_SCHED_INIT_DFA_BUBBLES’ can be used to initialize or create the nop insns.

You should remember that the scheduler does not insert the nop insns. It is not wise because of the following optimizations. The scheduler only considers such possibility to improve the result schedule. The nop insns should be inserted lately, e.g. on the final phase.

`rtx TARGET_SCHED_DFA_BUBBLE (int index)` [Target Hook]

This hook ‘FIRST_CYCLE_MULTIPASS_SCHEDULING’ is used to insert nop operations for better insn scheduling when DFA-based scheduler makes multipass insn scheduling (see also description of hook ‘TARGET_SCHED_INIT_DFA_BUBBLES’). This hook returns a nop insn with given *index*. The indexes start with zero. The hook should return NULL if there are no more nop insns with indexes greater than given index.

Macros in the following table are generated by the program ‘genattr’ and can be useful for writing the hooks.

TRADITIONAL_PIPELINE_INTERFACE

The macro definition is generated if there is a traditional pipeline description in ‘.md’ file. You should also remember that to simplify the insn scheduler sources an empty traditional pipeline description interface is generated even if there is no a traditional pipeline description in the ‘.md’ file. The macro can be used to distinguish the two types of the traditional interface.

DFA_PIPELINE_INTERFACE

The macro definition is generated if there is an automaton pipeline description in ‘.md’ file. You should also remember that to simplify the insn scheduler sources an empty automaton pipeline description interface is generated even if there is no an automaton pipeline description in the ‘.md’ file. The macro can be used to distinguish the two types of the automaton interface.

MAX_DFA_ISSUE_RATE

The macro definition is generated in the automaton based pipeline description interface. Its value is calculated from the automaton based pipeline description and is equal to maximal number of all insns described in constructions ‘define_insn_reservation’ which can be issued on the same processor cycle.

10.18 Dividing the Output into Sections (Texts, Data, . . .)

An object file is divided into sections containing different types of data. In the most common case, there are three sections: the *text section*, which holds instructions and read-only data; the *data section*, which holds initialized writable data; and the *bss section*, which holds uninitialized data. Some systems have other kinds of sections.

The compiler must tell the assembler when to switch sections. These macros control what commands to output to tell the assembler this. You can also define additional sections.

TEXT_SECTION_ASM_OP

A C expression whose value is a string, including spacing, containing the assembler operation that should precede instructions and read-only data. Normally “\t.text” is right.

TEXT_SECTION

A C statement that switches to the default section containing instructions. Normally this is not needed, as simply defining TEXT_SECTION_ASM_OP is enough. The MIPS port uses this to sort all functions after all data declarations.

HOT_TEXT_SECTION_NAME

If defined, a C string constant for the name of the section containing most frequently executed functions of the program. If not defined, GCC will provide a default definition if the target supports named sections.

UNLIKELY_EXECUTED_TEXT_SECTION_NAME

If defined, a C string constant for the name of the section containing unlikely executed functions in the program.

DATA_SECTION_ASM_OP

A C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as writable initialized data. Normally “\t.data” is right.

READONLY_DATA_SECTION_ASM_OP

A C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as read-only initialized data.

READONLY_DATA_SECTION

A macro naming a function to call to switch to the proper section for read-only data. The default is to use READONLY_DATA_SECTION_ASM_OP if defined, else fall back to `text_section`.

The most common definition will be `data_section`, if the target does not have a special read-only data section, and does not put data in the text section.

SHARED_SECTION_ASM_OP

If defined, a C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as shared data. If not defined, `DATA_SECTION_ASM_OP` will be used.

BSS_SECTION_ASM_OP

If defined, a C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as uninitialized global data. If not defined, and neither `ASM_OUTPUT_BSS` nor `ASM_OUTPUT_ALIGNED_BSS` are defined, uninitialized global data will be output in the data section if `'-fno-common'` is passed, otherwise `ASM_OUTPUT_COMMON` will be used.

SHARED_BSS_SECTION_ASM_OP

If defined, a C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as uninitialized global shared data. If not defined, and `BSS_SECTION_ASM_OP` is, the latter will be used.

INIT_SECTION_ASM_OP

If defined, a C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as initialization code. If not defined, GCC will assume such a section does not exist.

FINI_SECTION_ASM_OP

If defined, a C expression whose value is a string, including spacing, containing the assembler operation to identify the following data as finalization code. If not defined, GCC will assume such a section does not exist.

CRT_CALL_STATIC_FUNCTION (*section_op*, *function*)

If defined, an ASM statement that switches to a different section via *section_op*, calls *function*, and switches back to the text section. This is used in `'crtstuff.c'` if `INIT_SECTION_ASM_OP` or `FINI_SECTION_ASM_OP` to calls to initialization and finalization functions from the `init` and `fini` sections. By default, this macro uses a simple function call. Some ports need hand-crafted assembly code to avoid dependencies on registers initialized in the function prologue or to ensure that constant pools don't end up too far way in the text section.

FORCE_CODE_SECTION_ALIGN

If defined, an ASM statement that aligns a code section to some arbitrary boundary. This is used to force all fragments of the `.init` and `.fini` sections to have to same alignment and thus prevent the linker from having to add any padding.

EXTRA_SECTIONS

A list of names for sections other than the standard two, which are `in_text` and `in_data`. You need not define this macro on a system with no other sections (that GCC needs to use).

EXTRA_SECTION_FUNCTIONS

One or more functions to be defined in ‘`varasm.c`’. These functions should do jobs analogous to those of `text_section` and `data_section`, for your additional sections. Do not define this macro if you do not define `EXTRA_SECTIONS`.

JUMP_TABLES_IN_TEXT_SECTION

Define this macro to be an expression with a nonzero value if jump tables (for `tablejump` insns) should be output in the text section, along with the assembler instructions. Otherwise, the readonly data section is used.

This macro is irrelevant if there is no separate readonly data section.

void TARGET_ASM_SELECT_SECTION (*tree exp*, *int reloc*, [Target Hook]
unsigned HOST_WIDE_INT align)

Switches to the appropriate section for output of *exp*. You can assume that *exp* is either a `VAR_DECL` node or a constant of some sort. *reloc* indicates whether the initial value of *exp* requires link-time relocations. Bit 0 is set when variable contains local relocations only, while bit 1 is set for global relocations. Select the section by calling `data_section` or one of the alternatives for other sections. *align* is the constant alignment in bits.

The default version of this function takes care of putting read-only variables in `readonly_data_section`.

void TARGET_ASM_UNIQUE_SECTION (*tree decl*, *int reloc*) [Target Hook]

Build up a unique section name, expressed as a `STRING_CST` node, and assign it to ‘`DECL_SECTION_NAME (decl)`’. As with `TARGET_ASM_SELECT_SECTION`, *reloc* indicates whether the initial value of *exp* requires link-time relocations.

The default version of this function appends the symbol name to the ELF section name that would normally be used for the symbol. For example, the function `foo` would be placed in `.text.foo`. Whatever the actual target object format, this is often good enough.

void TARGET_ASM_SELECT_RTX_SECTION (*enum machine_mode mode*, *rtx x*, [Target Hook]
unsigned HOST_WIDE_INT align)

Switches to the appropriate section for output of constant pool entry *x* in *mode*. You can assume that *x* is some kind of constant in RTL. The argument *mode* is redundant except in the case of a `const_int` rtx. Select the section by calling `readonly_data_section` or one of the alternatives for other sections. *align* is the constant alignment in bits.

The default version of this function takes care of putting symbolic constants in `flag_pic` mode in `data_section` and everything else in `readonly_data_section`.

void TARGET_ENCODE_SECTION_INFO (*tree decl*, *int new_decl_p*) [Target Hook]

Define this hook if references to a symbol or a constant must be treated differently depending on something about the variable or function named by the symbol (such as what section it is in).

The hook is executed under two circumstances. One is immediately after the rtl for *decl* that represents a variable or a function has been created and stored in `DECL_RTL(decl)`. The value of the rtl will be a `mem` whose address is a `symbol_ref`. The

other is immediately after the rtl for *decl* that represents a constant has been created and stored in `TREE_CST_RTL` (*decl*). The macro is called once for each distinct constant in a source file.

The *new_decl_p* argument will be true if this is the first time that `ENCODE_SECTION_INFO` has been invoked on this decl. It will be false for subsequent invocations, which will happen for duplicate declarations. Whether or not anything must be done for the duplicate declaration depends on whether the hook examines `DECL_ATTRIBUTES`.

The usual thing for this hook to do is to record a flag in the `symbol_ref` (such as `SYMBOL_REF_FLAG`) or to store a modified name string in the `symbol_ref` (if one bit is not enough information).

`const char *TARGET_STRIP_NAME_ENCODING (const char *name)` [Target Hook]
Decode *name* and return the real name part, sans the characters that `TARGET_ENCODE_SECTION_INFO` may have added.

`bool TARGET_IN_SMALL_DATA_P (tree exp)` [Target Hook]
Returns true if *exp* should be placed into a “small data” section. The default version of this hook always returns false.

Target Hook `bool TARGET_HAVE_SRODATA_SECTION` [Variable]
Contains the value true if the target places read-only “small data” into a separate section. The default value is false.

`bool TARGET_BINDS_LOCAL_P (tree exp)` [Target Hook]
Returns true if *exp* names an object for which name resolution rules must resolve to the current “module” (dynamic shared library or executable image).
The default version of this hook implements the name resolution rules for ELF, which has a looser model of global name binding than other currently supported object file formats.

Target Hook `bool TARGET_HAVE_TLS` [Variable]
Contains the value true if the target supports thread-local storage. The default value is false.

10.19 Position Independent Code

This section describes macros that help implement generation of position independent code. Simply defining these macros is not enough to generate valid PIC; you must also add support to the macros `GO_IF_LEGITIMATE_ADDRESS` and `PRINT_OPERAND_ADDRESS`, as well as `LEGITIMIZE_ADDRESS`. You must modify the definition of ‘`movsi`’ to do something appropriate when the source operand contains a symbolic address. You may also need to alter the handling of switch statements so that they use relative addresses.

`PIC_OFFSET_TABLE_REGNUM`

The register number of the register used to address a table of static data addresses in memory. In some cases this register is defined by a processor’s “application binary interface” (ABI). When this macro is defined, RTL is generated for this register once, as with the stack pointer and frame pointer registers. If

this macro is not defined, it is up to the machine-dependent files to allocate such a register (if necessary). Note that this register must be fixed when in use (e.g. when `flag_pic` is true).

`PIC_OFFSET_TABLE_REG_CALL_CLOBBERED`

Define this macro if the register defined by `PIC_OFFSET_TABLE_REGNUM` is clobbered by calls. Do not define this macro if `PIC_OFFSET_TABLE_REGNUM` is not defined.

`FINALIZE_PIC`

By generating position-independent code, when two different programs (A and B) share a common library (libC.a), the text of the library can be shared whether or not the library is linked at the same address for both programs. In some of these environments, position-independent code requires not only the use of different addressing modes, but also special code to enable the use of these addressing modes.

The `FINALIZE_PIC` macro serves as a hook to emit these special codes once the function is being compiled into assembly code, but not before. (It is not done before, because in the case of compiling an inline function, it would lead to multiple PIC prologues being included in functions which used inline functions and were compiled to assembly language.)

`LEGITIMATE_PIC_OPERAND_P (x)`

A C expression that is nonzero if `x` is a legitimate immediate operand on the target machine when generating position independent code. You can assume that `x` satisfies `CONSTANT_P`, so you need not check this. You can also assume `flag_pic` is true, so you need not check it either. You need not define this macro if all constants (including `SYMBOL_REF`) can be immediate operands when generating position independent code.

10.20 Defining the Output Assembler Language

This section describes macros whose principal purpose is to describe how to write instructions in assembler language—rather than what the instructions do.

10.20.1 The Overall Framework of an Assembler File

This describes the overall framework of an assembler file.

`ASM_FILE_START (stream)`

A C expression which outputs to the stdio stream `stream` some appropriate text to go at the start of an assembler file.

Normally this macro is defined to output a line containing `'#NO_APP'`, which is a comment that has no effect on most assemblers but tells the GNU assembler that it can save time by not checking for certain assembler constructs.

On systems that use SDB, it is necessary to output certain commands; see `'attasm.h'`.

ASM_FILE_END (*stream*)

A C expression which outputs to the stdio stream *stream* some appropriate text to go at the end of an assembler file.

If this macro is not defined, the default is to output nothing special at the end of the file. Most systems don't require any definition.

On systems that use SDB, it is necessary to output certain commands; see 'attasm.h'.

ASM_COMMENT_START

A C string constant describing how to begin a comment in the target assembler language. The compiler assumes that the comment will end at the end of the line.

ASM_APP_ON

A C string constant for text to be output before each `asm` statement or group of consecutive ones. Normally this is "#APP", which is a comment that has no effect on most assemblers but tells the GNU assembler that it must check the lines that follow for all valid assembler constructs.

ASM_APP_OFF

A C string constant for text to be output after each `asm` statement or group of consecutive ones. Normally this is "#NO_APP", which tells the GNU assembler to resume making the time-saving assumptions that are valid for ordinary compiler output.

ASM_OUTPUT_SOURCE_FILENAME (*stream*, *name*)

A C statement to output COFF information or DWARF debugging information which indicates that filename *name* is the current source file to the stdio stream *stream*.

This macro need not be defined if the standard form of output for the file format in use is appropriate.

OUTPUT_QUOTED_STRING (*stream*, *string*)

A C statement to output the string *string* to the stdio stream *stream*. If you do not call the function `output_quoted_string` in your config files, GCC will only call it to output filenames to the assembler source. So you can use it to canonicalize the format of the filename using this macro.

ASM_OUTPUT_SOURCE_LINE (*stream*, *line*)

A C statement to output DBX or SDB debugging information before code for line number *line* of the current source file to the stdio stream *stream*.

This macro need not be defined if the standard form of debugging information for the debugger in use is appropriate.

ASM_OUTPUT_IDENT (*stream*, *string*)

A C statement to output something to the assembler file to handle a '#ident' directive containing the text *string*. If this macro is not defined, nothing is output for a '#ident' directive.

OBJC_PROLOGUE

A C statement to output any assembler statements which are required to precede any Objective-C object definitions or message sending. The statement is executed only when compiling an Objective-C program.

```
void TARGET_ASM_NAMED_SECTION (const char *name, unsigned [Target Hook]
    int flags, unsigned int align)
```

Output assembly directives to switch to section *name*. The section should have attributes as specified by *flags*, which is a bit mask of the `SECTION_*` flags defined in `'output.h'`. If *align* is nonzero, it contains an alignment in bytes to be used for the section, otherwise some target default should be used. Only targets that must specify an alignment within the section directive need pay attention to *align* – we will still use `ASM_OUTPUT_ALIGN`.

```
bool TARGET_HAVE_NAMED_SECTIONS [Target Hook]
    This flag is true if the target supports TARGET_ASM_NAMED_SECTION.
```

```
unsigned int TARGET_SECTION_TYPE_FLAGS (tree decl, const [Target Hook]
    char *name, int reloc)
```

Choose a set of section attributes for use by `TARGET_ASM_NAMED_SECTION` based on a variable or function *decl*, a section name, and whether or not the declaration's initializer may contain runtime relocations. *decl* may be null, in which case read-write data should be assumed.

The default version of this function handles choosing code vs data, read-only vs read-write data, and `flag_pic`. You should only need to override this if your target has special flags that might be set via `__attribute__`.

10.20.2 Output of Data

```
const char * TARGET_ASM_BYTE_OP [Target Hook]
const char * TARGET_ASM_ALIGNED_HI_OP [Target Hook]
const char * TARGET_ASM_ALIGNED_SI_OP [Target Hook]
const char * TARGET_ASM_ALIGNED_DI_OP [Target Hook]
const char * TARGET_ASM_ALIGNED_TI_OP [Target Hook]
const char * TARGET_ASM_UNALIGNED_HI_OP [Target Hook]
const char * TARGET_ASM_UNALIGNED_SI_OP [Target Hook]
const char * TARGET_ASM_UNALIGNED_DI_OP [Target Hook]
const char * TARGET_ASM_UNALIGNED_TI_OP [Target Hook]
```

These hooks specify assembly directives for creating certain kinds of integer object. The `TARGET_ASM_BYTE_OP` directive creates a byte-sized object, the `TARGET_ASM_ALIGNED_HI_OP` one creates an aligned two-byte object, and so on. Any of the hooks may be NULL, indicating that no suitable directive is available.

The compiler will print these strings at the start of a new line, followed immediately by the object's initial value. In most cases, the string should contain a tab, a pseudo-op, and then another tab.

`bool TARGET_ASM_INTEGER (rtx x, unsigned int size, int aligned_p)` [Target Hook]

The `assemble_integer` function uses this hook to output an integer object. *x* is the object's value, *size* is its size in bytes and *aligned_p* indicates whether it is aligned. The function should return `true` if it was able to output the object. If it returns false, `assemble_integer` will try to split the object into smaller parts.

The default implementation of this hook will use the `TARGET_ASM_BYTE_OP` family of strings, returning `false` when the relevant string is `NULL`.

`OUTPUT_ADDR_CONST_EXTRA (stream, x, fail)`

A C statement to recognize *rtx* patterns that `output_addr_const` can't deal with, and output assembly code to *stream* corresponding to the pattern *x*. This may be used to allow machine-dependent UNSPECS to appear within constants. If `OUTPUT_ADDR_CONST_EXTRA` fails to recognize a pattern, it must `goto fail`, so that a standard error message is printed. If it prints an error message itself, by calling, for example, `output_operand_lossage`, it may just complete normally.

`ASM_OUTPUT_ASCII (stream, ptr, len)`

A C statement to output to the stdio stream *stream* an assembler instruction to assemble a string constant containing the *len* bytes at *ptr*. *ptr* will be a C expression of type `char *` and *len* a C expression of type `int`.

If the assembler has a `.ascii` pseudo-op as found in the Berkeley Unix assembler, do not define the macro `ASM_OUTPUT_ASCII`.

`ASM_OUTPUT_FDESC (stream, decl, n)`

A C statement to output word *n* of a function descriptor for *decl*. This must be defined if `TARGET_VTABLE_USES_DESCRIPTOR` is defined, and is otherwise unused.

`CONSTANT_POOL_BEFORE_FUNCTION`

You may define this macro as a C expression. You should define the expression to have a nonzero value if GCC should output the constant pool for a function before the code for the function, or a zero value if GCC should output the constant pool after the function. If you do not define this macro, the usual case, GCC will output the constant pool before the function.

`ASM_OUTPUT_POOL_PROLOGUE (file, funname, fundecl, size)`

A C statement to output assembler commands to define the start of the constant pool for a function. *funname* is a string giving the name of the function. Should the return type of the function be required, it can be obtained via *fundecl*. *size* is the size, in bytes, of the constant pool that will be written immediately after this call.

If no constant-pool prefix is required, the usual case, this macro need not be defined.

`ASM_OUTPUT_SPECIAL_POOL_ENTRY (file, x, mode, align, labelno, jumpto)`

A C statement (with or without semicolon) to output a constant in the constant pool, if it needs special treatment. (This macro need not do anything for RTL expressions that can be output normally.)

The argument *file* is the standard I/O stream to output the assembler code on. *x* is the RTL expression for the constant to output, and *mode* is the machine mode (in case *x* is a ‘`const_int`’). *align* is the required alignment for the value *x*; you should output an assembler directive to force this much alignment.

The argument *labelno* is a number to use in an internal label for the address of this pool entry. The definition of this macro is responsible for outputting the label definition at the proper place. Here is how to do this:

```
ASM_OUTPUT_INTERNAL_LABEL (file, "LC", labelno);
```

When you output a pool entry specially, you should end with a `goto` to the label *jumpno*. This will prevent the same pool entry from being output a second time in the usual manner.

You need not define this macro if it would do nothing.

CONSTANT_AFTER_FUNCTION_P (*exp*)

Define this macro as a C expression which is nonzero if the constant *exp*, of type `tree`, should be output after the code for a function. The compiler will normally output all constants before the function; you need not define this macro if this is OK.

ASM_OUTPUT_POOL_EPILOGUE (*file funname fundecl size*)

A C statement to output assembler commands to at the end of the constant pool for a function. *funname* is a string giving the name of the function. Should the return type of the function be required, you can obtain it via *fundecl*. *size* is the size, in bytes, of the constant pool that GCC wrote immediately before this call.

If no constant-pool epilogue is required, the usual case, you need not define this macro.

IS_ASM_LOGICAL_LINE_SEPARATOR (*C*)

Define this macro as a C expression which is nonzero if *C* is used as a logical line separator by the assembler.

If you do not define this macro, the default is that only the character ‘;’ is treated as a logical line separator.

```
const char * TARGET_ASM_OPEN_PAREN [Target Hook]
const char * TARGET_ASM_CLOSE_PAREN [Target Hook]
```

These target hooks are C string constants, describing the syntax in the assembler for grouping arithmetic expressions. If not overridden, they default to normal parentheses, which is correct for most assemblers.

These macros are provided by ‘`real.h`’ for writing the definitions of `ASM_OUTPUT_DOUBLE` and the like:

```
REAL_VALUE_TO_TARGET_SINGLE (x, l)
REAL_VALUE_TO_TARGET_DOUBLE (x, l)
REAL_VALUE_TO_TARGET_LONG_DOUBLE (x, l)
```

These translate *x*, of type `REAL_VALUE_TYPE`, to the target’s floating point representation, and store its bit pattern in the variable *l*. For `REAL_VALUE_TO_TARGET_SINGLE`, this variable should be a simple `long int`. For the others, it

should be an array of `long int`. The number of elements in this array is determined by the size of the desired target floating point data type: 32 bits of it go in each `long int` array element. Each array element holds 32 bits of the result, even if `long int` is wider than 32 bits on the host machine.

The array element values are designed so that you can print them out using `fprintf` in the order they should appear in the target machine's memory.

10.20.3 Output of Uninitialized Variables

Each of the macros in this section is used to do the whole job of outputting a single uninitialized variable.

ASM_OUTPUT_COMMON (*stream*, *name*, *size*, *rounded*)

A C statement (sans semicolon) to output to the stdio stream *stream* the assembler definition of a common-label named *name* whose size is *size* bytes. The variable *rounded* is the size rounded up to whatever alignment the caller wants.

Use the expression `assemble_name (stream, name)` to output the name itself; before and after that, output the additional assembler syntax for defining the name, and a newline.

This macro controls how the assembler definitions of uninitialized common global variables are output.

ASM_OUTPUT_ALIGNED_COMMON (*stream*, *name*, *size*, *alignment*)

Like `ASM_OUTPUT_COMMON` except takes the required alignment as a separate, explicit argument. If you define this macro, it is used in place of `ASM_OUTPUT_COMMON`, and gives you more flexibility in handling the required alignment of the variable. The alignment is specified as the number of bits.

ASM_OUTPUT_ALIGNED_DECL_COMMON (*stream*, *decl*, *name*, *size*, *alignment*)

Like `ASM_OUTPUT_ALIGNED_COMMON` except that *decl* of the variable to be output, if there is one, or `NULL_TREE` if there is no corresponding variable. If you define this macro, GCC will use it in place of both `ASM_OUTPUT_COMMON` and `ASM_OUTPUT_ALIGNED_COMMON`. Define this macro when you need to see the variable's decl in order to chose what to output.

ASM_OUTPUT_SHARED_COMMON (*stream*, *name*, *size*, *rounded*)

If defined, it is similar to `ASM_OUTPUT_COMMON`, except that it is used when *name* is shared. If not defined, `ASM_OUTPUT_COMMON` will be used.

ASM_OUTPUT_BSS (*stream*, *decl*, *name*, *size*, *rounded*)

A C statement (sans semicolon) to output to the stdio stream *stream* the assembler definition of uninitialized global *decl* named *name* whose size is *size* bytes. The variable *rounded* is the size rounded up to whatever alignment the caller wants.

Try to use function `asm_output_bss` defined in 'varasm.c' when defining this macro. If unable, use the expression `assemble_name (stream, name)` to output the name itself; before and after that, output the additional assembler syntax for defining the name, and a newline.

This macro controls how the assembler definitions of uninitialized global variables are output. This macro exists to properly support languages like C++ which do not have `common` data. However, this macro currently is not defined for all targets. If this macro and `ASM_OUTPUT_ALIGNED_BSS` are not defined then `ASM_OUTPUT_COMMON` or `ASM_OUTPUT_ALIGNED_COMMON` or `ASM_OUTPUT_ALIGNED_DECL_COMMON` is used.

`ASM_OUTPUT_ALIGNED_BSS` (*stream, decl, name, size, alignment*)

Like `ASM_OUTPUT_BSS` except takes the required alignment as a separate, explicit argument. If you define this macro, it is used in place of `ASM_OUTPUT_BSS`, and gives you more flexibility in handling the required alignment of the variable. The alignment is specified as the number of bits.

Try to use function `asm_output_aligned_bss` defined in file ‘`varasm.c`’ when defining this macro.

`ASM_OUTPUT_SHARED_BSS` (*stream, decl, name, size, rounded*)

If defined, it is similar to `ASM_OUTPUT_BSS`, except that it is used when *name* is shared. If not defined, `ASM_OUTPUT_BSS` will be used.

`ASM_OUTPUT_LOCAL` (*stream, name, size, rounded*)

A C statement (sans semicolon) to output to the stdio stream *stream* the assembler definition of a local-common-label named *name* whose size is *size* bytes. The variable *rounded* is the size rounded up to whatever alignment the caller wants.

Use the expression `assemble_name` (*stream, name*) to output the name itself; before and after that, output the additional assembler syntax for defining the name, and a newline.

This macro controls how the assembler definitions of uninitialized static variables are output.

`ASM_OUTPUT_ALIGNED_LOCAL` (*stream, name, size, alignment*)

Like `ASM_OUTPUT_LOCAL` except takes the required alignment as a separate, explicit argument. If you define this macro, it is used in place of `ASM_OUTPUT_LOCAL`, and gives you more flexibility in handling the required alignment of the variable. The alignment is specified as the number of bits.

`ASM_OUTPUT_ALIGNED_DECL_LOCAL` (*stream, decl, name, size, alignment*)

Like `ASM_OUTPUT_ALIGNED_DECL` except that *decl* of the variable to be output, if there is one, or `NULL_TREE` if there is no corresponding variable. If you define this macro, GCC will use it in place of both `ASM_OUTPUT_DECL` and `ASM_OUTPUT_ALIGNED_DECL`. Define this macro when you need to see the variable’s decl in order to chose what to output.

`ASM_OUTPUT_SHARED_LOCAL` (*stream, name, size, rounded*)

If defined, it is similar to `ASM_OUTPUT_LOCAL`, except that it is used when *name* is shared. If not defined, `ASM_OUTPUT_LOCAL` will be used.

10.20.4 Output and Generation of Labels

This is about outputting labels.

ASM_OUTPUT_LABEL (*stream*, *name*)

A C statement (sans semicolon) to output to the stdio stream *stream* the assembler definition of a label named *name*. Use the expression `assemble_name (stream, name)` to output the name itself; before and after that, output the additional assembler syntax for defining the name, and a newline. A default definition of this macro is provided which is correct for most systems.

SIZE_ASM_OP

A C string containing the appropriate assembler directive to specify the size of a symbol, without any arguments. On systems that use ELF, the default (in `'config/elfos.h'`) is `"\t.size\t"`; on other systems, the default is not to define this macro.

Define this macro only if it is correct to use the default definitions of `ASM_OUTPUT_SIZE_DIRECTIVE` and `ASM_OUTPUT_MEASURED_SIZE` for your system. If you need your own custom definitions of those macros, or if you do not need explicit symbol sizes at all, do not define this macro.

ASM_OUTPUT_SIZE_DIRECTIVE (*stream*, *name*, *size*)

A C statement (sans semicolon) to output to the stdio stream *stream* a directive telling the assembler that the size of the symbol *name* is *size*. *size* is a `HOST_WIDE_INT`. If you define `SIZE_ASM_OP`, a default definition of this macro is provided.

ASM_OUTPUT_MEASURED_SIZE (*stream*, *name*)

A C statement (sans semicolon) to output to the stdio stream *stream* a directive telling the assembler to calculate the size of the symbol *name* by subtracting its address from the current address.

If you define `SIZE_ASM_OP`, a default definition of this macro is provided. The default assumes that the assembler recognizes a special `'.'` symbol as referring to the current address, and can calculate the difference between this and another symbol. If your assembler does not recognize `'.'` or cannot do calculations with it, you will need to redefine `ASM_OUTPUT_MEASURED_SIZE` to use some other technique.

TYPE_ASM_OP

A C string containing the appropriate assembler directive to specify the type of a symbol, without any arguments. On systems that use ELF, the default (in `'config/elfos.h'`) is `"\t.type\t"`; on other systems, the default is not to define this macro.

Define this macro only if it is correct to use the default definition of `ASM_OUTPUT_TYPE_DIRECTIVE` for your system. If you need your own custom definition of this macro, or if you do not need explicit symbol types at all, do not define this macro.

TYPE_OPERAND_FMT

A C string which specifies (using `printf` syntax) the format of the second operand to `TYPE_ASM_OP`. On systems that use ELF, the default (in `'config/elfos.h'`) is `"%s"`; on other systems, the default is not to define this macro.

Define this macro only if it is correct to use the default definition of `ASM_OUTPUT_TYPE_DIRECTIVE` for your system. If you need your own custom definition of this macro, or if you do not need explicit symbol types at all, do not define this macro.

ASM_OUTPUT_TYPE_DIRECTIVE (*stream*, *type*)

A C statement (sans semicolon) to output to the stdio stream *stream* a directive telling the assembler that the type of the symbol *name* is *type*. *type* is a C string; currently, that string is always either "function" or "object", but you should not count on this.

If you define `TYPE_ASM_OP` and `TYPE_OPERAND_FMT`, a default definition of this macro is provided.

