

```

*****
43526 Sun Oct 28 20:56:06 2012
new/gcc/common.opt
Implement -fstrict-calling-conventions
Stock GCC is overly willing to violate the ABI when calling local functions,
such that it passes arguments in registers on i386. This hampers debugging
with anything other than a fully-aware DWARF debugger, and is generally not
something we desire.
Implement a flag which disables this behaviour, enabled by default. The flag is
global, though only effective on i386, to more easily allow its globalization
later which, given the odds, is likely to be necessary.
*****
1 ; Options for the language- and target-independent parts of the compiler.

3 ; Copyright (C) 2003, 2004, 2005, 2006, 2007, 2008, 2009
4 ; Free Software Foundation, Inc.
5 ;
6 ; This file is part of GCC.
7 ;
8 ; GCC is free software; you can redistribute it and/or modify it under
9 ; the terms of the GNU General Public License as published by the Free
10 ; Software Foundation; either version 3, or (at your option) any later
11 ; version.
12 ;
13 ; GCC is distributed in the hope that it will be useful, but WITHOUT ANY
14 ; WARRANTY; without even the implied warranty of MERCHANTABILITY or
15 ; FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License
16 ; for more details.
17 ;
18 ; You should have received a copy of the GNU General Public License
19 ; along with GCC; see the file COPYING3. If not see
20 ; <http://www.gnu.org/licenses/>.

22 ; See the GCC internals manual (options.texi) for a description of this file's f

24 ; Please try to keep this file in ASCII collating order.

26 -help
27 Common
28 Display this information

30 -help=
31 Common Report Joined
32 --help=<class> Display descriptions of a specific class of options. <class> is

34 -target-help
35 Common
36 Alias for --help=target

38 ;; The following three entries are to work around the gcc driver
39 ;; program's insatiable desire to turn options starting with a
40 ;; double dash (--) into options starting with a dash f (-f).
41 fhhelp
42 Common

44 fhhelp=
45 Common Joined

47 ftarget-help
48 Common

50 -param
51 Common Separate
52 --param <param>=<value> Set parameter <param> to value. See below for a complet

54 -version

```

```

55 Common

57 G
58 Common Joined Separate UInteger
59 -G<number> Put global and static data smaller than <number> bytes into a sp

61 O
62 Common JoinedOrMissing Optimization
63 -O<number> Set optimization level to <number>

65 Os
66 Common Optimization
67 Optimize for space rather than speed

69 W
70 Common RejectNegative
71 This switch is deprecated; use -Wextra instead

73 Waggregate-return
74 Common Var(warn_aggregate_return) Warning
75 Warn about returning structures, unions or arrays

77 Warray-bounds
78 Common Var(warn_array_bounds) Warning
79 Warn if an array is accessed out of bounds

81 Wattributes
82 Common Var(warn_attributes) Init(1) Warning
83 Warn about inappropriate attribute usage

85 Wcast-align
86 Common Var(warn_cast_align) Warning
87 Warn about pointer casts which increase alignment

89 Wdeprecated-declarations
90 Common Var(warn_deprecated_decl) Init(1) Warning
91 Warn about uses of __attribute__((deprecated)) declarations

93 Wdisabled-optimization
94 Common Var(warn_disabled_optimization) Warning
95 Warn when an optimization pass is disabled

97 Werror
98 Common Var(warnings_are_errors)
99 Treat all warnings as errors

101 Werror=
102 Common Joined
103 Treat specified warning as error

105 Wextra
106 Common Warning
107 Print extra (possibly unwanted) warnings

109 Wfatal-errors
110 Common Var(flag_fatal_errors)
111 Exit on the first error occurred

113 Wframe-larger-than=
114 Common RejectNegative Joined UInteger
115 -Wframe-larger-than=<number> Warn if a function's stack frame requires more than

117 Winline
118 Common Var(warn_inline) Warning
119 Warn when an inlined function cannot be inlined

```

121 Wlarger-than-  
 122 Common RejectNegative Joined UInteger Warning

124 Wlarger-than=  
 125 Common RejectNegative Joined UInteger Warning  
 126 -Wlarger-than=<number> Warn if an object is larger than <number> bytes

128 Wlogical-op  
 129 Common Warning Var(warn\_logical\_op)  
 130 Warn when a logical operator is suspiciously always evaluating to true or false

132 Wunsafe-loop-optimizations  
 133 Common Var(warn\_unsafe\_loop\_optimizations) Warning  
 134 Warn if the loop cannot be optimized due to nontrivial assumptions.

136 Wmissing-noreturn  
 137 Common Var(warn\_missing\_noreturn) Warning  
 138 Warn about functions which might be candidates for \_\_attribute\_\_((noreturn))

140 Wmudflap  
 141 Common Var(warn\_mudflap) Init(1) Warning  
 142 Warn about constructs not instrumented by -fmudflap

144 Woverflow  
 145 Common Var(warn\_overflow) Init(1) Warning  
 146 Warn about overflow in arithmetic expressions

148 Wpacked  
 149 Common Var(warn\_packed) Warning  
 150 Warn when the packed attribute has no effect on struct layout

152 Wpadded  
 153 Common Var(warn\_padded) Warning  
 154 Warn when padding is required to align structure members

156 Wshadow  
 157 Common Var(warn\_shadow) Warning  
 158 Warn when one local variable shadows another

160 Wstack-protector  
 161 Common Var(warn\_stack\_protect) Warning  
 162 Warn when not issuing stack smashing protection for some reason

164 Wstrict-aliasing  
 165 Common Warning  
 166 Warn about code which might break strict aliasing rules

168 Wstrict-aliasing=  
 169 Common Joined UInteger Var(warn\_strict\_aliasing) Init(-1) Warning  
 170 Warn about code which might break strict aliasing rules

172 Wstrict-overflow  
 173 Common Warning  
 174 Warn about optimizations that assume that signed overflow is undefined

176 Wstrict-overflow=  
 177 Common Joined UInteger Var(warn\_strict\_overflow) Init(-1) Warning  
 178 Warn about optimizations that assume that signed overflow is undefined

180 Wswitch  
 181 Common Var(warn\_switch) Warning  
 182 Warn about enumerated switches, with no default, missing a case

184 Wswitch-default  
 185 Common Var(warn\_switch\_default) Warning  
 186 Warn about enumerated switches missing a \"default:\" statement

188 Wswitch-enum  
 189 Common Var(warn\_switch\_enum) Warning  
 190 Warn about all enumerated switches missing a specific case

192 Wsystem-headers  
 193 Common Var(warn\_system\_headers) Warning  
 194 Do not suppress warnings from system headers

196 Wtype-limits  
 197 Common Var(warn\_type\_limits) Init(-1) Warning  
 198 Warn if a comparison is always true or always false due to the limited range of

200 Wuninitialized  
 201 Common Var(warn\_uninitialized) Warning  
 202 Warn about uninitialized automatic variables

204 Wunreachable-code  
 205 Common Var(warn\_notreached) Warning  
 206 Warn about code that will never be executed

208 Wunused  
 209 Common Var(warn\_unused) Init(0) Warning  
 210 Enable all -Wunused- warnings

212 Wunused-function  
 213 Common Var(warn\_unused\_function) Init(-1) Warning  
 214 Warn when a function is unused

216 Wunused-label  
 217 Common Var(warn\_unused\_label) Init(-1) Warning  
 218 Warn when a label is unused

220 Wunused-parameter  
 221 Common Var(warn\_unused\_parameter) Init(-1) Warning  
 222 Warn when a function parameter is unused

224 Wunused-value  
 225 Common Var(warn\_unused\_value) Init(-1) Warning  
 226 Warn when an expression value is unused

228 Wunused-variable  
 229 Common Var(warn\_unused\_variable) Init(-1) Warning  
 230 Warn when a variable is unused

232 Wcoverage-mismatch  
 233 Common RejectNegative Var(warn\_coverage\_mismatch) Warning  
 234 Warn instead of error in case profiles in -fprofile-use do not match

236 aux-info  
 237 Common Separate  
 238 -aux-info <file> Emit declaration information into <file>

240 aux-info=  
 241 Common Joined

243 auxbase  
 244 Common Separate

246 auxbase-strip  
 247 Common Separate

249 d  
 250 Common Joined  
 251 -d<letters> Enable dumps from specific passes of the compiler

```

253 dumpbase
254 Common Separate
255 -dumpbase <file>          Set the file basename to be used for dumps

257 ; The version of the C++ ABI in use.  The following values are allowed:
258 ;
259 ; 0: The version of the ABI believed most conformant with the C++ ABI
260 ; specification.  This ABI may change as bugs are discovered and fixed.
261 ; Therefore, 0 will not necessarily indicate the same ABI in different
262 ; versions of G++.
263 ;
264 ; 1: The version of the ABI first used in G++ 3.2.
265 ;
266 ; 2: The version of the ABI first used in G++ 3.4.
267 ;
268 ; Additional positive integers will be assigned as new versions of
269 ; the ABI become the default version of the ABI.
270 fabi-version=
271 Common Joined UInteger Var(flag_abi_version) Init(2)

273 falign-functions
274 Common Report Var(align_functions,0) Optimization UInteger
275 Align the start of functions

277 falign-functions=
278 Common RejectNegative Joined UInteger

280 falign-jumps
281 Common Report Var(align_jumps,0) Optimization UInteger
282 Align labels which are only reached by jumping

284 falign-jumps=
285 Common RejectNegative Joined UInteger

287 falign-labels
288 Common Report Var(align_labels,0) Optimization UInteger
289 Align all labels

291 falign-labels=
292 Common RejectNegative Joined UInteger

294 falign-loops
295 Common Report Var(align_loops) Optimization UInteger
296 Align the start of loops

298 falign-loops=
299 Common RejectNegative Joined UInteger

301 ; This flag is only tested if alias checking is enabled.
302 ; 0 if pointer arguments may alias each other.  True in C.
303 ; 1 if pointer arguments may not alias each other but may alias
304 ; global variables.
305 ; 2 if pointer arguments may not alias each other and may not
306 ; alias global variables.
307 ; 3 if pointer arguments may not alias anything.  True in Fortran.
308 ; Set by the front end.
309 fargument-alias
310 Common Report Var(flag_argument_noalias,0) Optimization
311 Specify that arguments may alias each other and globals

313 fargument-noalias
314 Common Report Var(flag_argument_noalias,1) VarExists Optimization
315 Assume arguments may alias globals but not each other

317 fargument-noalias-global
318 Common Report Var(flag_argument_noalias,2) VarExists Optimization

```

```

319 Assume arguments alias neither each other nor globals

321 fargument-noalias-anything
322 Common Report Var(flag_argument_noalias,3) VarExists Optimization
323 Assume arguments alias no other storage

325 fasynchronous-unwind-tables
326 Common Report Var(flag_asynchronous_unwind_tables) Optimization
327 Generate unwind tables that are exact at each instruction boundary

329 fauto-inc-dec
330 Common Report Var(flag_auto_inc_dec) Init(1)
331 Generate auto-inc/dec instructions

333 ; -fcheck-bounds causes gcc to generate array bounds checks.
334 ; For C, C++ and ObjC: defaults off.
335 ; For Java: defaults to on.
336 ; For Fortran: defaults to off.
337 fbounds-check
338 Common Report Var(flag_bounds_check)
339 Generate code to check bounds before indexing arrays

341 fbranch-count-reg
342 Common Report Var(flag_branch_on_count_reg) Init(1) Optimization
343 Replace add, compare, branch with branch on count register

345 fbranch-probabilities
346 Common Report Var(flag_branch_probabilities) Optimization
347 Use profiling information for branch probabilities

349 fbranch-target-load-optimize
350 Common Report Var(flag_branch_target_load_optimize) Optimization
351 Perform branch target load optimization before prologue / epilogue threading

353 fbranch-target-load-optimize2
354 Common Report Var(flag_branch_target_load_optimize2) Optimization
355 Perform branch target load optimization after prologue / epilogue threading

357 fbtr-bb-exclusive
358 Common Report Var(flag_btr_bb_exclusive) Optimization
359 Restrict target load migration not to re-use registers in any basic block

361 fcall-saved-
362 Common Joined RejectNegative
363 -fcall-saved-<register> Mark <register> as being preserved across functions

365 fcall-used-
366 Common Joined RejectNegative
367 -fcall-used-<register> Mark <register> as being corrupted by function calls

369 ; Nonzero for -fcaller-saves: allocate values in regs that need to
370 ; be saved across function calls, if that produces overall better code.
371 ; Optional now, so people can test it.
372 fcaller-saves
373 Common Report Var(flag_caller_saves) Optimization
374 Save registers around function calls

376 fcheck-data-deps
377 Common Report Var(flag_check_data_deps)
378 Compare the results of several data dependence analyzers.

380 fcommon
381 Common Report Var(flag_no_common,0) Optimization
382 Do not put uninitialized globals in the common section

384 fconserve-stack