ASM_DECLARE_FUNCTION_NAME (*stream*, *name*, *decl*)

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for declaring the name *name* of a function which is being defined. This macro is responsible for outputting the label definition (perhaps using `ASM_OUTPUT_LABEL`). The argument *decl* is the `FUNCTION_DECL` tree node representing the function.

If this macro is not defined, then the function name is defined in the usual manner as a label (by means of `ASM_OUTPUT_LABEL`).

You may wish to use `ASM_OUTPUT_TYPE_DIRECTIVE` in the definition of this macro.

ASM_DECLARE_FUNCTION_SIZE (*stream*, *name*, *decl*)

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for declaring the size of a function which is being defined. The argument *name* is the name of the function. The argument *decl* is the `FUNCTION_DECL` tree node representing the function.

If this macro is not defined, then the function size is not defined.

You may wish to use `ASM_OUTPUT_MEASURED_SIZE` in the definition of this macro.

ASM_DECLARE_OBJECT_NAME (*stream*, *name*, *decl*)

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for declaring the name *name* of an initialized variable which is being defined. This macro must output the label definition (perhaps using `ASM_OUTPUT_LABEL`). The argument *decl* is the `VAR_DECL` tree node representing the variable.

If this macro is not defined, then the variable name is defined in the usual manner as a label (by means of `ASM_OUTPUT_LABEL`).

You may wish to use `ASM_OUTPUT_TYPE_DIRECTIVE` and/or `ASM_OUTPUT_SIZE_DIRECTIVE` in the definition of this macro.

ASM_DECLARE_CONSTANT_NAME (*stream*, *name*, *exp*, *size*)

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for declaring the name *name* of a constant which is being defined. This macro is responsible for outputting the label definition (perhaps using

ASM_OUTPUT_LABEL). The argument *exp* is the value of the constant, and *size* is the size of the constant in bytes. *name* will be an internal label.

If this macro is not defined, then the *name* is defined in the usual manner as a label (by means of ASM_OUTPUT_LABEL).

You may wish to use ASM_OUTPUT_TYPE_DIRECTIVE in the definition of this macro.

ASM_DECLARE_REGISTER_GLOBAL (*stream, decl, regno, name*)

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for claiming a register *regno* for a global variable *decl* with name *name*.

If you don't define this macro, that is equivalent to defining it to do nothing.

ASM_FINISH_DECLARE_OBJECT (*stream, decl, toplevel, atend*)

A C statement (sans semicolon) to finish up declaring a variable name once the compiler has processed its initializer fully and thus has had a chance to determine the size of an array when controlled by an initializer. This is used on systems where it's necessary to declare something about the size of the object.

If you don't define this macro, that is equivalent to defining it to do nothing.

You may wish to use ASM_OUTPUT_SIZE_DIRECTIVE and/or ASM_OUTPUT_MEASURED_SIZE in the definition of this macro.

void TARGET_ASM_GLOBALIZE_LABEL (FILE **stream*, const char [Target Hook]
**name*)

This target hook is a function to output to the stdio stream *stream* some commands that will make the label *name* global; that is, available for reference from other files.

The default implementation relies on a proper definition of GLOBAL_ASM_OP.

ASM_WEAKEN_LABEL (*stream, name*)

A C statement (sans semicolon) to output to the stdio stream *stream* some commands that will make the label *name* weak; that is, available for reference from other files but only used if no other definition is available. Use the expression `assemble_name (stream, name)` to output the name itself; before and after that, output the additional assembler syntax for making that name weak, and a newline.

If you don't define this macro or ASM_WEAKEN_DECL, GCC will not support weak symbols and you should not define the SUPPORTS_WEAK macro.

ASM_WEAKEN_DECL (*stream, decl, name, value*)

Combines (and replaces) the function of ASM_WEAKEN_LABEL and ASM_OUTPUT_WEAK_ALIAS, allowing access to the associated function or variable *decl*. If *value* is not NULL, this C statement should output to the stdio stream *stream* assembler code which defines (equates) the weak symbol *name* to have the value *value*. If *value* is NULL, it should output commands to make *name* weak.

SUPPORTS_WEAK

A C expression which evaluates to true if the target supports weak symbols.

If you don't define this macro, 'defaults.h' provides a default definition. If either `ASM_WEAKEN_LABEL` or `ASM_WEAKEN_DECL` is defined, the default definition is '1'; otherwise, it is '0'. Define this macro if you want to control weak symbol support with a compiler flag such as '-melf'.

`MAKE_DECL_ONE_ONLY`

A C statement (sans semicolon) to mark *decl* to be emitted as a public symbol such that extra copies in multiple translation units will be discarded by the linker. Define this macro if your object file format provides support for this concept, such as the 'COMDAT' section flags in the Microsoft Windows PE/COFF format, and this support requires changes to *decl*, such as putting it in a separate section.

`SUPPORTS_ONE_ONLY`

A C expression which evaluates to true if the target supports one-only semantics.

If you don't define this macro, 'varasm.c' provides a default definition. If `MAKE_DECL_ONE_ONLY` is defined, the default definition is '1'; otherwise, it is '0'. Define this macro if you want to control one-only symbol support with a compiler flag, or if setting the `DECL_ONE_ONLY` flag is enough to mark a declaration to be emitted as one-only.

```
void TARGET_ASM_ASSEMBLE_VISIBILITY (tree decl,      [Target Hook]
                                     const char *visibility)
```

This target hook is a function to output to *asm_out_file* some commands that will make the symbol(s) associated with *decl* have hidden, protected or internal visibility as specified by *visibility*.

`ASM_OUTPUT_EXTERNAL (stream, decl, name)`

A C statement (sans semicolon) to output to the stdio stream *stream* any text necessary for declaring the name of an external symbol named *name* which is referenced in this compilation but not defined. The value of *decl* is the tree node for the declaration.

This macro need not be defined if it does not need to output anything. The GNU assembler and most Unix assemblers don't require anything.

`ASM_OUTPUT_EXTERNAL_LIBCALL (stream, symref)`

A C statement (sans semicolon) to output on *stream* an assembler pseudo-op to declare a library function name external. The name of the library function is given by *symref*, which has type `rtx` and is a `symbol_ref`.

This macro need not be defined if it does not need to output anything. The GNU assembler and most Unix assemblers don't require anything.

`ASM_OUTPUT_LABELREF (stream, name)`

A C statement (sans semicolon) to output to the stdio stream *stream* a reference in assembler syntax to a label named *name*. This should add '_' to the front of the name, if that is customary on your operating system, as it is in most Berkeley Unix systems. This macro is used in `assemble_name`.

ASM_OUTPUT_SYMBOL_REF (*stream*, *sym*)

A C statement (sans semicolon) to output a reference to `SYMBOL_REF` *sym*. If not defined, `assemble_name` will be used to output the name of the symbol. This macro may be used to modify the way a symbol is referenced depending on information encoded by `TARGET_ENCODE_SECTION_INFO`.

ASM_OUTPUT_LABEL_REF (*stream*, *buf*)

A C statement (sans semicolon) to output a reference to *buf*, the result of `ASM_GENERATE_INTERNAL_LABEL`. If not defined, `assemble_name` will be used to output the name of the symbol. This macro is not used by `output_asm_label`, or the `%l` specifier that calls it; the intention is that this macro should be set when it is necessary to output a label differently when its address is being taken.

ASM_OUTPUT_INTERNAL_LABEL (*stream*, *prefix*, *num*)

A C statement to output to the stdio stream *stream* a label whose name is made from the string *prefix* and the number *num*.

It is absolutely essential that these labels be distinct from the labels used for user-level functions and variables. Otherwise, certain programs will have name conflicts with internal labels.

It is desirable to exclude internal labels from the symbol table of the object file. Most assemblers have a naming convention for labels that should be excluded; on many systems, the letter ‘L’ at the beginning of a label has this effect. You should find out what convention your system uses, and follow it.

The usual definition of this macro is as follows:

```
fprintf (stream, "L%s%d:\n", prefix, num)
```

ASM_OUTPUT_DEBUG_LABEL (*stream*, *prefix*, *num*)

A C statement to output to the stdio stream *stream* a debug info label whose name is made from the string *prefix* and the number *num*. This is useful for VLIW targets, where debug info labels may need to be treated differently than branch target labels. On some systems, branch target labels must be at the beginning of instruction bundles, but debug info labels can occur in the middle of instruction bundles.

If this macro is not defined, then `ASM_OUTPUT_INTERNAL_LABEL` will be used.

ASM_GENERATE_INTERNAL_LABEL (*string*, *prefix*, *num*)

A C statement to store into the string *string* a label whose name is made from the string *prefix* and the number *num*.

This string, when output subsequently by `assemble_name`, should produce the output that `ASM_OUTPUT_INTERNAL_LABEL` would produce with the same *prefix* and *num*.

If the string begins with ‘*’, then `assemble_name` will output the rest of the string unchanged. It is often convenient for `ASM_GENERATE_INTERNAL_LABEL` to use ‘*’ in this way. If the string doesn’t start with ‘*’, then `ASM_OUTPUT_LABELREF` gets to output the string, and may change it. (Of course, `ASM_OUTPUT_LABELREF` is also part of your machine description, so you should know what it does on your machine.)

ASM_FORMAT_PRIVATE_NAME (*outvar*, *name*, *number*)

A C expression to assign to *outvar* (which is a variable of type `char *`) a newly allocated string made from the string *name* and the number *number*, with some suitable punctuation added. Use `alloca` to get space for the string.

The string will be used as an argument to `ASM_OUTPUT_LABELREF` to produce an assembler label for an internal static variable whose name is *name*. Therefore, the string must be such as to result in valid assembler code. The argument *number* is different each time this macro is executed; it prevents conflicts between similarly-named internal static variables in different scopes.

Ideally this string should not be a valid C identifier, to prevent any conflict with the user's own symbols. Most assemblers allow periods or percent signs in assembler symbols; putting at least one of these between the name and the number will suffice.

ASM_OUTPUT_DEF (*stream*, *name*, *value*)

A C statement to output to the stdio stream *stream* assembler code which defines (equates) the symbol *name* to have the value *value*.

If `SET_ASM_OP` is defined, a default definition is provided which is correct for most systems.

ASM_OUTPUT_DEF_FROM_DECLS (*stream*, *decl_of_name*, *decl_of_value*)

A C statement to output to the stdio stream *stream* assembler code which defines (equates) the symbol whose tree node is *decl_of_name* to have the value of the tree node *decl_of_value*. This macro will be used in preference to 'ASM_OUTPUT_DEF' if it is defined and if the tree nodes are available.

If `SET_ASM_OP` is defined, a default definition is provided which is correct for most systems.

ASM_OUTPUT_WEAK_ALIAS (*stream*, *name*, *value*)

A C statement to output to the stdio stream *stream* assembler code which defines (equates) the weak symbol *name* to have the value *value*. If *value* is `NULL`, it defines *name* as an undefined weak symbol.

Define this macro if the target only supports weak aliases; define `ASM_OUTPUT_DEF` instead if possible.

OBJC_GEN_METHOD_LABEL (*buf*, *is_inst*, *class_name*, *cat_name*, *sel_name*)

Define this macro to override the default assembler names used for Objective-C methods.

The default name is a unique method number followed by the name of the class (e.g. '`_1_Foo`'). For methods in categories, the name of the category is also included in the assembler name (e.g. '`_1_Foo_Bar`').

These names are safe on most systems, but make debugging difficult since the method's selector is not present in the name. Therefore, particular systems define other ways of computing names.

buf is an expression of type `char *` which gives you a buffer in which to store the name; its length is as long as *class_name*, *cat_name* and *sel_name* put together, plus 50 characters extra.

The argument *is_inst* specifies whether the method is an instance method or a class method; *class_name* is the name of the class; *cat_name* is the name of the category (or NULL if the method is not in a category); and *sel_name* is the name of the selector.

On systems where the assembler can handle quoted names, you can use this macro to provide more human-readable names.

ASM_DECLARE_CLASS_REFERENCE (*stream*, *name*)

A C statement (sans semicolon) to output to the stdio stream *stream* commands to declare that the label *name* is an Objective-C class reference. This is only needed for targets whose linkers have special support for NeXT-style runtimes.

ASM_DECLARE_UNRESOLVED_REFERENCE (*stream*, *name*)

A C statement (sans semicolon) to output to the stdio stream *stream* commands to declare that the label *name* is an unresolved Objective-C class reference. This is only needed for targets whose linkers have special support for NeXT-style runtimes.

10.20.5 How Initialization Functions Are Handled

The compiled code for certain languages includes *constructors* (also called *initialization routines*)—functions to initialize data in the program when the program is started. These functions need to be called before the program is “started”—that is to say, before `main` is called.

Compiling some languages generates *destructors* (also called *termination routines*) that should be called when the program terminates.

To make the initialization and termination functions work, the compiler must output something in the assembler code to cause those functions to be called at the appropriate time. When you port the compiler to a new system, you need to specify how to do this.

There are two major ways that GCC currently supports the execution of initialization and termination functions. Each way has two variants. Much of the structure is common to all four variations.

The linker must build two lists of these functions—a list of initialization functions, called `__CTOR_LIST__`, and a list of termination functions, called `__DTOR_LIST__`.

Each list always begins with an ignored function pointer (which may hold 0, `-1`, or a count of the function pointers after it, depending on the environment). This is followed by a series of zero or more function pointers to constructors (or destructors), followed by a function pointer containing zero.

Depending on the operating system and its executable file format, either `'crtstuff.c'` or `'libgcc2.c'` traverses these lists at startup time and exit time. Constructors are called in reverse order of the list; destructors in forward order.

The best way to handle static constructors works only for object file formats which provide arbitrarily-named sections. A section is set aside for a list of constructors, and another for a list of destructors. Traditionally these are called `'ctors'` and `'dtors'`. Each object file that defines an initialization function also puts a word in the constructor section

to point to that function. The linker accumulates all these words into one contiguous `.ctors` section. Termination functions are handled similarly.

This method will be chosen as the default by `target-def.h` if `TARGET_ASM_NAMED_SECTION` is defined. A target that does not support arbitrary sections, but does support special designated constructor and destructor sections may define `CTORS_SECTION_ASM_OP` and `DTORS_SECTION_ASM_OP` to achieve the same effect.

When arbitrary sections are available, there are two variants, depending upon how the code in `crtstuff.c` is called. On systems that support a `.init` section which is executed at program startup, parts of `crtstuff.c` are compiled into that section. The program is linked by the `gcc` driver like this:

```
ld -o output_file crti.o crtbegin.o ... -lgcc crtend.o crtn.o
```

The prologue of a function (`__init`) appears in the `.init` section of `crti.o`; the epilogue appears in `crtn.o`. Likewise for the function `__fini` in the `.fini` section. Normally these files are provided by the operating system or by the GNU C library, but are provided by GCC for a few targets.

The objects `crtbegin.o` and `crtend.o` are (for most targets) compiled from `crtstuff.c`. They contain, among other things, code fragments within the `.init` and `.fini` sections that branch to routines in the `.text` section. The linker will pull all parts of a section together, which results in a complete `__init` function that invokes the routines we need at startup.

To use this variant, you must define the `INIT_SECTION_ASM_OP` macro properly.

If no `init` section is available, when GCC compiles any function called `main` (or more accurately, any function designated as a program entry point by the language front end calling `expand_main_function`), it inserts a procedure call to `__main` as the first executable code after the function prologue. The `__main` function is defined in `libgcc2.c` and runs the global constructors.

In file formats that don't support arbitrary sections, there are again two variants. In the simplest variant, the GNU linker (GNU `ld`) and an `'a.out'` format must be used. In this case, `TARGET_ASM_CONSTRUCTOR` is defined to produce a `.stabs` entry of type `'N_SETT'`, referencing the name `__CTOR_LIST__`, and with the address of the void function containing the initialization code as its value. The GNU linker recognizes this as a request to add the value to a `set`; the values are accumulated, and are eventually placed in the executable as a vector in the format described above, with a leading (ignored) count and a trailing zero element. `TARGET_ASM_DESTRUCTOR` is handled similarly. Since no `init` section is available, the absence of `INIT_SECTION_ASM_OP` causes the compilation of `main` to call `__main` as above, starting the initialization process.

The last variant uses neither arbitrary sections nor the GNU linker. This is preferable when you want to do dynamic linking and when using file formats which the GNU linker does not support, such as `'ECOFF'`. In this case, `TARGET_HAVE_CTORS_DTORS` is false, initialization and termination functions are recognized simply by their names. This requires an extra program in the linkage step, called `collect2`. This program pretends to be the linker, for use with GCC; it does its job by running the ordinary linker, but also arranges to include the vectors of initialization and termination functions. These functions are called via `__main` as described above. In order to use this method, `use_collect2` must be defined in the target in `'config.gcc'`.

10.20.6 Macros Controlling Initialization Routines

Here are the macros that control how the compiler handles initialization and termination functions:

INIT_SECTION_ASM_OP

If defined, a C string constant, including spacing, for the assembler operation to identify the following data as initialization code. If not defined, GCC will assume such a section does not exist. When you are using special sections for initialization and termination functions, this macro also controls how ‘`crtstuff.c`’ and ‘`libgcc2.c`’ arrange to run the initialization functions.

HAS_INIT_SECTION

If defined, `main` will not call `__main` as described above. This macro should be defined for systems that control start-up code on a symbol-by-symbol basis, such as OSF/1, and should not be defined explicitly for systems that support `INIT_SECTION_ASM_OP`.

LD_INIT_SWITCH

If defined, a C string constant for a switch that tells the linker that the following symbol is an initialization routine.

LD_FINI_SWITCH

If defined, a C string constant for a switch that tells the linker that the following symbol is a finalization routine.

COLLECT_SHARED_INIT_FUNC (*stream*, *func*)

If defined, a C statement that will write a function that can be automatically called when a shared library is loaded. The function should call *func*, which takes no arguments. If not defined, and the object format requires an explicit initialization function, then a function called `_GLOBAL__DI` will be generated.

This function and the following one are used by `collect2` when linking a shared library that needs constructors or destructors, or has DWARF2 exception tables embedded in the code.

COLLECT_SHARED_FINI_FUNC (*stream*, *func*)

If defined, a C statement that will write a function that can be automatically called when a shared library is unloaded. The function should call *func*, which takes no arguments. If not defined, and the object format requires an explicit finalization function, then a function called `_GLOBAL__DD` will be generated.

INVOKE__main

If defined, `main` will call `__main` despite the presence of `INIT_SECTION_ASM_OP`. This macro should be defined for systems where the `init` section is not actually run automatically, but is still useful for collecting the lists of constructors and destructors.

SUPPORTS_INIT_PRIORITY

If nonzero, the C++ `init_priority` attribute is supported and the compiler should emit instructions to control the order of initialization of objects. If zero, the compiler will issue an error message upon encountering an `init_priority` attribute.

bool TARGET_HAVE_CTORS_DTORS [Target Hook]

This value is true if the target supports some “native” method of collecting constructors and destructors to be run at startup and exit. It is false if we must use `collect2`.

void TARGET_ASM_CONSTRUCTOR (*rtx symbol*, *int priority*) [Target Hook]

If defined, a function that outputs assembler code to arrange to call the function referenced by *symbol* at initialization time.

Assume that *symbol* is a `SYMBOL_REF` for a function taking no arguments and with no return value. If the target supports initialization priorities, *priority* is a value between 0 and `MAX_INIT_PRIORITY`; otherwise it must be `DEFAULT_INIT_PRIORITY`.

If this macro is not defined by the target, a suitable default will be chosen if (1) the target supports arbitrary section names, (2) the target defines `CTORS_SECTION_ASM_OP`, or (3) `USE_COLLECT2` is not defined.

void TARGET_ASM_DESTRUCTOR (*rtx symbol*, *int priority*) [Target Hook]

This is like `TARGET_ASM_CONSTRUCTOR` but used for termination functions rather than initialization functions.

If `TARGET_HAVE_CTORS_DTORS` is true, the initialization routine generated for the generated object file will have static linkage.

If your system uses `collect2` as the means of processing constructors, then that program normally uses `nm` to scan an object file for constructor functions to be called.

On certain kinds of systems, you can define these macros to make `collect2` work faster (and, in some cases, make it work at all):

OBJECT_FORMAT_COFF

Define this macro if the system uses COFF (Common Object File Format) object files, so that `collect2` can assume this format and scan object files directly for dynamic constructor/destructor functions.

OBJECT_FORMAT_ROSE

Define this macro if the system uses ROSE format object files, so that `collect2` can assume this format and scan object files directly for dynamic constructor/destructor functions.

These macros are effective only in a native compiler; `collect2` as part of a cross compiler always uses `nm` for the target machine.

REAL_NM_FILE_NAME

Define this macro as a C string constant containing the file name to use to execute `nm`. The default is to search the path normally for `nm`.

If your system supports shared libraries and has a program to list the dynamic dependencies of a given library or executable, you can define these macros to enable support for running initialization and termination functions in shared libraries:

LDD_SUFFIX

Define this macro to a C string constant containing the name of the program which lists dynamic dependencies, like `"ldd"` under SunOS 4.

PARSE_LDD_OUTPUT (*ptr*)

Define this macro to be C code that extracts filenames from the output of the program denoted by `LDD_SUFFIX`. *ptr* is a variable of type `char *` that points to the beginning of a line of output from `LDD_SUFFIX`. If the line lists a dynamic dependency, the code must advance *ptr* to the beginning of the filename on that line. Otherwise, it must set *ptr* to `NULL`.

10.20.7 Output of Assembler Instructions

This describes assembler instruction output.

REGISTER_NAMES

A C initializer containing the assembler's names for the machine registers, each one as a C string constant. This is what translates register numbers in the compiler into assembler language.

ADDITIONAL_REGISTER_NAMES

If defined, a C initializer for an array of structures containing a name and a register number. This macro defines additional names for hard registers, thus allowing the `asm` option in declarations to refer to registers using alternate names.

ASM_OUTPUT_OPCODE (*stream*, *ptr*)

Define this macro if you are using an unusual assembler that requires different names for the machine instructions.

The definition is a C statement or statements which output an assembler instruction opcode to the stdio stream *stream*. The macro-operand *ptr* is a variable of type `char *` which points to the opcode name in its "internal" form—the form that is written in the machine description. The definition should output the opcode name to *stream*, performing any translation you desire, and increment the variable *ptr* to point at the end of the opcode so that it will not be output twice.

In fact, your macro definition may process less than the entire opcode name, or more than the opcode name; but if you want to process text that includes '%'-sequences to substitute operands, you must take care of the substitution yourself. Just be sure to increment *ptr* over whatever text should not be output normally.

If you need to look at the operand values, they can be found as the elements of `recog_data.operand`.

If the macro definition does nothing, the instruction is output in the usual way.

FINAL_PRESCAN_INSN (*insn*, *opvec*, *noperands*)

If defined, a C statement to be executed just prior to the output of assembler code for *insn*, to modify the extracted operands so they will be output differently.

Here the argument *opvec* is the vector containing the operands extracted from *insn*, and *noperands* is the number of elements of the vector which contain meaningful data for this *insn*. The contents of this vector are what will be

used to convert the `insn` template into assembler code, so you can change the assembler output by changing the contents of the vector.

This macro is useful when various assembler syntaxes share a single file of instruction patterns; by defining this macro differently, you can cause a large class of instructions to be output differently (such as with rearranged operands). Naturally, variations in assembler syntax affecting individual `insn` patterns ought to be handled by writing conditional output routines in those patterns.

If this macro is not defined, it is equivalent to a null statement.

FINAL_PRESCAN_LABEL

If defined, `FINAL_PRESCAN_INSN` will be called on each `CODE_LABEL`. In that case, `opvec` will be a null pointer and `noperands` will be zero.

PRINT_OPERAND (*stream*, *x*, *code*)

A C compound statement to output to stdio stream *stream* the assembler syntax for an instruction operand *x*. *x* is an RTL expression.

code is a value that can be used to specify one of several ways of printing the operand. It is used when identical operands must be printed differently depending on the context. *code* comes from the ‘%’ specification that was used to request printing of the operand. If the specification was just ‘%*digit*’ then *code* is 0; if the specification was ‘%*ltr digit*’ then *code* is the ASCII code for *ltr*.

If *x* is a register, this macro should print the register’s name. The names can be found in an array `reg_names` whose type is `char *[]`. `reg_names` is initialized from `REGISTER_NAMES`.

When the machine description has a specification ‘%*punct*’ (a ‘%’ followed by a punctuation character), this macro is called with a null pointer for *x* and the punctuation character for *code*.

PRINT_OPERAND_PUNCT_VALID_P (*code*)

A C expression which evaluates to true if *code* is a valid punctuation character for use in the `PRINT_OPERAND` macro. If `PRINT_OPERAND_PUNCT_VALID_P` is not defined, it means that no punctuation characters (except for the standard one, ‘%’) are used in this way.

PRINT_OPERAND_ADDRESS (*stream*, *x*)

A C compound statement to output to stdio stream *stream* the assembler syntax for an instruction operand that is a memory reference whose address is *x*. *x* is an RTL expression.

On some machines, the syntax for a symbolic address depends on the section that the address refers to. On these machines, define the hook `TARGET_ENCODE_SECTION_INFO` to store the information into the `symbol_ref`, and then check for it here. See Section 10.20 [Assembler Format], page 279.

DBR_OUTPUT_SEQEND(*file*)

A C statement, to be executed after all slot-filler instructions have been output. If necessary, call `dbr_sequence_length` to determine the number of slots filled in a sequence (zero if not currently outputting a sequence), to decide how many no-ops to output, or whatever.

Don't define this macro if it has nothing to do, but it is helpful in reading assembly output if the extent of the delay sequence is made explicit (e.g. with white space).

Note that output routines for instructions with delay slots must be prepared to deal with not being output as part of a sequence (i.e. when the scheduling pass is not run, or when no slot fillers could be found.) The variable `final_sequence` is null when not processing a sequence, otherwise it contains the `sequence` rtx being output.

REGISTER_PREFIX

LOCAL_LABEL_PREFIX

USER_LABEL_PREFIX

IMMEDIATE_PREFIX

If defined, C string expressions to be used for the '%R', '%L', '%U', and '%I' options of `asm_fprintf` (see 'final.c'). These are useful when a single 'md' file must support multiple assembler formats. In that case, the various 'tm.h' files can define these macros differently.

ASM_FPRINTF_EXTENSIONS(*file*, *argptr*, *format*)

If defined this macro should expand to a series of `case` statements which will be parsed inside the `switch` statement of the `asm_fprintf` function. This allows targets to define extra printf formats which may be useful when generating their assembler statements. Note that upper case letters are reserved for future generic extensions to `asm_fprintf`, and so are not available to target specific code. The output file is given by the parameter *file*. The varargs input pointer is *argptr* and the rest of the format string, starting the character after the one that is being switched upon, is pointed to by *format*.

ASSEMBLER_DIALECT

If your target supports multiple dialects of assembler language (such as different opcodes), define this macro as a C expression that gives the numeric index of the assembler language dialect to use, with zero as the first variant.

If this macro is defined, you may use constructs of the form

```
{option0|option1|option2...}
```

in the output templates of patterns (see Section 9.5 [Output Template], page 123) or in the first argument of `asm_fprintf`. This construct outputs 'option0', 'option1', 'option2', etc., if the value of `ASSEMBLER_DIALECT` is zero, one, two, etc. Any special characters within these strings retain their usual meaning. If there are fewer alternatives within the braces than the value of `ASSEMBLER_DIALECT`, the construct outputs nothing.

If you do not define this macro, the characters '{', '|' and '}' do not have any special meaning when used in templates or operands to `asm_fprintf`.

Define the macros `REGISTER_PREFIX`, `LOCAL_LABEL_PREFIX`, `USER_LABEL_PREFIX` and `IMMEDIATE_PREFIX` if you can express the variations in assembler language syntax with that mechanism. Define `ASSEMBLER_DIALECT` and use the '{option0|option1}' syntax if the syntax variants are larger and involve such things as different opcodes or operand order.

ASM_OUTPUT_REG_PUSH (*stream*, *regno*)

A C expression to output to *stream* some assembler code which will push hard register number *regno* onto the stack. The code need not be optimal, since this macro is used only when profiling.

ASM_OUTPUT_REG_POP (*stream*, *regno*)

A C expression to output to *stream* some assembler code which will pop hard register number *regno* off of the stack. The code need not be optimal, since this macro is used only when profiling.

10.20.8 Output of Dispatch Tables

This concerns dispatch tables.

ASM_OUTPUT_ADDR_DIFF_ELT (*stream*, *body*, *value*, *rel*)

A C statement to output to the stdio stream *stream* an assembler pseudo-instruction to generate a difference between two labels. *value* and *rel* are the numbers of two internal labels. The definitions of these labels are output using **ASM_OUTPUT_INTERNAL_LABEL**, and they must be printed in the same way here. For example,

```
fprintf (stream, "\t.word L%d-L%d\n",
        value, rel)
```

You must provide this macro on machines where the addresses in a dispatch table are relative to the table's own address. If defined, GCC will also use this macro on all machines when producing PIC. *body* is the body of the **ADDR_DIFF_VEC**; it is provided so that the mode and flags can be read.

ASM_OUTPUT_ADDR_VEC_ELT (*stream*, *value*)

This macro should be provided on machines where the addresses in a dispatch table are absolute.

The definition should be a C statement to output to the stdio stream *stream* an assembler pseudo-instruction to generate a reference to a label. *value* is the number of an internal label whose definition is output using **ASM_OUTPUT_INTERNAL_LABEL**. For example,

```
fprintf (stream, "\t.word L%d\n", value)
```

ASM_OUTPUT_CASE_LABEL (*stream*, *prefix*, *num*, *table*)

Define this if the label before a jump-table needs to be output specially. The first three arguments are the same as for **ASM_OUTPUT_INTERNAL_LABEL**; the fourth argument is the jump-table which follows (a **jump_insn** containing an **addr_vec** or **addr_diff_vec**).