```

```

385 Common Var(flag_conserve_stack) Optimization
386 Do not perform optimizations increasing noticeably stack usage

388 fcprop-registers
389 Common Report Var(flag_cprop_registers) Optimization
390 Perform a register copy-propagation optimization pass

392 fcrossjumping
393 Common Report Var(flag_crossjumping) Optimization
394 Perform cross-jumping optimization

396 fcse-follow-jumps
397 Common Report Var(flag_cse_follow_jumps) Optimization
398 When running CSE, follow jumps to their targets

400 fcse-skip-blocks
401 Common Report Var(flag_cse_skip_blocks) Optimization
402 When running CSE, follow conditional jumps

404 fcx-limited-range
405 Common Report Var(flag_cx_limited_range) Optimization
406 Omit range reduction step when performing complex division

408 fcx-fortran-rules
409 Common Report Var(flag_cx_fortran_rules) Optimization
410 Complex multiplication and division follow Fortran rules

412 fdata-sections
413 Common Report Var(flag_data_sections) Optimization
414 Place data items into their own section

416 fdbg-cnt-list
417 Common Report
418 List all available debugging counters with their limits and counts.

420 fdbg-cnt=
421 Common RejectNegative Joined
422 -fdbg-cnt=<counter>:<limit>[,<counter>:<limit>,...] Set the debug counter lim

424 fdebug-prefix-map=
425 Common Joined RejectNegative
426 Map one directory name to another in debug information

428 ; Nonzero for -fdefer-pop: don't pop args after each function call
429 ; instead save them up to pop many calls' args with one insns.
430 fdefer-pop
431 Common Report Var(flag_defer_pop) Optimization
432 Defer popping functions args from stack until later

434 fdelayed-branch
435 Common Report Var(flag_delayed_branch) Optimization
436 Attempt to fill delay slots of branch instructions

438 fdelete-null-pointer-checks
439 Common Report Var(flag_delete_null_pointer_checks) Optimization
440 Delete useless null pointer checks

442 fdiagnostics-show-location=
443 Common Joined RejectNegative
444 -fdiagnostics-show-location=[once|every-line] How often to emit source locatio

446 fdiagnostics-show-option
447 Common
448 Amend appropriate diagnostic messages with the command line option that controls

450 fdump-
```

```

451 Common Joined RejectNegative
452 -fdump-<type> Dump various compiler internals to a file

454 fdump-noaddr
455 Common Report Var(flag_dump_noaddr)
456 Suppress output of addresses in debugging dumps

458 fdump-unnumbered
459 Common Report Var(flag_dump_unnumbered) VarExists
460 Suppress output of instruction numbers, line number notes and addresses in debug

462 fdwarf2-cfi-asm
463 Common Report Var(flag_dwarf2_cfi_asm) Init(HAVE_GAS_CFI_DIRECTIVE)
464 Enable CFI tables via GAS assembler directives.

466 fearly-inlining
467 Common Report Var(flag_early_inlining) Init(1) Optimization
468 Perform early inlining

470 feliminate-dwarf2-dups
471 Common Report Var(flag_eliminate_dwarf2_dups)
472 Perform DWARF2 duplicate elimination

474 feliminate-unused-debug-symbols
475 Common Report Var(flag_debug_only_used_symbols)
476 Perform unused type elimination in debug info

478 feliminate-unused-debug-types
479 Common Report Var(flag_eliminate_unused_debug_types) Init(1)
480 Perform unused type elimination in debug info

482 femit-class-debug-always
483 Common Report Var(flag_emit_class_debug_always) Init(0)
484 Do not suppress C++ class debug information.

486 fexceptions
487 Common Report Var(flag_exceptions) Optimization
488 Enable exception handling

490 fexpensive-optimizations
491 Common Report Var(flag_expensive_optimizations) Optimization
492 Perform a number of minor, expensive optimizations

494 ffast-math
495 Common

497 ffinite-math-only
498 Common Report Var(flag_finite_math_only) Optimization
499 Assume no NaNs or infinities are generated

501 ffixed-
502 Common Joined RejectNegative
503 -ffixed-<register> Mark <register> as being unavailable to the compiler

505 ffloat-store
506 Common Report Var(flag_float_store) Optimization
507 Don't allocate floats and doubles in extended-precision registers

509 fforce-addr
510 Common
511 Does nothing. Preserved for backward compatibility.

513 fforward-propagate
514 Common Report Var(flag_forward_propagate) Optimization
515 Perform a forward propagation pass on RTL
```

```

517 ; Nonzero means don't put addresses of constant functions in registers.
518 ; Used for compiling the Unix kernel, where strange substitutions are
519 ; done on the assembly output.
520 ffunction-cse
521 Common Report Var(flag_no_function_cse,0)
522 Allow function addresses to be held in registers

524 ffunction-sections
525 Common Report Var(flag_function_sections)
526 Place each function into its own section

528 fgcse
529 Common Report Var(flag_gcse) Optimization
530 Perform global common subexpression elimination

532 fgcse-lm
533 Common Report Var(flag_gcse_lm) Init(1) Optimization
534 Perform enhanced load motion during global common subexpression elimination

536 fgcse-sm
537 Common Report Var(flag_gcse_sm) Init(0) Optimization
538 Perform store motion after global common subexpression elimination

540 fgcse-las
541 Common Report Var(flag_gcse_las) Init(0) Optimization
542 Perform redundant load after store elimination in global common subexpression
543 elimination

545 fgcse-after-reload
546 Common Report Var(flag_gcse_after_reload) Optimization
547 Perform global common subexpression elimination after register allocation
548 has finished

550 ; This option is not documented yet as its semantics will change.
551 fgraphite
552 Common Report Var(flag_graphite)
553 Enable in and out of Graphite representation

555 floop-strip-mine
556 Common Report Var(flag_loop_strip_mine) Optimization
557 Enable Loop Strip Mining transformation

559 floop-interchange
560 Common Report Var(flag_loop_interchange) Optimization
561 Enable Loop Interchange transformation

563 floop-block
564 Common Report Var(flag_loop_block) Optimization
565 Enable Loop Blocking transformation

567 ; This option is not documented as it does not perform any useful optimization.
568 fgraphite-identity
569 Common Report Var(flag_graphite_identity) Optimization
570 Enable Graphite Identity transformation

572 fguess-branch-probability
573 Common Report Var(flag_guess_branch_prob) Optimization
574 Enable guessing of branch probabilities

576 ; Nonzero means ignore '#ident' directives. 0 means handle them.
577 ; Generate position-independent code for executables if possible
578 ; On SVR4 targets, it also controls whether or not to emit a
579 ; string identifying the compiler.
580 fident
581 Common Report Var(flag_no_ident,0)
582 Process #ident directives

```

```

584 fif-conversion
585 Common Report Var(flag_if_conversion) Optimization
586 Perform conversion of conditional jumps to branchless equivalents

588 fif-conversion2
589 Common Report Var(flag_if_conversion2) Optimization
590 Perform conversion of conditional jumps to conditional execution

592 ; -finhibit-size-directive inhibits output of .size for ELF.
593 ; This is used only for compiling crtstuff.c,
594 ; and it may be extended to other effects
595 ; needed for crtstuff.c on other systems.
596 finhibit-size-directive
597 Common Report Var(flag_inhibit_size_directive)
598 Do not generate .size directives

600 findirect-inlining
601 Common Report Var(flag_indirect_inlining)
602 Perform indirect inlining

604 ; Nonzero means that functions declared 'inline' will be treated
605 ; as 'static'. Prevents generation of zillions of copies of unused
606 ; static inline functions; instead, 'inlines' are written out
607 ; only when actually used. Used in conjunction with -g. Also
608 ; does the right thing with #pragma interface.
609 finline
610 Common Report Var(flag_no_inline,0) Init(0)
611 Pay attention to the "\"inline\" keyword

613 finline-small-functions
614 Common Report Var(flag_inline_small_functions) Optimization
615 Integrate simple functions into their callers when code size is known to not gro

617 finline-functions
618 Common Report Var(flag_inline_functions) Optimization
619 Integrate simple functions into their callers

621 finline-functions-called-once
622 Common Report Var(flag_inline_functions_called_once) Init(1) Optimization
623 Integrate functions called once into their callers

625 finline-limit=
626 Common RejectNegative Joined UInteger

628 finline-limit=
629 Common RejectNegative Joined UInteger
630 -finline-limit=<number> Limit the size of inlined functions to <number>

632 finstrument-functions
633 Common Report Var(flag_instrument_function_entry_exit)
634 Instrument function entry and exit with profiling calls

636 finstrument-functions-exclude-function-list=
637 Common RejectNegative Joined
638 -finstrument-functions-exclude-function-list=name,... Do not instrument listed

640 finstrument-functions-exclude-file-list=
641 Common RejectNegative Joined
642 -finstrument-functions-exclude-file-list=filename,... Do not instrument functio

644 fipa-cp
645 Common Report Var(flag_ipa_cp) Optimization
646 Perform Interprocedural constant propagation

648 fipa-cp-clone

```

```

649 Common Report Var(flag_ipa_cp_clone) Optimization
650 Perform cloning to make Interprocedural constant propagation stronger

652 fipa-pure-const
653 Common Report Var(flag_ipa_pure_const) Init(0) Optimization
654 Discover pure and const functions

656 fipa-pta
657 Common Report Var(flag_ipa_pta) Init(0) Optimization
658 Perform interprocedural points-to analysis

660 fipa-reference
661 Common Report Var(flag_ipa_reference) Init(0) Optimization
662 Discover readonly and non addressable static variables

664 fipa-type-escape
665 Common Report Var(flag_ipa_type_escape) Init(0) Optimization
666 Type based escape and alias analysis

668 fipa-matrix-reorg
669 Common Report Var(flag_ipa_matrix_reorg) Optimization
670 Perform matrix layout flattening and transposing based
671 on profiling information.

673 fipa-struct-reorg
674 Common Report Var(flag_ipa_struct_reorg)
675 Perform structure layout optimizations based
676 on profiling information.

678 fira-algorithm=
679 Common Joined RejectNegative
680 -fira-algorithm=[CB|priority] Set the used IRA algorithm

682 fira-region=
683 Common Joined RejectNegative
684 -fira-region=[one|all|mixed] Set regions for IRA

686 fira-coalesce
687 Common Report Var(flag_ira_coalesce) Init(0)
688 Do optimistic coalescing.

690 fira-share-save-slots
691 Common Report Var(flag_ira_share_save_slots) Init(1)
692 Share slots for saving different hard registers.

694 fira-share-spill-slots
695 Common Report Var(flag_ira_share_spill_slots) Init(1)
696 Share stack slots for spilled pseudo-registers.

698 fira-verbose=
699 Common RejectNegative Joined UInteger
700 -fira-verbose=<number> Control IRA's level of diagnostic messages.

702 fivopts
703 Common Report Var(flag_ivopts) Init(1) Optimization
704 Optimize induction variables on trees

706 fjump-tables
707 Common Var(flag_jump_tables) Init(1) Optimization
708 Use jump tables for sufficiently large switch statements

710 fkeep-inline-functions
711 Common Report Var(flag_keep_inline_functions)
712 Generate code for functions even if they are fully inlined

714 fkeep-static-consts

```

```

715 Common Report Var(flag_keep_static_consts) Init(1)
716 Emit static const variables even if they are not used

718 fleading-underscore
719 Common Report Var(flag_leading_underscore) Init(-1)
720 Give external symbols a leading underscore

722 floop-optimize
723 Common
724 Does nothing. Preserved for backward compatibility.

726 fmath-errno
727 Common Report Var(flag_errno_math) Init(1) Optimization
728 Set errno after built-in math functions

730 fmem-report
731 Common Report Var(mem_report)
732 Report on permanent memory allocation

734 ; This will attempt to merge constant section constants, if 1 only
735 ; string constants and constants from constant pool, if 2 also constant
736 ; variables.
737 fmerge-all-constants
738 Common Report Var(flag_merge_constants,2) Init(1) Optimization
739 Attempt to merge identical constants and constant variables

741 fmerge-constants
742 Common Report Var(flag_merge_constants,1) VarExists Optimization
743 Attempt to merge identical constants across compilation units

745 fmerge-debug-strings
746 Common Report Var(flag_merge_debug_strings) Init(1)
747 Attempt to merge identical debug strings across compilation units

749 fmessage-length=
750 Common RejectNegative Joined UInteger
751 -fmessage-length=<number> Limit diagnostics to <number> characters per lin

753 fmodulo-sched
754 Common Report Var(flag_modulo_sched) Optimization
755 Perform SMS based modulo scheduling before the first scheduling pass

757 fmodulo-sched-allow-regmoves
758 Common Report Var(flag_modulo_sched_allow_regmoves)
759 Perform SMS based modulo scheduling with register moves allowed

761 fmove-loop-invariants
762 Common Report Var(flag_move_loop_invariants) Init(1) Optimization
763 Move loop invariant computations out of loops

765 fmodflap
766 Common RejectNegative Report Var(flag_modflap)
767 Add modflap bounds-checking instrumentation for single-threaded program

769 fmodflapth
770 Common RejectNegative Report VarExists Var(flag_modflap,2)
771 Add modflap bounds-checking instrumentation for multi-threaded program

773 fmodflapir
774 Common RejectNegative Report Var(flag_modflap_ignore_reads)
775 Ignore read operations when inserting modflap instrumentation

777 fdce
778 Common Var(flag_dce) Init(1) Optimization
779 Use the RTL dead code elimination pass

```

```

781 fdse
782 Common Var(flag_dse) Init(1) Optimization
783 Use the RTL dead store elimination pass

785 freschedule-modulo-scheduled-loops
786 Common Report Var(flag_resched_modulo_sched) Optimization
787 Enable/Disable the traditional scheduling in loops that already passed modulo sc

789 fnon-call-exceptions
790 Common Report Var(flag_non_call_exceptions) Optimization
791 Support synchronous non-call exceptions

793 fomit-frame-pointer
794 Common Report Var(flag_omit_frame_pointer) Optimization
795 When possible do not generate stack frames

797 foptimize-register-move
798 Common Report Var(flag_regmove) Optimization
799 Do the full register move optimization pass

801 foptimize-sibling-calls
802 Common Report Var(flag_optimize_sibling_calls) Optimization
803 Optimize sibling and tail recursive calls

805 fpre-ipa-mem-report
806 Common Report Var(pre_ipa_mem_report)
807 Report on memory allocation before interprocedural optimization

809 fpost-ipa-mem-report
810 Common Report Var(post_ipa_mem_report)
811 Report on memory allocation before interprocedural optimization

813 fpack-struct
814 Common Report Var(flag_pack_struct) Optimization
815 Pack structure members together without holes

817 fpack-struct=
818 Common RejectNegative Joined UInteger Optimization
819 -fpack-struct=<number> Set initial maximum structure member alignment

821 fpcc-struct-return
822 Common Report Var(flag_pcc_struct_return,1) VarExists
823 Return small aggregates in memory, not registers

825 fpeel-loops
826 Common Report Var(flag_peel_loops) Optimization
827 Perform loop peeling

829 fpeephole
830 Common Report Var(flag_no_peephole,0) Optimization
831 Enable machine specific peephole optimizations

833 fpeephole2
834 Common Report Var(flag_peephole2) Optimization
835 Enable an RTL peephole pass before sched2

837 fPIC
838 Common Report Var(flag_pic,2)
839 Generate position-independent code if possible (large mode)

841 fPIE
842 Common Report Var(flag_pie,2)
843 Generate position-independent code for executables if possible (large mode)

845 fpic
846 Common Report Var(flag_pic,1) VarExists