This feature is used on system V to output a **swbeg** statement for the table.

If this macro is not defined, these labels are output with **ASM_OUTPUT_INTERNAL_LABEL**.

ASM_OUTPUT_CASE_END (*stream*, *num*, *table*)

Define this if something special must be output at the end of a jump-table. The definition should be a C statement to be executed after the assembler code for

the table is written. It should write the appropriate code to stdio stream *stream*. The argument *table* is the jump-table insn, and *num* is the label-number of the preceding label.

If this macro is not defined, nothing special is output at the end of the jump-table.

10.20.9 Assembler Commands for Exception Regions

This describes commands marking the start and the end of an exception region.

EH_FRAME_SECTION_NAME

If defined, a C string constant for the name of the section containing exception handling frame unwind information. If not defined, GCC will provide a default definition if the target supports named sections. ‘*crtstuff.c*’ uses this macro to switch to the appropriate section.

You should define this symbol if your target supports DWARF 2 frame unwind information and the default definition does not work.

EH_FRAME_IN_DATA_SECTION

If defined, DWARF 2 frame unwind information will be placed in the data section even though the target supports named sections. This might be necessary, for instance, if the system linker does garbage collection and sections cannot be marked as not to be collected.

Do not define this macro unless `TARGET_ASM_NAMED_SECTION` is also defined.

MASK_RETURN_ADDR

An rtx used to mask the return address found via `RETURN_ADDR_RTX`, so that it does not contain any extraneous set bits in it.

DWARF2_UNWIND_INFO

Define this macro to 0 if your target supports DWARF 2 frame unwind information, but it does not yet work with exception handling. Otherwise, if your target supports this information (if it defines ‘`INCOMING_RETURN_ADDR_RTX`’ and either ‘`UNALIGNED_INT_ASM_OP`’ or ‘`OBJECT_FORMAT_ELF`’), GCC will provide a default definition of 1.

If this macro is defined to 1, the DWARF 2 unwinder will be the default exception handling mechanism; otherwise, `setjmp/longjmp` will be used by default.

If this macro is defined to anything, the DWARF 2 unwinder will be used instead of inline unwinders and `__unwind_function` in the non-`setjmp` case.

DWARF_CIE_DATA_ALIGNMENT

This macro need only be defined if the target might save registers in the function prologue at an offset to the stack pointer that is not aligned to `UNITS_PER_WORD`. The definition should be the negative minimum alignment if `STACK_GROWS_DOWNWARD` is defined, and the positive minimum alignment otherwise. See Section 10.21.5 [SDB and DWARF], page 308. Only applicable if the target supports DWARF 2 frame unwind information.

`void TARGET_ASM_EXCEPTION_SECTION ()` [Target Hook]

If defined, a function that switches to the section in which the main exception table is to be placed (see Section 10.18 [Sections], page 275). The default is a function that switches to a section named `.gcc_except_table` on machines that support named sections via `TARGET_ASM_NAMED_SECTION`, otherwise if `'-fpic'` or `'-fPIC'` is in effect, the `data_section`, otherwise the `readonly_data_section`.

`void TARGET_ASM_EH_FRAME_SECTION ()` [Target Hook]

If defined, a function that switches to the section in which the DWARF 2 frame unwind information to be placed (see Section 10.18 [Sections], page 275). The default is a function that outputs a standard GAS section directive, if `EH_FRAME_SECTION_NAME` is defined, or else a data section directive followed by a synthetic label.

Target Hook `bool TARGET_TERMINATE_DW2_EH_FRAME_INFO` [Variable]

Contains the value true if the target should add a zero word onto the end of a Dwarf-2 frame info section when used for exception handling. Default value is false if `EH_FRAME_SECTION_NAME` is defined, and true otherwise.

10.20.10 Assembler Commands for Alignment

This describes commands for alignment.

`JUMP_ALIGN (label)`

The alignment (log base 2) to put in front of *label*, which is a common destination of jumps and has no fallthru incoming edge.

This macro need not be defined if you don't want any special alignment to be done at such a time. Most machine descriptions do not currently define the macro.

Unless it's necessary to inspect the *label* parameter, it is better to set the variable *align_jumps* in the target's `OVERWRITE_OPTIONS`. Otherwise, you should try to honor the user's selection in *align_jumps* in a `JUMP_ALIGN` implementation.

`LABEL_ALIGN_AFTER_BARRIER (label)`

The alignment (log base 2) to put in front of *label*, which follows a `BARRIER`.

This macro need not be defined if you don't want any special alignment to be done at such a time. Most machine descriptions do not currently define the macro.

`LABEL_ALIGN_AFTER_BARRIER_MAX_SKIP`

The maximum number of bytes to skip when applying `LABEL_ALIGN_AFTER_BARRIER`. This works only if `ASM_OUTPUT_MAX_SKIP_ALIGN` is defined.

`LOOP_ALIGN (label)`

The alignment (log base 2) to put in front of *label*, which follows a `NOTE_INSN_LOOP_BEG` note.

This macro need not be defined if you don't want any special alignment to be done at such a time. Most machine descriptions do not currently define the macro.

Unless it's necessary to inspect the *label* parameter, it is better to set the variable `align_loops` in the target's `OVERRIDE_OPTIONS`. Otherwise, you should try to honor the user's selection in `align_loops` in a `LOOP_ALIGN` implementation.

`LOOP_ALIGN_MAX_SKIP`

The maximum number of bytes to skip when applying `LOOP_ALIGN`. This works only if `ASM_OUTPUT_MAX_SKIP_ALIGN` is defined.

`LABEL_ALIGN (label)`

The alignment (log base 2) to put in front of *label*. If `LABEL_ALIGN_AFTER_BARRIER / LOOP_ALIGN` specify a different alignment, the maximum of the specified values is used.

Unless it's necessary to inspect the *label* parameter, it is better to set the variable `align_labels` in the target's `OVERRIDE_OPTIONS`. Otherwise, you should try to honor the user's selection in `align_labels` in a `LABEL_ALIGN` implementation.

`LABEL_ALIGN_MAX_SKIP`

The maximum number of bytes to skip when applying `LABEL_ALIGN`. This works only if `ASM_OUTPUT_MAX_SKIP_ALIGN` is defined.

`ASM_OUTPUT_SKIP (stream, nbytes)`

A C statement to output to the stdio stream *stream* an assembler instruction to advance the location counter by *nbytes* bytes. Those bytes should be zero when loaded. *nbytes* will be a C expression of type `int`.

`ASM_NO_SKIP_IN_TEXT`

Define this macro if `ASM_OUTPUT_SKIP` should not be used in the text section because it fails to put zeros in the bytes that are skipped. This is true on many Unix systems, where the pseudo-op to skip bytes produces no-op instructions rather than zeros when used in the text section.

`ASM_OUTPUT_ALIGN (stream, power)`

A C statement to output to the stdio stream *stream* an assembler command to advance the location counter to a multiple of 2 to the *power* bytes. *power* will be a C expression of type `int`.

`ASM_OUTPUT_ALIGN_WITH_NOP (stream, power)`

Like `ASM_OUTPUT_ALIGN`, except that the “nop” instruction is used for padding, if necessary.

`ASM_OUTPUT_MAX_SKIP_ALIGN (stream, power, max_skip)`

A C statement to output to the stdio stream *stream* an assembler command to advance the location counter to a multiple of 2 to the *power* bytes, but only if *max_skip* or fewer bytes are needed to satisfy the alignment request. *power* and *max_skip* will be a C expression of type `int`.

10.21 Controlling Debugging Information Format

This describes how to specify debugging information.

10.21.1 Macros Affecting All Debugging Formats

These macros affect all debugging formats.

DBX_REGISTER_NUMBER (*regno*)

A C expression that returns the DBX register number for the compiler register number *regno*. In the default macro provided, the value of this expression will be *regno* itself. But sometimes there are some registers that the compiler knows about and DBX does not, or vice versa. In such cases, some register may need to have one number in the compiler and another for DBX.

If two registers have consecutive numbers inside GCC, and they can be used as a pair to hold a multiword value, then they *must* have consecutive numbers after renumbering with DBX_REGISTER_NUMBER. Otherwise, debuggers will be unable to access such a pair, because they expect register pairs to be consecutive in their own numbering scheme.

If you find yourself defining DBX_REGISTER_NUMBER in way that does not preserve register pairs, then what you must do instead is redefine the actual register numbering scheme.

DEBUGGER_AUTO_OFFSET (*x*)

A C expression that returns the integer offset value for an automatic variable having address *x* (an RTL expression). The default computation assumes that *x* is based on the frame-pointer and gives the offset from the frame-pointer. This is required for targets that produce debugging output for DBX or COFF-style debugging output for SDB and allow the frame-pointer to be eliminated when the ‘-g’ options is used.

DEBUGGER_ARG_OFFSET (*offset*, *x*)

A C expression that returns the integer offset value for an argument having address *x* (an RTL expression). The nominal offset is *offset*.

PREFERRED_DEBUGGING_TYPE

A C expression that returns the type of debugging output GCC should produce when the user specifies just ‘-g’. Define this if you have arranged for GCC to support more than one format of debugging output. Currently, the allowable values are DBX_DEBUG, SDB_DEBUG, DWARF_DEBUG, DWARF2_DEBUG, XCOFF_DEBUG, VMS_DEBUG, and VMS_AND_DWARF2_DEBUG.

When the user specifies ‘-gddb’, GCC normally also uses the value of this macro to select the debugging output format, but with two exceptions. If DWARF2_DEBUGGING_INFO is defined and LINKER_DOES_NOT_WORK_WITH_DWARF2 is not defined, GCC uses the value DWARF2_DEBUG. Otherwise, if DBX_DEBUGGING_INFO is defined, GCC uses DBX_DEBUG.

The value of this macro only affects the default debugging output; the user can always get a specific type of output by using `-gstabs`, `-gcoff`, `-gdwarf-1`, `-gdwarf-2`, `-gxcoff`, or `-gvms`.

10.21.2 Specific Options for DBX Output

These are specific options for DBX output.

DBX_DEBUGGING_INFO

Define this macro if GCC should produce debugging output for DBX in response to the `-g` option.

XCOFF_DEBUGGING_INFO

Define this macro if GCC should produce XCOFF format debugging output in response to the `-g` option. This is a variant of DBX format.

DEFAULT_GDB_EXTENSIONS

Define this macro to control whether GCC should by default generate GDB's extended version of DBX debugging information (assuming DBX-format debugging information is enabled at all). If you don't define the macro, the default is 1: always generate the extended information if there is any occasion to.

DEBUG_SYMS_TEXT

Define this macro if all `.stabs` commands should be output while in the text section.

ASM_STABS_OP

A C string constant, including spacing, naming the assembler pseudo op to use instead of `"\t.stabs\t"` to define an ordinary debugging symbol. If you don't define this macro, `"\t.stabs\t"` is used. This macro applies only to DBX debugging information format.

ASM_STABD_OP

A C string constant, including spacing, naming the assembler pseudo op to use instead of `"\t.stabd\t"` to define a debugging symbol whose value is the current location. If you don't define this macro, `"\t.stabd\t"` is used. This macro applies only to DBX debugging information format.

ASM_STABN_OP

A C string constant, including spacing, naming the assembler pseudo op to use instead of `"\t.stabn\t"` to define a debugging symbol with no name. If you don't define this macro, `"\t.stabn\t"` is used. This macro applies only to DBX debugging information format.

DBX_NO_XREFS

Define this macro if DBX on your system does not support the construct `'xstagname'`. On some systems, this construct is used to describe a forward reference to a structure named *tagname*. On other systems, this construct is not supported at all.

DBX_CONTIN_LENGTH

A symbol name in DBX-format debugging information is normally continued (split into two separate `.stabs` directives) when it exceeds a certain length (by default, 80 characters). On some operating systems, DBX requires this splitting; on others, splitting must not be done. You can inhibit splitting by defining this macro with the value zero. You can override the default splitting-length by defining this macro as an expression for the length you desire.

DBX_CONTIN_CHAR

Normally continuation is indicated by adding a `\` character to the end of a `.stabs` string when a continuation follows. To use a different character instead, define this macro as a character constant for the character you want to use. Do not define this macro if backslash is correct for your system.

DBX_STATIC_STAB_DATA_SECTION

Define this macro if it is necessary to go to the data section before outputting the `.stabs` pseudo-op for a non-global static variable.

DBX_TYPE_DECL_STABS_CODE

The value to use in the “code” field of the `.stabs` directive for a typedef. The default is `N_LSYM`.

DBX_STATIC_CONST_VAR_CODE

The value to use in the “code” field of the `.stabs` directive for a static variable located in the text section. DBX format does not provide any “right” way to do this. The default is `N_FUN`.

DBX_REGPARAM_STABS_CODE

The value to use in the “code” field of the `.stabs` directive for a parameter passed in registers. DBX format does not provide any “right” way to do this. The default is `N_RSYM`.

DBX_REGPARAM_STABS_LETTER

The letter to use in DBX symbol data to identify a symbol as a parameter passed in registers. DBX format does not customarily provide any way to do this. The default is `'P'`.

DBX_MEMPARAM_STABS_LETTER

The letter to use in DBX symbol data to identify a symbol as a stack parameter. The default is `'p'`.

DBX_FUNCTION_FIRST

Define this macro if the DBX information for a function and its arguments should precede the assembler code for the function. Normally, in DBX format, the debugging information entirely follows the assembler code.

DBX_LBRAC_FIRST

Define this macro if the `N_LBRAC` symbol for a block should precede the debugging information for variables and functions defined in that block. Normally, in DBX format, the `N_LBRAC` symbol comes first.

DBX_BLOCKS_FUNCTION_RELATIVE

Define this macro if the value of a symbol describing the scope of a block (`N_LBRAC` or `N_RBRAC`) should be relative to the start of the enclosing function. Normally, GCC uses an absolute address.

DBX_USE_BINCL

Define this macro if GCC should generate `N_BINCL` and `N_EINCL` stabs for included header files, as on Sun systems. This macro also directs GCC to output a type number as a pair of a file number and a type number within the file. Normally, GCC does not generate `N_BINCL` or `N_EINCL` stabs, and it outputs a single number for a type number.

10.21.3 Open-Ended Hooks for DBX Format

These are hooks for DBX format.

DBX_OUTPUT_LBRAC (*stream*, *name*)

Define this macro to say how to output to *stream* the debugging information for the start of a scope level for variable names. The argument *name* is the name of an assembler symbol (for use with `assemble_name`) whose value is the address where the scope begins.

DBX_OUTPUT_RBRAC (*stream*, *name*)

Like `DBX_OUTPUT_LBRAC`, but for the end of a scope level.

DBX_OUTPUT_NFUN (*stream*, *lscope_label*, *decl*)

Define this macro if the target machine requires special handling to output an `N_FUN` entry for the function *decl*.

DBX_OUTPUT_ENUM (*stream*, *type*)

Define this macro if the target machine requires special handling to output an enumeration type. The definition should be a C statement (sans semicolon) to output the appropriate information to *stream* for the type *type*.

DBX_OUTPUT_FUNCTION_END (*stream*, *function*)

Define this macro if the target machine requires special output at the end of the debugging information for a function. The definition should be a C statement (sans semicolon) to output the appropriate information to *stream*. *function* is the `FUNCTION_DECL` node for the function.

DBX_OUTPUT_STANDARD_TYPES (*syms*)

Define this macro if you need to control the order of output of the standard data types at the beginning of compilation. The argument *syms* is a `tree` which is a chain of all the predefined global symbols, including names of data types.

Normally, DBX output starts with definitions of the types for integers and characters, followed by all the other predefined types of the particular language in no particular order.

On some machines, it is necessary to output different particular types first. To do this, define `DBX_OUTPUT_STANDARD_TYPES` to output those symbols in the

necessary order. Any predefined types that you don't explicitly output will be output afterward in no particular order.

Be careful not to define this macro so that it works only for C. There are no global variables to access most of the built-in types, because another language may have another set of types. The way to output a particular type is to look through `syms` to see if you can find it. Here is an example:

```

{
  tree decl;
  for (decl = syms; decl; decl = TREE_CHAIN (decl))
    if (!strcmp (IDENTIFIER_POINTER (DECL_NAME (decl)),
                "long int"))
      dbxout_symbol (decl);
  ...
}

```

This does nothing if the expected type does not exist.

See the function `init_decl_processing` in 'c-decl.c' to find the names to use for all the built-in C types.

Here is another way of finding a particular type:

```

{
  tree decl;
  for (decl = syms; decl; decl = TREE_CHAIN (decl))
    if (TREE_CODE (decl) == TYPE_DECL
        && (TREE_CODE (TREE_TYPE (decl))
            == INTEGER_CST)
        && TYPE_PRECISION (TREE_TYPE (decl)) == 16
        && TYPE_UNSIGNED (TREE_TYPE (decl)))
      /* This must be unsigned short. */
      dbxout_symbol (decl);
  ...
}

```

`NO_DBX_FUNCTION_END`

Some stabs encapsulation formats (in particular ECOFF), cannot handle the `.stabs "",N_FUN,,0,0,Lscope-function-1` gdb dbx extension construct. On those machines, define this macro to turn this feature off without disturbing the rest of the gdb extensions.

10.21.4 File Names in DBX Format

This describes file names in DBX format.

`DBX_WORKING_DIRECTORY`

Define this if DBX wants to have the current directory recorded in each object file.

Note that the working directory is always recorded if GDB extensions are enabled.

`DBX_OUTPUT_MAIN_SOURCE_FILENAME (stream, name)`

A C statement to output DBX debugging information to the stdio stream `stream` which indicates that file `name` is the main source file—the file specified as the input file for compilation. This macro is called only once, at the beginning of compilation.

This macro need not be defined if the standard form of output for DBX debugging information is appropriate.

DBX_OUTPUT_MAIN_SOURCE_DIRECTORY (*stream*, *name*)

A C statement to output DBX debugging information to the stdio stream *stream* which indicates that the current directory during compilation is named *name*.

This macro need not be defined if the standard form of output for DBX debugging information is appropriate.

DBX_OUTPUT_MAIN_SOURCE_FILE_END (*stream*, *name*)

A C statement to output DBX debugging information at the end of compilation of the main source file *name*.

If you don't define this macro, nothing special is output at the end of compilation, which is correct for most machines.

DBX_OUTPUT_SOURCE_FILENAME (*stream*, *name*)

A C statement to output DBX debugging information to the stdio stream *stream* which indicates that file *name* is the current source file. This output is generated each time input shifts to a different source file as a result of '#include', the end of an included file, or a '#line' command.

This macro need not be defined if the standard form of output for DBX debugging information is appropriate.

10.21.5 Macros for SDB and DWARF Output

Here are macros for SDB and DWARF output.

SDB_DEBUGGING_INFO

Define this macro if GCC should produce COFF-style debugging output for SDB in response to the '-g' option.

DWARF_DEBUGGING_INFO

Define this macro if GCC should produce dwarf format debugging output in response to the '-g' option.

DWARF2_DEBUGGING_INFO

Define this macro if GCC should produce dwarf version 2 format debugging output in response to the '-g' option.

To support optional call frame debugging information, you must also define INCOMING_RETURN_ADDR_RTX and either set RTX_FRAME_RELATED_P on the prologue insns if you use RTL for the prologue, or call dwarf2out_def_cfa and dwarf2out_reg_save as appropriate from TARGET_ASM_FUNCTION_PROLOGUE if you don't.

DWARF2_FRAME_INFO

Define this macro to a nonzero value if GCC should always output Dwarf 2 frame information. If DWARF2_UNWIND_INFO (see Section 10.20.9 [Exception Region Output], page 300) is nonzero, GCC will output this information not matter how you define DWARF2_FRAME_INFO.

LINKER_DOES_NOT_WORK_WITH_DWARF2

Define this macro if the linker does not work with Dwarf version 2. Normally, if the user specifies only `-ggdb` GCC will use Dwarf version 2 if available; this macro disables this. See the description of the `PREFERRED_DEBUGGING_TYPE` macro for more details.

DWARF2_GENERATE_TEXT_SECTION_LABEL

By default, the Dwarf 2 debugging information generator will generate a label to mark the beginning of the text section. If it is better simply to use the name of the text section itself, rather than an explicit label, to indicate the beginning of the text section, define this macro to zero.

DWARF2_ASM_LINE_DEBUG_INFO

Define this macro to be a nonzero value if the assembler can generate Dwarf 2 line debug info sections. This will result in much more compact line number tables, and hence is desirable if it works.

ASM_OUTPUT_DWARF_DELTA (*stream, size, label1, label2*)

A C statement to issue assembly directives that create a difference between the two given labels, using an integer of the given size.

ASM_OUTPUT_DWARF_OFFSET (*stream, size, label*)

A C statement to issue assembly directives that create a section-relative reference to the given label, using an integer of the given size.

ASM_OUTPUT_DWARF_PCREL (*stream, size, label*)

A C statement to issue assembly directives that create a self-relative reference to the given label, using an integer of the given size.

PUT_SDB...

Define these macros to override the assembler syntax for the special SDB assembler directives. See `'sdbout.c'` for a list of these macros and their arguments. If the standard syntax is used, you need not define them yourself.

SDB_DELIM

Some assemblers do not support a semicolon as a delimiter, even between SDB assembler directives. In that case, define this macro to be the delimiter to use (usually `'\n'`). It is not necessary to define a new set of `PUT_SDB_op` macros if this is the only change required.

SDB_GENERATE_FAKE

Define this macro to override the usual method of constructing a dummy name for anonymous structure and union types. See `'sdbout.c'` for more information.

SDB_ALLOW_UNKNOWN_REFERENCES

Define this macro to allow references to unknown structure, union, or enumeration tags to be emitted. Standard COFF does not allow handling of unknown references, MIPS ECOFF has support for it.

SDB_ALLOW_FORWARD_REFERENCES

Define this macro to allow references to structure, union, or enumeration tags that have not yet been seen to be handled. Some assemblers choke if forward tags are used, while some require it.

10.21.6 Macros for VMS Debug Format

Here are macros for VMS debug format.

VMS_DEBUGGING_INFO

Define this macro if GCC should produce debugging output for VMS in response to the ‘-g’ option. The default behavior for VMS is to generate minimal debug info for a traceback in the absence of ‘-g’ unless explicitly overridden with ‘-g0’. This behavior is controlled by `OPTIMIZATION_OPTIONS` and `OVERRIDE_OPTIONS`.

10.22 Cross Compilation and Floating Point

While all modern machines use two's-complement representation for integers, there are a variety of representations for floating point numbers. This means that in a cross-compiler the representation of floating point numbers in the compiled program may be different from that used in the machine doing the compilation.

Because different representation systems may offer different amounts of range and precision, all floating point constants must be represented in the target machine's format. Therefore, the cross compiler cannot safely use the host machine's floating point arithmetic; it must emulate the target's arithmetic. To ensure consistency, GCC always uses emulation to work with floating point values, even when the host and target floating point formats are identical.

The following macros are provided by ‘`real.h`’ for the compiler to use. All parts of the compiler which generate or optimize floating-point calculations must use these macros. They may evaluate their operands more than once, so operands must not have side effects.

REAL_VALUE_TYPE [Macro]

The C data type to be used to hold a floating point value in the target machine's format. Typically this is a `struct` containing an array of `HOST_WIDE_INT`, but all code should treat it as an opaque quantity.

`int REAL_VALUES_EQUAL (REAL_VALUE_TYPE x, REAL_VALUE_TYPE y)` [Macro]

Compares for equality the two values, `x` and `y`. If the target floating point format supports negative zeroes and/or NaNs, ‘`REAL_VALUES_EQUAL (-0.0, 0.0)`’ is true, and ‘`REAL_VALUES_EQUAL (NaN, NaN)`’ is false.

`int REAL_VALUES_LESS (REAL_VALUE_TYPE x, REAL_VALUE_TYPE y)` [Macro]

Tests whether `x` is less than `y`.

`HOST_WIDE_INT REAL_VALUE_FIX (REAL_VALUE_TYPE x)` [Macro]

Truncates `x` to a signed integer, rounding toward zero.

`unsigned HOST_WIDE_INT REAL_VALUE_UNSIGNED_FIX (REAL_VALUE_TYPE x)` [Macro]

Truncates `x` to an unsigned integer, rounding toward zero. If `x` is negative, returns zero.

- `REAL_VALUE_TYPE REAL_VALUE_ATOF (const char *string, enum machine_mode mode)` [Macro]
 Converts *string* into a floating point number in the target machine's representation for mode *mode*. This routine can handle both decimal and hexadecimal floating point constants, using the syntax defined by the C language for both.
- `int REAL_VALUE_NEGATIVE (REAL_VALUE_TYPE x)` [Macro]
 Returns 1 if *x* is negative (including negative zero), 0 otherwise.
- `int REAL_VALUE_ISINF (REAL_VALUE_TYPE x)` [Macro]
 Determines whether *x* represents infinity (positive or negative).
- `int REAL_VALUE_ISNAN (REAL_VALUE_TYPE x)` [Macro]
 Determines whether *x* represents a "NaN" (not-a-number).
- `void REAL_ARITHMETIC (REAL_VALUE_TYPE output, enum tree_code code, REAL_VALUE_TYPE x, REAL_VALUE_TYPE y)` [Macro]
 Calculates an arithmetic operation on the two floating point values *x* and *y*, storing the result in *output* (which must be a variable).
 The operation to be performed is specified by *code*. Only the following codes are supported: PLUS_EXPR, MINUS_EXPR, MULT_EXPR, RDIV_EXPR, MAX_EXPR, MIN_EXPR.
 If REAL_ARITHMETIC is asked to evaluate division by zero and the target's floating point format cannot represent infinity, it will call `abort`. Callers should check for this situation first, using MODE_HAS_INFINITIES. See Section 10.5 [Storage Layout], page 205.
- `REAL_VALUE_TYPE REAL_VALUE_NEGATE (REAL_VALUE_TYPE x)` [Macro]
 Returns the negative of the floating point value *x*.
- `REAL_VALUE_TYPE REAL_VALUE_ABS (REAL_VALUE_TYPE x)` [Macro]
 Returns the absolute value of *x*.
- `REAL_VALUE_TYPE REAL_VALUE_TRUNCATE (REAL_VALUE_TYPE mode, enum machine_mode x)` [Macro]
 Truncates the floating point value *x* to fit in *mode*. The return value is still a full-size REAL_VALUE_TYPE, but it has an appropriate bit pattern to be output as a floating constant whose precision accords with mode *mode*.
- `void REAL_VALUE_TO_INT (HOST_WIDE_INT low, HOST_WIDE_INT high, REAL_VALUE_TYPE x)` [Macro]
 Converts a floating point value *x* into a double-precision integer which is then stored into *low* and *high*. If the value is not integral, it is truncated.
- `void REAL_VALUE_FROM_INT (REAL_VALUE_TYPE x, HOST_WIDE_INT low, HOST_WIDE_INT high, enum machine_mode mode)` [Macro]
 Converts a double-precision integer found in *low* and *high*, into a floating point value which is then stored into *x*. The value is truncated to fit in mode *mode*.

10.23 Mode Switching Instructions

The following macros control mode switching optimizations:

OPTIMIZE_MODE_SWITCHING (*entity*)

Define this macro if the port needs extra instructions inserted for mode switching in an optimizing compilation.

For an example, the SH4 can perform both single and double precision floating point operations, but to perform a single precision operation, the FPSCR PR bit has to be cleared, while for a double precision operation, this bit has to be set. Changing the PR bit requires a general purpose register as a scratch register, hence these FPSCR sets have to be inserted before reload, i.e. you can't put this into instruction emitting or `MACHINE_DEPENDENT_REORG`.

You can have multiple entities that are mode-switched, and select at run time which entities actually need it. `OPTIMIZE_MODE_SWITCHING` should return nonzero for any *entity* that needs mode-switching. If you define this macro, you also have to define `NUM_MODES_FOR_MODE_SWITCHING`, `MODE_NEEDED`, `MODE_PRIORITY_TO_MODE` and `EMIT_MODE_SET`. `NORMAL_MODE` is optional.

NUM_MODES_FOR_MODE_SWITCHING

If you define `OPTIMIZE_MODE_SWITCHING`, you have to define this as initializer for an array of integers. Each initializer element *N* refers to an entity that needs mode switching, and specifies the number of different modes that might need to be set for this entity. The position of the initializer in the initializer - starting counting at zero - determines the integer that is used to refer to the mode-switched entity in question. In macros that take mode arguments / yield a mode result, modes are represented as numbers $0 \dots N - 1$. *N* is used to specify that no mode switch is needed / supplied.

MODE_NEEDED (*entity*, *insn*)

entity is an integer specifying a mode-switched entity. If `OPTIMIZE_MODE_SWITCHING` is defined, you must define this macro to return an integer value not larger than the corresponding element in `NUM_MODES_FOR_MODE_SWITCHING`, to denote the mode that *entity* must be switched into prior to the execution of *insn*.

NORMAL_MODE (*entity*)

If this macro is defined, it is evaluated for every *entity* that needs mode switching. It should evaluate to an integer, which is a mode that *entity* is assumed to be switched to at function entry and exit.

MODE_PRIORITY_TO_MODE (*entity*, *n*)

This macro specifies the order in which modes for *entity* are processed. 0 is the highest priority, `NUM_MODES_FOR_MODE_SWITCHING[entity] - 1` the lowest. The value of the macro should be an integer designating a mode for *entity*. For any fixed *entity*, `mode_priority_to_mode (entity, n)` shall be a bijection in $0 \dots \text{num_modes_for_mode_switching}[entity] - 1$.

`EMIT_MODE_SET` (*entity*, *mode*, *hard_regs_live*)

Generate one or more insns to set *entity* to *mode*. *hard_reg_live* is the set of hard registers live at the point where the insn(s) are to be inserted.

10.24 Defining target-specific uses of `__attribute__`

Target-specific attributes may be defined for functions, data and types. These are described using the following target hooks; they also need to be documented in ‘`extend.texi`’.