```

```

847 Generate position-independent code if possible (small mode)

849 fpie
850 Common Report Var(flag_pie,1) VarExists
851 Generate position-independent code for executables if possible (small mode)

853 fpredictive-commoning
854 Common Report Var(flag_predictive_commoning) Optimization
855 Run predictive commoning optimization.

857 fprefetch-loop-arrays
858 Common Report Var(flag_prefetch_loop_arrays) Optimization
859 Generate prefetch instructions, if available, for arrays in loops

861 fprofile
862 Common Report Var(profile_flag)
863 Enable basic program profiling code

865 fprofile-arcs
866 Common Report Var(profile_arc_flag)
867 Insert arc-based program profiling code

869 fprofile-dir=
870 Common Joined RejectNegative
871 Set the top-level directory for storing the profile data.
872 The default is 'pwd'.

874 fprofile-correction
875 Common Report Var(flag_profile_correction)
876 Enable correction of flow inconsistent profile data input

878 fprofile-generate
879 Common
880 Enable common options for generating profile info for profile feedback directed

882 fprofile-generate=
883 Common Joined RejectNegative
884 Enable common options for generating profile info for profile feedback directed

886 fprofile-use
887 Common Var(flag_profile_use)
888 Enable common options for performing profile feedback directed optimizations

890 fprofile-use=
891 Common Joined RejectNegative
892 Enable common options for performing profile feedback directed optimizations, an

894 fprofile-values
895 Common Report Var(flag_profile_values)
896 Insert code to profile values of expressions

898 frandom-seed
899 Common

901 frandom-seed=
902 Common Joined RejectNegative
903 -frandom-seed=<string> Make compile reproducible using <string>

905 ; This switch causes the command line that was used to create an
906 ; object file to be recorded into the object file. The exact format
907 ; of this recording is target and binary file format dependent.
908 ; It is related to the -fverbose-asm switch, but that switch only
909 ; records information in the assembler output file as comments, so
910 ; they never reach the object file.
911 frecord-gcc-switches
912 Common Report Var(flag_record_gcc_switches)

```

```

913 Record gcc command line switches in the object file.

915 freg-struct-return
916 Common Report Var(flag_pcc_struct_return,0) VarExists Optimization
917 Return small aggregates in registers

919 fregmove
920 Common Report Var(flag_regmove) Optimization
921 Enables a register move optimization

923 frename-registers
924 Common Report Var(flag_rename_registers) Init(2) Optimization
925 Perform a register renaming optimization pass

927 freorder-blocks
928 Common Report Var(flag_reorder_blocks) Optimization
929 Reorder basic blocks to improve code placement

931 freorder-blocks-and-partition
932 Common Report Var(flag_reorder_blocks_and_partition) Optimization
933 Reorder basic blocks and partition into hot and cold sections

935 freorder-functions
936 Common Report Var(flag_reorder_functions) Optimization
937 Reorder functions to improve code placement

939 frerun-cse-after-loop
940 Common Report Var(flag_rerun_cse_after_loop) Init(2) Optimization
941 Add a common subexpression elimination pass after loop optimizations

943 frerun-loop-opt
944 Common
945 Does nothing. Preserved for backward compatibility.

947 frounding-math
948 Common Report Var(flag_rounding_math) Optimization
949 Disable optimizations that assume default FP rounding behavior

951 fsched-interblock
952 Common Report Var(flag_schedule_interblock) Init(1) Optimization
953 Enable scheduling across basic blocks

955 fsched-spec
956 Common Report Var(flag_schedule_speculative) Init(1) Optimization
957 Allow speculative motion of non-loads

959 fsched-spec-load
960 Common Report Var(flag_schedule_speculative_load) Optimization
961 Allow speculative motion of some loads

963 fsched-spec-load-dangerous
964 Common Report Var(flag_schedule_speculative_load_dangerous) Optimization
965 Allow speculative motion of more loads

967 fsched-verbose=
968 Common RejectNegative Joined
969 -fsched-verbose=<number> Set the verbosity level of the scheduler

971 fsched2-use-superblocks
972 Common Report Var(flag_sched2_use_superblocks) Optimization
973 If scheduling post reload, do superblock scheduling

975 fsched2-use-traces
976 Common Report Var(flag_sched2_use_traces) Optimization
977 If scheduling post reload, do trace scheduling

```

```

979 fschedule-insns
980 Common Report Var(flag_schedule_insns) Optimization
981 Reschedule instructions before register allocation

983 fschedule-insns2
984 Common Report Var(flag_schedule_insns_after_reload) Optimization
985 Reschedule instructions after register allocation

987 ; This flag should be on when a target implements non-trivial
988 ; scheduling hooks, maybe saving some information for its own sake.
989 ; On IA64, for example, this is used for correct bundling.
990 fselective-scheduling
991 Common Report Var(flag_selective_scheduling) Optimization
992 Schedule instructions using selective scheduling algorithm

994 fselective-scheduling2
995 Common Report Var(flag_selective_scheduling2) Optimization
996 Run selective scheduling after reload

998 fsel-sched-pipelining
999 Common Report Var(flag_sel_sched_pipelining) Init(0) Optimization
1000 Perform software pipelining of inner loops during selective scheduling

1002 fsel-sched-pipelining-outer-loops
1003 Common Report Var(flag_sel_sched_pipelining_outer_loops) Init(0) Optimization
1004 Perform software pipelining of outer loops during selective scheduling

1006 fsel-sched-reschedule-pipelined
1007 Common Report Var(flag_sel_sched_reschedule_pipelined) Init(0) Optimization
1008 Reschedule pipelined regions without pipelining

1010 ; sched_stalled_insns means that insns can be moved prematurely from the queue
1011 ; of stalled insns into the ready list.
1012 fsched-stalled-insns
1013 Common Report Var(flag_sched_stalled_insns) Optimization UInteger
1014 Allow premature scheduling of queued insns

1016 fsched-stalled-insns=
1017 Common RejectNegative Joined UInteger
1018 -fsched-stalled-insns=<number> Set number of queued insns that can be premature

1020 ; sched_stalled_insns_dep controls how many recently scheduled cycles will
1021 ; be examined for a dependency on a stalled insn that is candidate for
1022 ; premature removal from the queue of stalled insns into the ready list (has
1023 ; an effect only if the flag 'sched_stalled_insns' is set).
1024 fsched-stalled-insns-dep
1025 Common Report Var(flag_sched_stalled_insns_dep,1) Init(1) Optimization UInteger
1026 Set dependence distance checking in premature scheduling of queued insns

1028 fsched-stalled-insns-dep=
1029 Common RejectNegative Joined UInteger
1030 -fsched-stalled-insns-dep=<number> Set dependence distance checking in prem

1032 fsection-anchors
1033 Common Report Var(flag_section_anchors) Optimization
1034 Access data in the same section from shared anchor points

1036 frtl-abstract-sequences
1037 Common Report Var(flag_rtl_seqabstr) Optimization
1038 Perform sequence abstraction optimization on RTL

1040 fsee
1041 Common Report Var(flag_see) Init(0)
1042 Eliminate redundant sign extensions using LCM.

1044 fshow-column

```



```

1045 Common C ObjC C++ ObjC++ Report Var(flag_show_column) Init(0)
1046 Show column numbers in diagnostics, when available. Default off

1048 fsignaling-nans
1049 Common Report Var(flag_signaling_nans) Optimization
1050 Disable optimizations observable by IEEE signaling NaNs

1052 fsigned-zeros
1053 Common Report Var(flag_signed_zeros) Init(1) Optimization
1054 Disable floating point optimizations that ignore the IEEE signedness of zero

1056 fsingle-precision-constant
1057 Common Report Var(flag_single_precision_constant) Optimization
1058 Convert floating point constants to single precision constants

1060 fsplit-ivs-in-unroller
1061 Common Report Var(flag_split_ivs_in_unroller) Init(1) Optimization
1062 Split lifetimes of induction variables when loops are unrolled

1064 fsplit-wide-types
1065 Common Report Var(flag_split_wide_types) Optimization
1066 Split wide types into independent registers

1068 fvariable-expansion-in-unroller
1069 Common Report Var(flag_variable_expansion_in_unroller) Optimization
1070 Apply variable expansion when loops are unrolled

1072 fstack-check=
1073 Common Report RejectNegative Joined
1074 -fstack-check=[no|generic|specific]      Insert stack checking code into the prog

1076 fstack-check
1077 Common Report
1078 Insert stack checking code into the program. Same as -fstack-check=specific

1080 fstack-limit
1081 Common

1083 fstack-limit-register=
1084 Common Report RejectNegative Joined
1085 -fstack-limit-register=<register>      Trap if the stack goes past <register>

1087 fstack-limit-symbol=
1088 Common Report RejectNegative Joined
1089 -fstack-limit-symbol=<name>          Trap if the stack goes past symbol <name>

1091 fstack-protector
1092 Common Report Var(flag_stack_protect, 1)
1093 Use propolice as a stack protection method

1095 fstack-protector-all
1096 Common Report RejectNegative Var(flag_stack_protect, 2) VarExists
1097 Use a stack protection method for every function

1099 fstrength-reduce
1100 Common
1101 Does nothing. Preserved for backward compatibility.

1103 ; Nonzero if we should do (language-dependent) alias analysis.
1104 ; Typically, this analysis will assume that expressions of certain
1105 ; types do not alias expressions of certain other types. Only used
1106 ; if alias analysis (in general) is enabled.
1107 fstrict-aliasing
1108 Common Report Var(flag_strict_aliasing) Optimization
1109 Assume strict aliasing rules apply

```

```

1111 fstrict-calling-conventions
1112 Common Report Var(flag_strict_calling_conventions) Init(1)
1113 Use strict ABI calling conventions even for static functions

1115 #endif /* ! codereview */
1116 fstrict-overflow
1117 Common Report Var(flag_strict_overflow)
1118 Treat signed overflow as undefined

1120 fsyntax-only
1121 Common Report Var(flag_syntax_only)
1122 Check for syntax errors, then stop

1124 ftest-coverage
1125 Common Report Var(flag_test_coverage)
1126 Create data files needed by "gcov"

1128 fthread-jumps
1129 Common Report Var(flag_thread_jumps) Optimization
1130 Perform jump threading optimizations

1132 ftime-report
1133 Common Report Var(time_report)
1134 Report the time taken by each compiler pass

1136 ftls-model=
1137 Common Joined RejectNegative
1138 -ftls-model=[global-dynamic|local-dynamic|initial-exec|local-exec]      Set the

1140 ftoplevel-reorder
1141 Common Report Var(flag_toplevel_reorder) Init(2) Optimization
1142 Reorder top level functions, variables, and asms

1144 ftracer
1145 Common Report Var(flag_tracer)
1146 Perform superblock formation via tail duplication

1148 ; Zero means that floating-point math operations cannot generate a
1149 ; (user-visible) trap. This is the case, for example, in nonstop
1150 ; IEEE 754 arithmetic.
1151 ftrapping-math
1152 Common Report Var(flag_trapping_math) Init(1) Optimization
1153 Assume floating-point operations can trap

1155 ftrapv
1156 Common Report Var(flag_trapv) Optimization
1157 Trap for signed overflow in addition, subtraction and multiplication

1159 ftree-ccp
1160 Common Report Var(flag_tree_ccp) Optimization
1161 Enable SSA-CCP optimization on trees

1163 ftree-store-ccp
1164 Common
1165 Does nothing. Preserved for backward compatibility.

1167 ftree-ch
1168 Common Report Var(flag_tree_ch) Optimization
1169 Enable loop header copying on trees

1171 ftree-copyrename
1172 Common Report Var(flag_tree_copyrename) Optimization
1173 Replace SSA temporaries with better names in copies

1175 ftree-copy-prop
1176 Common Report Var(flag_tree_copy_prop) Optimization

```

```

1177 Enable copy propagation on trees

1179 ftree-store-copy-prop
1180 Common
1181 Does nothing. Preserved for backward compatibility.

1183 ftree-cselim
1184 Common Report Var(flag_tree_cselim) Init(2) Optimization
1185 Transform condition stores into unconditional ones

1187 ftree-switch-conversion
1188 Common Report Var(flag_tree_switch_conversion) Optimization
1189 Perform conversions of switch initializations.

1191 ftree-dce
1192 Common Report Var(flag_tree_dce) Optimization
1193 Enable SSA dead code elimination optimization on trees

1195 ftree-dominator-opts
1196 Common Report Var(flag_tree_dom) Optimization
1197 Enable dominator optimizations

1199 ftree-dse
1200 Common Report Var(flag_tree_dse) Optimization
1201 Enable dead store elimination

1203 ftree-fre
1204 Common Report Var(flag_tree_fre) Optimization
1205 Enable Full Redundancy Elimination (FRE) on trees

1207 ftree-loop-distribution
1208 Common Report Var(flag_tree_loop_distribution) Optimization
1209 Enable loop distribution on trees

1211 ftree-loop-im
1212 Common Report Var(flag_tree_loop_im) Init(1) Optimization
1213 Enable loop invariant motion on trees

1215 ftree-loop-linear
1216 Common Report Var(flag_tree_loop_linear) Optimization
1217 Enable linear loop transforms on trees

1219 ftree-loop-ivcanon
1220 Common Report Var(flag_tree_loop_ivcanon) Init(1) Optimization
1221 Create canonical induction variables in loops

1223 ftree-loop-optimize
1224 Common Report Var(flag_tree_loop_optimize) Init(1) Optimization
1225 Enable loop optimizations on tree level

1227 ftree-parallelize-loops=
1228 Common Report Joined UInteger Var(flag_tree_parallelize_loops) Init(1)
1229 Enable automatic parallelization of loops

1231 ftree-pre
1232 Common Report Var(flag_tree_pre) Optimization
1233 Enable SSA-PRE optimization on trees

1235 ftree-reassoc
1236 Common Report Var(flag_tree_reassoc) Init(1) Optimization
1237 Enable reassociation on tree level

1239 ftree-salias
1240 Common
1241 Does nothing. Preserved for backward compatibility.

```

```

1243 ftree-sink
1244 Common Report Var(flag_tree_sink) Optimization
1245 Enable SSA code sinking on trees

1247 ftree-sra
1248 Common Report Var(flag_tree_sra) Optimization
1249 Perform scalar replacement of aggregates

1251 ftree-ter
1252 Common Report Var(flag_tree_ter) Init(1) Optimization
1253 Replace temporary expressions in the SSA->normal pass

1255 ftree-lrs
1256 Common Report Var(flag_tree_live_range_split) Optimization
1257 Perform live range splitting during the SSA->normal pass

1259 ftree-vrp
1260 Common Report Var(flag_tree_vrp) Init(0) Optimization
1261 Perform Value Range Propagation on trees

1263 funit-at-a-time
1264 Common Report Var(flag_unit_at_a_time) Init(1) Optimization
1265 Compile whole compilation unit at a time

1267 funroll-loops
1268 Common Report Var(flag_unroll_loops) Optimization
1269 Perform loop unrolling when iteration count is known

1271 funroll-all-loops
1272 Common Report Var(flag_unroll_all_loops) Optimization
1273 Perform loop unrolling for all loops

1275 ; Nonzero means that loop optimizer may assume that the induction variables
1276 ; that control loops do not overflow and that the loops with nontrivial
1277 ; exit condition are not infinite
1278 funsafe-loop-optimizations
1279 Common Report Var(flag_unsafe_loop_optimizations) Optimization
1280 Allow loop optimizations to assume that the loops behave in normal way

1282 fassociative-math
1283 Common Report Var(flag_associative_math)
1284 Allow optimization for floating-point arithmetic which may change the
1285 result of the operation due to rounding.

1287 freciprocal-math
1288 Common Report Var(flag_reciprocal_math)
1289 Same as -fassociative-math for expressions which include division.

1291 ; Nonzero means that unsafe floating-point math optimizations are allowed
1292 ; for the sake of speed. IEEE compliance is not guaranteed, and operations
1293 ; are allowed to assume that their arguments and results are "normal"
1294 ; (e.g., nonnegative for SQRT).
1295 funsafe-math-optimizations
1296 Common Report Var(flag_unsafe_math_optimizations) Optimization
1297 Allow math optimizations that may violate IEEE or ISO standards

1299 funswitch-loops
1300 Common Report Var(flag_unswitch_loops) Optimization
1301 Perform loop unswitching

1303 funwind-tables
1304 Common Report Var(flag_unwind_tables) Optimization
1305 Just generate unwind tables for exception handling

1307 fvar-tracking
1308 Common Report Var(flag_var_tracking) VarExists Optimization

```

```

1309 Perform variable tracking

1311 fvar-tracking-uninit
1312 Common Report Var(flag_var_tracking_uninit) Optimization
1313 Perform variable tracking and also tag variables that are uninitialized

1315 ftree-vectorize
1316 Common Report Var(flag_tree_vectorize) Optimization
1317 Enable loop vectorization on trees

1319 fvect-cost-model
1320 Common Report Var(flag_vect_cost_model) Optimization
1321 Enable use of cost model in vectorization

1323 ftree-vect-loop-version
1324 Common Report Var(flag_tree_vect_loop_version) Init(1) Optimization
1325 Enable loop versioning when doing loop vectorization on trees

1327 ftree-vectorizer-verbose=
1328 Common RejectNegative Joined
1329 -ftree-vectorizer-verbose=<number>      Set the verbosity level of the vectorize

1331 ftree-scev-cprop
1332 Common Report Var(flag_tree_scev_cprop) Init(1) Optimization
1333 Enable copy propagation of scalar-evolution information.

1335 ; -fverbose-asm causes extra commentary information to be produced in
1336 ; the generated assembly code (to make it more readable). This option
1337 ; is generally only of use to those who actually need to read the
1338 ; generated assembly code (perhaps while debugging the compiler itself).
1339 ; -fno-verbose-asm, the default, causes the extra information
1340 ; to not be added and is useful when comparing two assembler files.
1341 fverbose-asm
1342 Common Report Var(flag_verbose_asm)
1343 Add extra commentary to assembler output

1345 fvisibility=
1346 Common Joined RejectNegative
1347 -fvisibility=[default|internal|hidden|protected]      Set the default symbol v

1350 fvpt
1351 Common Report Var(flag_value_profile_transformations) Optimization
1352 Use expression value profiles in optimizations

1354 fweb
1355 Common Report Var(flag_web) Init(2) Optimization
1356 Construct webs and split unrelated uses of single variable

1358 ftree-builtin-call-dce
1359 Common Report Var(flag_tree_builtin_call_dce) Init(0) Optimization
1360 Enable conditional dead code elimination for builtin calls

1362 fwhole-program
1363 Common Report Var(flag_whole_program) Init(0) Optimization
1364 Perform whole program optimizations

1366 fwrapv
1367 Common Report Var(flag_wrapv) Optimization
1368 Assume signed arithmetic overflow wraps around

1370 fzero-initialized-in-bss
1371 Common Report Var(flag_zero_initialized_in_bss) Init(1)
1372 Put zero initialized data in the bss section

1374 g

```

```

1375 Common JoinedOrMissing
1376 Generate debug information in default format

1378 gcoff
1379 Common JoinedOrMissing Negative(gdwarf-2)
1380 Generate debug information in COFF format

1382 gdwarf-2
1383 Common JoinedOrMissing Negative(gstabs)
1384 Generate debug information in DWARF v2 format

1386 ggdb
1387 Common JoinedOrMissing
1388 Generate debug information in default extended format

1390 gstabs
1391 Common JoinedOrMissing Negative(gstabs+)
1392 Generate debug information in STABS format

1394 gstabs+
1395 Common JoinedOrMissing Negative(gvms)
1396 Generate debug information in extended STABS format

1398 gvms
1399 Common JoinedOrMissing Negative(gxcoff)
1400 Generate debug information in VMS format

1402 gxcoff
1403 Common JoinedOrMissing Negative(gxcoff+)
1404 Generate debug information in XCOFF format

1406 gxcoff+
1407 Common JoinedOrMissing Negative(gcoff)
1408 Generate debug information in extended XCOFF format

1410 o
1411 Common Joined Separate
1412 -o <file>      Place output into <file>

1414 p
1415 Common Var(profile_flag)
1416 Enable function profiling

1418 pedantic
1419 Common Var(pedantic)
1420 Issue warnings needed for strict compliance to the standard

1422 pedantic-errors
1423 Common
1424 Like -pedantic but issue them as errors

1426 quiet
1427 Common Var(quiet_flag)
1428 Do not display functions compiled or elapsed time

1430 version
1431 Common Var(version_flag)
1432 Display the compiler's version

1434 w
1435 Common Var(inhibit_warnings)
1436 Suppress warnings

1438 shared
1439 Common RejectNegative Negative(pie)
1440 Create a shared library

```

```
1442 pie
1443 Common RejectNegative Negative(shared)
1444 Create a position independent executable

1446 ; This comment is to ensure we retain the blank line above.
```

new/gcc/config/i386/i386.c

1

```
*****
964570 Sun Oct 28 20:56:07 2012
new/gcc/config/i386/i386.c
Implement -fstrict-calling-conventions
Stock GCC is overly willing to violate the ABI when calling local functions,
such that it passes arguments in registers on i386. This hampers debugging
with anything other than a fully-aware DWARF debugger, and is generally not
something we desire.
Implement a flag which disables this behaviour, enabled by default. The flag is
global, though only effective on i386, to more easily allow its globalization
later which, given the odds, is likely to be necessary.
*****
_____unchanged_portion_omitted_____
```

new/gcc/config/i386/i386.c

2

```
4346 /* Return the regparm value for a function with the indicated TYPE and DECL.
4347 DECL may be NULL when calling function indirectly
4348 or considering a libcall. */

4350 static int
4351 ix86_function_regparm (const_tree type, const_tree decl)
4352 {
4353     tree attr;
4354     int regparm;

4356     static bool error_issued;

4358     if (TARGET_64BIT)
4359         return (ix86_function_type_abi (type) == SYSV_ABI
4360                 ? X86_64_REGPARAM_MAX : X64_REGPARAM_MAX);

4362     regparm = ix86_regparm;
4363     attr = lookup_attribute ("regparm", TYPE_ATTRIBUTES (type));
4364     if (attr)
4365     {
4366         regparm
4367             = TREE_INT_CST_LOW (TREE_VALUE (TREE_VALUE (attr)));

4369         if (decl && TREE_CODE (decl) == FUNCTION_DECL)
4370         {
4371             /* We can't use regparm(3) for nested functions because
4372              these pass static chain pointer in %ecx register. */
4373             if (!error_issued && regparm == 3
4374                 && decl_function_context (decl)
4375                 && !DECL_NO_STATIC_CHAIN (decl))
4376             {
4377                 error ("nested functions are limited to 2 register parameters");
4378                 error_issued = true;
4379                 return 0;
4380             }
4381         }

4383         return regparm;
4384     }

4386     if (lookup_attribute ("fastcall", TYPE_ATTRIBUTES (type)))
4387         return 2;

4389     /* Use register calling convention for local functions when possible. */
4390     if (decl
4391         && TREE_CODE (decl) == FUNCTION_DECL
4392         && optimize
4393         && (TARGET_64BIT || !flag_strict_calling_conventions)
4394         #endif /* ! codereview */
4395         && !profile_flag)
4396     {
4397         /* FIXME: remove this CONST_CAST when cgraph.[ch] is constified. */
4398         struct cgraph_local_info *i = cgraph_local_info (CONST_CAST_TREE(decl));
4399         if (i && i->local)
4400         {
4401             int local_regparm, globals = 0, regno;
4402             struct function *f;

4404             /* Make sure no regparm register is taken by a
4405              fixed register variable. */
4406             for (local_regparm = 0; local_regparm < REGPARAM_MAX; local_regparm++)
4407                 if (fixed_regs[local_regparm])
4408                     break;

4410             /* We can't use regparm(3) for nested functions as these use
4411              static chain pointer in third argument. */
```

```

4412     if (local_regparm == 3
4413         && decl_function_context (decl)
4414         && !DECL_NO_STATIC_CHAIN (decl))
4415         local_regparm = 2;

4417     /* If the function realigns its stackpointer, the prologue will
4418        clobber %ecx. If we've already generated code for the callee,
4419        the callee DECL_STRUCT_FUNCTION is gone, so we fall back to
4420        scanning the attributes for the self-realigning property. */
4421     f = DECL_STRUCT_FUNCTION (decl);
4422     /* Since current internal arg pointer won't conflict with
4423        parameter passing regs, so no need to change stack
4424        realignment and adjust regparm number.

4426        Each fixed register usage increases register pressure,
4427        so less registers should be used for argument passing.
4428        This functionality can be overridden by an explicit
4429        regparm value. */
4430     for (regno = 0; regno <= DI_REG; regno++)
4431         if (fixed_regs[regno])
4432             globals++;

4434     local_regparm
4435     = globals < local_regparm ? local_regparm - globals : 0;

4437     if (local_regparm > regparm)
4438         regparm = local_regparm;
4439 }
4440 }

4442 return regparm;
4443 }

4445 /* Return 1 or 2, if we can pass up to SSE_REGPARAM_MAX SFmode (1) and
4446    DFmode (2) arguments in SSE registers for a function with the
4447    indicated TYPE and DECL. DECL may be NULL when calling function
4448    indirectly or considering a libcall. Otherwise return 0. */

4450 static int
4451 ix86_function_sseregparm (const_tree type, const_tree decl, bool warn)
4452 {
4453     gcc_assert (!TARGET_64BIT);

4455     /* Use SSE registers to pass SFmode and DFmode arguments if requested
4456        by the sseregparm attribute. */
4457     if (TARGET_SSEREGPARM
4458         || (type && lookup_attribute ("sseregparm", TYPE_ATTRIBUTES (type))))
4459     {
4460         if (!TARGET_SSE)
4461         {
4462             if (warn)
4463             {
4464                 if (decl)
4465                     error ("Calling %qD with attribute sseregparm without "
4466                          "SSE/SSE2 enabled", decl);
4467                 else
4468                     error ("Calling %qT with attribute sseregparm without "
4469                          "SSE/SSE2 enabled", type);
4470             }
4471             return 0;
4472         }
4474         return 2;
4475     }

4477     /* For local functions, pass up to SSE_REGPARAM_MAX SFmode

```

```

4478     (and DFmode for SSE2) arguments in SSE registers. */
4479     if (decl && TARGET_SSE_MATH && optimize && !profile_flag &&
4480         (TARGET_64BIT || !flag_strict_calling_conventions))
4481     if (decl && TARGET_SSE_MATH && optimize && !profile_flag)
4482     {
4483         /* FIXME: remove this CONST_CAST when cgraph.[ch] is constified. */
4484         struct cgraph_local_info *i = cgraph_local_info (CONST_CAST_TREE(decl));
4485         if (i && i->local)
4486             return TARGET_SSE2 ? 2 : 1;
4488     }
4489     return 0;

```

unchanged\_portion\_omitted

\*\*\*\*\*

651723 Sun Oct 28 20:56:09 2012

new/gcc/doc/invoke.texi

Implement `-fstrict-calling-conventions`

Stock GCC is overly willing to violate the ABI when calling local functions, such that it passes arguments in registers on i386. This hampers debugging with anything other than a fully-aware DWARF debugger, and is generally not something we desire.

Implement a flag which disables this behaviour, enabled by default. The flag is global, though only effective on i386, to more easily allow its globalization later which, given the odds, is likely to be necessary.

\*\*\*\*\*

unchanged portion omitted

2811 The option `@option{-Wextra}` also prints warning messages for the  
2812 following cases:

2814 `@itemize @bullet`

2816 `@item`

2817 A pointer is compared against integer zero with `@samp{<}`, `@samp{<=}`,  
2818 `@samp{>}`, or `@samp{>=}`.

2820 `@item`

2821 (C++ only) An enumerator and a non-enumerator both appear in a  
2822 conditional expression.

2824 `@item`

2825 (C++ only) Ambiguous virtual bases.

2827 `@item`

2828 (C++ only) Subscripting an array which has been declared `@samp{register}`.

2830 `@item`

2831 (C++ only) Taking the address of a variable which has been declared  
2832 `@samp{register}`.

2834 `@item`

2835 (C++ only) A base class is not initialized in a derived class' copy  
2836 constructor.

2838 `@end itemize`

2840 `@item -Wchar-subscripts`

2841 `@opindex Wchar-subscripts`

2842 `@opindex Wno-char-subscripts`

2843 Warn if an array subscript has type `@code{char}`. This is a common cause

2844 of error, as programmers often forget that this type is signed on some  
2845 machines.

2846 This warning is enabled by `@option{-Wall}`.

2848 `@item -Wcomment`

2849 `@opindex Wcomment`

2850 `@opindex Wno-comment`

2851 Warn whenever a comment-start sequence `@samp{/**}` appears in a `@samp{/**}`

2852 comment, or whenever a Backslash-Newline appears in a `@samp{/**}` comment.