`const struct attribute_spec * TARGET_ATTRIBUTE_TABLE` [Target Hook]

If defined, this target hook points to an array of ‘`struct attribute_spec`’ (defined in ‘`tree.h`’) specifying the machine specific attributes for this target and some of the restrictions on the entities to which these attributes are applied and the arguments they take.

`int TARGET_COMP_TYPE_ATTRIBUTES` (*tree type1*, *tree type2*) [Target Hook]

If defined, this target hook is a function which returns zero if the attributes on *type1* and *type2* are incompatible, one if they are compatible, and two if they are nearly compatible (which causes a warning to be generated). If this is not defined, machine-specific attributes are supposed always to be compatible.

`void TARGET_SET_DEFAULT_TYPE_ATTRIBUTES` (*tree type*) [Target Hook]

If defined, this target hook is a function which assigns default attributes to newly defined *type*.

`tree TARGET_MERGE_TYPE_ATTRIBUTES` (*tree type1*, *tree type2*) [Target Hook]

Define this target hook if the merging of type attributes needs special handling. If defined, the result is a list of the combined `TYPE_ATTRIBUTES` of *type1* and *type2*. It is assumed that `comptypes` has already been called and returned 1. This function may call `merge_attributes` to handle machine-independent merging.

`tree TARGET_MERGE_DECL_ATTRIBUTES` (*tree olddecl*, *tree newdecl*) [Target Hook]

Define this target hook if the merging of decl attributes needs special handling. If defined, the result is a list of the combined `DECL_ATTRIBUTES` of *olddecl* and *newdecl*. *newdecl* is a duplicate declaration of *olddecl*. Examples of when this is needed are when one attribute overrides another, or when an attribute is nullified by a subsequent definition. This function may call `merge_attributes` to handle machine-independent merging.

If the only target-specific handling you require is ‘`dllimport`’ for Windows targets, you should define the macro `TARGET_DLLIMPORT_DECL_ATTRIBUTES`. This links in a function called `merge_dllimport_decl_attributes` which can then be defined as the expansion of `TARGET_MERGE_DECL_ATTRIBUTES`. This is done in ‘`i386/cygwin.h`’ and ‘`i386/i386.c`’, for example.

`void TARGET_INSERT_ATTRIBUTES` (*tree node*, *tree *attr_ptr*) [Target Hook]

Define this target hook if you want to be able to add attributes to a decl when it is being created. This is normally useful for back ends which wish to implement a

pragma by using the attributes which correspond to the pragma's effect. The *node* argument is the decl which is being created. The *attr_ptr* argument is a pointer to the attribute list for this decl. The list itself should not be modified, since it may be shared with other decls, but attributes may be chained on the head of the list and **attr_ptr* modified to point to the new attributes, or a copy of the list may be made if further changes are needed.

bool TARGET_FUNCTION_ATTRIBUTE_INLINABLE_P (*tree fndecl*) [Target Hook]
 This target hook returns **true** if it is ok to inline *fndecl* into the current function, despite its having target-specific attributes, **false** otherwise. By default, if a function has a target specific attribute attached to it, it will not be inlined.

10.25 Defining coprocessor specifics for MIPS targets.

The MIPS specification allows MIPS implementations to have as many as 4 coprocessors, each with as many as 32 private registers. gcc supports accessing these registers and transferring values between the registers and memory using asm-ized variables. For example:

```
register unsigned int cp0count asm ("c0r1");
unsigned int d;

d = cp0count + 3;
```

("c0r1" is the default name of register 1 in coprocessor 0; alternate names may be added as described below, or the default names may be overridden entirely in SUBTARGET_CONDITIONAL_REGISTER_USAGE.)

Coprocessor registers are assumed to be epilogue-used; sets to them will be preserved even if it does not appear that the register is used again later in the function.

Another note: according to the MIPS spec, coprocessor 1 (if present) is the FPU. One accesses COP1 registers through standard mips floating-point support; they are not included in this mechanism.

There is one macro used in defining the MIPS coprocessor interface which you may want to override in subtargets; it is described below.

ALL_COP_ADDITIONAL_REGISTER_NAMES

A comma-separated list (with leading comma) of pairs describing the alternate names of coprocessor registers. The format of each entry should be

```
{ alternatename, register_number}
```

Default: empty.

10.26 Miscellaneous Parameters

Here are several miscellaneous parameters.

PREDICATE_CODES

Define this if you have defined special-purpose predicates in the file '*machine.c*'. This macro is called within an initializer of an array of structures. The first field in the structure is the name of a predicate and the second field is an array of rtl codes. For each predicate, list all rtl codes that can be in expressions

matched by the predicate. The list should have a trailing comma. Here is an example of two entries in the list for a typical RISC machine:

```
#define PREDICATE_CODES \
    {"gen_reg_rtx_operand", {SUBREG, REG}}, \
    {"reg_or_short_cint_operand", {SUBREG, REG, CONST_INT}},
```

Defining this macro does not affect the generated code (however, incorrect definitions that omit an rtl code that may be matched by the predicate can cause the compiler to malfunction). Instead, it allows the table built by ‘genrecog’ to be more compact and efficient, thus speeding up the compiler. The most important predicates to include in the list specified by this macro are those used in the most insn patterns.

For each predicate function named in PREDICATE_CODES, a declaration will be generated in ‘insn-codes.h’.

SPECIAL_MODE_PREDICATES

Define this if you have special predicates that know special things about modes. Genrecog will warn about certain forms of `match_operand` without a mode; if the operand predicate is listed in SPECIAL_MODE_PREDICATES, the warning will be suppressed.

Here is an example from the IA-32 port (`ext_register_operand` specially checks for HImode or SImode in preparation for a byte extraction from %ah etc.).

```
#define SPECIAL_MODE_PREDICATES \
    "ext_register_operand",
```

CASE_VECTOR_MODE

An alias for a machine mode name. This is the machine mode that elements of a jump-table should have.

CASE_VECTOR_SHORTEN_MODE (*min_offset*, *max_offset*, *body*)

Optional: return the preferred mode for an `addr_diff_vec` when the minimum and maximum offset are known. If you define this, it enables extra code in branch shortening to deal with `addr_diff_vec`. To make this work, you also have to define `INSN_ALIGN` and make the alignment for `addr_diff_vec` explicit. The *body* argument is provided so that the `offset_unsigned` and `scale` flags can be updated.

CASE_VECTOR_PC_RELATIVE

Define this macro to be a C expression to indicate when jump-tables should contain relative addresses. If jump-tables never contain relative addresses, then you need not define this macro.

CASE_DROPS_THROUGH

Define this if control falls through a `case` insn when the index value is out of range. This means the specified default-label is actually ignored by the `case` insn proper.

CASE_VALUES_THRESHOLD

Define this to be the smallest number of different values for which it is best to use a jump-table instead of a tree of conditional branches. The default is

four for machines with a `casesi` instruction and five otherwise. This is best for most machines.

WORD_REGISTER_OPERATIONS

Define this macro if operations between registers with integral mode smaller than a word are always performed on the entire register. Most RISC machines have this property and most CISC machines do not.

LOAD_EXTEND_OP (*mode*)

Define this macro to be a C expression indicating when insns that read memory in *mode*, an integral mode narrower than a word, set the bits outside of *mode* to be either the sign-extension or the zero-extension of the data read. Return `SIGN_EXTEND` for values of *mode* for which the insn sign-extends, `ZERO_EXTEND` for which it zero-extends, and `NIL` for other modes.

This macro is not called with *mode* non-integral or with a width greater than or equal to `BITS_PER_WORD`, so you may return any value in this case. Do not define this macro if it would always return `NIL`. On machines where this macro is defined, you will normally define it as the constant `SIGN_EXTEND` or `ZERO_EXTEND`.

SHORT_IMMEDIATES_SIGN_EXTEND

Define this macro if loading short immediate values into registers sign extends.

FIXUNS_TRUNC_LIKE_FIX_TRUNC

Define this macro if the same instructions that convert a floating point number to a signed fixed point number also convert validly to an unsigned one.

MOVE_MAX The maximum number of bytes that a single instruction can move quickly between memory and registers or between two memory locations.

MAX_MOVE_MAX

The maximum number of bytes that a single instruction can move quickly between memory and registers or between two memory locations. If this is undefined, the default is `MOVE_MAX`. Otherwise, it is the constant value that is the largest value that `MOVE_MAX` can have at run-time.

SHIFT_COUNT_TRUNCATED

A C expression that is nonzero if on this machine the number of bits actually used for the count of a shift operation is equal to the number of bits needed to represent the size of the object being shifted. When this macro is nonzero, the compiler will assume that it is safe to omit a sign-extend, zero-extend, and certain bitwise ‘and’ instructions that truncates the count of a shift operation. On machines that have instructions that act on bit-fields at variable positions, which may include ‘bit test’ instructions, a nonzero `SHIFT_COUNT_TRUNCATED` also enables deletion of truncations of the values that serve as arguments to bit-field instructions.

If both types of instructions truncate the count (for shifts) and position (for bit-field operations), or if no variable-position bit-field instructions exist, you should define this macro.

However, on some machines, such as the 80386 and the 680x0, truncation only applies to shift operations and not the (real or pretended) bit-field operations.

Define `SHIFT_COUNT_TRUNCATED` to be zero on such machines. Instead, add patterns to the ‘`md`’ file that include the implied truncation of the shift instructions.

You need not define this macro if it would always have the value of zero.

`TRULY_NOOP_TRUNCATION` (*outprec*, *inprec*)

A C expression which is nonzero if on this machine it is safe to “convert” an integer of *inprec* bits to one of *outprec* bits (where *outprec* is smaller than *inprec*) by merely operating on it as if it had only *outprec* bits.

On many machines, this expression can be 1.

When `TRULY_NOOP_TRUNCATION` returns 1 for a pair of sizes for modes for which `MODES_TIEABLE_P` is 0, suboptimal code can result. If this is the case, making `TRULY_NOOP_TRUNCATION` return 0 in such cases may improve things.

`STORE_FLAG_VALUE`

A C expression describing the value returned by a comparison operator with an integral mode and stored by a store-flag instruction (‘`scond`’) when the condition is true. This description must apply to *all* the ‘`scond`’ patterns and all the comparison operators whose results have a `MODE_INT` mode.

A value of 1 or -1 means that the instruction implementing the comparison operator returns exactly 1 or -1 when the comparison is true and 0 when the comparison is false. Otherwise, the value indicates which bits of the result are guaranteed to be 1 when the comparison is true. This value is interpreted in the mode of the comparison operation, which is given by the mode of the first operand in the ‘`scond`’ pattern. Either the low bit or the sign bit of `STORE_FLAG_VALUE` be on. Presently, only those bits are used by the compiler.

If `STORE_FLAG_VALUE` is neither 1 or -1 , the compiler will generate code that depends only on the specified bits. It can also replace comparison operators with equivalent operations if they cause the required bits to be set, even if the remaining bits are undefined. For example, on a machine whose comparison operators return an `SI` mode value and where `STORE_FLAG_VALUE` is defined as ‘`0x80000000`’, saying that just the sign bit is relevant, the expression

```
(ne:SI (and:SI x (const_int power-of-2)) (const_int 0))
```

can be converted to

```
(ashift:SI x (const_int n))
```

where *n* is the appropriate shift count to move the bit being tested into the sign bit.

There is no way to describe a machine that always sets the low-order bit for a true value, but does not guarantee the value of any other bits, but we do not know of any machine that has such an instruction. If you are trying to port GCC to such a machine, include an instruction to perform a logical-and of the result with 1 in the pattern for the comparison operators and let us know at gcc@gcc.gnu.org.

Often, a machine will have multiple instructions that obtain a value from a comparison (or the condition codes). Here are rules to guide the choice of value for `STORE_FLAG_VALUE`, and hence the instructions to be used:

- Use the shortest sequence that yields a valid definition for `STORE_FLAG_VALUE`. It is more efficient for the compiler to “normalize” the value (convert it to, e.g., 1 or 0) than for the comparison operators to do so because there may be opportunities to combine the normalization with other operations.
- For equal-length sequences, use a value of 1 or -1 , with -1 being slightly preferred on machines with expensive jumps and 1 preferred on other machines.
- As a second choice, choose a value of ‘0x80000001’ if instructions exist that set both the sign and low-order bits but do not define the others.
- Otherwise, use a value of ‘0x80000000’.

Many machines can produce both the value chosen for `STORE_FLAG_VALUE` and its negation in the same number of instructions. On those machines, you should also define a pattern for those cases, e.g., one matching

```
(set A (neg:m (ne:m B C)))
```

Some machines can also perform `and` or `plus` operations on condition code values with less instructions than the corresponding ‘`scond`’ insn followed by `and` or `plus`. On those machines, define the appropriate patterns. Use the names `incsc` and `decsc`, respectively, for the patterns which perform `plus` or `minus` operations on condition code values. See ‘`rs6000.md`’ for some examples. The GNU Superoptimizer can be used to find such instruction sequences on other machines.

You need not define `STORE_FLAG_VALUE` if the machine has no store-flag instructions.

`FLOAT_STORE_FLAG_VALUE` (*mode*)

A C expression that gives a nonzero `REAL_VALUE_TYPE` value that is returned when comparison operators with floating-point results are true. Define this macro on machine that have comparison operations that return floating-point values. If there are no such operations, do not define this macro.

`Pmode`

An alias for the machine mode for pointers. On most machines, define this to be the integer mode corresponding to the width of a hardware pointer; `SImode` on 32-bit machine or `DImode` on 64-bit machines. On some machines you must define this to be one of the partial integer modes, such as `PSImode`.

The width of `Pmode` must be at least as large as the value of `POINTER_SIZE`. If it is not equal, you must define the macro `POINTERS_EXTEND_UNSIGNED` to specify how pointers are extended to `Pmode`.

`FUNCTION_MODE`

An alias for the machine mode used for memory references to functions being called, in `call` RTL expressions. On most machines this should be `QImode`.

`INTEGRATE_THRESHOLD` (*decl*)

A C expression for the maximum number of instructions above which the function *decl* should not be inlined. *decl* is a `FUNCTION_DECL` node.

The default definition of this macro is 64 plus 8 times the number of arguments that the function accepts. Some people think a larger threshold should be used on RISC machines.

STDC_0_IN_SYSTEM_HEADERS

In normal operation, the preprocessor expands `__STDC__` to the constant 1, to signify that GCC conforms to ISO Standard C. On some hosts, like Solaris, the system compiler uses a different convention, where `__STDC__` is normally 0, but is 1 if the user specifies strict conformance to the C Standard.

Defining `STDC_0_IN_SYSTEM_HEADERS` makes GNU CPP follow the host convention when processing system header files, but when processing user files `__STDC__` will always expand to 1.

NO_IMPLICIT_EXTERN_C

Define this macro if the system header files support C++ as well as C. This macro inhibits the usual method of using system header files in C++, which is to pretend that the file's contents are enclosed in `'extern "C" {...}'`.

HANDLE_PRAGMA (*getc*, *ungetc*, *name*)

This macro is no longer supported. You must use `REGISTER_TARGET_PRAGMAS` instead.

REGISTER_TARGET_PRAGMAS (*pfile*)

Define this macro if you want to implement any target-specific pragmas. If defined, it is a C expression which makes a series of calls to `cpp_register_pragma` for each pragma, with *pfile* passed as the first argument to these functions. The macro may also do any setup required for the pragmas.

The primary reason to define this macro is to provide compatibility with other compilers for the same target. In general, we discourage definition of target-specific pragmas for GCC.

If the pragma can be implemented by attributes then you should consider defining the target hook `'TARGET_INSERT_ATTRIBUTES'` as well.

Preprocessor macros that appear on pragma lines are not expanded. All `'#pragma'` directives that do not match any registered pragma are silently ignored, unless the user specifies `'-Wunknown-pragmas'`.

```
void cpp_register_pragma (cpp_reader *pfile, const [Function]
                        char *space, const char *name, void (*callback) (cpp_reader
                        *))
```

Each call to `cpp_register_pragma` establishes one pragma. The *callback* routine will be called when the preprocessor encounters a pragma of the form

```
#pragma [space] name ...
```

space is the case-sensitive namespace of the pragma, or NULL to put the pragma in the global namespace. The callback routine receives *pfile* as its first argument, which can be passed on to `cpplib`'s functions if necessary. You can lex tokens after the *name* by calling `c_lex`. Tokens that are not read by the callback will be silently ignored. The end of the line is indicated by a token of type `CPP_EOF`.

For an example use of this routine, see ‘c4x.h’ and the callback routines defined in ‘c4x-c.c’.

Note that the use of `c_lex` is specific to the C and C++ compilers. It will not work in the Java or Fortran compilers, or any other language compilers for that matter. Thus if `c_lex` is going to be called from target-specific code, it must only be done so when building the C and C++ compilers. This can be done by defining the variables `c_target_objs` and `cxx_target_objs` in the target entry in the ‘config.gcc’ file. These variables should name the target-specific, language-specific object file which contains the code that uses `c_lex`. Note it will also be necessary to add a rule to the makefile fragment pointed to by `tmake_file` that shows how to build this object file.

HANDLE_SYSV_PRAGMA

Define this macro (to a value of 1) if you want the System V style pragmas ‘`#pragma pack(<n>)`’ and ‘`#pragma weak <name> [=<value>]`’ to be supported by gcc.

The pack pragma specifies the maximum alignment (in bytes) of fields within a structure, in much the same way as the ‘`__aligned__`’ and ‘`__packed__`’ `__attribute__`s do. A pack value of zero resets the behavior to the default.

A subtlety for Microsoft Visual C/C++ style bit-field packing (e.g. `-mms-bitfields`) for targets that support it: When a bit-field is inserted into a packed record, the whole size of the underlying type is used by one or more same-size adjacent bit-fields (that is, if its `long:3`, 32 bits is used in the record, and any additional adjacent long bit-fields are packed into the same chunk of 32 bits. However, if the size changes, a new field of that size is allocated).

If both MS bit-fields and ‘`__attribute__((packed))`’ are used, the latter will take precedence. If ‘`__attribute__((packed))`’ is used on a single field when MS bit-fields are in use, it will take precedence for that field, but the alignment of the rest of the structure may affect its placement.

The weak pragma only works if `SUPPORTS_WEAK` and `ASM_WEAKEN_LABEL` are defined. If enabled it allows the creation of specifically named weak labels, optionally with a value.

HANDLE_PRAGMA_PACK_PUSH_POP

Define this macro (to a value of 1) if you want to support the Win32 style pragmas ‘`#pragma pack(push,n)`’ and ‘`#pragma pack(pop)`’. The ‘`pack(push,n)`’ pragma specifies the maximum alignment (in bytes) of fields within a structure, in much the same way as the ‘`__aligned__`’ and ‘`__packed__`’ `__attribute__`s do. A pack value of zero resets the behavior to the default. Successive invocations of this pragma cause the previous values to be stacked, so that invocations of ‘`#pragma pack(pop)`’ will return to the previous value.

DOLLARS_IN_IDENTIFIERS

Define this macro to control use of the character ‘\$’ in identifier names. 0 means ‘\$’ is not allowed by default; 1 means it is allowed. 1 is the default; there is no need to define this macro in that case. This macro controls the compiler proper; it does not affect the preprocessor.

NO_DOLLAR_IN_LABEL

Define this macro if the assembler does not accept the character '\$' in label names. By default constructors and destructors in G++ have '\$' in the identifiers. If this macro is defined, '.' is used instead.

NO_DOT_IN_LABEL

Define this macro if the assembler does not accept the character '.' in label names. By default constructors and destructors in G++ have names that use '.'. If this macro is defined, these names are rewritten to avoid '.'.

DEFAULT_MAIN_RETURN

Define this macro if the target system expects every program's `main` function to return a standard "success" value by default (if no other value is explicitly returned).

The definition should be a C statement (sans semicolon) to generate the appropriate rtl instructions. It is used only when compiling the end of `main`.

NEED_ATEXIT

Define this if the target system lacks the function `atexit` from the ISO C standard. If this macro is defined, a default definition will be provided to support C++. If `ON_EXIT` is not defined, a default `exit` function will also be provided.

ON_EXIT

Define this macro if the target has another way to implement `atexit` functionality without replacing `exit`. For instance, SunOS 4 has a similar `on_exit` library function.

The definition should be a functional macro which can be used just like the `atexit` function.

EXIT_BODY

Define this if your `exit` function needs to do something besides calling an external function `_cleanup` before terminating with `_exit`. The `EXIT_BODY` macro is only needed if `NEED_ATEXIT` is defined and `ON_EXIT` is not defined.

INSN_SETS_ARE_DELAYED (*insn*)

Define this macro as a C expression that is nonzero if it is safe for the delay slot scheduler to place instructions in the delay slot of *insn*, even if they appear to use a resource set or clobbered in *insn*. *insn* is always a `jump_insn` or an `insn`; GCC knows that every `call_insn` has this behavior. On machines where some `insn` or `jump_insn` is really a function call and hence has this behavior, you should define this macro.

You need not define this macro if it would always return zero.

INSN_REFERENCES_ARE_DELAYED (*insn*)

Define this macro as a C expression that is nonzero if it is safe for the delay slot scheduler to place instructions in the delay slot of *insn*, even if they appear to set or clobber a resource referenced in *insn*. *insn* is always a `jump_insn` or an `insn`. On machines where some `insn` or `jump_insn` is really a function call and its operands are registers whose use is actually in the subroutine it calls, you should define this macro. Doing so allows the delay slot scheduler to move

instructions which copy arguments into the argument registers into the delay slot of *insn*.

You need not define this macro if it would always return zero.

MACHINE_DEPENDENT_REORG (*insn*)

In rare cases, correct code generation requires extra machine dependent processing between the second jump optimization pass and delayed branch scheduling. On those machines, define this macro as a C statement to act on the code starting at *insn*.

MULTIPLE_SYMBOL_SPACES

Define this macro if in some cases global symbols from one translation unit may not be bound to undefined symbols in another translation unit without user intervention. For instance, under Microsoft Windows symbols must be explicitly imported from shared libraries (DLLs).

MD_ASM_CLOBBERS (*clobbers*)

A C statement that adds to *clobbers* `STRING_CST` trees for any hard regs the port wishes to automatically clobber for all asms.

MAX_INTEGER_COMPUTATION_MODE

Define this to the largest integer machine mode which can be used for operations other than load, store and copy operations.

You need only define this macro if the target holds values larger than `word_mode` in general purpose registers. Most targets should not define this macro.

MATH_LIBRARY

Define this macro as a C string constant for the linker argument to link in the system math library, or "" if the target does not have a separate math library.

You need only define this macro if the default of `"-lm"` is wrong.

LIBRARY_PATH_ENV

Define this macro as a C string constant for the environment variable that specifies where the linker should look for libraries.

You need only define this macro if the default of `"LIBRARY_PATH"` is wrong.

TARGET_HAS_F_SETLKW

Define this macro if the target supports file locking with `fcntl / F_SETLKW`. Note that this functionality is part of POSIX. Defining `TARGET_HAS_F_SETLKW` will enable the test coverage code to use file locking when exiting a program, which avoids race conditions if the program has forked.

MAX_CONDITIONAL_EXECUTE

A C expression for the maximum number of instructions to execute via conditional execution instructions instead of a branch. A value of `BRANCH_COST+1` is the default if the machine does not use `cc0`, and 1 if it does use `cc0`.

IFCVT_MODIFY_TESTS(*ce_info*, *true_expr*, *false_expr*)

Used if the target needs to perform machine-dependent modifications on the conditionals used for turning basic blocks into conditionally executed code. *ce_info* points to a data structure, `struct ce_if_block`, which contains information about the currently processed blocks. *true_expr* and *false_expr* are the

tests that are used for converting the then-block and the else-block, respectively. Set either *true_expr* or *false_expr* to a null pointer if the tests cannot be converted.

`IFCVT_MODIFY_MULTIPLE_TESTS(ce_info, bb, true_expr, false_expr)`

Like `IFCVT_MODIFY_TESTS`, but used when converting more complicated if-statements into conditions combined by `and` and `or` operations. *bb* contains the basic block that contains the test that is currently being processed and about to be turned into a condition.

`IFCVT_MODIFY_INSN(ce_info, pattern, insn)`

A C expression to modify the *PATTERN* of an *INSN* that is to be converted to conditional execution format. *ce_info* points to a data structure, `struct ce_if_block`, which contains information about the currently processed blocks.

`IFCVT_MODIFY_FINAL(ce_info)`

A C expression to perform any final machine dependent modifications in converting code to conditional execution. The involved basic blocks can be found in the `struct ce_if_block` structure that is pointed to by *ce_info*.

`IFCVT_MODIFY_CANCEL(ce_info)`

A C expression to cancel any machine dependent modifications in converting code to conditional execution. The involved basic blocks can be found in the `struct ce_if_block` structure that is pointed to by *ce_info*.

`IFCVT_INIT_EXTRA_FIELDS(ce_info)`

A C expression to initialize any extra fields in a `struct ce_if_block` structure, which are defined by the `IFCVT_EXTRA_FIELDS` macro.

`IFCVT_EXTRA_FIELDS`

If defined, it should expand to a set of field declarations that will be added to the `struct ce_if_block` structure. These should be initialized by the `IFCVT_INIT_EXTRA_FIELDS` macro.

`void TARGET_INIT_BUILTINS ()` [Target Hook]

Define this hook if you have any machine-specific built-in functions that need to be defined. It should be a function that performs the necessary setup.

Machine specific built-in functions can be useful to expand special machine instructions that would otherwise not normally be generated because they have no equivalent in the source language (for example, SIMD vector instructions or prefetch instructions).

To create a built-in function, call the function `builtin_function` which is defined by the language front end. You can use any type nodes set up by `build_common_tree_nodes` and `build_common_tree_nodes_2`; only language front ends that use those two functions will call ‘`TARGET_INIT_BUILTINS`’.

`rtx TARGET_EXPAND_BUILTIN (tree exp, rtx target, rtx subtarget, enum machine_mode mode, int ignore)` [Target Hook]

Expand a call to a machine specific built-in function that was set up by ‘`TARGET_INIT_BUILTINS`’. *exp* is the expression for the function call; the result

should go to *target* if that is convenient, and have mode *mode* if that is convenient. *subtarget* may be used as the target for computing one of *exp*'s operands. *ignore* is nonzero if the value is to be ignored. This function should return the result of the call to the built-in function.

MD_CAN_REDIRECT_BRANCH(*branch1*, *branch2*)

Take a branch insn in *branch1* and another in *branch2*. Return true if redirecting *branch1* to the destination of *branch2* is possible.

On some targets, branches may have a limited range. Optimizing the filling of delay slots can result in branches being redirected, and this may in turn cause a branch offset to overflow.

ALLOCATE_INITIAL_VALUE(*hard_reg*)

When the initial value of a hard register has been copied in a pseudo register, it is often not necessary to actually allocate another register to this pseudo register, because the original hard register or a stack slot it has been saved into can be used. `ALLOCATE_INITIAL_VALUE`, if defined, is called at the start of register allocation once for each hard register that had its initial value copied by using `get_func_hard_reg_initial_val` or `get_hard_reg_initial_val`. Possible values are `NULL_RTX`, if you don't want to do any special allocation, a `REG` rtx—that would typically be the hard register itself, if it is known not to be clobbered—or a `MEM`. If you are returning a `MEM`, this is only a hint for the allocator; it might decide to use another register anyways. You may use `current_function_leaf_function` in the definition of the macro, functions that use `REG_N_SETS`, to determine if the hard register in question will not be clobbered.

TARGET_OBJECT_SUFFIX

Define this macro to be a C string representing the suffix for object files on your target machine. If you do not define this macro, GCC will use `'.o'` as the suffix for object files.

TARGET_EXECUTABLE_SUFFIX

Define this macro to be a C string representing the suffix to be automatically added to executable files on your target machine. If you do not define this macro, GCC will use the null string as the suffix for executable files.

COLLECT_EXPORT_LIST

If defined, `collect2` will scan the individual object files specified on its command line and create an export list for the linker. Define this macro for systems like AIX, where the linker discards object files that are not referenced from `main` and uses export lists.

MODIFY_JNI_METHOD_CALL (*mdecl*)

Define this macro to a C expression representing a variant of the method call *mdecl*, if Java Native Interface (JNI) methods must be invoked differently from other methods on your target. For example, on 32-bit Windows, JNI methods must be invoked using the `stdcall` calling convention and this macro is then defined as this expression:

```
build_type_attribute_variant (mdecl,  
                              build_tree_list  
                              (get_identifier ("stdcall"),  
                               NULL))
```

`bool TARGET_CANNOT_MODIFY_JUMPS_P (void)` [Target Hook]

This target hook returns `true` past the point in which new jump instructions could be created. On machines that require a register for every jump such as the SHmedia ISA of SH5, this point would typically be `reload`, so this target hook should be defined to a function such as:

```
static bool  
cannot_modify_jumps_past_reload_p ()  
{  
    return (reload_completed || reload_in_progress);  
}
```


11 Host Configuration Headers

Host configuration headers contain macro definitions that describe the machine and system on which the compiler is running. They are usually unnecessary. Most of the things GCC needs to know about the host system can be deduced by the `configure` script.

If your host does need a special configuration header, it should be named `'xm-machine.h'`, where *machine* is a short mnemonic for the machine. Here are some macros which this header can define.

VMS Define this macro if the host system is VMS.

FATAL_EXIT_CODE

A C expression for the status code to be returned when the compiler exits after serious errors. The default is the system-provided macro `'EXIT_FAILURE'`, or `'1'` if the system doesn't define that macro. Define this macro only if these defaults are incorrect.

SUCCESS_EXIT_CODE

A C expression for the status code to be returned when the compiler exits without serious errors. (Warnings are not serious errors.) The default is the system-provided macro `'EXIT_SUCCESS'`, or `'0'` if the system doesn't define that macro. Define this macro only if these defaults are incorrect.

USE_C_ALLOCA

Define this macro if GCC should use the C implementation of `alloca` provided by `'libiberty.a'`. This only affects how some parts of the compiler itself allocate memory. It does not change code generation.

When GCC is built with a compiler other than itself, the C `alloca` is always used. This is because most other implementations have serious bugs. You should define this macro only on a system where no stack-based `alloca` can possibly work. For instance, if a system has a small limit on the size of the stack, GCC's builtin `alloca` will not work reliably.

HAVE_DOS_BASED_FILE_SYSTEM

Define this macro if the host file system obeys the semantics defined by MS-DOS instead of Unix. DOS file systems are case insensitive, file specifications may begin with a drive letter, and both forward slash and backslash (`'/'` and `'\'`) are directory separators. If you define this macro, you probably need to define the next three macros too.

PATH_SEPARATOR

If defined, this macro should expand to a character constant specifying the separator for elements of search paths. The default value is a colon (`':'`). DOS-based systems usually use semicolon (`';'`).

DIR_SEPARATOR

DIR_SEPARATOR_2

If defined, these macros expand to character constants specifying separators for directory names within a file specification. They are used somewhat inconsistently throughout the compiler. If your system behaves like Unix (only forward

slash separates pathnames), define neither of them. If your system behaves like DOS (both forward and backward slash can be used), define `DIR_SEPARATOR` to `'/'` and `DIR_SEPARATOR_2` to `'\'`.

HOST_OBJECT_SUFFIX

Define this macro to be a C string representing the suffix for object files on your host machine. If you do not define this macro, GCC will use `'.o'` as the suffix for object files.

HOST_EXECUTABLE_SUFFIX

Define this macro to be a C string representing the suffix for executable files on your host machine. If you do not define this macro, GCC will use the null string as the suffix for executable files.

HOST_BIT_BUCKET

A pathname defined by the host operating system, which can be opened as a file and written to, but all the information written is discarded. This is commonly known as a *bit bucket* or *null device*. If you do not define this macro, GCC will use `'/dev/null'` as the bit bucket. If the host does not support a bit bucket, define this macro to an invalid filename.

COLLECT2_HOST_INITIALIZATION

If defined, a C statement (sans semicolon) that performs host-dependent initialization when `collect2` is being initialized.

GCC_DRIVER_HOST_INITIALIZATION

If defined, a C statement (sans semicolon) that performs host-dependent initialization when a compilation driver is being initialized.

UPDATE_PATH_HOST_CANONICALIZE (*path*)

If defined, a C statement (sans semicolon) that performs host-dependent canonicalization when a path used in a compilation driver or preprocessor is canonicalized. *path* is a malloc-ed path to be canonicalized. If the C statement does canonicalize *path* into a different buffer, the old path should be freed and the new buffer should have been allocated with malloc.

DUMPFILFORMAT

Define this macro to be a C string representing the format to use for constructing the index part of debugging dump file names. The resultant string must fit in fifteen bytes. The full filename will be the concatenation of: the prefix of the assembler file name, the string resulting from applying this format to an index number, and a string unique to each dump file kind, e.g. `'rtl'`.

If you do not define this macro, GCC will use `'.%02d.'`. You should define this macro if using the default will create an invalid file name.

SMALL_ARG_MAX

Define this macro if the host system has a small limit on the total size of an argument vector. This causes the driver to take more care not to pass unnecessary arguments to subprocesses.

In addition, if `configure` generates an incorrect definition of any of the macros in `'auto-host.h'`, you can override that definition in a host configuration header. If you need to do this, first see if it is possible to fix `configure`.

If you need to define only a few of these macros, and they have simple definitions, consider using the `xm_defines` variable in your `'config.gcc'` entry instead of creating a host configuration header. See Section 5.3.2.2 [System Config], page 14.

12 Makefile Fragments

When you configure GCC using the ‘`configure`’ script, it will construct the file ‘`Makefile`’ from the template file ‘`Makefile.in`’. When it does this, it can incorporate makefile fragments from the ‘`config`’ directory. These are used to set Makefile parameters that are not amenable to being calculated by `autoconf`. The list of fragments to incorporate is set by ‘`config.gcc`’; See Section 5.3.2.2 [System Config], page 14.

Fragments are named either ‘`t-target`’ or ‘`x-host`’, depending on whether they are relevant to configuring GCC to produce code for a particular target, or to configuring GCC to run on a particular host. Here *target* and *host* are mnemonics which usually have some relationship to the canonical system name, but no formal connection.

If these files do not exist, it means nothing needs to be added for a given target or host. Most targets need a few ‘`t-target`’ fragments, but needing ‘`x-host`’ fragments is rare.

12.1 Target Makefile Fragments

Target makefile fragments can set these Makefile variables.

LIBGCC2_CFLAGS

Compiler flags to use when compiling ‘`libgcc2.c`’.

LIB2FUNCS_EXTRA

A list of source file names to be compiled or assembled and inserted into ‘`libgcc.a`’.

Floating Point Emulation

To have GCC include software floating point libraries in ‘`libgcc.a`’ define `FPBIT` and `DPBIT` along with a few rules as follows:

```
# We want fine grained libraries, so use the new code
# to build the floating point emulation libraries.
FPBIT = fp-bit.c
DPBIT = dp-bit.c
```

```
fp-bit.c: $(srcdir)/config/fp-bit.c
    echo '#define FLOAT' > fp-bit.c
    cat $(srcdir)/config/fp-bit.c >> fp-bit.c
```

```
dp-bit.c: $(srcdir)/config/fp-bit.c
    cat $(srcdir)/config/fp-bit.c > dp-bit.c
```

You may need to provide additional `#defines` at the beginning of ‘`fp-bit.c`’ and ‘`dp-bit.c`’ to control target endianness and other options.

CRTSTUFF_T_CFLAGS

Special flags used when compiling ‘`crtstuff.c`’. See Section 10.20.5 [Initialization], page 292.

CRTSTUFF_T_CFLAGS_S

Special flags used when compiling ‘`crtstuff.c`’ for shared linking. Used if you use ‘`crtbeginS.o`’ and ‘`crtendS.o`’ in `EXTRA-PARTS`. See Section 10.20.5 [Initialization], page 292.

MULTILIB_OPTIONS

For some targets, invoking GCC in different ways produces objects that can not be linked together. For example, for some targets GCC produces both big and little endian code. For these targets, you must arrange for multiple versions of `'libgcc.a'` to be compiled, one for each set of incompatible options. When GCC invokes the linker, it arranges to link in the right version of `'libgcc.a'`, based on the command line options used.

The `MULTILIB_OPTIONS` macro lists the set of options for which special versions of `'libgcc.a'` must be built. Write options that are mutually incompatible side by side, separated by a slash. Write options that may be used together separated by a space. The build procedure will build all combinations of compatible options.

For example, if you set `MULTILIB_OPTIONS` to `'m68000/m68020 msoft-float'`, `'Makefile'` will build special versions of `'libgcc.a'` using the following sets of options: `'-m68000'`, `'-m68020'`, `'-msoft-float'`, `'-m68000 -msoft-float'`, and `'-m68020 -msoft-float'`.

MULTILIB_DIRNAMES

If `MULTILIB_OPTIONS` is used, this variable specifies the directory names that should be used to hold the various libraries. Write one element in `MULTILIB_DIRNAMES` for each element in `MULTILIB_OPTIONS`. If `MULTILIB_DIRNAMES` is not used, the default value will be `MULTILIB_OPTIONS`, with all slashes treated as spaces.

For example, if `MULTILIB_OPTIONS` is set to `'m68000/m68020 msoft-float'`, then the default value of `MULTILIB_DIRNAMES` is `'m68000 m68020 msoft-float'`. You may specify a different value if you desire a different set of directory names.

MULTILIB_MATCHES

Sometimes the same option may be written in two different ways. If an option is listed in `MULTILIB_OPTIONS`, GCC needs to know about any synonyms. In that case, set `MULTILIB_MATCHES` to a list of items of the form `'option=option'` to describe all relevant synonyms. For example, `'m68000=mc68000 m68020=mc68020'`.

MULTILIB_EXCEPTIONS

Sometimes when there are multiple sets of `MULTILIB_OPTIONS` being specified, there are combinations that should not be built. In that case, set `MULTILIB_EXCEPTIONS` to be all of the switch exceptions in shell case syntax that should not be built.

For example the ARM processor cannot execute both hardware floating point instructions and the reduced size THUMB instructions at the same time, so there is no need to build libraries with both of these options enabled. Therefore `MULTILIB_EXCEPTIONS` is set to:

```
*mthumb/*mhard-float*
```

MULTILIB_EXTRA_OPTS

Sometimes it is desirable that when building multiple versions of `'libgcc.a'` certain options should always be passed on to the compiler. In that case, set `MULTILIB_EXTRA_OPTS` to be the list of options to be used for all builds. If you set this, you should probably set `CRTSTUFF_T_CFLAGS` to a dash followed by it.

SPECS Unfortunately, setting `MULTILIB_EXTRA_OPTS` is not enough, since it does not affect the build of target libraries, at least not the build of the default multilib. One possible work-around is to use `DRIVER_SELF_SPECS` to bring options from the `'specs'` file as if they had been passed in the compiler driver command line. However, you don't want to be adding these options after the toolchain is installed, so you can instead tweak the `'specs'` file that will be used during the toolchain build, while you still install the original, built-in `'specs'`. The trick is to set `SPECS` to some other filename (say `'specs.install'`), that will then be created out of the built-in specs, and introduce a `'Makefile'` rule to generate the `'specs'` file that's going to be used at build time out of your `'specs.install'`.

12.2 Host Makefile Fragments

The use of `'x-host'` fragments is discouraged. You should do so only if there is no other mechanism to get the behavior desired. Host fragments should never forcibly override variables set by the configure script, as they may have been adjusted by the user.

Variables provided for host fragments to set include:

X_CFLAGS

X_CPPFLAGS

These are extra flags to pass to the C compiler and preprocessor, respectively. They are used both when building GCC, and when compiling things with the just-built GCC.

XCFLAGS These are extra flags to use when building the compiler. They are not used when compiling `'libgcc.a'`. However, they *are* used when recompiling the compiler with itself in later stages of a bootstrap.

BOOT_LDFLAGS

Flags to be passed to the linker when recompiling the compiler with itself in later stages of a bootstrap. You might need to use this if, for instance, one of the front ends needs more text space than the linker provides by default.

EXTRA_PROGRAMS

A list of additional programs required to use the compiler on this host, which should be compiled with GCC and installed alongside the front ends. If you set this variable, you must also provide rules to build the extra programs.

13 `collect2`

GCC uses a utility called `collect2` on nearly all systems to arrange to call various initialization 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 GNU CC). 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 file names:

- `'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 passes of the compiler. This includes directories that you specify with `-B`.

Cross-compilers search a little differently:

- `'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 file name 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 above for `ld`.

14 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 GCC, nothing special need be done).

`GPLUSPLUS_INCLUDE_DIR` means the same thing for native and cross. It is where `g++` looks first for header files. The C++ library installs only target independent header files in that directory.

`LOCAL_INCLUDE_DIR` is used only by native compilers. GCC doesn't install anything there. It is normally `'/usr/local/include'`. This is where local additions to a packaged system should place header files.

`CROSS_INCLUDE_DIR` is used only by cross compilers. 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.

15 Memory Management and Type Information

GCC uses some fairly sophisticated memory management techniques, which involve determining information about GCC's data structures from GCC's source code and using this information to perform garbage collection.

A full C parser would be too overcomplicated for this task, so a limited subset of C is interpreted and special markers are used to determine what parts of the source to look at. The parser can also detect simple typedefs of the form `typedef struct ID1 *ID2;` and `typedef int ID3;`, and these don't need to be specially marked.

The two forms that do need to be marked are:

```
struct ID1 GTY(([options]))
{
    [fields]
};

typedef struct ID2 GTY(([options]))
{
    [fields]
} ID3;
```

15.1 The Inside of a GTY(())

Sometimes the C code is not enough to fully describe the type structure. Extra information can be provided by using more GTY markers. These markers can be placed:

- In a structure definition, before the open brace;
- In a global variable declaration, after the keyword `static` or `extern`; and
- In a structure field definition, before the name of the field.

The format of a marker is

```
GTY (([name] ([param]), [name] ([param]) ...))
```

The parameter is either a string or a type name.

When the parameter is a string, often it is a fragment of C code. Three special escapes may be available:

<code>%h</code>	This expands to an expression that evaluates to the current structure.
<code>%1</code>	This expands to an expression that evaluates to the structure that immediately contains the current structure.
<code>%0</code>	This expands to an expression that evaluates to the outermost structure that contains the current structure.
<code>%a</code>	This expands to the string of the form <code>[i1][i2]...</code> that indexes the array item currently being marked. For instance, if the field being marked is <code>foo</code> , then <code>%1.foo%a</code> is the same as <code>%h</code> .

The available options are:

length

There are two places the type machinery will need to be explicitly told the length of an array. The first case is when a structure ends in a variable-length array, like this:

```
struct rtvec_def GTY(()) {
    int num_elem;          /* number of elements */
    rtx GTY ((length ("%h.num_elem"))) elem[1];
};
```

In this case, the `length` option is used to override the specified array length (which should usually be 1). The parameter of the option is a fragment of C code that calculates the length.

The second case is when a structure or a global variable contains a pointer to an array, like this:

```
tree *
    GTY ((length ("%h.regno_pointer_align_length"))) regno_decl;
```

In this case, `regno_decl` has been allocated by writing something like

```
x->regno_decl =
    gcc_alloc (x->regno_pointer_align_length * sizeof (tree));
```

and the `length` provides the length of the field.

This second use of `length` also works on global variables, like:

```
static GTY((length ("reg_base_value_size")))
    rtx *reg_base_value;
```

skip

If `skip` is applied to a field, the type machinery will ignore it. This is somewhat dangerous; the only safe use is in a union when one field really isn't ever used.

desc**tag****default**

The type machinery needs to be told which field of a union is currently active. This is done by giving each field a constant `tag` value, and then specifying a discriminator using `desc`. For example,

```
struct tree_binding GTY(())
{
    struct tree_common common;
    union tree_binding_u {
        tree GTY ((tag ("0"))) scope;
        struct cp_binding_level * GTY ((tag ("1"))) level;
    } GTY ((desc ("BINDING_HAS_LEVEL_P ((tree)&%0)"))) scope;
    tree value;
};
```

In the `desc` option, the “current structure” is the union that it discriminates. Use `%1` to mean the structure containing it. (There are no escapes available to the `tag` option, since it's supposed to be a constant.)

Each `tag` should be different. If no `tag` is matched, the field marked with `default` is used if there is one, otherwise no field in the union will be marked.

`param_is`
`use_param`

Sometimes it's convenient to define some data structure to work on generic pointers (that is, PTR) and then use it with a specific type. `param_is` specifies the real type pointed to, and `use_param` says where in the generic data structure that type should be put.

For instance, to have a `htab_t` that points to trees, one should write

```
htab_t GTY ((param_is (union tree_node))) ict;
```

`paramn_is`
`use_paramn`

In more complicated cases, the data structure might need to work on several different types, which might not necessarily all be pointers. For this, `param1_is` through `param9_is` may be used to specify the real type of a field identified by `use_param1` through `use_param9`.

`use_params`

When a structure contains another structure that is parameterized, there's no need to do anything special, the inner structure inherits the parameters of the outer one. When a structure contains a pointer to a parameterized structure, the type machinery won't automatically detect this (it could, it just doesn't yet), so it's necessary to tell it that the pointed-to structure should use the same parameters as the outer structure. This is done by marking the pointer with the `use_params` option.

`deletable`

`deletable`, when applied to a global variable, indicates that when garbage collection runs, there's no need to mark anything pointed to by this variable, it can just be set to NULL instead. This is used to keep a list of free structures around for re-use.

`if_marked`

Suppose you want some kinds of object to be unique, and so you put them in a hash table. If garbage collection marks the hash table, these objects will never be freed, even if the last other reference to them goes away. GGC has special handling to deal with this: if you use the `if_marked` option on a global hash table, GGC will call the routine whose name is the parameter to the option on each hash table entry. If the routine returns nonzero, the hash table entry will be marked as usual. If the routine returns zero, the hash table entry will be deleted.

The routine `ggc_marked_p` can be used to determine if an element has been marked already; in fact, the usual case is to use `if_marked ("ggc_marked_p")`.

`maybe_undef`

When applied to a field, `maybe_undef` indicates that it's OK if the structure that this field points to is never defined, so long as this field is always NULL. This is used to avoid requiring backends to define certain optional structures. It doesn't work with language frontends.

`special`

The `special` option is used for those bizarre cases that are just too hard to deal with otherwise. Don't use it for new code.

15.2 Marking Roots for the Garbage Collector

In addition to keeping track of types, the type machinery also locates the global variables that the garbage collector starts at. There are two syntaxes it accepts to indicate a root:

1. `extern GTY ([[options]]) [type] ID;`
2. `static GTY ([[options]]) [type] ID;`

These are the only syntaxes that are accepted. In particular, if you want to mark a variable that is only declared as

```
int ID;
```

or similar, you should either make it `static` or you should create a `extern` declaration in a header file somewhere.

15.3 Source Files Containing Type Information

Whenever you add `GTY` markers to a new source file, there are three things you need to do:

1. You need to add the file to the list of source files the type machinery scans. There are three cases:
 - a. For a back-end file, this is usually done automatically; if not, you should add it to `target_gtfiles` in the appropriate port's entries in `'config.gcc'`.
 - b. For files shared by all front ends, this is done by adding the filename to the `GTFILES` variable in `'Makefile.in'`.
 - c. For any other file used by a front end, this is done by adding the filename to the `gtfiles` variable defined in `'config-lang.in'`. For C, the file is `'c-config-lang.in'`. This list should include all files that have `GTY` macros in them that are used in that front end, other than those defined in the previous list items. For example, it is common for front end writers to use `'c-common.c'` and other files from the C front end, and these should be included in the `'gtfiles'` variable for such front ends.
2. If the file was a header file, you'll need to check that it's included in the right place to be visible to the generated files. For a back-end header file, this should be done automatically. For a front-end header file, it needs to be included by the same file that includes `'gtype-lang.h'`. For other header files, it needs to be included in `'gtype-desc.c'`, which is a generated file, so add it to `ifiles` in `open_base_file` in `'gengtype.c'`.

For source files that aren't header files, the machinery will generate a header file that should be included in the source file you just changed. The file will be called `'gt-path.h'` where `path` is the pathname relative to the `'gcc'` directory with slashes replaced by `-`, so for example the header file to be included in `'objc/objc-parse.c'` is called `'gt-objc-objc-parse.c'`. The generated header file should be included after everything else in the source file. Don't forget to mention this file as a dependency in the `'Makefile'`!

3. If a new `gt-path.h` file is needed, you need to arrange to add a `Makefile` rule that will ensure this file can be built. This is done by making it a dependency of `s-gtype`, like this:

```
gt-path.h : s-gtype ; @true
```

For language frontends, there is another file that needs to be included somewhere. It will be called `gtype-lang.h`, where *lang* is the name of the subdirectory the language is contained in. It will need `Makefile` rules just like the other generated files.

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 Compiler Collection 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.

The GNU Project and GNU/Linux

The GNU Project was launched in 1984 to develop a complete Unix-like operating system which is free software: the GNU system. (GNU is a recursive acronym for “GNU’s Not Unix”; it is pronounced “guh-NEW”.) Variants of the GNU operating system, which use the kernel Linux, are now widely used; though these systems are often referred to as “Linux”, they are more accurately called GNU/Linux systems.

For more information, see:

<http://www.gnu.org/>

<http://www.gnu.org/gnu/linux-and-gnu.html>

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:
 - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b. 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.
 - c. 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 to give the program's name and a brief idea of what it does.
Copyright (C) year 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 this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type 'show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type 'show c' for details.
```

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ‘show w’ and ‘show c’; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; 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.

GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 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.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none. The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and

that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements.”

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called

an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document 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.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation 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. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C) year your name.  
Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.2  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled ‘‘GNU  
Free Documentation License’’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with  
the Front-Cover Texts being list, and with the Back-Cover Texts  
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Contributors to GCC

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@redhat.com or gerald@pfeifer.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.
- John David Anglin for threading-related fixes and improvements to libstdc++-v3, and the HP-UX port.
- James van Artsdalen wrote the code that makes efficient use of the Intel 80387 register stack.
- Alasdair Baird for various bug fixes.
- Giovanni Bajo for analyzing lots of complicated C++ problem reports.
- Gerald Baumgartner added the signature extension to the C++ front end.
- Godmar Back for his Java improvements and encouragement.
- Scott Bambrough for help porting the Java compiler.
- Wolfgang Bangerth for processing tons of bug reports.
- Jon Beniston for his Windows port of Java.
- Daniel Berlin for better DWARF2 support, faster/better optimizations, improved alias analysis, plus migrating us to Bugzilla.
- Geoff Berry for his Java object serialization work and various patches.
- Eric Blake for helping to make GCJ and libgcj conform to the specifications.
- Segher Boessenkool for various fixes.
- Hans-J. Boehm for his garbage collector, IA-64 libffi port, and other Java work.
- Neil Booth for work on cpplib, lang hooks, debug hooks and other miscellaneous clean-ups.
- Eric Botcazou for fixing middle- and backend bugs left and right.
- Per Bothner for his direction via the steering committee and various improvements to our infrastructure for supporting new languages. Chill front end implementation. Initial implementations of cpplib, fix-header, config.guess, libio, and past C++ library (libg++) maintainer. Dreaming up, designing and implementing much of GCJ.
- Devon Bowen helped port GCC to the Tahoe.
- Don Bowman for mips-vxworks contributions.
- Dave Brolley for work on cpplib and Chill.
- Robert Brown implemented the support for Encore 32000 systems.
- Christian Bruel for improvements to local store elimination.
- Herman A.J. ten Brugge for various fixes.
- Joerg Brunsmann for Java compiler hacking and help with the GCJ FAQ.
- Joe Buck for his direction via the steering committee.
- Craig Burley for leadership of the Fortran effort.

- Stephan Buys for contributing Doxygen notes for libstdc++.
- Paolo Carlini for libstdc++ work: lots of efficiency improvements to the string class, hard detective work on the frustrating localization issues, and keeping up with the problem reports.
- John Carr for his alias work, SPARC hacking, infrastructure improvements, previous contributions to the steering committee, loop optimizations, etc.
- Stephane Carrez for 68HC11 and 68HC12 ports.
- Steve Chamberlain for support for the Renesas SH and H8 processors and the PicoJava processor, and for GCJ config fixes.
- Glenn Chambers for help with the GCJ FAQ.
- John-Marc Chandonia for various libgcj patches.
- Scott Christley for his Objective-C contributions.
- Eric Christopher for his Java porting help and clean-ups.
- Branko Cibej for more warning contributions.
- The GNU Classpath project for all of their merged runtime code.
- Nick Clifton for arm, mcore, fr30, v850, m32r work, ‘--help’, and other random hacking.
- Michael Cook for libstdc++ cleanup patches to reduce warnings.
- Ralf Corsepius 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.
- Russell Davidson for fstream and stringstream fixes in libstdc++.
- Mo DeJong for GCJ and libgcj bug fixes.
- DJ Delorie for the DJGPP port, build and libiberty maintenance, and various bug fixes.
- Gabriel Dos Reis for contributions to g++, contributions and maintenance of GCC diagnostics infrastructure, libstdc++-v3, including valarray<>, complex<>, maintaining the numerics library (including that pesky <limits> :-)) and keeping up-to-date anything to do with numbers.
- Ulrich Drepper for his work on glibc, testing of GCC using glibc, ISO C99 support, CFG dumping support, etc., plus support of the C++ runtime libraries including for all kinds of C interface issues, contributing and maintaining complex<>, sanity checking and disbursement, configuration architecture, libio maintenance, and early math work.
- Zdenek Dvorak for a new loop unroller and various fixes.
- 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, help cleaning up Haifa loop changes, doing the entire AIX port of libstdc++ with his bare hands, and for ensuring GCC properly keeps working on AIX.

- Kevin Ediger for the floating point formatting of `num_put::do_put` in `libstdc++`.
- Phil Edwards for `libstdc++` work including configuration hackery, documentation maintainer, chief breaker of the web pages, the occasional `iostream` bug fix, and work on shared library symbol versioning.
- Paul Eggert for random hacking all over GCC.
- Mark Elbrecht for various DJGPP improvements, and for `libstdc++` configuration support for locales and `fstream`-related fixes.
- Vadim Egorov for `libstdc++` fixes in strings, `streambufs`, and `iostreams`.
- Christian Ehrhardt for dealing with bug reports.
- Ben Elliston for his work to move the Objective-C runtime into its own subdirectory and for his work on `autoconf`.
- Marc Espie for OpenBSD support.
- Doug Evans for much of the global optimization framework, `arc`, `m32r`, and SPARC work.
- Christopher Faylor for his work on the Cygwin port and for caring and feeding the `gcc.gnu.org` box and saving its users tons of spam.
- Fred Fish for BeOS support and Ada fixes.
- Ivan Fontes Garcia for the Portugese translation of the GCJ FAQ.
- Peter Gerwinski for various bug fixes and the Pascal front end.
- Kaveh Ghazi for his direction via the steering committee, amazing work to make ‘`-Wall`’ useful, and continuously testing GCC on a plethora of platforms.
- John Gilmore for a donation to the FSF earmarked improving GNU Java.
- 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, improved leaf function register allocation, and his direction via the steering committee.
- Anthony Green for his ‘`-Os`’ contributions and Java front end work.
- Stu Grossman for `gdb` hacking, allowing GCJ developers to debug our code.
- 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 bug fixes.
- Andrew Haley for his amazing Java compiler and library efforts.
- 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.
- Dara Hazeghi for wading through myriads of target-specific bug reports.
- Kate Hedstrom for staking the `g77` folks with an initial testsuite.

- Richard Henderson for his ongoing SPARC, alpha, ia32, and ia64 work, loop opts, and generally fixing lots of old problems we've ignored for years, flow rewrite and lots of further stuff, including reviewing tons of patches.
- Aldy Hernandez for working on the PowerPC port, SIMD support, and various fixes.
- Nobuyuki Hikichi of Software Research Associates, Tokyo, contributed the support for the Sony NEWS machine.
- Kazu Hirata for caring and feeding the Renesas H8/300 port and various fixes.
- Manfred Hollstein for his ongoing work to keep the m88k alive, lots of testing and bug fixing, particularly of our configury code.
- Steve Holmgren for MachTen patches.
- Jan Hubicka for his x86 port improvements.
- Falk Hueffner for working on C and optimization bug reports.
- Christian Iseli for various bug fixes.
- Kamil Iskra for general m68k hacking.
- Lee Iverson for random fixes and MIPS testing.
- Andreas Jaeger for testing and benchmarking of GCC and various bug fixes.
- Jakub Jelinek for his SPARC work and sibling call optimizations as well as lots of bug fixes and test cases, and for improving the Java build system.
- Janis Johnson for ia64 testing and fixes, her quality improvement sidetracks, and web page maintenance.
- Kean Johnston for SCO OpenServer support and various fixes.
- Tim Josling for the sample language treelang based originally on Richard Kenner's "toy" language".
- Nicolai Josuttis for additional libstdc++ documentation.
- Klaus Kaempf for his ongoing work to make alpha-vms a viable target.
- David Kashtan of SRI adapted GCC to VMS.
- Ryszard Kabatek for many, many libstdc++ bug fixes and optimizations of strings, especially member functions, and for auto_ptr fixes.
- Geoffrey Keating for his ongoing work to make the PPC work for GNU/Linux and his automatic regression tester.
- Brendan Kehoe for his ongoing work with g++ and for a lot of early work in just about every part of libstdc++.
- Oliver M. Kellogg of Deutsche Aerospace contributed the port to the MIL-STD-1750A.
- 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.
- Mumit Khan for various contributions to the Cygwin and Mingw32 ports and maintaining binary releases for Windows hosts, and for massive libstdc++ porting work to Cygwin/Mingw32.

- Robin Kirkham for cpu32 support.
- Mark Klein for PA improvements.
- Thomas Koenig for various bug fixes.
- Bruce Korb for the new and improved fixincludes code.
- Benjamin Kosnik for his g++ work and for leading the libstdc++-v3 effort.
- Charles LaBrec contributed the support for the Integrated Solutions 68020 system.
- Jeff Law for his direction via the steering committee, coordinating the entire egcs project and GCC 2.95, rolling out snapshots and releases, handling merges from GCC2, reviewing tons of patches that might have fallen through the cracks else, and random but extensive hacking.
- Marc Lehmann for his direction via the steering committee and helping with analysis and improvements of x86 performance.
- Ted Lemon wrote parts of the RTL reader and printer.
- Kriang Lerdsuwanakij for C++ improvements including template as template parameter support, and many C++ fixes.
- Warren Levy for tremendous work on libgcj (Java Runtime Library) and random work on the Java front end.
- Alain Lichnewsky ported GCC to the MIPS CPU.
- Oskar Liljeblad for hacking on AWT and his many Java bug reports and patches.
- Robert Lipe for OpenServer support, new testsuites, testing, etc.
- Weiwen Liu for testing and various bug fixes.
- Dave Love for his ongoing work with the Fortran front end and runtime libraries.
- Martin von Löwis for internal consistency checking infrastructure, various C++ improvements including namespace support, and tons of assistance with libstdc++/compiler merges.
- H.J. Lu for his previous contributions to the steering committee, many x86 bug reports, prototype patches, and keeping the GNU/Linux ports working.
- Greg McGary for random fixes and (someday) bounded pointers.
- Andrew MacLeod for his ongoing work in building a real EH system, various code generation improvements, work on the global optimizer, etc.
- Vladimir Makarov for hacking some ugly i960 problems, PowerPC hacking improvements to compile-time performance, overall knowledge and direction in the area of instruction scheduling, and design and implementation of the automaton based instruction scheduler.
- Bob Manson for his behind the scenes work on dejagnu.
- Philip Martin for lots of libstdc++ string and vector iterator fixes and improvements, and string clean up and testsuites.
- All of the Mauve project contributors, for Java test code.
- Bryce McKinlay for numerous GCJ and libgcj fixes and improvements.
- Adam Megacz for his work on the Windows port of GCJ.
- Michael Meissner for LRS framework, ia32, m32r, v850, m88k, MIPS, powerpc, haifa, ECOFF debug support, and other assorted hacking.

- Jason Merrill for his direction via the steering committee and leading the g++ effort.
- David Miller for his 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.
- Alfred Minarik for libstdc++ string and ios bug fixes, and turning the entire libstdc++ testsuite namespace-compatible.
- Mark Mitchell for his direction via the steering committee, mountains of C++ work, load/store hoisting out of loops, alias analysis improvements, ISO C restrict support, and serving as release manager for GCC 3.x.
- Alan Modra for various GNU/Linux bits and testing.
- Toon Moene for his direction via the steering committee, Fortran maintenance, and his ongoing work to make us make Fortran run fast.
- 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. Doing all this work on scrap paper and the backs of envelopes would have been... difficult.
- Catherine Moore for fixing various ugly problems we have sent her way, including the haifa bug which was killing the Alpha & PowerPC Linux kernels.
- Mike Moreton for his various Java patches.
- 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 C99 support.
- Bill Moyer for his behind the scenes work on various issues.
- Philippe De Muyter for his work on the m68k port.
- Joseph S. Myers for his work on the PDP-11 port, format checking and ISO C99 support, and continuous emphasis on (and contributions to) documentation.
- Nathan Myers for his work on libstdc++-v3: architecture and authorship through the first three snapshots, including implementation of locale infrastructure, string, shadow C headers, and the initial project documentation (DESIGN, CHECKLIST, and so forth). Later, more work on MT-safe string and shadow headers.
- Felix Natter for documentation on porting libstdc++.
- Nathanael Nerode for cleaning up the configuration/build process.
- NeXT, Inc. donated the front end that supports the Objective-C language.
- Hans-Peter Nilsson for the CRIS and MMIX ports, improvements to the search engine setup, various documentation fixes and other small fixes.
- Geoff Noer for this work on getting cygwin native builds working.
- Diego Novillo for his SPEC performance tracking web pages and assorted fixes in the middle end and various back ends.
- David O'Brien for the FreeBSD/alpha, FreeBSD/AMD x86-64, FreeBSD/ARM, FreeBSD/PowerPC, and FreeBSD/SPARC64 ports and related infrastructure improvements.

- Alexandre Oliva for various build infrastructure improvements, scripts and amazing testing work, including keeping libtool issues sane and happy.
- Melissa O'Neill for various NeXT fixes.
- Rainer Orth for random MIPS work, including improvements to our o32 ABI support, improvements to dejagnu's MIPS support, Java configuration clean-ups and porting work, etc.
- Hartmut Penner for work on the s390 port.
- Paul Petersen wrote the machine description for the Alliant FX/8.
- Alexandre Petit-Bianco for implementing much of the Java compiler and continued Java maintainership.
- Matthias Pfaller for major improvements to the NS32k port.
- Gerald Pfeifer for his direction via the steering committee, pointing out lots of problems we need to solve, maintenance of the web pages, and taking care of documentation maintenance in general.
- Andrew Pinski for processing bug reports by the dozen.
- 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.
- Rolf W. Rasmussen for hacking on AWT.
- David Reese of Sun Microsystems contributed to the Solaris on PowerPC port.
- Volker Reichelt for keeping up with the problem reports.
- Joern Rennecke for maintaining the sh port, loop, regmove & reload hacking.
- Loren J. Rittle for improvements to libstdc++-v3 including the FreeBSD port, threading fixes, thread-related configury changes, critical threading documentation, and solutions to really tricky I/O problems, as well as keeping GCC properly working on FreeBSD and continuous testing.
- Craig Rodrigues for processing tons of bug reports.
- 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.
- Chip Salzenberg for libstdc++ patches and improvements to locales, traits, Makefiles, libio, libtool hackery, and "long long" support.
- Juha Sarlin for improvements to the H8 code generator.
- Greg Satz assisted in making GCC work on HP-UX for the 9000 series 300.
- Roger Sayle for improvements to constant folding and GCC's RTL optimizers as well as for fixing numerous bugs.
- Bradley Schatz for his work on the GCJ FAQ.
- 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 as well as serving as release manager for GCC 2.95.3.

- Peter Schmid for constant testing of libstdc++ – especially application testing, going above and beyond what was requested for the release criteria – and libstdc++ header file tweaks.
- Jason Schroeder for jcf-dump patches.
- Andreas Schwab for his work on the m68k port.
- 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 RMS 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 & m68k ports.
- Kenny Simpson for prompting libstdc++ fixes due to defect reports from the LWG (thereby keeping us in line with updates from the ISO).
- Franz Sirl for his ongoing work with making the PPC port stable for linux.
- Andrey Slepuhin for assorted AIX hacking.
- Christopher Smith did the port for Convex machines.
- Danny Smith for his major efforts on the Mingw (and Cygwin) ports.
- Randy Smith finished the Sun FPA support.
- Scott Snyder for queue, iterator, istream, and string fixes and libstdc++ testsuite entries.
- Brad Spencer for contributions to the GLIBCPP_FORCE_NEW technique.
- 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 mips16 related fixes/improvements.
- Jonathan Stone wrote the machine description for the Pyramid computer.
- Graham Stott for various infrastructure improvements.
- John Stracke for his Java HTTP protocol fixes.
- Mike Stump for his Elxsi port, g++ contributions over the years and more recently his vxworks contributions
- Jeff Sturm for Java porting help, bug fixes, and encouragement.
- Shigeya Suzuki for this fixes for the bsdi platforms.
- Ian Lance Taylor for his mips16 work, general configury hacking, fixincludes, etc.
- Holger Teutsch provided the support for the Clipper CPU.
- Gary Thomas for his ongoing work to make the PPC work for GNU/Linux.
- Philipp Thomas for random bug fixes throughout the compiler
- Jason Thorpe for thread support in libstdc++ on NetBSD.
- Kresten Krab Thorup wrote the run time support for the Objective-C language and the fantastic Java bytecode interpreter.
- Michael Tiemann for random bug fixes, the first instruction scheduler, initial C++ support, function integration, NS32k, SPARC and M88k machine description work, delay slot scheduling.

- Andreas Tobler for his work porting libgcj to Darwin.
- 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 Tromey for internationalization support and for his many Java contributions and libgcj maintainership.
- Lassi Tuura for improvements to config.guess to determine HP processor types.
- Petter Urkedal for libstdc++ CXXFLAGS, math, and algorithms fixes.
- Brent Verner for work with the libstdc++ cshadow files and their associated configure steps.
- Todd Vierling for contributions for NetBSD ports.
- Jonathan Wakely for contributing libstdc++ Doxygen notes and XHTML guidance.
- Dean Wakerley for converting the install documentation from HTML to texinfo in time for GCC 3.0.
- Krister Walfridsson for random bug fixes.
- Stephen M. Webb for time and effort on making libstdc++ shadow files work with the tricky Solaris 8+ headers, and for pushing the build-time header tree.
- 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.
- Ulrich Weigand for work on the s390 port.
- Zack Weinberg for major work on cplib and various other bug fixes.
- Matt Welsh for help with Linux Threads support in GCJ.
- Urban Widmark for help fixing java.io.
- Mark Wielaard for new Java library code and his work integrating with Classpath.
- Dale Wiles helped port GCC to the Tahoe.
- Bob Wilson from Tensilica, Inc. for the Xtensa port.
- Jim Wilson for his direction via the steering committee, 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.
- Gilles Zunino for help porting Java to Irix.

In addition to the above, all of which also contributed time and energy in testing GCC, we would like to thank the following for their contributions to testing:

- Michael Abd-El-Malek
- Thomas Arend

- Bonzo Armstrong
- Steven Ashe
- Chris Baldwin
- David Billingham
- Jim Blandy
- Stephane Bortzmeyer
- Horst von Brand
- Frank Braun
- Rodney Brown
- Sidney Cadot
- Bradford Castalia
- Ralph Doncaster
- Richard Emberson
- Levente Farkas
- Graham Fawcett
- Robert A. French
- Jörgen Freyh
- Mark K. Gardner
- Charles-Antoine Gauthier
- Yung Shing Gene
- David Gilbert
- Simon Gornall
- Fred Gray
- John Griffin
- Patrik Hagglund
- Phil Hargett
- Amancio Hasty
- Bryan W. Headley
- Kevin B. Hendricks
- Joep Jansen
- Christian Joensson
- David Kidd
- Tobias Kuipers
- Anand Krishnaswamy
- llewelly
- Damon Love
- Brad Lucier
- Matthias Klose
- Martin Knoblauch

- Jesse Macnish
- Stefan Morrell
- Anon A. Mous
- Matthias Mueller
- Pekka Nikander
- Jon Olson
- Magnus Persson
- Chris Pollard
- Richard Polton
- David Rees
- Paul Reilly
- Tom Reilly
- Torsten Rueger
- Danny Sadinoff
- Marc Schifer
- David Schuler
- Vin Shelton
- Tim Souder
- Adam Sulmicki
- George Talbot
- Gregory Warnes
- David E. Young
- And many others

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.

Option Index

GCC's command line options are indexed here without any initial '-' or '--'. Where an option has both positive and negative forms (such as '*-foption*' and '*-fno-option*'), relevant entries in the manual are indexed under the most appropriate form; it may sometimes be useful to look up both forms.

D

dB	36
dc	34
dd	36
de	33
dE	35
df	34
dg	36
dG	34
di	32
dj	33
dk	36
dl	35
dL	34
dN	35
dr	32
dR	36

ds	34
dS	35
dt	34
dW	33
dX	33

F

fnew-ra	35
frerun-cse-after-loop	34
fssa	33
fssa-ccp	33
fssa-dce	33
fthread-jumps	33

M

msoft-float	8
-------------	---

Index

- !**
 ‘!’ in constraint 131
- #**
 ‘#’ in constraint 132
 # in template 124
 #pragma 319, 320
- %**
 ‘%’ in constraint 132
 % in GTY option 339
 ‘%’ in template 123
- &**
 ‘&’ in constraint 131
- (**
 (nil) 72
- ***
 ‘*’ in constraint 132
 * in template 125
- +**
 ‘+’ in constraint 131
- /**
 ‘/c’ in RTL dump 79
 ‘/f’ in RTL dump 79
 ‘/i’ in RTL dump 79
 ‘/j’ in RTL dump 80
 ‘/s’ in RTL dump 79
 ‘/u’ in RTL dump 80
 ‘/v’ in RTL dump 80
- <**
 ‘<’ in constraint 127
- =**
 ‘=’ in constraint 131
- >**
 ‘>’ in constraint 127
- ?**
 ‘?’ in constraint 131
- __builtin_args_info 254
 __builtin_classify_type 254
 __builtin_next_arg 254
 __builtin_saveregs 253
 __CTOR_LIST__ 292
 __DTOR_LIST__ 292
 __main 335
- **
 \ 124
- 0**
 ‘0’ in constraint 127
- A**
 abort 5
 abs 93
 abs and attributes 177
 absence_set 188
 absm2 instruction pattern 147
 absolute value 93
 access to operands 74
 accessors 74
 ACCUMULATE_OUTGOING_ARGS 240
 ACCUMULATE_OUTGOING_ARGS and stack frames
 250
 ADA_LONG_TYPE_SIZE 213
 ADDITIONAL_REGISTER_NAMES 296
 addm3 instruction pattern 146
 addr_diff_vec 102
 addr_diff_vec, length of 181
 ADDR_EXPR 61
 addr_vec 102
 addr_vec, length of 181
 address constraints 128
 ADDRESS_COST 267
 address_operand 128
 addressing modes 260
 addressof 90
 ADJUST_FIELD_ALIGN 207
 ADJUST_INSN_LENGTH 181
 aggregates as return values 247
 ALL_COP_ADDITIONAL_REGISTER_NAMES 314
 ALL_REGS 223
 ALLOCATE_INITIAL_VALUE 324
 allocate_stack instruction pattern 155

ALLOCATE_TRAMPOLINE	257	asm_operands, usage	104
alternate entry points	106	ASM_OUTPUT_ADDR_DIFF_ELT	299
analysis, data flow	34	ASM_OUTPUT_ADDR_VEC_ELT	299
and	92	ASM_OUTPUT_ALIGN	302
and and attributes	176	ASM_OUTPUT_ALIGN_WITH_NOP	302
and, canonicalization of	163	ASM_OUTPUT_ALIGNED_BSS	285
andm3 instruction pattern	146	ASM_OUTPUT_ALIGNED_COMMON	284
APPLY_RESULT_SIZE	246	ASM_OUTPUT_ALIGNED_DECL_COMMON	284
ARG_POINTER_CFA_OFFSET	232	ASM_OUTPUT_ALIGNED_DECL_LOCAL	285
ARG_POINTER_REGNUM	236	ASM_OUTPUT_ALIGNED_LOCAL	285
ARG_POINTER_REGNUM and virtual registers	87	ASM_OUTPUT_ASCII	282