2853 This warning is enabled by `@option{-Wall}`.

2855 `@item -Wformat`

2856 `@opindex Wformat`

2857 `@opindex Wno-format`

2858 `@opindex ffreestanding`

2859 `@opindex fno-builtin`

2860 Check calls to `@code{printf}` and `@code{scanf}`, etc., to make sure that

2861 the arguments supplied have types appropriate to the format string

2862 specified, and that the conversions specified in the format string make

2863 sense. This includes standard functions, and others specified by format  
2864 attributes (`@pxref{Function Attributes}`), in the `@code{printf}`,  
2865 `@code{scanf}`, `@code{strftime}` and `@code{strfmon}` (an X/Open extension,  
2866 not in the C standard) families (or other target-specific families).  
2867 Which functions are checked without format attributes having been  
2868 specified depends on the standard version selected, and such checks of  
2869 functions without the attribute specified are disabled by  
2870 `@option{-ffreestanding}` or `@option{-fno-builtin}`.

2872 The formats are checked against the format features supported by GNU  
2873 libc version 2.2. These include all ISO C90 and C99 features, as well  
2874 as features from the Single Unix Specification and some BSD and GNU  
2875 extensions. Other library implementations may not support all these  
2876 features; GCC does not support warning about features that go beyond a  
2877 particular library's limitations. However, if `@option{-pedantic}` is used  
2878 with `@option{-Wformat}`, warnings will be given about format features not  
2879 in the selected standard version (but not for `@code{strfmon}` formats,  
2880 since those are not in any version of the C standard). `@xref{C Dialect`  
2881 Options,,Options Controlling C Dialect}.

2883 Since `@option{-Wformat}` also checks for null format arguments for  
2884 several functions, `@option{-Wformat}` also implies `@option{-Wnonnull}`.

2886 `@option{-Wformat}` is included in `@option{-Wall}`. For more control over some  
2887 aspects of format checking, the options `@option{-Wformat-y2k}`,  
2888 `@option{-Wno-format-extra-args}`, `@option{-Wno-format-zero-length}`,  
2889 `@option{-Wformat-nonliteral}`, `@option{-Wformat-security}`, and  
2890 `@option{-Wformat=2}` are available, but are not included in `@option{-Wall}`.

2892 `@item -Wformat-y2k`

2893 `@opindex Wformat-y2k`

2894 `@opindex Wno-format-y2k`

2895 If `@option{-Wformat}` is specified, also warn about `@code{strftime}`

2896 formats which may yield only a two-digit year.

2898 `@item -Wno-format-contains-nul`

2899 `@opindex Wno-format-contains-nul`

2900 `@opindex Wformat-contains-nul`

2901 If `@option{-Wformat}` is specified, do not warn about format strings that

2902 contain NUL bytes.

2904 `@item -Wno-format-extra-args`

2905 `@opindex Wno-format-extra-args`

2906 `@opindex Wformat-extra-args`

2907 If `@option{-Wformat}` is specified, do not warn about excess arguments to a

2908 `@code{printf}` or `@code{scanf}` format function. The C standard specifies

2909 that such arguments are ignored.

2911 Where the unused arguments lie between used arguments that are

2912 specified with `@samp{}` operand number specifications, normally

2913 warnings are still given, since the implementation could not know what

2914 type to pass to `@code{va_arg}` to skip the unused arguments. However,

2915 in the case of `@code{scanf}` formats, this option will suppress the

2916 warning if the unused arguments are all pointers, since the Single

2917 Unix Specification says that such unused arguments are allowed.

2919 `@item -Wno-format-zero-length @r{(C and Objective-C only)}`

2920 `@opindex Wno-format-zero-length`

2921 `@opindex Wformat-zero-length`

2922 If `@option{-Wformat}` is specified, do not warn about zero-length formats.

2923 The C standard specifies that zero-length formats are allowed.

2925 `@item -Wformat-nonliteral`

2926 `@opindex Wformat-nonliteral`

2927 `@opindex Wno-format-nonliteral`

2928 If `@option{-Wformat}` is specified, also warn if the format string is not a

2929 string literal and so cannot be checked, unless the format function  
 2930 takes its format arguments as a `{va_list}`.

2932 @item -Wformat-security  
 2933 @opindex Wformat-security  
 2934 @opindex Wno-format-security  
 2935 If @option{-Wformat} is specified, also warn about uses of format  
 2936 functions that represent possible security problems. At present, this  
 2937 warns about calls to `{printf}` and `{scanf}` functions where the  
 2938 format string is not a string literal and there are no format arguments,  
 2939 as in `{printf (foo);}`. This may be a security hole if the format  
 2940 string came from untrusted input and contains `{%n}`. (This is  
 2941 currently a subset of what @option{-Wformat-nonliteral} warns about, but  
 2942 in future warnings may be added to @option{-Wformat-security} that are not  
 2943 included in @option{-Wformat-nonliteral}.)

2945 @item -Wformat=2  
 2946 @opindex Wformat=2  
 2947 @opindex Wno-format=2  
 2948 Enable @option{-Wformat} plus format checks not included in  
 2949 @option{-Wformat}. Currently equivalent to `{@samp{-Wformat`  
 2950 `-Wformat-nonliteral -Wformat-security -Wformat-y2k}`.

2952 @item -Wnonnull @r{(C and Objective-C only)}  
 2953 @opindex Wnonnull  
 2954 @opindex Wno-nonnull  
 2955 Warn about passing a null pointer for arguments marked as  
 2956 requiring a non-null value by the `{nonnull}` function attribute.

2958 @option{-Wnonnull} is included in @option{-Wall} and @option{-Wformat}. It  
 2959 can be disabled with the @option{-Wno-nonnull} option.

2961 @item -Winit-self @r{(C, C++, Objective-C and Objective-C++ only)}  
 2962 @opindex Winit-self  
 2963 @opindex Wno-init-self  
 2964 Warn about uninitialized variables which are initialized with themselves.  
 2965 Note this option can only be used with the @option{-Wuninitialized} option.

2967 For example, GCC will warn about `{i}` being uninitialized in the  
 2968 following snippet only when @option{-Winit-self} has been specified:  
 2969 @smallexample  
 2970 @group  
 2971 int f()  
 2972 @{  
 2973 int i = i;  
 2974 return i;  
 2975 @}  
 2976 @end group  
 2977 @end smallexample

2979 @item -Wimplicit-int @r{(C and Objective-C only)}  
 2980 @opindex Wimplicit-int  
 2981 @opindex Wno-implicit-int  
 2982 Warn when a declaration does not specify a type.  
 2983 This warning is enabled by @option{-Wall}.

2985 @item -Wimplicit-function-declaration @r{(C and Objective-C only)}  
 2986 @opindex Wimplicit-function-declaration  
 2987 @opindex Wno-implicit-function-declaration  
 2988 Give a warning whenever a function is used before being declared. In  
 2989 C99 mode (@option{-std=c99} or @option{-std=gnu99}), this warning is  
 2990 enabled by default and it is made into an error by  
 2991 @option{-pedantic-errors}. This warning is also enabled by  
 2992 @option{-Wall}.

2994 @item -Wimplicit

2995 @opindex Wimplicit  
 2996 @opindex Wno-implicit  
 2997 Same as @option{-Wimplicit-int} and @option{-Wimplicit-function-declaration}.  
 2998 This warning is enabled by @option{-Wall}.

3000 @item -Wignored-qualifiers @r{(C and C++ only)}  
 3001 @opindex Wignored-qualifiers  
 3002 @opindex Wno-ignored-qualifiers  
 3003 Warn if the return type of a function has a type qualifier  
 3004 such as `{const}`. For ISO C such a type qualifier has no effect,  
 3005 since the value returned by a function is not an lvalue.  
 3006 For C++, the warning is only emitted for scalar types or `{void}`.  
 3007 ISO C prohibits qualified `{void}` return types on function  
 3008 definitions, so such return types always receive a warning  
 3009 even without this option.

3011 This warning is also enabled by @option{-Wextra}.

3013 @item -Wmain  
 3014 @opindex Wmain  
 3015 @opindex Wno-main  
 3016 Warn if the type of `{main}` is suspicious. `{main}` should be  
 3017 a function with external linkage, returning int, taking either zero  
 3018 arguments, two, or three arguments of appropriate types. This warning  
 3019 is enabled by default in C++ and is enabled by either @option{-Wall}  
 3020 or @option{-pedantic}.

3022 @item -Wmissing-braces  
 3023 @opindex Wmissing-braces  
 3024 @opindex Wno-missing-braces  
 3025 Warn if an aggregate or union initializer is not fully bracketed. In  
 3026 the following example, the initializer for `{a}` is not fully  
 3027 bracketed, but that for `{b}` is fully bracketed.

3029 @smallexample  
 3030 int a[2][2] = @{{ 0, 1, 2, 3 @};  
 3031 int b[2][2] = @{{ @{ 0, 1 @}, @{ 2, 3 @} @};  
 3032 @end smallexample

3034 This warning is enabled by @option{-Wall}.

3036 @item -Wmissing-include-dirs @r{(C, C++, Objective-C and Objective-C++ only)}  
 3037 @opindex Wmissing-include-dirs  
 3038 @opindex Wno-missing-include-dirs  
 3039 Warn if a user-supplied include directory does not exist.

3041 @item -Wparentheses  
 3042 @opindex Wparentheses  
 3043 @opindex Wno-parentheses  
 3044 Warn if parentheses are omitted in certain contexts, such  
 3045 as when there is an assignment in a context where a truth value  
 3046 is expected, or when operators are nested whose precedence people  
 3047 often get confused about.

3049 Also warn if a comparison like `{x<=y<=z}` appears; this is  
 3050 equivalent to `{(x<=y ? 1 : 0) <= z}`, which is a different  
 3051 interpretation from that of ordinary mathematical notation.

3053 Also warn about constructions where there may be confusion to which  
 3054 `{if}` statement an `{else}` branch belongs. Here is an example of  
 3055 such a case:

3057 @smallexample  
 3058 @group  
 3059 @  
 3060 if (a)



```

3061     if (b)
3062         foo ();
3063     else
3064         bar ();
3065 }
3066 @end group
3067 @end smallexample

```

3069 In C/C++, every `@code{else}` branch belongs to the innermost possible `@code{if}` statement, which in this example is `@code{if (b)}`. This is often not what the programmer expected, as illustrated in the above example by indentation the programmer chose. When there is the potential for this confusion, GCC will issue a warning when this flag is specified. To eliminate the warning, add explicit braces around the innermost `@code{if}` statement so there is no way the `@code{else}` could belong to the enclosing `@code{if}`. The resulting code would look like this:

```

3079 @smallexample
3080 @group
3081 @{
3082     if (a)
3083     @{
3084         if (b)
3085             foo ();
3086         else
3087             bar ();
3088     @}
3089 @}
3090 @end group
3091 @end smallexample

```

3093 This warning is enabled by `@option{-Wall}`.

```

3095 @item -Wsequence-point
3096 @opindex Wsequence-point
3097 @opindex Wno-sequence-point
3098 Warn about code that may have undefined semantics because of violations
3099 of sequence point rules in the C and C++ standards.

```

3101 The C and C++ standards defines the order in which expressions in a C/C++ program are evaluated in terms of `@dfn{sequence points}`, which represent a partial ordering between the execution of parts of the program: those executed before the sequence point, and those executed after it. These occur after the evaluation of a full expression (one which is not part of a larger expression), after the evaluation of the first operand of a `@code{&&}`, `@code{|||}`, `@code{?:}` or `@code{,}` (comma) operator, before a function is called (but after the evaluation of its arguments and the expression denoting the called function), and in certain other places. Other than as expressed by the sequence point rules, the order of evaluation of subexpressions of an expression is not specified. All these rules describe only a partial order rather than a total order, since, for example, if two functions are called within one expression with no sequence point between them, the order in which the functions are called is not specified. However, the standards committee have ruled that function calls do not overlap.

3118 It is not specified when between sequence points modifications to the values of objects take effect. Programs whose behavior depends on this have undefined behavior; the C and C++ standards specify that "Between the previous and next sequence point an object shall have its stored value modified at most once by the evaluation of an expression. Furthermore, the prior value shall be read only to determine the value to be stored." If a program breaks these rules, the results on any particular implementation are entirely unpredictable.

3127 Examples of code with undefined behavior are `@code{a = a++;}`, `@code{a[n] = b[n++]}` and `@code{a[a[i++] = i]}`. Some more complicated cases are not diagnosed by this option, and it may give an occasional false positive result, but in general it has been found fairly effective at detecting this sort of problem in programs.

3133 The standard is worded confusingly, therefore there is some debate over the precise meaning of the sequence point rules in subtle cases. 3135 Links to discussions of the problem, including proposed formal definitions, may be found on the GCC readings page, at `@w{@uref{http://gcc.gnu.org/readings.html}}`.

3139 This warning is enabled by `@option{-Wall}` for C and C++.

```

3141 @item -Wreturn-type
3142 @opindex Wreturn-type
3143 @opindex Wno-return-type
3144 Warn whenever a function is defined with a return-type that defaults
3145 to @code{int}. Also warn about any @code{return} statement with no
3146 return-value in a function whose return-type is not @code{void}
3147 (falling off the end of the function body is considered returning
3148 without a value), and about a @code{return} statement with a
3149 expression in a function whose return-type is @code{void}.

```

3151 For C++, a function without return type always produces a diagnostic message, even when `@option{-Wno-return-type}` is specified. The only exceptions are `@samp{main}` and functions defined in system headers.

3155 This warning is enabled by `@option{-Wall}`.

```

3157 @item -Wswitch
3158 @opindex Wswitch
3159 @opindex Wno-switch
3160 Warn whenever a @code{switch} statement has an index of enumerated type
3161 and lacks a @code{case} for one or more of the named codes of that
3162 enumeration. (The presence of a @code{default} label prevents this
3163 warning.) @code{case} labels outside the enumeration range also
3164 provoke warnings when this option is used.
3165 This warning is enabled by @option{-Wall}.

```

```

3167 @item -Wswitch-default
3168 @opindex Wswitch-default
3169 @opindex Wno-switch-default
3170 Warn whenever a @code{switch} statement does not have a @code{default}
3171 case.

```

```

3173 @item -Wswitch-enum
3174 @opindex Wswitch-enum
3175 @opindex Wno-switch-enum
3176 Warn whenever a @code{switch} statement has an index of enumerated type
3177 and lacks a @code{case} for one or more of the named codes of that
3178 enumeration. @code{case} labels outside the enumeration range also
3179 provoke warnings when this option is used.