arg_pointer_rtx	237	ASM_OUTPUT_BSS	284
ARGS_GROW_DOWNWARD	231	ASM_OUTPUT_CASE_END	299
argument passing	7	ASM_OUTPUT_CASE_LABEL	299
arguments in registers	242	ASM_OUTPUT_COMMON	284
arguments on stack	239	ASM_OUTPUT_DEBUG_LABEL	290
arithmetic libraries	8	ASM_OUTPUT_DEF	291
arithmetic shift	92	ASM_OUTPUT_DEF_FROM_DECLS	291
arithmetic simplifications	32	ASM_OUTPUT_DWARF_DELTA	309
arithmetic, in RTL	90	ASM_OUTPUT_DWARF_OFFSET	309
ARITHMETIC_TYPE_P	43	ASM_OUTPUT_DWARF_PCREL	309
array	42	ASM_OUTPUT_EXTERNAL	289
ARRAY_REF	61	ASM_OUTPUT_EXTERNAL_LIBCALL	289
ARRAY_TYPE	42	ASM_OUTPUT_FDESC	282
ashift	92	ASM_OUTPUT_IDENT	280
ashift and attributes	177	ASM_OUTPUT_INTERNAL_LABEL	290
ashiftrt	93	ASM_OUTPUT_LABEL	286
ashiftrt and attributes	177	ASM_OUTPUT_LABEL_REF	290
ashlm3 instruction pattern	147	ASM_OUTPUT_LABELREF	289
ashrm3 instruction pattern	147	ASM_OUTPUT_LOCAL	285
ASM_APP_OFF	280	ASM_OUTPUT_MAX_SKIP_ALIGN	302
ASM_APP_ON	280	ASM_OUTPUT_MEASURED_SIZE	286
ASM_CLOBBERS	55	ASM_OUTPUT_OPCODE	296
ASM_COMMENT_START	280	ASM_OUTPUT_POOL_EPILOGUE	283
ASM_CV_QUAL	55	ASM_OUTPUT_POOL_PROLOGUE	282
ASM_DECLARE_CLASS_REFERENCE	292	ASM_OUTPUT_REG_POP	299
ASM_DECLARE_CONSTANT_NAME	287	ASM_OUTPUT_REG_PUSH	298
ASM_DECLARE_FUNCTION_NAME	287	ASM_OUTPUT_SHARED_BSS	285
ASM_DECLARE_FUNCTION_SIZE	287	ASM_OUTPUT_SHARED_COMMON	284
ASM_DECLARE_OBJECT_NAME	287	ASM_OUTPUT_SHARED_LOCAL	285
ASM_DECLARE_REGISTER_GLOBAL	288	ASM_OUTPUT_SIZE_DIRECTIVE	286
ASM_DECLARE_UNRESOLVED_REFERENCE	292	ASM_OUTPUT_SKIP	302
ASM_FILE_END	279	ASM_OUTPUT_SOURCE_FILENAME	280
ASM_FILE_START	279	ASM_OUTPUT_SOURCE_LINE	280
ASM_FINAL_SPEC	195	ASM_OUTPUT_SPECIAL_POOL_ENTRY	282
ASM_FINISH_DECLARE_OBJECT	288	ASM_OUTPUT_SYMBOL_REF	289
ASM_FORMAT_PRIVATE_NAME	290	ASM_OUTPUT_TYPE_DIRECTIVE	287
asm_fprintf	298	ASM_OUTPUT_WEAK_ALIAS	291
ASM_FPRINTF_EXTENSIONS	298	ASM_OUTPUTS	55
ASM_GENERATE_INTERNAL_LABEL	290	ASM_PREFERRED_EH_DATA_FORMAT	234
asm_input	102	ASM_SPEC	195
asm_input and '/v'	76	ASM_STABD_OP	304
ASM_INPUTS	55	ASM_STABN_OP	304
ASM_MAYBE_OUTPUT_ENCODED_ADDR_RTX	234	ASM_STABS_OP	304
ASM_NO_SKIP_IN_TEXT	302	ASM_STMT	55
asm_noperands	108	ASM_STRING	55
asm_operands and '/v'	76	ASM_WEAKEN_DECL	288
asm_operands, RTL sharing	114	ASM_WEAKEN_LABEL	288

- assemble_name 286
 - assembler format 279
 - assembler instructions in RTL 104
 - ASSEMBLER_DIALECT 298
 - assigning attribute values to insns 178
 - assignment operator 52
 - asterisk in template 125
 - attr 178, 179
 - attr_flag 177
 - attribute expressions 176
 - attribute specifications 180
 - attribute specifications example 180
 - attributes 61
 - attributes, defining 175
 - attributes, target-specific 313
 - autoincrement addressing, availability 5
 - autoincrement/decrement addressing 126
 - autoincrement/decrement analysis 34
 - automata_option 188
 - automaton based pipeline description ... 183, 186, 190
 - automaton based scheduler 183
 - AVOID_CCMODE_COPIES 222
- B**
- backslash 124
 - barrier 106
 - barrier and '/f' 77
 - barrier and '/i' 77
 - barrier and '/v' 75
 - BASE_REG_CLASS 224
 - basic block reordering 36
 - basic blocks 34
 - bcond instruction pattern 151
 - bcopy, implicit usage 260
 - BIGGEST_ALIGNMENT 207
 - BIGGEST_FIELD_ALIGNMENT 207
 - BImode 81
 - BIND_EXPR 61
 - BINFO_TYPE 48
 - bit-fields 95
 - BIT_AND_EXPR 61
 - BIT_IOR_EXPR 61
 - BIT_NOT_EXPR 61
 - BIT_XOR_EXPR 61
 - BITFIELD_NBYTES_LIMITED 210
 - BITS_BIG_ENDIAN 205
 - BITS_BIG_ENDIAN, effect on sign_extract 95
 - BITS_PER_UNIT 205
 - BITS_PER_WORD 205
 - bitwise complement 92
 - bitwise exclusive-or 92
 - bitwise inclusive-or 92
 - bitwise logical-and 92
 - BLKmode 82
 - BLKmode, and function return values 113
 - bool 278, 301
 - BOOL_TYPE_SIZE 214
 - BOOLEAN_TYPE 42
 - branch shortening 36
 - BRANCH_COST 269
 - break_out_memory_refs 262
 - BREAK_STMT 55
 - BSS_SECTION_ASM_OP 276
 - builtin_longjmp instruction pattern 157
 - BUILTIN_SETJMP_FRAME_VALUE 232
 - builtin_setjmp_receiver instruction pattern 156
 - builtin_setjmp_setup instruction pattern ... 156
 - byte_mode 84
 - BYTES_BIG_ENDIAN 205
 - BYTES_BIG_ENDIAN, effect on subreg 88
 - bzero, implicit usage 260
- C**
- C statements for assembler output 124
 - C/C++ Internal Representation 39
 - C4X_FLOAT_FORMAT 211
 - call 79, 99
 - call instruction pattern 151
 - call usage 113
 - call-clobbered register 218, 219
 - call-saved register 218, 219
 - call-used register 218, 219
 - CALL_EXPR 61
 - call_insn 105
 - call_insn and '/f' 77
 - call_insn and '/i' 77
 - call_insn and '/j' 78
 - call_insn and '/s' 75, 77
 - call_insn and '/u' 75
 - call_insn and '/v' 75
 - CALL_INSN_FUNCTION_USAGE 105
 - call_pop instruction pattern 151
 - CALL_POPS_ARGS 241
 - CALL_REALLY_USED_REGISTERS 218
 - CALL_USED_REGISTERS 218
 - call_used_regs 219
 - call_value instruction pattern 151
 - call_value_pop instruction pattern 151
 - CALLER_SAVE_PROFITABLE 248
 - calling conventions 230
 - calling functions in RTL 113
 - CAN_DEBUG_WITHOUT_FP 203
 - CAN_ELIMINATE 239
 - canadian 11
 - canonicalization of instructions 163
 - CANONICALIZE_COMPARISON 265
 - canonicalize_funcptr_for_compare instruction pattern 154
 - CASE_DROPS_THROUGH 315
 - CASE_VALUES_THRESHOLD 315
 - CASE_VECTOR_MODE 315
 - CASE_VECTOR_PC_RELATIVE 315

CASE_VECTOR_SHORTEN_MODE	315	COMPOUND_BODY	55
casesi instruction pattern	153	COMPOUND_EXPR	61
cc_status	264	COMPOUND_LITERAL_EXPR	61
CC_STATUS_MDEP	264	COMPOUND_LITERAL_EXPR_DECL	69
CC_STATUS_MDEP_INIT	264	COMPOUND_LITERAL_EXPR_DECL_STMT	69
cc0	89	COMPOUND_STMT	55
cc0, RTL sharing	114	computing the length of an insn	180
cc0_rtx	89	concat and '/u'	77
CC1_SPEC	195	cond	94
CC1PLUS_SPEC	195	cond and attributes	176
CCmode	82	cond_exec	101
CDImode	82	COND_EXPR	61
change_address	144	condition code register	89
char	278	condition code status	264
CHAR_TYPE_SIZE	214	condition codes	93
check_stack instruction pattern	155	conditional constant propagation	33
CHImode	82	Conditional Constant Propagation, SSA based	33
class	48	conditional execution	190
class definitions, register	223	CONDITIONAL_REGISTER_USAGE	219
class preference constraints	131	conditional_trap instruction pattern	158
CLASS_LIKELY_SPILLED_P	228	conditions, in patterns	118
CLASS_MAX_NREGS	228	configuration file	327
CLASS_TYPE_P	43	configure terms	11
classes of RTX codes	72	CONJ_EXPR	61
CLASSTYPE_DECLARED_CLASS	48	const and '/i'	77
CLASSTYPE_HAS_MUTABLE	49	CONST_COSTS	267
CLASSTYPE_NON_POD_P	49	CONST_DECL	49
CLEANUP_DECL	55	const_double	84
CLEANUP_EXPR	55	const_double, RTL sharing	114
CLEANUP_POINT_EXPR	61	CONST_DOUBLE_CHAIN	85
CLEANUP_STMT	55	CONST_DOUBLE_LOW	85
CLEAR_BY_PIECES_P	270	CONST_DOUBLE_MEM	85
CLEAR_INSN_CACHE	258	CONST_DOUBLE_OK_FOR_LETTER_P	229
CLEAR_RATIO	270	const_int	84
clobber	99	const_int and attribute tests	176
clrstrm instruction pattern	148	const_int and attributes	176
cmpm instruction pattern	148	const_int, RTL sharing	114
cmpstrm instruction pattern	149	CONST_OK_FOR_LETTER_P	229
code generation RTL sequences	164	CONST_OR_PURE_CALL_P	75
code motion	34	const_string	85
code_label	106	const_string and attributes	176
code_label and '/i'	76, 77	const_true_rtx	84
code_label and '/v'	75	const_vector	84
CODE_LABEL_NUMBER	106	const_vector, RTL sharing	114
codes, RTL expression	71	const0_rtx	84
COImode	82	CONST0_RTX	85
COLLECT_EXPORT_LIST	324	const1_rtx	84
COLLECT2_HOST_INITIALIZATION	328	CONST1_RTX	85
combiner pass	88	const2_rtx	84
common subexpression elimination	33	CONST2_RTX	85
compare	91	constant attributes	182
compare, canonicalization of	163	constant definitions	191
compiler passes and files	31	constant folding	32
complement, bitwise	92	constant propagation	33, 34
COMPLEX_CST	61	CONSTANT_ADDRESS_P	260
COMPLEX_EXPR	61	CONSTANT_AFTER_FUNCTION_P	283
COMPLEX_TYPE	42	CONSTANT_ALIGNMENT	208
COMPONENT_REF	61		

CONSTANT_P 260
 CONSTANT_POOL_ADDRESS_P 75
 CONSTANT_POOL_BEFORE_FUNCTION 282
 constants in constraints 127
 constm1_rtx 84
 constraint modifier characters 131
 constraint, matching 128
 constraints 126
 constraints, machine specific 132
 constructor 52
 CONSTRUCTOR 61
 constructors, automatic calls 335
 constructors, output of 292
 container 41
 CONTINUE_STMT 55
 contributors 363
 controlling register usage 219
 controlling the compilation driver 193
 conventions, run-time 7
 conversions 96
 CONVERT_EXPR 61
 copy constructor 52
 copy propagation 34
 copy_rtx 263
 copy_rtx_if_shared 114
 cosm2 instruction pattern 147
 costs of instructions 267
 COSTS_N_INSNS 267
 CP_INTEGRAL_TYPE 43
 cp_namespace_decls 47
 CP_TYPE_CONST_NON_VOLATILE_P 42
 CP_TYPE_CONST_P 42
 CP_TYPE_QUALS 42
 CP_TYPE_RESTRICT_P 42
 CP_TYPE_VOLATILE_P 42
 CPLUSPLUS_CPP_SPEC 195
 CPP_PREDEFINES 201
 cpp_register_pragma 319
 CPP_SPEC 195
 CQImode 82
 cross compilation and floating point 310
 CRT_CALL_STATIC_FUNCTION 276
 CRTSTUFF_T_CFLAGS 331
 CRTSTUFF_T_CFLAGS_S 331
 CSImode 82
 CTImode 82
 CUMULATIVE_ARGS 243
 current_function_epilogue_delay_list 251
 current_function_is_leaf 222
 current_function_outgoing_args_size 240
 current_function_pops_args 250
 current_function_pretend_args_size 250
 current_function_uses_only_leaf_regs 222
 current_insn_predicate 190

D

data bypass 186, 188, 190
 data dependence delays 183
 data flow analysis 34
 data structures 204
 DATA_ALIGNMENT 208
 data_section 276
 DATA_SECTION_ASM_OP 275
 DBR_OUTPUT_SEQEND 297
 dbr_sequence_length 297
 DBX_BLOCKS_FUNCTION_RELATIVE 305
 DBX_CONTIN_CHAR 305
 DBX_CONTIN_LENGTH 304
 DBX_DEBUGGING_INFO 304
 DBX_FUNCTION_FIRST 305
 DBX_LBRAC_FIRST 305
 DBX_MEMPARG_STABS_LETTER 305
 DBX_NO_XREFS 304
 DBX_OUTPUT_ENUM 306
 DBX_OUTPUT_FUNCTION_END 306
 DBX_OUTPUT_LBRAC 306
 DBX_OUTPUT_MAIN_SOURCE_DIRECTORY 308
 DBX_OUTPUT_MAIN_SOURCE_FILE_END 308
 DBX_OUTPUT_MAIN_SOURCE_FILENAME 307
 DBX_OUTPUT_NFUN 306
 DBX_OUTPUT_RBRAC 306
 DBX_OUTPUT_SOURCE_FILENAME 308
 DBX_OUTPUT_STANDARD_TYPES 306
 DBX_REGISTER_NUMBER 303
 DBX_REGPARG_STABS_CODE 305
 DBX_REGPARG_STABS_LETTER 305
 DBX_STATIC_CONST_VAR_CODE 305
 DBX_STATIC_STAB_DATA_SECTION 305
 DBX_TYPE_DECL_STABS_CODE 305
 DBX_USE_BINCL 306
 DBX_WORKING_DIRECTORY 307
 DCE, SSA based 33
 DCmode 82
 De Morgan's law 163
 dead code 33
 dead code elimination 33
 dead_or_set_p 172
 DEBUG_SYMS_TEXT 304
 DEBUGGER_ARG_OFFSET 303
 DEBUGGER_AUTO_OFFSET 303
 debugging information generation 36
 DECL_ALIGN 49
 DECL_ANTICIPATED 53
 DECL_ARGUMENTS 55
 DECL_ARRAY_DELETE_OPERATOR_P 55
 DECL_ARTIFICIAL 50, 52, 55
 DECL_ASSEMBLER_NAME 52, 53
 DECL_ATTRIBUTES 61
 DECL_BASE_CONSTRUCTOR_P 54
 DECL_CLASS_SCOPE_P 50
 DECL_COMPLETE_CONSTRUCTOR_P 54
 DECL_COMPLETE_DESTRUCTOR_P 54
 DECL_CONST_MEMFUNC_P 53

DECL_CONSTRUCTOR_P	52, 53	define_insn_reservation	186
DECL_CONTEXT	47	define_peephole	171
DECL_CONV_FN_P	52, 54	define_peephole2	174
DECL_COPY_CONSTRUCTOR_P	54	define_query_cpu_unit	186
DECL_DESTRUCTOR_P	52, 54	define_reservation	188
DECL_EXTERN_C_FUNCTION_P	53	define_split	167
DECL_EXTERNAL	49, 53	defining attributes and their values	175
DECL_FUNCTION_MEMBER_P	52, 53	defining jump instruction patterns	160
DECL_FUNCTION_SCOPE_P	50	defining looping instruction patterns	161
DECL_GLOBAL_CTOR_P	52, 54	defining peephole optimizers	171
DECL_GLOBAL_DTOR_P	52, 54	defining RTL sequences for code generation	164
DECL_INITIAL	49	delay slots, defining	182
DECL_LINKONCE_P	52, 53	DELAY_SLOTS_FOR_EPILOGUE	251
DECL_LOCAL_FUNCTION_P	53	delayed branch scheduling	36
DECL_MAIN_P	52	deletable	341
DECL_NAME	47, 49, 52	Dependent Patterns	159
DECL_NAMESPACE_ALIAS	47	desc	340
DECL_NAMESPACE_SCOPE_P	50	destructor	52
DECL_NAMESPACE_STD_P	47	destructors, output of	292
DECL_NON_THUNK_FUNCTION_P	54	deterministic finite state automaton	183, 188
DECL_NONCONVERTING_P	54	DFA_PIPELINE_INTERFACE	274
DECL_NONSTATIC_MEMBER_FUNCTION_P	53	DFmode	81
DECL_OVERLOADED_OPERATOR_P	52, 54	digits in constraint	127
DECL_RESULT	55	DImode	81
DECL_SIZE	49	DIR_SEPARATOR	327
DECL_SOURCE_FILE	49	DIR_SEPARATOR_2	327
DECL_SOURCE_LINE	49	directory options .md	171
DECL_STATIC_FUNCTION_P	53	disabling certain registers	219
DECL_STMT	55	dispatch table	299
DECL_STMT_DECL	55	div	92
DECL_THUNK_P	54	div and attributes	177
DECL_VOLATILE_MEMFUNC_P	53	DIVDI3_LIBCALL	259
declaration	49	division	92
declarations, RTL	97	divm3 instruction pattern	146
DECLARE_LIBRARY_RENAMES	259	divmodm4 instruction pattern	147
decrement_and_branch_until_zero instruction		DIVSI3_LIBCALL	258
pattern	153	DO_BODY	55
default	340	DO_COND	55
DEFAULT_CALLER_SAVES	248	DO_STMT	55
DEFAULT_GDB_EXTENSIONS	304	DOLLARS_IN_IDENTIFIERS	320
DEFAULT_MAIN_RETURN	321	doloop_begin instruction pattern	154
DEFAULT_PCC_STRUCT_RETURN	247	doloop_end instruction pattern	153
DEFAULT_RTX_COSTS	267	DONE	165
DEFAULT_SHORT_ENUMS	215	DONT_REDUCE_ADDR	269
DEFAULT_SIGNED_CHAR	214	DOUBLE_TYPE_SIZE	214
define_asm_attributes	179	driver	193
define_attr	175	DRIVER_SELF_SPECS	194
define_automaton	186	DUMPFILe_FORMAT	328
define_bypass	188	DWARF_CIE_DATA_ALIGNMENT	300
define_cond_exec	190	DWARF_DEBUGGING_INFO	308
define_constants	191	DWARF_FRAME_REGISTERS	237
define_cpu_unit	186	DWARF2_ASM_LINE_DEBUG_INFO	309
define_delay	182	DWARF2_DEBUGGING_INFO	308
define_expand	164	DWARF2_FRAME_INFO	308
define_function_unit	184	DWARF2_GENERATE_TEXT_SECTION_LABEL	309
define_insn	117	DWARF2_UNWIND_INFO	300
define_insn example	118	DYNAMIC_CHAIN_ADDRESS	231
define_insn_and_split	169		

E

- 'E' in constraint 127
 - earlyclobber operand 131
 - EDOM, implicit usage 259
 - EH_FRAME_IN_DATA_SECTION 300
 - EH_FRAME_SECTION_NAME 300
 - eh_return instruction pattern 157
 - EH_RETURN_DATA_REGNO 233
 - EH_RETURN_HANDLER_RTX 233
 - EH_RETURN_STACKADJ_RTX 233
 - EH_USES 251
 - ELIGIBLE_FOR_EPILOGUE_DELAY 251
 - ELIMINABLE_REGS 238
 - ELSE_CLAUSE 55
 - EMIT_MODE_SET 312
 - EMPTY_CLASS_EXPR 55
 - EMPTY_FIELD_BOUNDARY 208
 - ENDFILE_SPEC 196
 - endianness 5
 - enum machine_mode 81
 - enum reg_class 224
 - ENUMERAL_TYPE 42
 - epilogue 248
 - epilogue instruction pattern 157
 - EPILOGUE_USES 251
 - eq 94
 - eq and attributes 177
 - eq_attr 177
 - EQ_EXPR 61
 - equal 94
 - errno, implicit usage 259
 - escape sequences 217
 - exception handling 233
 - exception_receiver instruction pattern 156
 - exclamation point 131
 - exclusion_set 188
 - exclusive-or, bitwise 92
 - EXIT_BODY 321
 - EXIT_EXPR 61
 - EXIT_IGNORE_STACK 251
 - EXPAND_BUILTIN_SAVEREGS 254
 - expander definitions 164
 - expm2 instruction pattern 147
 - expr_list 113
 - EXPR_STMT 55
 - EXPR_STMT_EXPR 55
 - expression 61
 - expression codes 71
 - extendmn2 instruction pattern 149
 - extensible constraints 128
 - extern int target_flags 201
 - EXTRA_ADDRESS_CONSTRAINT 230
 - EXTRA_CC_MODES 265
 - EXTRA_CONSTRAINT 229
 - EXTRA_MEMORY_CONSTRAINT 229
 - EXTRA_SECTION_FUNCTIONS 276
 - EXTRA_SECTIONS 276
 - EXTRA_SPECS 196
 - extv instruction pattern 150
 - extzv instruction pattern 150
- F**
- 'F' in constraint 127
 - FAIL 165
 - FATAL_EXIT_CODE 327
 - FDL, GNU Free Documentation License 355
 - features, optional, in system conventions 201
 - ffs 93
 - ffsm2 instruction pattern 148
 - FIELD_DECL 49
 - FILE_STMT 55
 - FILE_STMT_FILENAME 55
 - files and passes of the compiler 31
 - files, generated 342
 - final pass 36
 - FINAL_PRESCAN_INSN 296
 - FINAL_PRESCAN_LABEL 297
 - FINAL_REG_PARM_STACK_SPACE 240
 - final_scan_insn 251
 - final_sequence 298
 - FINALIZE_PIC 279
 - FIND_BASE_TERM 262
 - FINI_SECTION_ASM_OP 276
 - finite state automaton minimization 188
 - FIRST_INSN_ADDRESS 181
 - FIRST_PARM_OFFSET 231
 - FIRST_PARM_OFFSET and virtual registers 87
 - FIRST_PSEUDO_REGISTER 218
 - FIRST_STACK_REG 223
 - FIRST_VIRTUAL_REGISTER 87
 - fix 97
 - FIX_TRUNC_EXPR 61
 - fix_truncmn2 instruction pattern 149
 - fixed register 218
 - FIXED_REGISTERS 218
 - fixed_regs 219
 - fixmn2 instruction pattern 149
 - FIXUNS_TRUNC_LIKE_FIX_TRUNC 316
 - fixuns_truncmn2 instruction pattern 149
 - fixunsmn2 instruction pattern 149
 - flags in RTL expression 75
 - float 97
 - FLOAT_EXPR 61
 - float_extend 96
 - FLOAT_LIB_COMPARE_RETURNS_BOOL (*mode*,
 comparison) 259
 - FLOAT_STORE_FLAG_VALUE 318
 - float_truncate 97
 - FLOAT_TYPE_SIZE 214
 - FLOAT_WORDS_BIG_ENDIAN 205
 - FLOAT_WORDS_BIG_ENDIAN, (lack of) effect on
 subreg 88
 - floating point and cross compilation 310
 - Floating Point Emulation 331
 - floatmn2 instruction pattern 149

floatunsmn2 instruction pattern	149
FOR_BODY	55
FOR_COND	55
FOR_EXPR	55
FOR_INIT_STMT	55
FOR_STMT	55
FORCE_CODE_SECTION_ALIGN	276
FORCE_PREFERRED_STACK_BOUNDARY_IN_MAIN	207
force_reg	144
frame layout	230
FRAME_GROWS_DOWNWARD	231
FRAME_GROWS_DOWNWARD and virtual registers	87
frame_pointer_needed	249
FRAME_POINTER_REGNUM	236
FRAME_POINTER_REGNUM and virtual registers	87
FRAME_POINTER_REQUIRED	238
frame_pointer_rtx	237
frame_related	79
frame_related, in insn, call_insn, jump_insn, barrier, and set	77
frame_related, in mem	76
frame_related, in reg	76
frame_related, in symbol_ref	78
ftruncm2 instruction pattern	149
function	52
function body	55
function call conventions	7
function entry and exit	248
function units, for scheduling	184
function-call insns	113
FUNCTION_ARG	242
FUNCTION_ARG_ADVANCE	244
FUNCTION_ARG_BOUNDARY	245
FUNCTION_ARG_CALLEE_COPIES	243
FUNCTION_ARG_PADDING	244
FUNCTION_ARG_PARTIAL_NREGS	243
FUNCTION_ARG_PASS_BY_REFERENCE	243
FUNCTION_ARG_REGNO_P	245
FUNCTION_BOUNDARY	207
FUNCTION_DECL	52
FUNCTION_INCOMING_ARG	242
FUNCTION_MODE	318
FUNCTION_OK_FOR_SIBCALL	253
FUNCTION_OUTGOING_VALUE	246
FUNCTION_PROFILER	252
FUNCTION_TYPE	42
FUNCTION_VALUE	245
FUNCTION_VALUE_REGNO_P	246
functions, leaf	222
fundamental type	42
G	
'g' in constraint	127
'G' in constraint	127
GCC and portability	5
GCC_DRIVER_HOST_INITIALIZATION	328
GCOV_TYPE_SIZE	215
ge	94
ge and attributes	177
GE_EXPR	61
GEN_ERRNO_RTX	259
gencodes	32
genconfig	37
general_operand	120
GENERAL_REGS	223
generated files	342
generating assembler output	124
generating insns	119
genflags	32
get_attr	177
get_attr_length	181
GET_CLASS_NARROWEST_MODE	84
GET_CODE	71
get_frame_size	238
get_insns	105
get_last_insn	105
GET_MODE	83
GET_MODE_ALIGNMENT	84
GET_MODE_BITSIZE	83
GET_MODE_CLASS	83
GET_MODE_MASK	83
GET_MODE_NAME	83
GET_MODE_NUNITS	84
GET_MODE_SIZE	83
GET_MODE_UNIT_SIZE	84
GET_MODE_WIDER_MODE	83
GET_RTX_CLASS	72
GET_RTX_FORMAT	73
GET_RTX_LENGTH	73
geu	94
geu and attributes	177
GGC	339
global common subexpression elimination	34
global register allocation	35
GLOBAL_INIT_PRIORITY	52, 55
GO_IF_LEGITIMATE_ADDRESS	261
GO_IF_MODE_DEPENDENT_ADDRESS	263
GOTO_DESTINATION	55
GOTO_FAKE_P	55
GOTO_STMT	55
graph coloring register allocation	35
greater than	94
gt	94
gt and attributes	177
GT_EXPR	61
gtu	94
gtu and attributes	177
GTY	339

H

'H' in constraint 127
HANDLE_PRAGMA 319
HANDLE_PRAGMA_PACK_PUSH_POP 320
HANDLE_SYSV_PRAGMA 320
HANDLER 55
HANDLER_BODY 55
HANDLER_PARMS 55
hard registers 86
HARD_FRAME_POINTER_REGNUM 236
HARD_REGNO_CALL_PART_CLOBBERED 218
HARD_REGNO_CALLER_SAVE_MODE 248
HARD_REGNO_MODE_OK 220
HARD_REGNO_NREGS 220
HAS_INIT_SECTION 294
HAVE_DOS_BASED_FILE_SYSTEM 327
HAVE_POST_DECREMENT 260
HAVE_POST_INCREMENT 260
HAVE_POST_MODIFY_DISP 260
HAVE_POST_MODIFY_REG 260
HAVE_PRE_DECREMENT 260
HAVE_PRE_INCREMENT 260
HAVE_PRE_MODIFY_DISP 260
HAVE_PRE_MODIFY_REG 260
HCmode 82
HFmode 81
high 86
HImode 81
HImode, in insn 108
host makefile fragment 333
HOST_BIT_BUCKET 328
HOST_EXECUTABLE_SUFFIX 328
HOST_OBJECT_SUFFIX 328
HOT_TEXT_SECTION_NAME 275

I

'i' in constraint 127
'I' in constraint 127
IBM_FLOAT_FORMAT 211
identifier 41
IDENTIFIER_LENGTH 41
IDENTIFIER_NODE 41
IDENTIFIER_OPNAME_P 41
IDENTIFIER_POINTER 41
IDENTIFIER_TYPENAME_P 41
IEEE_FLOAT_FORMAT 211
if conversion 35
IF_COND 55
if_marked 341
IF_STMT 55
if_then_else 94
if_then_else and attributes 176
if_then_else usage 98
IFCVT_EXTRA_FIELDS 323
IFCVT_INIT_EXTRA_FIELDS 323
IFCVT_MODIFY_CANCEL 323
IFCVT_MODIFY_FINAL 323

IFCVT_MODIFY_INSN 323
IFCVT_MODIFY_MULTIPLE_TESTS 323
IFCVT_MODIFY_TESTS 322
IMAGPART_EXPR 61
immediate_operand 120
IMMEDIATE_PREFIX 298
in_data 276
in_struct 79
in_struct, in code_label and note 76
in_struct, in insn 75
in_struct, in insn and jump_insn and call_insn
..... 75
in_struct, in insn, jump_insn and call_insn
..... 77
in_struct, in label_ref 75
in_struct, in mem 76
in_struct, in reg 76
in_struct, in subreg 78
in_text 276
include 170
INCLUDE_DEFAULTS 199
inclusive-or, bitwise 92
INCOMING_FRAME_SP_OFFSET 232
INCOMING_REGNO 219
INCOMING_RETURN_ADDR_RTX 232
INDEX_REG_CLASS 225
indirect_jump instruction pattern 153
INDIRECT_REF 61
INIT_CUMULATIVE_ARGS 244
INIT_CUMULATIVE_INCOMING_ARGS 244
INIT_CUMULATIVE_LIBCALL_ARGS 244
INIT_ENVIRONMENT 198
INIT_EXPANDERS 204
INIT_EXPR 61
init_machine_status 204
INIT_SECTION_ASM_OP 276, 294
INIT_TARGET_OPTABS 259
INITIAL_ELIMINATION_OFFSET 239
INITIAL_FRAME_POINTER_OFFSET 238
initialization routines 292
INITIALIZE_TRAMPOLINE 256
inline on rtx, automatic 32
inline on trees, automatic 32
inlining 314
insn 105
insn and '/f' 77
insn and '/i' 77
insn and '/j' 78
insn and '/s' 75, 77
insn and '/u' 75
insn and '/v' 75
insn attributes 175
insn canonicalization 163
insn includes 170
insn lengths, computing 180
insn splitting 167
insn-attr.h 175
INSN_ANNULLED_BRANCH_P 75

INSN_CACHE_DEPTH	257
INSN_CACHE_LINE_WIDTH	257
INSN_CACHE_SIZE	257
INSN_CODE	108
INSN_DEAD_CODE_P	75
INSN_DELETED_P	75
INSN_FROM_TARGET_P	75
insn_list	113
insn_list and ‘/i’	77
INSN_REFERENCES_ARE_DELAYED	321
INSN_SETS_ARE_DELAYED	321
INSN_UID	104
insns	104
insns, generating	119
insns, recognizing	119
instruction attributes	175
instruction combination	34
instruction latency time	183, 186, 188, 190
instruction patterns	117
instruction recognizer	37
instruction scheduling	35, 36
instruction splitting	167
insv instruction pattern	150
INT_TYPE_SIZE	213
INTEGER_CST	61
INTEGER_TYPE	42
INTEGRATE_THRESHOLD	318
integrated	79
integrated, in insn, call_insn, jump_insn, barrier, code_label, insn_list, const, and note	77
integrated, in reg	76
integrated, in symbol_ref	78
Interdependence of Patterns	159
interfacing to GCC output	7
interlock delays	183, 190
INTMAX_TYPE	216
introduction	1
INVOKE_main	294
ior	92
ior and attributes	176
ior, canonicalization of	163
ior3 instruction pattern	146
IS_ASM_LOGICAL_LINE_SEPARATOR	283

J

jump	80
jump instruction pattern	151
jump instruction patterns	160
jump instructions and set	98
jump optimization	33
jump threading	33
jump, in call_insn	78
jump, in insn	78
jump, in mem	76
JUMP_ALIGN	301
jump_insn	105

jump_insn and ‘/f’	77
jump_insn and ‘/i’	77
jump_insn and ‘/s’	75, 77
jump_insn and ‘/u’	75
jump_insn and ‘/v’	75
JUMP_LABEL	105
JUMP_TABLES_IN_TEXT_SECTION	277

L

LABEL_ALIGN	302
LABEL_ALIGN_AFTER_BARRIER	301
LABEL_ALIGN_AFTER_BARRIER_MAX_SKIP	301
LABEL_ALIGN_MAX_SKIP	302
LABEL_ALT_ENTRY_P	106
LABEL_DECL	49
LABEL_KIND	106
LABEL_NUSES	106
LABEL_OUTSIDE_LOOP_P	75
LABEL_PRESERVE_P	76
label_ref	85
label_ref and ‘/s’	75
label_ref and ‘/v’	76
label_ref, RTL sharing	114
LABEL_REF_NONLOCAL_P	76
LABEL_STMT	55
LABEL_STMT_LABEL	55
large return values	247
LARGEST_EXPONENT_IS_NORMAL	212
LAST_STACK_REG	223
LAST_VIRTUAL_REGISTER	87
LD_FINI_SWITCH	294
LD_INIT_SWITCH	294
LDD_SUFFIX	295
le	94
le and attributes	177
LE_EXPR	61
leaf functions	222
leaf_function_p	152
LEAF_REG_REMAP	222
LEAF_REGISTERS	222
left rotate	93
left shift	92
LEGITIMATE_CONSTANT_P	264
LEGITIMATE_PIC_OPERAND_P	279
LEGITIMIZE_ADDRESS	262
LEGITIMIZE_RELOAD_ADDRESS	263
length	339
less than	94
less than or equal	94
leu	94
leu and attributes	177
LIB_SPEC	195
LIB2FUNCS_EXTRA	331
LIBCALL_VALUE	246
‘libgcc.a’	258
LIBGCC_NEEDS_DOUBLE	260
LIBGCC_SPEC	195

- LIBGCC2_CFLAGS 331
 - LIBGCC2_WORDS_BIG_ENDIAN 205
 - library subroutine names 258
 - LIBRARY_PATH_ENV 322
 - LIMIT_RELOAD_CLASS 226
 - LINK_COMMAND_SPEC 197
 - LINK_ELIMINATE_DUPLICATE_LDIRECTORIES 197
 - LINK_GCC_C_SEQUENCE_SPEC 197
 - LINK_LIBGCC_SPECIAL 197
 - LINK_LIBGCC_SPECIAL_1 197
 - LINK_SPEC 195
 - linkage 52
 - LINKER_DOES_NOT_WORK_WITH_DWARF2 308
 - list 41
 - lo_sum 90
 - load address instruction 128
 - LOAD_ARGS_REVERSED 245
 - LOAD_EXTEND_OP 316
 - load_multiple instruction pattern 145
 - local register allocation 35
 - LOCAL_ALIGNMENT 208
 - LOCAL_CLASS_P 48
 - LOCAL_INCLUDE_DIR 198
 - LOCAL_LABEL_PREFIX 298
 - LOCAL_REGNO 219
 - LOG_LINKS 108
 - logical-and, bitwise 92
 - logm2 instruction pattern 147
 - LONG_DOUBLE_TYPE_SIZE 214
 - LONG_LONG_TYPE_SIZE 214
 - LONG_TYPE_SIZE 213
 - longjmp and automatic variables 7
 - loop optimization 34
 - LOOP_ALIGN 301
 - LOOP_ALIGN_MAX_SKIP 302
 - LOOP_EXPR 61
 - looping instruction patterns 161
 - LSHIFT_EXPR 61
 - lshiftrt 93
 - lshiftrt and attributes 177
 - lshrm3 instruction pattern 147
 - lt 94
 - lt and attributes 177
 - LT_EXPR 61
 - ltu 94
- M**
- 'm' in constraint 126
 - machine attributes 313
 - machine description macros 193
 - machine descriptions 117
 - machine mode conversions 96
 - machine modes 81
 - machine specific constraints 132
 - MACHINE_DEPENDENT_REORG 322
 - machine_mode 267
 - macros, target description 193
 - MAKE_DECL_ONE_ONLY (*decl*) 289
 - make_safe_from 166
 - makefile fragment 331
 - makefile targets 15
 - marking roots 342
 - MASK_RETURN_ADDR 300
 - match_dup 120, 174
 - match_dup and attributes 180
 - match_insn 123
 - match_insn2 123
 - match_op_dup 122
 - match_operand 119
 - match_operand and attributes 176