```

```

3181 @item -Wsync-nand @r{(C and C++ only)}
3182 @opindex Wsync-nand
3183 @opindex Wno-sync-nand
3184 Warn when @code{__sync_fetch_and_nand} and @code{__sync_nand_and_fetch}
3185 built-in functions are used. These functions changed semantics in GCC 4.4.

```

```

3187 @item -Wtrigraphs
3188 @opindex Wtrigraphs
3189 @opindex Wno-trigraphs
3190 Warn if any trigraphs are encountered that might change the meaning of
3191 the program (trigraphs within comments are not warned about).
3192 This warning is enabled by @option{-Wall}.

```

3194 @item -Wunused-function  
 3195 @opindex Wunused-function  
 3196 @opindex Wno-unused-function  
 3197 Warn whenever a static function is declared but not defined or a  
 3198 non-inline static function is unused.  
 3199 This warning is enabled by @option{-Wall}.

3201 @item -Wunused-label  
 3202 @opindex Wunused-label  
 3203 @opindex Wno-unused-label  
 3204 Warn whenever a label is declared but not used.  
 3205 This warning is enabled by @option{-Wall}.

3207 To suppress this warning use the @samp{unused} attribute  
 3208 (@pxref{Variable Attributes}).

3210 @item -Wunused-parameter  
 3211 @opindex Wunused-parameter  
 3212 @opindex Wno-unused-parameter  
 3213 Warn whenever a function parameter is unused aside from its declaration.

3215 To suppress this warning use the @samp{unused} attribute  
 3216 (@pxref{Variable Attributes}).

3218 @item -Wunused-variable  
 3219 @opindex Wunused-variable  
 3220 @opindex Wno-unused-variable  
 3221 Warn whenever a local variable or non-constant static variable is unused  
 3222 aside from its declaration.  
 3223 This warning is enabled by @option{-Wall}.

3225 To suppress this warning use the @samp{unused} attribute  
 3226 (@pxref{Variable Attributes}).

3228 @item -Wunused-value  
 3229 @opindex Wunused-value  
 3230 @opindex Wno-unused-value  
 3231 Warn whenever a statement computes a result that is explicitly not  
 3232 used. To suppress this warning cast the unused expression to  
 3233 @samp{void}. This includes an expression-statement or the left-hand  
 3234 side of a comma expression that contains no side effects. For example,  
 3235 an expression such as @samp{x[i,j]} will cause a warning, while  
 3236 @samp{x[(void)i,j]} will not.

3238 This warning is enabled by @option{-Wall}.

3240 @item -Wunused  
 3241 @opindex Wunused  
 3242 @opindex Wno-unused  
 3243 All the above @option{-Wunused} options combined.

3245 In order to get a warning about an unused function parameter, you must  
 3246 either specify @samp{-Wextra -Wunused} (note that @samp{-Wall} implies  
 3247 @samp{-Wunused}), or separately specify @option{-Wunused-parameter}.

3249 @item -Wuninitialized  
 3250 @opindex Wuninitialized  
 3251 @opindex Wno-uninitialized  
 3252 Warn if an automatic variable is used without first being initialized  
 3253 or if a variable may be clobbered by a @code{setjmp} call. In C++,  
 3254 warn if a non-static reference or non-static @samp{const} member  
 3255 appears in a class without constructors.

3257 If you want to warn about code which uses the uninitialized value of the  
 3258 variable in its own initializer, use the @option{-Winit-self} option.

3260 These warnings occur for individual uninitialized or clobbered  
 3261 elements of structure, union or array variables as well as for  
 3262 variables which are uninitialized or clobbered as a whole. They do  
 3263 not occur for variables or elements declared @code{volatile}. Because  
 3264 these warnings depend on optimization, the exact variables or elements  
 3265 for which there are warnings will depend on the precise optimization  
 3266 options and version of GCC used.

3268 Note that there may be no warning about a variable that is used only  
 3269 to compute a value that itself is never used, because such  
 3270 computations may be deleted by data flow analysis before the warnings  
 3271 are printed.

3273 These warnings are made optional because GCC is not smart  
 3274 enough to see all the reasons why the code might be correct  
 3275 despite appearing to have an error. Here is one example of how  
 3276 this can happen:

```
3278 @smallexample
3279 @group
3280 @{
3281   int x;
3282   switch (y)
3283   @{
3284     case 1: x = 1;
3285     break;
3286     case 2: x = 4;
3287     break;
3288     case 3: x = 5;
3289   @}
3290   foo (x);
3291 @}
3292 @end group
3293 @end smallexample
```

3295 @noindent  
 3296 If the value of @code{y} is always 1, 2 or 3, then @code{x} is  
 3297 always initialized, but GCC doesn't know this. Here is  
 3298 another common case:

```
3300 @smallexample
3301 @{
3302   int save_y;
3303   if (change_y) save_y = y, y = new_y;
3304   @dots{}
3305   if (change_y) y = save_y;
3306 @}
3307 @end smallexample
```

3309 @noindent  
 3310 This has no bug because @code{save\_y} is used only if it is set.

3312 @cindex @code{longjmp} warnings  
 3313 This option also warns when a non-volatile automatic variable might be  
 3314 changed by a call to @code{longjmp}. These warnings as well are possible  
 3315 only in optimizing compilation.

3317 The compiler sees only the calls to @code{setjmp}. It cannot know  
 3318 where @code{longjmp} will be called; in fact, a signal handler could  
 3319 call it at any point in the code. As a result, you may get a warning  
 3320 even when there is in fact no problem because @code{longjmp} cannot  
 3321 in fact be called at the place which would cause a problem.

3323 Some spurious warnings can be avoided if you declare all the functions  
 3324 you use that never return as @code{noreturn}. @xref{Function

3325 Attributes}.

3327 This warning is enabled by @option{-Wall} or @option{-Wextra}.

3329 @item -Wunknown-pragmas  
 3330 @opindex Wunknown-pragmas  
 3331 @opindex Wno-unknown-pragmas  
 3332 @cindex warning for unknown pragmas  
 3333 @cindex unknown pragmas, warning  
 3334 @cindex pragmas, warning of unknown  
 3335 Warn when a #pragma directive is encountered which is not understood by GCC. If this command line option is used, warnings will even be issued for unknown pragmas in system header files. This is not the case if the warnings were only enabled by the @option{-Wall} command line option.

3340 @item -Wno-pragmas  
 3341 @opindex Wno-pragmas  
 3342 @opindex Wpragmas  
 3343 Do not warn about misuses of pragmas, such as incorrect parameters, invalid syntax, or conflicts between pragmas. See also @samp{-Wunknown-pragmas}.

3347 @item -Wstrict-aliasing  
 3348 @opindex Wstrict-aliasing  
 3349 @opindex Wno-strict-aliasing  
 3350 This option is only active when @option{-fstrict-aliasing} is active.  
 3351 It warns about code which might break the strict aliasing rules that the compiler is using for optimization. The warning does not catch all cases, but does attempt to catch the more common pitfalls. It is included in @option{-Wall}.

3355 It is equivalent to @option{-Wstrict-aliasing=3}

3357 @item -Wstrict-aliasing=n  
 3358 @opindex Wstrict-aliasing=n  
 3359 @opindex Wno-strict-aliasing=n  
 3360 This option is only active when @option{-fstrict-aliasing} is active.  
 3361 It warns about code which might break the strict aliasing rules that the compiler is using for optimization.  
 3362 Higher levels correspond to higher accuracy (fewer false positives).  
 3363 Higher levels also correspond to more effort, similar to the way -O works.  
 3364 @option{-Wstrict-aliasing} is equivalent to @option{-Wstrict-aliasing=n}, with n=3.

3368 Level 1: Most aggressive, quick, least accurate.  
 3369 Possibly useful when higher levels  
 3370 do not warn but -fstrict-aliasing still breaks the code, as it has very few false negatives. However, it has many false positives.  
 3371 Warns for all pointer conversions between possibly incompatible types, even if never dereferenced. Runs in the frontend only.

3375 Level 2: Aggressive, quick, not too precise.  
 3376 May still have many false positives (not as many as level 1 though), and few false negatives (but possibly more than level 1).  
 3377 Unlike level 1, it only warns when an address is taken. Warns about incomplete types. Runs in the frontend only.

3381 Level 3 (default for @option{-Wstrict-aliasing}):  
 3382 Should have very few false positives and few false negatives. Slightly slower than levels 1 or 2 when optimization is enabled.  
 3383 Takes care of the common punn+dereference pattern in the frontend:  
 3384 @code{\*(int\*)&some\_float}.

3386 If optimization is enabled, it also runs in the backend, where it deals with multiple statement cases using flow-sensitive points-to information.  
 3387 Only warns when the converted pointer is dereferenced.  
 3388 Does not warn about incomplete types.

3391 @item -Wstrict-overflow  
 3392 @itemx -Wstrict-overflow=@var{n}  
 3393 @opindex Wstrict-overflow  
 3394 @opindex Wno-strict-overflow  
 3395 This option is only active when @option{-fstrict-overflow} is active.  
 3396 It warns about cases where the compiler optimizes based on the assumption that signed overflow does not occur. Note that it does not warn about all cases where the code might overflow: it only warns about cases where the compiler implements some optimization. Thus this warning depends on the optimization level.

3402 An optimization which assumes that signed overflow does not occur is perfectly safe if the values of the variables involved are such that overflow never does, in fact, occur. Therefore this warning can easily give a false positive: a warning about code which is not actually a problem. To help focus on important issues, several warning levels are defined. No warnings are issued for the use of undefined signed overflow when estimating how many iterations a loop will require, in particular when determining whether a loop will be executed at all.

3412 @table @gcctabopt  
 3413 @item -Wstrict-overflow=1  
 3414 Warn about cases which are both questionable and easy to avoid. For example: @code{x + 1 > x}; with @option{-fstrict-overflow}, the compiler will simplify this to @code{1}. This level of @option{-Wstrict-overflow} is enabled by @option{-Wall}; higher levels are not, and must be explicitly requested.

3420 @item -Wstrict-overflow=2  
 3421 Also warn about other cases where a comparison is simplified to a constant. For example: @code{abs(x) >= 0}. This can only be simplified when @option{-fstrict-overflow} is in effect, because @code{abs(INT\_MIN)} overflows to @code{INT\_MIN}, which is less than zero. @option{-Wstrict-overflow} (with no level) is the same as @option{-Wstrict-overflow=2}.

3428 @item -Wstrict-overflow=3  
 3429 Also warn about other cases where a comparison is simplified. For example: @code{x + 1 > 1} will be simplified to @code{x > 0}.

3432 @item -Wstrict-overflow=4  
 3433 Also warn about other simplifications not covered by the above cases.  
 3434 For example: @code{(x \* 10) / 5} will be simplified to @code{x \* 2}.

3436 @item -Wstrict-overflow=5  
 3437 Also warn about cases where the compiler reduces the magnitude of a constant involved in a comparison. For example: @code{x + 2 > y} will be simplified to @code{x + 1 >= y}. This is reported only at the highest warning level because this simplification applies to many comparisons, so this warning level will give a very large number of false positives.

3443 @end table

3445 @item -Warray-bounds  
 3446 @opindex Wno-array-bounds  
 3447 @opindex Warray-bounds  
 3448 This option is only active when @option{-ftree-vcv} is active  
 3449 (default for -O2 and above). It warns about subscripts to arrays that are always out of bounds. This warning is enabled by @option{-Wall}.

3452 @item -Wno-div-by-zero  
 3453 @opindex Wno-div-by-zero  
 3454 @opindex Wdiv-by-zero  
 3455 Do not warn about compile-time integer division by zero. Floating point division by zero is not warned about, as it can be a legitimate way of

3457 obtaining infinities and NaNs.

3459 @item -Wsystem-headers  
 3460 @opindex Wsystem-headers  
 3461 @opindex Wno-system-headers  
 3462 @cindex warnings from system headers  
 3463 @cindex system headers, warnings from  
 3464 Print warning messages for constructs found in system header files.  
 3465 Warnings from system headers are normally suppressed, on the assumption  
 3466 that they usually do not indicate real problems and would only make the  
 3467 compiler output harder to read. Using this command line option tells  
 3468 GCC to emit warnings from system headers as if they occurred in user  
 3469 code. However, note that using @option{-Wall} in conjunction with this  
 3470 option will @emph{not} warn about unknown pragmas in system  
 3471 headers---for that, @option{-Wunknown-pragmas} must also be used.

3473 @item -Wfloat-equal  
 3474 @opindex Wfloat-equal  
 3475 @opindex Wno-float-equal  
 3476 Warn if floating point values are used in equality comparisons.

3478 The idea behind this is that sometimes it is convenient (for the  
 3479 programmer) to consider floating-point values as approximations to  
 3480 infinitely precise real numbers. If you are doing this, then you need  
 3481 to compute (by analyzing the code, or in some other way) the maximum or  
 3482 likely maximum error that the computation introduces, and allow for it  
 3483 when performing comparisons (and when producing output, but that's a  
 3484 different problem). In particular, instead of testing for equality, you  
 3485 would check to see whether the two values have ranges that overlap; and  
 3486 this is done with the relational operators, so equality comparisons are  
 3487 probably mistaken.

3489 @item -Wtraditional @r{(C and Objective-C only)}  
 3490 @opindex Wtraditional  
 3491 @opindex Wno-traditional  
 3492 Warn about certain constructs that behave differently in traditional and  
 3493 ISO C@. Also warn about ISO C constructs that have no traditional C  
 3494 equivalent, and/or problematic constructs which should be avoided.

3496 @itemize @bullet  
 3497 @item  
 3498 Macro parameters that appear within string literals in the macro body.  
 3499 In traditional C macro replacement takes place within string literals,  
 3500 but does not in ISO C@.

3502 @item  
 3503 In traditional C, some preprocessor directives did not exist.  
 3504 Traditional preprocessors would only consider a line to be a directive  
 3505 if the @samp{#} appeared in column 1 on the line. Therefore  
 3506 @option{-Wtraditional} warns about directives that traditional C  
 3507 understands but would ignore because the @samp{#} does not appear as the  
 3508 first character on the line. It also suggests you hide directives like  
 3509 @samp{#pragma} not understood by traditional C by indenting them. Some  
 3510 traditional implementations would not recognize @samp{#elif}, so it  
 3511 suggests avoiding it altogether.