 - match_operator 121
 - match_par_dup 123
 - match_parallel 122
 - match_scratch 120, 174
 - matching constraint 128
 - matching operands 124
 - math libraries 8
 - math, in RTL 90
 - MATH_LIBRARY 322
 - MAX_BITS_PER_WORD 205
 - MAX_CONDITIONAL_EXECUTE 322
 - MAX_DFA_ISSUE_RATE 274
 - MAX_FIXED_MODE_SIZE 210
 - MAX_INTEGER_COMPUTATION_MODE 322
 - MAX_LONG_DOUBLE_TYPE_SIZE 214
 - MAX_LONG_TYPE_SIZE 214
 - MAX_MOVE_MAX 316
 - MAX_OFFILE_ALIGNMENT 208
 - MAX_REGS_PER_ADDRESS 261
 - MAX_WCHAR_TYPE_SIZE 215
 - maxm3 instruction pattern 146
 - MAYBE_REG_PARM_STACK_SPACE 240
 - maybe_undef 341
 - mcount 252
 - MD_ASM_CLOBBERS 322
 - MD_CAN_REDIRECT_BRANCH 324
 - MD_EXEC_PREFIX 198
 - MD_FALLBACK_FRAME_STATE_FOR 234
 - MD_HANDLE_UNWABI 234
 - MD_STARTFILE_PREFIX 198
 - MD_STARTFILE_PREFIX_1 198
 - mem 90
 - mem and '/f' 76
 - mem and '/j' 76
 - mem and '/s' 76
 - mem and '/u' 77
 - mem and '/v' 76
 - mem, RTL sharing 114
 - MEM_IN_STRUCT_P 76
 - MEM_KEEP_ALIAS_SET_P 76
 - MEM_SCALAR_P 76
 - MEM_VOLATILE_P 76
 - MEMBER_TYPE_FORCES_BLK 210
 - memcpy, implicit usage 260
 - memmove, implicit usage 260

memory reference, nonoffsettable	130
memory references in constraints	126
MEMORY_MOVE_COST	268
memset, implicit usage	260
METHOD_TYPE	42
MIN_UNITS_PER_WORD	206
MINIMUM_ATOMIC_ALIGNMENT	207
minm3 instruction pattern	146
minus	90
minus and attributes	177
minus, canonicalization of	163
MINUS_EXPR	61
MIPS coprocessor-definition macros	314
mod	92
mod and attributes	177
MODDI3_LIBCALL	259
mode classes	82
mode switching	312
MODE_BASE_REG_CLASS	225
MODE_CC	83
MODE_COMPLEX_FLOAT	83
MODE_COMPLEX_INT	83
MODE_FLOAT	83
MODE_FUNCTION	83
MODE_HAS_INFINITIES	212
MODE_HAS_NANS	211
MODE_HAS_SIGN_DEPENDENT_ROUNDING	212
MODE_HAS_SIGNED_ZEROS	212
MODE_INT	82
MODE_NEEDED	312
MODE_PARTIAL_INT	82
MODE_PRIORITY_TO_MODE	312
MODE_RANDOM	83
MODES_TIEABLE_P	221
modifiers in constraints	131
MODIFY_EXPR	61
MODIFY_JNI_METHOD_CALL	324
MODIFY_TARGET_NAME	198
modm3 instruction pattern	146
MODSI3_LIBCALL	258
MOVE_BY_PIECES_P	270
MOVE_MAX	316
MOVE_MAX_PIECES	270
MOVE_RATIO	270
movm instruction pattern	144
movmodecc instruction pattern	150
movstrictm instruction pattern	145
movstrm instruction pattern	148
MULDI3_LIBCALL	258
mulhisi3 instruction pattern	146
mulm3 instruction pattern	146
mulqihi3 instruction pattern	146
MULSI3_LIBCALL	258
mulside3 instruction pattern	146
mult	91
mult and attributes	177
mult, canonicalization of	163
MULT_EXPR	61
MULTILIB_DEFAULTS	197
MULTILIB_DIRNAMES	332
MULTILIB_EXCEPTIONS	332
MULTILIB_EXTRA_OPTS	332
MULTILIB_MATCHES	332
MULTILIB_OPTIONS	331
multiple alternative constraints	130
MULTIPLE_SYMBOL_SPACES	322
multiplication	91
MUST_PASS_IN_STACK	242
MUST_PASS_IN_STACK, and FUNCTION_ARG	242
N	
'n' in constraint	127
N_REG_CLASSES	224
name	41
named patterns and conditions	118
names, pattern	144
namespace	47
namespace, class, scope	46
NAMESPACE_DECL	47, 49
ne	94
ne and attributes	177
NE_EXPR	61
NEED_ATEXIT	321
neg	91
neg and attributes	177
neg, canonicalization of	163
NEGATE_EXPR	61
negm2 instruction pattern	147
nested functions, trampolines for	256
next_cc0_user	160
NEXT_INSN	104
NEXT_OBJC_RUNTIME	260
nil	72
NO_DBX_FUNCTION_END	307
NO_DOLLAR_IN_LABEL	320
NO_DOT_IN_LABEL	321
NO_FUNCTION_CSE	271
NO_IMPLICIT_EXTERN_C	319
no_new_pseudos	145
NO_PROFILE_COUNTERS	253
NO_RECURSIVE_FUNCTION_CSE	271
NO_REGS	223
NON_SAVING_SETJMP	219
nondeterministic finite state automaton	188
nonlocal_goto instruction pattern	156
nonlocal_goto_receiver instruction pattern	156
nonoffsettable memory reference	130
nop instruction pattern	152
NOP_EXPR	61
NORMAL_MODE	312
not	92
not and attributes	176
not equal	94
not, canonicalization of	163

note 107
 note and ‘/i’ 76, 77
 note and ‘/v’ 75
 NOTE_INSN_BLOCK_BEG 107
 NOTE_INSN_BLOCK_END 107
 NOTE_INSN_DELETED 107
 NOTE_INSN_DELETED_LABEL 107
 NOTE_INSN_EH_REGION_BEG 107
 NOTE_INSN_EH_REGION_END 107
 NOTE_INSN_FUNCTION_END 107
 NOTE_INSN_LOOP_BEG 107
 NOTE_INSN_LOOP_CONT 107
 NOTE_INSN_LOOP_END 107
 NOTE_INSN_LOOP_VTOP 107
 NOTE_INSN_SETJMP 108
 NOTE_LINE_NUMBER 107
 NOTE_SOURCE_FILE 107
 NOTICE_UPDATE_CC 264
 NUM_MACHINE_MODES 83
 NUM_MODES_FOR_MODE_SWITCHING 312

O

‘o’ in constraint 126
 OBJC_GEN_METHOD_LABEL 291
 OBJC_PROLOGUE 280
 OBJECT_FORMAT_COFF 295
 OBJECT_FORMAT_ROSE 295
 OFFSET_TYPE 42
 offsettable address 126
 OImode 81
 old pipeline description 184, 190
 ON_EXIT 321
 one_cmplm2 instruction pattern 148
 operand access 74
 operand constraints 126
 operand substitution 123
 operands 118
 OPTIMIZATION_OPTIONS 203
 OPTIMIZE_MODE_SWITCHING 312
 optional hardware or system features 201
 options, directory search 171
 order of register allocation 220
 ORDER_REGS_FOR_LOCAL_ALLOC 220
 Ordering of Patterns 158
 other register constraints 128
 OUTGOING_REG_PARM_STACK_SPACE 240
 OUTGOING_REGNO 219
 output of assembler code 279
 output statements 124
 output templates 123
 OUTPUT_ADDR_CONST_EXTRA 282
 output_asm_insn 125
 OUTPUT_QUOTED_STRING 280
 OVERLOAD 52
 OVERRIDE_OPTIONS 203
 OVL_CURRENT 52
 OVL_NEXT 52

P

‘p’ in constraint 128
 PAD_VARARGS_DOWN 245
 parallel 100
 param_is 340
 parameters, miscellaneous 314
 paramn_is 341
 PARM_BOUNDARY 207
 PARM_DECL 49
 PARSE_LDD_OUTPUT 295
 parsing pass 31
 passes and files of the compiler 31
 passing arguments 7
 PATH_SEPARATOR 327
 PATTERN 108
 pattern conditions 118
 pattern names 144
 Pattern Ordering 158
 patterns 117
 pc 90
 pc and attributes 181
 pc, RTL sharing 114
 pc_rtx 90
 PCC_BITFIELD_TYPE_MATTERS 209
 PCC_STATIC_STRUCT_RETURN 248
 PDI mode 81
 peephole optimization 36
 peephole optimization, RTL representation ... 101
 peephole optimizer definitions 171
 per-function data 204
 percent sign 123
 PIC 278
 PIC_OFFSET_TABLE_REG_CALL_CLOBBERED 279
 PIC_OFFSET_TABLE_REGNUM 278
 pipeline hazard recognizer 183, 186, 190
 plus 90
 plus and attributes 177
 plus, canonicalization of 163
 PLUS_EXPR 61
 Pmode 318
 pointer 42
 POINTER_SIZE 206
 POINTER_TYPE 42
 POINTERS_EXTEND_UNSIGNED 206
 portability 5
 position independent code 278
 post_dec 103
 post_inc 103
 post_modify 103
 pragma 319, 320
 pre_dec 102
 PRE_GCC3_DWARF_FRAME_REGISTERS 237
 pre_inc 103
 pre_modify 103
 predefined macros 200
 PREDICATE_CODES 314
 predication 190
 PREFERRED_DEBUGGING_TYPE 303

PREFERRED_OUTPUT_RELOAD_CLASS	226
PREFERRED_RELOAD_CLASS	225
PREFERRED_STACK_BOUNDARY	207
prefetch	102
prefetch instruction pattern	158
presence_set	188
PRETEND_OUTGOING_VARARGS_NAMED	255
prev_active_insn	172
prev_cc0_setter	160
PREV_INSN	104
PRINT_OPERAND	297
PRINT_OPERAND_ADDRESS	297
PRINT_OPERAND_PUNCT_VALID_P	297
probe instruction pattern	155
processor functional units	183, 186, 190
processor pipeline description	183
product	91
PROFILE_BEFORE_PROLOGUE	253
PROFILE_HOOK	253
profiling, code generation	252
program counter	90
prologue	248
prologue instruction pattern	157
PROMOTE_FOR_CALL_ONLY	206
PROMOTE_FUNCTION_ARGS	206
PROMOTE_FUNCTION_RETURN	206
PROMOTE_MODE	206
PROMOTE_PROTOTYPES	239
pseudo registers	86
PSImode	81
PTRDIFF_TYPE	215
PTRMEM_CST	61
PTRMEM_CST_CLASS	61
PTRMEM_CST_MEMBER	61
push address instruction	128
PUSH_ARGS	239
push_reload	263
PUSH_ROUNDING	239
PUSH_ROUNDING, interaction with PREFERRED_STACK_BOUNDARY	207
pushm instruction pattern	146
PUT_CODE	71
PUT_MODE	83
PUT_REG_NOTE_KIND	109
PUT_SDB	309

Q

QCmode	82
QFmode	81
QImode	81
QImode, in insn	108
qualified type	42
querying function unit reservations	186
question mark	131
quotient	92

R

'r' in constraint	127
RDIV_EXPR	61
READONLY_DATA_SECTION	275
READONLY_DATA_SECTION_ASM_OP	275
REAL_ARITHMETIC	311
REAL_CST	61
REAL_NM_FILE_NAME	295
REAL_TYPE	42
REAL_VALUE_ABS	311
REAL_VALUE_ATOF	311
REAL_VALUE_FIX	310
REAL_VALUE_FROM_INT	311
REAL_VALUE_ISINF	311
REAL_VALUE_ISNAN	311
REAL_VALUE_NEGATE	311
REAL_VALUE_NEGATIVE	311
REAL_VALUE_TO_INT	311
REAL_VALUE_TO_TARGET_DOUBLE	283
REAL_VALUE_TO_TARGET_LONG_DOUBLE	283
REAL_VALUE_TO_TARGET_SINGLE	283
REAL_VALUE_TRUNCATE	311
REAL_VALUE_TYPE	310
REAL_VALUE_UNSIGNED_FIX	310
REAL_VALUES_EQUAL	310
REAL_VALUES_LESS	310
REALPART_EXPR	61
recog_data.operand	296
recognizing insns	119
RECORD_TYPE	42, 48
reference	42
REFERENCE_TYPE	42
reg	86
reg and '/f'	76
reg and '/i'	76
reg and '/s'	76
reg and '/u'	77
reg and '/v'	76
reg, RTL sharing	114
REG_ALLOC_ORDER	220
REG_BR_PRED	112
REG_BR_PROB	112
REG_CC_SETTER	112
REG_CC_USER	112
REG_CLASS_CONTENTS	224
REG_CLASS_FROM_LETTER	225
REG_CLASS_NAMES	224
REG_DEAD	109
REG_DEP_ANTI	112
REG_DEP_OUTPUT	112
REG_EQUAL	110
REG_EQUIV	110
REG_FRAME_RELATED_EXPR	112
REG_FUNCTION_VALUE_P	76
REG_INC	109
REG_LABEL	110
reg_label and '/v'	76
REG_LIBCALL	112

- REG_LOOP_TEST_P 76
- REG_MODE_OK_FOR_BASE_P 262
- reg_names 297
- REG_NO_CONFLICT 110
- REG_NONNEG 109
- REG_NOTE_KIND 109
- REG_NOTES 108
- REG_OK_FOR_BASE_P 262
- REG_OK_FOR_INDEX_P 262
- REG_OK_STRICT 261
- REG_PARM_STACK_SPACE 240
- REG_PARM_STACK_SPACE, and FUNCTION_ARG ... 242
- REG_POINTER 76
- REG_RETVAL 111
- REG_UNUSED 109
- REG_USERVAR_P 76
- REG_WAS_0 111
- register allocation 35
- register allocation order 220
- register class definitions 223
- register class preference constraints 131
- register class preference pass 35
- register movement 35
- register pairs 221
- Register Transfer Language (RTL) 71
- register usage 218
- register use analysis 33
- register-to-stack conversion 36
- REGISTER_MOVE_COST 268
- REGISTER_NAMES 296
- register_operand 120
- REGISTER_PREFIX 298
- REGISTER_TARGET_PRAGMAS 319
- registers arguments 242
- registers in constraints 127
- REGNO_MODE_OK_FOR_BASE_P 225
- REGNO_OK_FOR_BASE_P 225
- REGNO_OK_FOR_INDEX_P 225
- REGNO_REG_CLASS 224
- regs_ever_live 249
- regular expressions 183, 186
- relative costs 267
- RELATIVE_PREFIX_NOT_LINKDIR 197
- reload pass 88
- reload_completed 152
- reload_in instruction pattern 145
- reload_in_progress 144
- reload_out instruction pattern 145
- reloading 35
- remainder 92
- reordering, block 36
- representation of RTL 71
- reservation delays 183
- rest_of_compilation 31
- rest_of_decl_compilation 31
- restore_stack_block instruction pattern ... 154
- restore_stack_function instruction pattern
..... 154
- restore_stack_nonlocal instruction pattern
..... 154
- RESULT_DECL 49
- return 98
- return instruction pattern 152
- return values in registers 245
- RETURN_ADDR_IN_PREVIOUS_FRAME 232
- RETURN_ADDR_RTX 232
- RETURN_ADDRESS_POINTER_REGNUM 237
- RETURN_EXPR 55
- RETURN_IN_MEMORY 247
- RETURN_INIT 55
- RETURN_POPS_ARGS 241
- RETURN_STMT 55
- returning aggregate values 247
- returning structures and unions 7
- REVERSE_CONDEEXEC_PREDICATES_P 266
- REVERSE_CONDITION (*code, mode*) 266
- REVERSIBLE_CC_MODE 266
- right rotate 93
- right shift 93
- RISC 183, 188
- roots, marking 342
- rotate 93
- rotatert 93
- rotlm3 instruction pattern 147
- rotrm3 instruction pattern 147
- ROUND_TOWARDS_ZERO 212
- ROUND_TYPE_ALIGN 210
- ROUND_TYPE_SIZE 210
- ROUND_TYPE_SIZE_UNIT 210
- RSHIFT_EXPR 61
- RTL addition 90
- RTL addition with signed saturation 91
- RTL addition with unsigned saturation 91
- RTL classes 72
- RTL comparison 91
- RTL comparison operations 93
- RTL constant expression types 84
- RTL constants 84
- RTL declarations 97
- RTL difference 90
- RTL expression 71
- RTL expressions for arithmetic 90
- RTL format 73
- RTL format characters 73
- RTL function-call insns 113
- RTL generation 32
- RTL insn template 119
- RTL integers 71
- RTL memory expressions 86
- RTL object types 71
- RTL postdecrement 102
- RTL postincrement 102
- RTL predecrement 102
- RTL preincrement 102
- RTL register expressions 86
- RTL representation 71

RTL side effect expressions	97	SET_SRC	98
RTL strings	71	SETUP_FRAME_ADDRESSES	231
RTL structure sharing assumptions	114	SETUP_INCOMING_VARARGS	255
RTL subtraction	90	SFmode	81
RTL sum	90	SHARED_BSS_SECTION_ASM_OP	276
RTL vectors	71	SHARED_SECTION_ASM_OP	275
RTX (See RTL)	71	sharing of RTL components	114
RTX codes, classes of	72	shift	92
RTX_COSTS	267	SHIFT_COUNT_TRUNCATED	316
RTX_FRAME_RELATED_P	77	SHORT_IMMEDIATES_SIGN_EXTEND	316
RTX_INTEGRATED_P	77	SHORT_TYPE_SIZE	213
RTX_UNCHANGING_P	77	sibcall_epilogue instruction pattern	157
run-time conventions	7	sibling call optimization	32
run-time target specification	200	SIBLING_CALL_P	78
S			
's' in constraint	127	sign_extend	96
same_type_p	43	sign_extract	95
save_stack_block instruction pattern	154	sign_extract, canonicalization of	164
save_stack_function instruction pattern	154	signed division	92
save_stack_nonlocal instruction pattern	154	signed maximum	92
saveable_obstack	261	signed minimum	92
scalars, returned as values	245	SImode	81
SCHED_GROUP_P	77	simple constraints	126
scheduling, delayed branch	36	simplifications, arithmetic	32
scheduling, instruction	35, 36	Single Static Assignment optimizations	33
SCmode	82	sinm2 instruction pattern	147
scond instruction pattern	150	SIZE_ASM_OP	286
SCOPE_BEGIN_P	55	SIZE_TYPE	215
SCOPE_END_P	55	skip	340
SCOPE_NULLIFIED_P	55	SLOW_BYTE_ACCESS	269
SCOPE_STMT	55	SLOW_UNALIGNED_ACCESS	269
scratch	89	SMALL_ARG_MAX	328
scratch operands	89	SMALL_REGISTER_CLASSES	228
scratch, RTL sharing	114	SMALL_STACK	233
SDB_ALLOW_FORWARD_REFERENCES	309	smax	92
SDB_ALLOW_UNKNOWN_REFERENCES	309	smaxm3 instruction pattern	146
SDB_DEBUGGING_INFO	308	smin	92
SDB_DELIM	309	sminm3 instruction pattern	146
SDB_GENERATE_FAKE	309	smulm3_highpart instruction pattern	146
search options	171	special	341
SECONDARY_INPUT_RELOAD_CLASS	226	SPECIAL_MODE_PREDICATES	315
SECONDARY_MEMORY_NEEDED	227	SPECS	332
SECONDARY_MEMORY_NEEDED_MODE	227	speed of instructions	267
SECONDARY_MEMORY_NEEDED_RTX	227	splitting instructions	167
SECONDARY_OUTPUT_RELOAD_CLASS	226	sqrt	93
SECONDARY_RELOAD_CLASS	226	sqrtm2 instruction pattern	147
SELECT_CC_MODE	265	square root	93
sequence	101	ss_minus	91
set	98	ss_plus	91
set and '/f'	77	ss_truncate	96
SET_ASM_OP	291	SSA Conditional Constant Propagation	33
set_attr	178	SSA DCE	33
set_attr_alternative	179	SSA optimizations	33
SET_DEST	98	stack arguments	239
SET_IS_RETURN_P	78	stack frame layout	230
SET_LABEL_KIND	106	STACK_BOUNDARY	207
		STACK_CHECK_BUILTIN	235
		STACK_CHECK_FIXED_FRAME_SIZE	235
		STACK_CHECK_MAX_FRAME_SIZE	235

- STACK_CHECK_MAX_VAR_SIZE 236
 - STACK_CHECK_PROBE_INTERVAL 235
 - STACK_CHECK_PROBE_LOAD 235
 - STACK_CHECK_PROTECT 235
 - STACK_DYNAMIC_OFFSET 231
 - STACK_DYNAMIC_OFFSET and virtual registers ... 87
 - STACK_GROWS_DOWNWARD 230
 - STACK_PARAMS_IN_REG_PARM_AREA 240
 - STACK_POINTER_OFFSET 231
 - STACK_POINTER_OFFSET and virtual registers ... 87
 - STACK_POINTER_REGNUM 236
 - STACK_POINTER_REGNUM and virtual registers ... 87
 - stack_pointer_rtx 237
 - STACK_PUSH_CODE 230
 - STACK_REGS 223
 - STACK_SAVEAREA_MODE 211
 - STACK_SIZE_MODE 211
 - standard pattern names 144
 - STANDARD_EXEC_PREFIX 198
 - STANDARD_INCLUDE_COMPONENT 199
 - STANDARD_INCLUDE_DIR 199
 - STANDARD_STARTFILE_PREFIX 198
 - STARTFILE_SPEC 196
 - STARTING_FRAME_OFFSET 231
 - STARTING_FRAME_OFFSET and virtual registers .. 87
 - statements 55
 - STATIC_CHAIN 237
 - STATIC_CHAIN_INCOMING 237
 - STATIC_CHAIN_INCOMING_REGNUM 237
 - STATIC_CHAIN_REGNUM 237
 - 'stdarg.h' and register arguments 242
 - STDC_0_IN_SYSTEM_HEADERS 319
 - STMT_EXPR 61
 - STMT_IS_FULL_EXPR_P 56
 - STMT_LINENO 56
 - storage layout 205
 - STORE_FLAG_VALUE 317
 - 'store_multiple' instruction pattern 146
 - strcpy 208
 - strength-reduction 34
 - STRICT_ALIGNMENT 208
 - STRICT_ARGUMENT_NAMING 255
 - strict_low_part 97
 - strict_memory_address_p 263
 - STRING_CST 61
 - STRING_POOL_ADDRESS_P 78
 - strlenm instruction pattern 149
 - STRUCT_VALUE 247
 - STRUCT_VALUE_INCOMING 247
 - STRUCT_VALUE_INCOMING_REGNUM 247
 - STRUCT_VALUE_REGNUM 247
 - structure value address 247
 - STRUCTURE_SIZE_BOUNDARY 208
 - structures, returning 7
 - subm3 instruction pattern 146
 - SUBOBJECT 55
 - SUBOBJECT_CLEANUP 55
 - subreg 88
 - subreg and '/s' 78
 - subreg and '/u' 78
 - subreg and '/u' and '/v' 78
 - subreg, in strict_low_part 97
 - subreg, special reload handling 88
 - SUBREG_BYTE 89
 - SUBREG_PROMOTED_UNSIGNED_P 78
 - SUBREG_PROMOTED_UNSIGNED_SET 78
 - SUBREG_PROMOTED_VAR_P 78
 - SUBREG_REG 89
 - SUCCESS_EXIT_CODE 327
 - SUPPORTS_INIT_PRIORITY 294
 - SUPPORTS_ONE_ONLY 289
 - SUPPORTS_WEAK 288
 - SWITCH_BODY 55
 - SWITCH_COND 55
 - SWITCH_CURTAILS_COMPILATION 194
 - SWITCH_STMT 55
 - SWITCH_TAKES_ARG 193
 - SWITCHES_NEED_SPACES 194
 - symbol_ref 85
 - symbol_ref and '/f' 78
 - symbol_ref and '/i' 78
 - symbol_ref and '/u' 75
 - symbol_ref and '/v' 78
 - symbol_ref, RTL sharing 114
 - SYMBOL_REF_FLAG 78
 - SYMBOL_REF_FLAG, in
TARGET_ENCODE_SECTION_INFO 278
 - SYMBOL_REF_USED 78
 - SYMBOL_REF_WEAK 78
 - symbolic label 114
 - SYSTEM_INCLUDE_DIR 199
- ## T
- 't-target' 331
 - tablejump instruction pattern 153
 - tag 340
 - tagging insns 178
 - tail calls 253
 - tail recursion optimization 32
 - target attributes 313
 - target description macros 193
 - target functions 193
 - target hooks 193
 - target makefile fragment 331
 - target specifications 200
 - target-parameter-dependent code 32
 - TARGET_ASM_ALIGNED_DI_OP 281
 - TARGET_ASM_ALIGNED_HI_OP 281
 - TARGET_ASM_ALIGNED_SI_OP 281
 - TARGET_ASM_ALIGNED_TI_OP 281
 - TARGET_ASM_ASSEMBLE_VISIBILITY 289
 - TARGET_ASM_BYTE_OP 281
 - TARGET_ASM_CLOSE_PAREN 283
 - TARGET_ASM_CONSTRUCTOR 295
 - TARGET_ASM_DESTRUCTOR 295

TARGET_ASM_EH_FRAME_SECTION	301	TARGET_OPTIONS	202
TARGET_ASM_EXCEPTION_SECTION	301	TARGET_OS_CPP_BUILTINS	201
TARGET_ASM_FUNCTION_BEGIN_EPILOGUE	249	TARGET_PTRMEMFUNC_VBIT_LOCATION	216
TARGET_ASM_FUNCTION_END_PROLOGUE	249	TARGET_SCHED_ADJUST_COST	271
TARGET_ASM_FUNCTION_EPILOGUE	249	TARGET_SCHED_ADJUST_PRIORITY	272
TARGET_ASM_FUNCTION_EPILOGUE and trampolines	257	TARGET_SCHED_DFA_BUBBLE	274
TARGET_ASM_FUNCTION_PROLOGUE	248	TARGET_SCHED_DFA_POST_CYCLE_INSN	273
TARGET_ASM_FUNCTION_PROLOGUE and trampolines	257	TARGET_SCHED_DFA_PRE_CYCLE_INSN	273
TARGET_ASM_GLOBALIZE_LABEL	288	TARGET_SCHED_FINISH	273
TARGET_ASM_INTEGER	282	TARGET_SCHED_FIRST_CYCLE_MULTIPASS_DFA_ LOOKAHEAD	273
TARGET_ASM_NAMED_SECTION	281	TARGET_SCHED_INIT	272
TARGET_ASM_OPEN_PAREN	283	TARGET_SCHED_INIT_DFA_BUBBLES	274
TARGET_ASM_OUTPUT_MI_THUNK	251, 252	TARGET_SCHED_INIT_DFA_POST_CYCLE_INSN ...	273
TARGET_ASM_OUTPUT_MI_VCALL_THUNK	252	TARGET_SCHED_INIT_DFA_PRE_CYCLE_INSN ...	273
TARGET_ASM_SELECT_RTX_SECTION	277	TARGET_SCHED_ISSUE_RATE	271
TARGET_ASM_SELECT_SECTION	277	TARGET_SCHED_REORDER	272
TARGET_ASM_UNALIGNED_DI_OP	281	TARGET_SCHED_REORDER2	272
TARGET_ASM_UNALIGNED_HI_OP	281	TARGET_SCHED_USE_DFA_PIPELINE_INTERFACE	273
TARGET_ASM_UNALIGNED_SI_OP	281	TARGET_SCHED_VARIABLE_ISSUE	271
TARGET_ASM_UNALIGNED_TI_OP	281	TARGET_SECTION_TYPE_FLAGS	281
TARGET_ASM_UNIQUE_SECTION	277	TARGET_SET_DEFAULT_TYPE_ATTRIBUTES	313
TARGET_ATTRIBUTE_TABLE	313	TARGET_SWITCHES	202
TARGET_BELL	217	TARGET_TAB	217
TARGET_BINDS_LOCAL_P	278	TARGET_VERSION	203
TARGET_BS	217	TARGET_VT	217
TARGET_CANNOT_MODIFY_JUMPS_P	325	TARGET_VTABLE_DATA_ENTRY_DISTANCE	217
TARGET_COMP_TYPE_ATTRIBUTES	313	TARGET_VTABLE_ENTRY_ALIGN	217
TARGET_CPU_CPP_BUILTINS	200	TARGET_VTABLE_USES_DESCRIPTOR	217
TARGET_CR	217	targetm	193
TARGET_DLLIMPORT_DECL_ATTRIBUTES	313	targets, makefile	15
TARGET_EDOM	259	TCmode	82
TARGET_ENCODE_SECTION_INFO	277	TEMPLATE_DECL	49
TARGET_ENCODE_SECTION_INFO and address validation	261	termination routines	292
TARGET_ENCODE_SECTION_INFO usage	297	text_section	276
TARGET_ESC	217	TEXT_SECTION	275
TARGET_EXECUTABLE_SUFFIX	324	TEXT_SECTION_ASM_OP	275
TARGET_EXPAND_BUILTIN	323	TFmode	82
TARGET_FF	217	THEN_CLAUSE	55
TARGET_FIXED_CONDITION_CODE_REGS	266	THREAD_MODEL_SPEC	196
TARGET_FLOAT_FORMAT	211	THROW_EXPR	61
TARGET_FLT_EVAL_METHOD	214	THUNK_DECL	49
TARGET_FUNCTION_ATTRIBUTE_INLINABLE_P ...	314	THUNK_DELTA	49
TARGET_HAS_F_SETLKW	322	TImode	81
TARGET_HAVE_CTORS_DTORS	295	TImode, in insn	108
TARGET_HAVE_NAMED_SECTIONS	281	'tm.h' macros	193
TARGET_IN_SMALL_DATA_P	278	top level of compiler	31
TARGET_INIT_BUILTINS	323	TQFmode	81
TARGET_INSERT_ATTRIBUTES	313	TRADITIONAL_PIPELINE_INTERFACE	274
TARGET_MEM_FUNCTIONS	260	TRAMPOLINE_ADJUST_ADDRESS	256
TARGET_MERGE_DECL_ATTRIBUTES	313	TRAMPOLINE_ALIGNMENT	256
TARGET_MERGE_TYPE_ATTRIBUTES	313	TRAMPOLINE_SECTION	256
TARGET_MS_BITFIELD_LAYOUT_P	213	TRAMPOLINE_SIZE	256
TARGET_NEWLINE	217	TRAMPOLINE_TEMPLATE	256
TARGET_OBJECT_SUFFIX	324	trampolines for nested functions	256
TARGET_OPTION_TRANSLATE_TABLE	194	TRANSFER_FROM_TRAMPOLINE	258
		trap instruction pattern	158

- tree 39, 40
 - Tree optimization 32
 - TREE_CODE 39
 - tree_int_cst_equal 61
 - TREE_INT_CST_HIGH 61
 - TREE_INT_CST_LOW 61
 - tree_int_cst_lt 61
 - TREE_LIST 41
 - TREE_OPERAND 61
 - TREE_PUBLIC 52, 53
 - TREE_PURPOSE 41
 - TREE_STRING_LENGTH 61
 - TREE_STRING_POINTER 61
 - TREE_TYPE 42, 49, 55, 61
 - TREE_VALUE 41
 - TREE_VEC 41
 - TREE_VEC_ELT 41
 - TREE_VEC_LENGTH 41
 - TREE_VIA_PRIVATE 48
 - TREE_VIA_PROTECTED 48
 - TREE_VIA_PUBLIC 48
 - Trees 39
 - TRULY_NOOP_TRUNCATION 317
 - TRUNC_DIV_EXPR 61
 - TRUNC_MOD_EXPR 61
 - truncate 96
 - truncm2 instruction pattern 149
 - TRUTH_AND_EXPR 61
 - TRUTH_ANDIF_EXPR 61
 - TRUTH_NOT_EXPR 61
 - TRUTH_OR_EXPR 61
 - TRUTH_ORIF_EXPR 61
 - TRUTH_XOR_EXPR 61
 - TRY_BLOCK 55
 - TRY_HANDLERS 55
 - TRY_STMTS 55
 - tstm instruction pattern 148
 - type 42
 - type declaration 49
 - TYPE_ALIGN 42
 - TYPE_ARG_TYPES 42
 - TYPE_ASM_OP 286
 - TYPE_ATTRIBUTES 61
 - TYPE_BINFO 48
 - TYPE_BUILT_IN 43
 - TYPE_CONTEXT 42
 - TYPE_DECL 49
 - TYPE_FIELDS 42, 48
 - TYPE_HAS_ARRAY_NEW_OPERATOR 49
 - TYPE_HAS_DEFAULT_CONSTRUCTOR 49
 - TYPE_HAS_MUTABLE_P 49
 - TYPE_HAS_NEW_OPERATOR 49
 - TYPE_MAIN_VARIANT 42
 - TYPE_MAX_VALUE 42
 - TYPE_METHOD_BASETYPE 42
 - TYPE_METHODS 48
 - TYPE_MIN_VALUE 42
 - TYPE_NAME 42, 43
 - TYPE_NOTHROW_P 55
 - TYPE_OFFSET_BASETYPE 42
 - TYPE_OPERAND_FMT 286
 - TYPE_OVERLOADS_ARRAY_REF 49
 - TYPE_OVERLOADS_ARROW 49
 - TYPE_OVERLOADS_CALL_EXPR 49
 - TYPE_POLYMORPHIC_P 49
 - TYPE_PRECISION 42
 - TYPE_PTR_P 43
 - TYPE_PTRFN_P 43
 - TYPE_PTRMEM_P 42, 43
 - TYPE_PTROB_P 43
 - TYPE_PTROBV_P 42
 - TYPE_QUAL_CONST 42
 - TYPE_QUAL_RESTRICT 42
 - TYPE_QUAL_VOLATILE 42
 - TYPE_RAISES_EXCEPTIONS 55
 - TYPE_SIZE 42
 - TYPE_UNQUALIFIED 42
 - TYPE_VFIELD 48
 - TYPENAME_TYPE 42
 - TYPENAME_TYPE_FULLNAME 42
 - TYPEOF_TYPE 42
- ## U
- udiv 92
 - UDIVDI3_LIBCALL 259
 - udivm3 instruction pattern 146
 - udivmodm4 instruction pattern 147
 - UDIVSI3_LIBCALL 258
 - UINTMAX_TYPE 216
 - umax 92
 - umaxm3 instruction pattern 146
 - umin 92
 - uminm3 instruction pattern 146
 - umod 92
 - UMODDI3_LIBCALL 259
 - umodm3 instruction pattern 146
 - UMODSI3_LIBCALL 258
 - umulhisi3 instruction pattern 146
 - umulm3_highpart instruction pattern 146
 - umulqih3 instruction pattern 146
 - umulsidi3 instruction pattern 146
 - unchanging 80
 - unchanging, in call_insn 75
 - unchanging, in jump_insn, call_insn and insn
..... 75
 - unchanging, in reg and mem 77
 - unchanging, in subreg 78
 - unchanging, in symbol_ref 75
 - UNION_TYPE 42, 48
 - unions, returning 7
 - UNITS_PER_WORD 205
 - UNKNOWN_FLOAT_FORMAT 211
 - UNKNOWN_TYPE 42
 - UNLIKELY_EXECUTED_TEXT_SECTION_NAME 275
 - unreachable code 33

unshare_all_rtl	114
unsigned division	92
unsigned greater than	94
unsigned less than	94
unsigned minimum and maximum	92
unsigned_fix	97
unsigned_float	97
unspec	102
unspec_volatile	102
untyped_call instruction pattern	152
untyped_return instruction pattern	152
UPDATE_PATH_HOST_CANONICALIZE (<i>path</i>)	328
us_minus	91
us_plus	91
us_truncate	96
use	100
USE_C_ALLOCA	327
USE_LOAD_POST_DECREMENT	270
USE_LOAD_POST_INCREMENT	270
USE_LOAD_PRE_DECREMENT	270
USE_LOAD_PRE_INCREMENT	270
use_param	340
use_paramn	341
use_params	341
USE_STORE_POST_DECREMENT	270
USE_STORE_POST_INCREMENT	270
USE_STORE_PRE_DECREMENT	271
USE_STORE_PRE_INCREMENT	271
used	80
used, in symbol_ref	78
USER_LABEL_PREFIX	298
USING_DECL	49
USING_STMT	55

V

'V' in constraint	127
VA_ARG_EXPR	61
values, returned by functions	245
VAR_DECL	49, 61
varargs implementation	253
variable	49
VAX_FLOAT_FORMAT	211
vec_concat	95
vec_const	95
vec_duplicate	96
vec_merge	95
vec_select	95
vector	41
vector operations	95
VECTOR_CST	61
VECTOR_MODE_SUPPORTED_P	210
VIRTUAL_INCOMING_ARGS_REGNUM	87
VIRTUAL_OUTGOING_ARGS_REGNUM	87
VIRTUAL_STACK_DYNAMIC_REGNUM	87
VIRTUAL_STACK_VARS_REGNUM	87
VLIW	183, 188
VMS	327

VMS_DEBUGGING_INFO	310
VOID_TYPE	42
VOIDmode	82
volatil	80
volatil, in insn, call_insn, jump_insn, code_label, barrier, and note	75
volatil, in label_ref and reg_label	76
volatil, in mem, asm_operands, and asm_input	76
volatil, in reg	76
volatil, in subreg	78
volatil, in symbol_ref	78
volatile memory references	80
voting between constraint alternatives	131
VTABLE_REF	61

W

WCHAR_TYPE	215
WCHAR_TYPE_SIZE	215
which_alternative	125
WHILE_BODY	55
WHILE_COND	55
WHILE_STMT	55
WIDEST_HARDWARE_FP_SIZE	214
WINT_TYPE	215
word_mode	84
WORD_REGISTER_OPERATIONS	316
WORD_SWITCH_TAKES_ARG	193
WORDS_BIG_ENDIAN	205
WORDS_BIG_ENDIAN, effect on subreg	88

X

'X' in constraint	127
'x-host'	333
XCmode	82
XCOFF_DEBUGGING_INFO	304
XEXP	74
XFmode	81
XINT	74
'xm-machine.h'	327
xor	92
xor, canonicalization of	164
xorm3 instruction pattern	146
XSTR	74
XVEC	74
XVECEXP	74
XVECLEN	74
XWINT	74

Z

zero_extend	96
zero_extendmn2 instruction pattern	149
zero_extract	95
zero_extract, canonicalization of	164