3513 @item  
 3514 A function-like macro that appears without arguments.

3516 @item  
 3517 The unary plus operator.

3519 @item  
 3520 The @samp{U} integer constant suffix, or the @samp{F} or @samp{L} floating point  
 3521 constant suffixes. (Traditional C does support the @samp{L} suffix on integer  
 3522 constants.) Note, these suffixes appear in macros defined in the system

3523 headers of most modern systems, e.g.@: the @samp{\_MIN}/@samp{\_MAX} macros in @co  
 3524 Use of these macros in user code might normally lead to spurious  
 3525 warnings, however GCC's integrated preprocessor has enough context to  
 3526 avoid warning in these cases.

3528 @item  
 3529 A function declared external in one block and then used after the end of  
 3530 the block.

3532 @item  
 3533 A @code{switch} statement has an operand of type @code{long}.

3535 @item  
 3536 A non-@code{static} function declaration follows a @code{static} one.  
 3537 This construct is not accepted by some traditional C compilers.

3539 @item  
 3540 The ISO type of an integer constant has a different width or  
 3541 signedness from its traditional type. This warning is only issued if  
 3542 the base of the constant is ten. I.e.@: hexadecimal or octal values, which  
 3543 typically represent bit patterns, are not warned about.

3545 @item  
 3546 Usage of ISO string concatenation is detected.

3548 @item  
 3549 Initialization of automatic aggregates.

3551 @item  
 3552 Identifier conflicts with labels. Traditional C lacks a separate  
 3553 namespace for labels.

3555 @item  
 3556 Initialization of unions. If the initializer is zero, the warning is  
 3557 omitted. This is done under the assumption that the zero initializer in  
 3558 user code appears conditioned on e.g.@: @code{\_\_STDC\_\_} to avoid missing  
 3559 initializer warnings and relies on default initialization to zero in the  
 3560 traditional C case.

3562 @item  
 3563 Conversions by prototypes between fixed/floating point values and vice  
 3564 versa. The absence of these prototypes when compiling with traditional  
 3565 C would cause serious problems. This is a subset of the possible  
 3566 conversion warnings, for the full set use @option{-Wtraditional-conversion}.

3568 @item  
 3569 Use of ISO C style function definitions. This warning intentionally is  
 3570 @emph{not} issued for prototype declarations or variadic functions  
 3571 because these ISO C features will appear in your code when using  
 3572 libliberty's traditional C compatibility macros, @code{PARAMS} and  
 3573 @code{VPARAMS}. This warning is also bypassed for nested functions  
 3574 because that feature is already a GCC extension and thus not relevant to  
 3575 traditional C compatibility.  
 3576 @end itemize

3578 @item -Wtraditional-conversion @r{(C and Objective-C only)}  
 3579 @opindex Wtraditional-conversion  
 3580 @opindex Wno-traditional-conversion  
 3581 Warn if a prototype causes a type conversion that is different from what  
 3582 would happen to the same argument in the absence of a prototype. This  
 3583 includes conversions of fixed point to floating and vice versa, and  
 3584 conversions changing the width or signedness of a fixed point argument  
 3585 except when the same as the default promotion.

3587 @item -Wdeclaration-after-statement @r{(C and Objective-C only)}  
 3588 @opindex Wdeclaration-after-statement

3589 @opindex Wno-declaration-after-statement  
 3590 Warn when a declaration is found after a statement in a block. This  
 3591 construct, known from C++, was introduced with ISO C99 and is by default  
 3592 allowed in GCC@. It is not supported by ISO C90 and was not supported by  
 3593 GCC versions before GCC 3.0. @xref{Mixed Declarations}.

3595 @item -Wundef  
 3596 @opindex Wundef  
 3597 @opindex Wno-undef  
 3598 Warn if an undefined identifier is evaluated in an @samp{#if} directive.

3600 @item -Wno-endif-labels  
 3601 @opindex Wno-endif-labels  
 3602 @opindex Wendif-labels  
 3603 Do not warn whenever an @samp{#else} or an @samp{#endif} are followed by text.

3605 @item -Wshadow  
 3606 @opindex Wshadow  
 3607 @opindex Wno-shadow  
 3608 Warn whenever a local variable shadows another local variable, parameter or  
 3609 global variable or whenever a built-in function is shadowed.

3611 @item -Wlarger-than=@var{len}  
 3612 @opindex Wlarger-than=@var{len}  
 3613 @opindex Wlarger-than=@var{len}  
 3614 Warn whenever an object of larger than @var{len} bytes is defined.

3616 @item -Wframe-larger-than=@var{len}  
 3617 @opindex Wframe-larger-than  
 3618 Warn if the size of a function frame is larger than @var{len} bytes.  
 3619 The computation done to determine the stack frame size is approximate  
 3620 and not conservative.  
 3621 The actual requirements may be somewhat greater than @var{len}  
 3622 even if you do not get a warning. In addition, any space allocated  
 3623 via @code{alloca}, variable-length arrays, or related constructs  
 3624 is not included by the compiler when determining  
 3625 whether or not to issue a warning.

3627 @item -Wunsafe-loop-optimizations  
 3628 @opindex Wunsafe-loop-optimizations  
 3629 @opindex Wno-unsafe-loop-optimizations  
 3630 Warn if the loop cannot be optimized because the compiler could not  
 3631 assume anything on the bounds of the loop indices. With  
 3632 @option{-funsafe-loop-optimizations} warn if the compiler made  
 3633 such assumptions.

3635 @item -Wno-pedantic-ms-format @r{(MinGW targets only)}  
 3636 @opindex Wno-pedantic-ms-format  
 3637 @opindex Wpedantic-ms-format  
 3638 Disables the warnings about non-ISO @code{printf} / @code{scanf} format  
 3639 width specifiers @code{I32}, @code{I64}, and @code{I} used on Windows targets  
 3640 depending on the MS runtime, when you are using the options @option{-Wformat}  
 3641 and @option{-pedantic} without gnu-extensions.

3643 @item -Wpointer-arith  
 3644 @opindex Wpointer-arith  
 3645 @opindex Wno-pointer-arith  
 3646 Warn about anything that depends on the 'size of' a function type or  
 3647 of @code{void}. GNU C assigns these types a size of 1, for  
 3648 convenience in calculations with @code{void \*} pointers and pointers  
 3649 to functions. In C++, warn also when an arithmetic operation involves  
 3650 @code{NULL}. This warning is also enabled by @option{-pedantic}.

3652 @item -Wtype-limits  
 3653 @opindex Wtype-limits  
 3654 @opindex Wno-type-limits

3655 Warn if a comparison is always true or always false due to the limited  
 3656 range of the data type, but do not warn for constant expressions. For  
 3657 example, warn if an unsigned variable is compared against zero with  
 3658 @samp{@<} or @samp{@>=}. This warning is also enabled by  
 3659 @option{-Wextra}.

3661 @item -Wbad-function-cast @r{(C and Objective-C only)}  
 3662 @opindex Wbad-function-cast  
 3663 @opindex Wno-bad-function-cast  
 3664 Warn whenever a function call is cast to a non-matching type.  
 3665 For example, warn if @code{int malloc()} is cast to @code{anything \*}.

3667 @item -Wc++-compat @r{(C and Objective-C only)}  
 3668 Warn about ISO C constructs that are outside of the common subset of  
 3669 ISO C and ISO C++, e.g.@: request for implicit conversion from  
 3670 @code{void \*} to a pointer to non-@code{void} type.

3672 @item -Wc++0x-compat @r{(C++ and Objective-C++ only)}  
 3673 Warn about C++ constructs whose meaning differs between ISO C++ 1998 and  
 3674 ISO C++ 200x, e.g., identifiers in ISO C++ 1998 that will become keywords  
 3675 in ISO C++ 200x. This warning is enabled by @option{-Wall}.

3677 @item -Wcast-qual  
 3678 @opindex Wcast-qual  
 3679 @opindex Wno-cast-qual  
 3680 Warn whenever a pointer is cast so as to remove a type qualifier from  
 3681 the target type. For example, warn if a @code{const char \*} is cast  
 3682 to an ordinary @code{char \*}.

3684 @item -Wcast-align  
 3685 @opindex Wcast-align  
 3686 @opindex Wno-cast-align  
 3687 Warn whenever a pointer is cast such that the required alignment of the  
 3688 target is increased. For example, warn if a @code{char \*} is cast to  
 3689 an @code{int \*} on machines where integers can only be accessed at  
 3690 two- or four-byte boundaries.

3692 @item -Wwrite-strings  
 3693 @opindex Wwrite-strings  
 3694 @opindex Wno-write-strings  
 3695 When compiling C, give string constants the type @code{const  
 3696 char[@var{length}]} so that copying the address of one into a  
 3697 non-@code{const} @code{char \*} pointer will get a warning. These  
 3698 warnings will help you find at compile time code that can try to write  
 3699 into a string constant, but only if you have been very careful about  
 3700 using @code{const} in declarations and prototypes. Otherwise, it will  
 3701 just be a nuisance. This is why we did not make @option{-Wall} request  
 3702 these warnings.

3704 When compiling C++, warn about the deprecated conversion from string  
 3705 literals to @code{char \*}. This warning is enabled by default for C++  
 3706 programs.

3708 @item -Wclobbered  
 3709 @opindex Wclobbered  
 3710 @opindex Wno-clobbered  
 3711 Warn for variables that might be changed by @samp{longjmp} or  
 3712 @samp{vfork}. This warning is also enabled by @option{-Wextra}.

3714 @item -Wconversion  
 3715 @opindex Wconversion  
 3716 @opindex Wno-conversion  
 3717 Warn for implicit conversions that may alter a value. This includes  
 3718 conversions between real and integer, like @code{abs(x)} when  
 3719 @code{x} is @code{double}; conversions between signed and unsigned,  
 3720 like @code{unsigned ui = -1}; and conversions to smaller types, like

3721 @code{sqrtf (M\_PI)}. Do not warn for explicit casts like @code{abs  
3722 ((int) x)} and @code{ui = (unsigned) -1}, or if the value is not  
3723 changed by the conversion like in @code{abs (2.0)}. Warnings about  
3724 conversions between signed and unsigned integers can be disabled by  
3725 using @option{-Wno-sign-conversion}.

3727 For C++, also warn for conversions between @code{NULL} and non-pointer  
3728 types; confusing overload resolution for user-defined conversions; and  
3729 conversions that will never use a type conversion operator:  
3730 conversions to @code{void}, the same type, a base class or a reference  
3731 to them. Warnings about conversions between signed and unsigned  
3732 integers are disabled by default in C++ unless  
3733 @option{-Wsign-conversion} is explicitly enabled.

3735 @item -Wempty-body  
3736 @opindex Wempty-body  
3737 @opindex Wno-empty-body  
3738 Warn if an empty body occurs in an @samp{if}, @samp{else} or @samp{do  
3739 while} statement. This warning is also enabled by @option{-Wextra}.

3741 @item -Wenum-compare @r{(C++ and Objective-C++ only)}  
3742 @opindex Wenum-compare  
3743 @opindex Wno-enum-compare  
3744 Warn about a comparison between values of different enum types. This  
3745 warning is enabled by default.

3747 @item -Wsign-compare  
3748 @opindex Wsign-compare  
3749 @opindex Wno-sign-compare  
3750 @cindex warning for comparison of signed and unsigned values  
3751 @cindex comparison of signed and unsigned values, warning  
3752 @cindex signed and unsigned values, comparison warning  
3753 Warn when a comparison between signed and unsigned values could produce  
3754 an incorrect result when the signed value is converted to unsigned.  
3755 This warning is also enabled by @option{-Wextra}; to get the other warnings  
3756 of @option{-Wextra} without this warning, use @samp{-Wextra -Wno-sign-compare}.

3758 @item -Wsign-conversion  
3759 @opindex Wsign-conversion  
3760 @opindex Wno-sign-conversion  
3761 Warn for implicit conversions that may change the sign of an integer  
3762 value, like assigning a signed integer expression to an unsigned  
3763 integer variable. An explicit cast silences the warning. In C, this  
3764 option is enabled also by @option{-Wconversion}.

3766 @item -Waddress  
3767 @opindex Waddress  
3768 @opindex Wno-address  
3769 Warn about suspicious uses of memory addresses. These include using  
3770 the address of a function in a conditional expression, such as  
3771 @code{void func(void); if (func)}, and comparisons against the memory  
3772 address of a string literal, such as @code{if (x == "abc")}. Such  
3773 uses typically indicate a programmer error: the address of a function  
3774 always evaluates to true, so their use in a conditional usually  
3775 indicate that the programmer forgot the parentheses in a function  
3776 call; and comparisons against string literals result in unspecified  
3777 behavior and are not portable in C, so they usually indicate that the  
3778 programmer intended to use @code{strcmp}. This warning is enabled by  
3779 @option{-Wall}.

3781 @item -Wlogical-op  
3782 @opindex Wlogical-op  
3783 @opindex Wno-logical-op  
3784 Warn about suspicious uses of logical operators in expressions.  
3785 This includes using logical operators in contexts where a  
3786 bit-wise operator is likely to be expected.

3788 @item -Waggregate-return  
3789 @opindex Waggregate-return  
3790 @opindex Wno-aggregate-return  
3791 Warn if any functions that return structures or unions are defined or  
3792 called. (In languages where you can return an array, this also elicits  
3793 a warning.)

3795 @item -Wno-attributes  
3796 @opindex Wno-attributes  
3797 @opindex Wattributes  
3798 Do not warn if an unexpected @code{\_\_attribute\_\_} is used, such as  
3799 unrecognized attributes, function attributes applied to variables,  
3800 etc. This will not stop errors for incorrect use of supported  
3801 attributes.

3803 @item -Wno-builtin-macro-redefined  
3804 @opindex Wno-builtin-macro-redefined  
3805 @opindex Wbuiltin-macro-redefined  
3806 Do not warn if certain built-in macros are redefined. This suppresses  
3807 warnings for redefinition of @code{\_\_TIMESTAMP\_\_}, @code{\_\_TIME\_\_},  
3808 @code{\_\_DATE\_\_}, @code{\_\_FILE\_\_}, and @code{\_\_BASE\_FILE\_\_}.

3810 @item -Wstrict-prototypes @r{(C and Objective-C only)}  
3811 @opindex Wstrict-prototypes  
3812 @opindex Wno-strict-prototypes  
3813 Warn if a function is declared or defined without specifying the  
3814 argument types. (An old-style function definition is permitted without  
3815 a warning if preceded by a declaration which specifies the argument  
3816 types.)

3818 @item -Wold-style-declaration @r{(C and Objective-C only)}  
3819 @opindex Wold-style-declaration  
3820 @opindex Wno-old-style-declaration  
3821 Warn for obsolescent usages, according to the C Standard, in a  
3822 declaration. For example, warn if storage-class specifiers like  
3823 @code{static} are not the first things in a declaration. This warning  
3824 is also enabled by @option{-Wextra}.

3826 @item -Wold-style-definition @r{(C and Objective-C only)}  
3827 @opindex Wold-style-definition  
3828 @opindex Wno-old-style-definition  
3829 Warn if an old-style function definition is used. A warning is given  
3830 even if there is a previous prototype.

3832 @item -Wmissing-parameter-type @r{(C and Objective-C only)}  
3833 @opindex Wmissing-parameter-type  
3834 @opindex Wno-missing-parameter-type  
3835 A function parameter is declared without a type specifier in K&R-style  
3836 functions:

3838 @smalllexample  
3839 void foo(bar) @{ @}  
3840 @end smalllexample

3842 This warning is also enabled by @option{-Wextra}.

3844 @item -Wmissing-prototypes @r{(C and Objective-C only)}  
3845 @opindex Wmissing-prototypes  
3846 @opindex Wno-missing-prototypes  
3847 Warn if a global function is defined without a previous prototype  
3848 declaration. This warning is issued even if the definition itself  
3849 provides a prototype. The aim is to detect global functions that fail  
3850 to be declared in header files.

3852 @item -Wmissing-declarations

3853 @opindex Wmissing-declarations  
3854 @opindex Wno-missing-declarations  
3855 Warn if a global function is defined without a previous declaration.  
3856 Do so even if the definition itself provides a prototype.  
3857 Use this option to detect global functions that are not declared in  
3858 header files. In C++, no warnings are issued for function templates,  
3859 or for inline functions, or for functions in anonymous namespaces.

3861 @item -Wmissing-field-initializers  
3862 @opindex Wmissing-field-initializers  
3863 @opindex Wno-missing-field-initializers  
3864 @opindex W  
3865 @opindex Wextra  
3866 @opindex Wno-extra  
3867 Warn if a structure's initializer has some fields missing. For  
3868 example, the following code would cause such a warning, because  
3869 @code{x.h} is implicitly zero:

```
3871 @smallexample
3872 struct s @ { int f, g, h; @ };
3873 struct s x = @ { 3, 4 @ };
3874 @end smallexample
```

3876 This option does not warn about designated initializers, so the following  
3877 modification would not trigger a warning:

```
3879 @smallexample
3880 struct s @ { int f, g, h; @ };
3881 struct s x = @ { .f = 3, .g = 4 @ };
3882 @end smallexample
```

3884 This warning is included in @option{-Wextra}. To get other @option{-Wextra}  
3885 warnings without this one, use @samp{-Wextra -Wno-missing-field-initializers}.

3887 @item -Wmissing-noreturn  
3888 @opindex Wmissing-noreturn  
3889 @opindex Wno-missing-noreturn  
3890 Warn about functions which might be candidates for attribute @code{noreturn}.  
3891 Note these are only possible candidates, not absolute ones. Care should  
3892 be taken to manually verify functions actually do not ever return before  
3893 adding the @code{noreturn} attribute, otherwise subtle code generation  
3894 bugs could be introduced. You will not get a warning for @code{main} in  
3895 hosted C environments.

3897 @item -Wmissing-format-attribute  
3898 @opindex Wmissing-format-attribute  
3899 @opindex Wno-missing-format-attribute  
3900 @opindex Wformat  
3901 @opindex Wno-format  
3902 Warn about function pointers which might be candidates for @code{format}  
3903 attributes. Note these are only possible candidates, not absolute ones.  
3904 GCC will guess that function pointers with @code{format} attributes that  
3905 are used in assignment, initialization, parameter passing or return  
3906 statements should have a corresponding @code{format} attribute in the  
3907 resulting type. I.e.@: the left-hand side of the assignment or  
3908 initialization, the type of the parameter variable, or the return type  
3909 of the containing function respectively should also have a @code{format}  
3910 attribute to avoid the warning.

3912 GCC will also warn about function definitions which might be  
3913 candidates for @code{format} attributes. Again, these are only  
3914 possible candidates. GCC will guess that @code{format} attributes  
3915 might be appropriate for any function that calls a function like  
3916 @code{vprintf} or @code{vscanf}, but this might not always be the  
3917 case, and some functions for which @code{format} attributes are  
3918 appropriate may not be detected.

3920 @item -Wno-multichar  
3921 @opindex Wno-multichar  
3922 @opindex Wmultichar  
3923 Do not warn if a multicharacter constant (@samp{'FOOF'}) is used.  
3924 Usually they indicate a typo in the user's code, as they have  
3925 implementation-defined values, and should not be used in portable code.

3927 @item -Wnormalized=<none|id|nfc|nfkc>  
3928 @opindex Wnormalized=  
3929 @cindex NFC  
3930 @cindex NFKC  
3931 @cindex character set, input normalization  
3932 In ISO C and ISO C++, two identifiers are different if they are  
3933 different sequences of characters. However, sometimes when characters  
3934 outside the basic ASCII character set are used, you can have two  
3935 different character sequences that look the same. To avoid confusion,  
3936 the ISO 10646 standard sets out some @dfn{normalization rules} which  
3937 when applied ensure that two sequences that look the same are turned into  
3938 the same sequence. GCC can warn you if you are using identifiers which  
3939 have not been normalized; this option controls that warning.

3941 There are four levels of warning that GCC supports. The default is  
3942 @option{-Wnormalized=nfc}, which warns about any identifier which is  
3943 not in the ISO 10646 'C' normalized form, @dfn{NFC}. NFC is the  
3944 recommended form for most uses.

3946 Unfortunately, there are some characters which ISO C and ISO C++ allow  
3947 in identifiers that when turned into NFC aren't allowable as  
3948 identifiers. That is, there's no way to use these symbols in portable  
3949 ISO C or C++ and have all your identifiers in NFC@.  
3950 @option{-Wnormalized=id} suppresses the warning for these characters.  
3951 It is hoped that future versions of the standards involved will correct  
3952 this, which is why this option is not the default.

3954 You can switch the warning off for all characters by writing  
3955 @option{-Wnormalized=none}. You would only want to do this if you  
3956 were using some other normalization scheme (like 'D'), because  
3957 otherwise you can easily create bugs that are literally impossible to see.

3959 Some characters in ISO 10646 have distinct meanings but look identical  
3960 in some fonts or display methodologies, especially once formatting has  
3961 been applied. For instance @code{\u207F}, 'SUPERSCRIPT LATIN SMALL  
3962 LETTER N', will display just like a regular @code{n} which has been  
3963 placed in a superscript. ISO 10646 defines the @dfn{NFKC}  
3964 normalization scheme to convert all these into a standard form as  
3965 well, and GCC will warn if your code is not in NFKC if you use  
3966 @option{-Wnormalized=nfkc}. This warning is comparable to warning  
3967 about every identifier that contains the letter O because it might be  
3968 confused with the digit 0, and so is not the default, but may be  
3969 useful as a local coding convention if the programming environment is  
3970 unable to be fixed to display these characters distinctly.

3972 @item -Wno-deprecated  
3973 @opindex Wno-deprecated  
3974 @opindex Wdeprecated  
3975 Do not warn about usage of deprecated features. @xref{Deprecated Features}.

3977 @item -Wno-deprecated-declarations  
3978 @opindex Wno-deprecated-declarations  
3979 @opindex Wdeprecated-declarations  
3980 Do not warn about uses of functions (@pxref{Function Attributes}),  
3981 variables (@pxref{Variable Attributes}), and types (@pxref{Type  
3982 Attributes}) marked as deprecated by using the @code{deprecated}  
3983 attribute.

```

3985 @item -Wno-overflow
3986 @opindex Wno-overflow
3987 @opindex Woverflow
3988 Do not warn about compile-time overflow in constant expressions.

3990 @item -Woverride-init @r{(C and Objective-C only)}
3991 @opindex Woverride-init
3992 @opindex Wno-override-init
3993 @opindex W
3994 @opindex Wextra
3995 @opindex Wno-extra
3996 Warn if an initialized field without side effects is overridden when
3997 using designated initializers (@pxref{Designated Inits, , Designated
3998 Initializers}).

4000 This warning is included in @option{-Wextra}. To get other
4001 @option{-Wextra} warnings without this one, use @samp{-Wextra
4002 -Wno-override-init}.

4004 @item -Wpacked
4005 @opindex Wpacked
4006 @opindex Wno-packed
4007 Warn if a structure is given the packed attribute, but the packed
4008 attribute has no effect on the layout or size of the structure.
4009 Such structures may be mis-aligned for little benefit. For
4010 instance, in this code, the variable @code{f.x} in @code{struct bar}
4011 will be misaligned even though @code{struct bar} does not itself
4012 have the packed attribute:

4014 @smallexample
4015 @group
4016 struct foo @{
4017     int x;
4018     char a, b, c, d;
4019 } __attribute__((packed));
4020 struct bar @{
4021     char z;
4022     struct foo f;
4023 };
4024 @end group
4025 @end smallexample

4027 @item -Wpacked-bitfield-compat
4028 @opindex Wpacked-bitfield-compat
4029 @opindex Wno-packed-bitfield-compat
4030 The 4.1, 4.2 and 4.3 series of GCC ignore the @code{packed} attribute
4031 on bit-fields of type @code{char}. This has been fixed in GCC 4.4 but
4032 the change can lead to differences in the structure layout. GCC
4033 informs you when the offset of such a field has changed in GCC 4.4.
4034 For example there is no longer a 4-bit padding between field @code{a}
4035 and @code{b} in this structure:

4037 @smallexample
4038 struct foo
4039 @{
4040     char a:4;
4041     char b:8;
4042 } __attribute__((packed));
4043 @end smallexample

4045 This warning is enabled by default. Use
4046 @option{-Wno-packed-bitfield-compat} to disable this warning.

4048 @item -Wpadded
4049 @opindex Wpadded
4050 @opindex Wno-padded

```

```

4051 Warn if padding is included in a structure, either to align an element
4052 of the structure or to align the whole structure. Sometimes when this
4053 happens it is possible to rearrange the fields of the structure to
4054 reduce the padding and so make the structure smaller.

4056 @item -Wredundant-decls
4057 @opindex Wredundant-decls
4058 @opindex Wno-redundant-decls
4059 Warn if anything is declared more than once in the same scope, even in
4060 cases where multiple declaration is valid and changes nothing.

4062 @item -Wnested-externs @r{(C and Objective-C only)}
4063 @opindex Wnested-externs
4064 @opindex Wno-nested-externs
4065 Warn if an @code{extern} declaration is encountered within a function.

4067 @item -Wunreachable-code
4068 @opindex Wunreachable-code
4069 @opindex Wno-unreachable-code
4070 Warn if the compiler detects that code will never be executed.

4072 This option is intended to warn when the compiler detects that at
4073 least a whole line of source code will never be executed, because
4074 some condition is never satisfied or because it is after a
4075 procedure that never returns.

4077 It is possible for this option to produce a warning even though there
4078 are circumstances under which part of the affected line can be executed,
4079 so care should be taken when removing apparently-unreachable code.

4081 For instance, when a function is inlined, a warning may mean that the
4082 line is unreachable in only one inlined copy of the function.

4084 This option is not made part of @option{-Wall} because in a debugging
4085 version of a program there is often substantial code which checks
4086 correct functioning of the program and is, hopefully, unreachable
4087 because the program does work. Another common use of unreachable
4088 code is to provide behavior which is selectable at compile-time.

4090 @item -Winline
4091 @opindex Winline
4092 @opindex Wno-inline
4093 Warn if a function can not be inlined and it was declared as inline.
4094 Even with this option, the compiler will not warn about failures to
4095 inline functions declared in system headers.

4097 The compiler uses a variety of heuristics to determine whether or not
4098 to inline a function. For example, the compiler takes into account
4099 the size of the function being inlined and the amount of inlining
4100 that has already been done in the current function. Therefore,
4101 seemingly insignificant changes in the source program can cause the
4102 warnings produced by @option{-Winline} to appear or disappear.

4104 @item -Wno-invalid-offsetof @r{(C++ and Objective-C++ only)}
4105 @opindex Wno-invalid-offsetof
4106 @opindex Winvalid-offsetof
4107 Suppress warnings from applying the @samp{offsetof} macro to a non-POD
4108 type. According to the 1998 ISO C++ standard, applying @samp{offsetof}
4109 to a non-POD type is undefined. In existing C++ implementations,
4110 however, @samp{offsetof} typically gives meaningful results even when
4111 applied to certain kinds of non-POD types. (Such as a simple
4112 @samp{struct} that fails to be a POD type only by virtue of having a
4113 constructor.) This flag is for users who are aware that they are
4114 writing nonportable code and who have deliberately chosen to ignore the
4115 warning about it.

```



4117 The restrictions on @samp{offsetof} may be relaxed in a future version  
 4118 of the C++ standard.

4120 @item -Wno-int-to-pointer-cast @r{(C and Objective-C only)}  
 4121 @opindex Wno-int-to-pointer-cast  
 4122 @opindex Wint-to-pointer-cast  
 4123 Suppress warnings from casts to pointer type of an integer of a  
 4124 different size.

4126 @item -Wno-pointer-to-int-cast @r{(C and Objective-C only)}  
 4127 @opindex Wno-pointer-to-int-cast  
 4128 @opindex Wpointer-to-int-cast  
 4129 Suppress warnings from casts from a pointer to an integer type of a  
 4130 different size.

4132 @item -Winvalid-pch  
 4133 @opindex Winvalid-pch  
 4134 @opindex Wno-invalid-pch  
 4135 Warn if a precompiled header (@pxref{Precompiled Headers}) is found in  
 4136 the search path but can't be used.

4138 @item -Wlong-long  
 4139 @opindex Wlong-long  
 4140 @opindex Wno-long-long  
 4141 Warn if @samp{long long} type is used. This is default. To inhibit  
 4142 the warning messages, use @option{-Wno-long-long}. Flags  
 4143 @option{-Wlong-long} and @option{-Wno-long-long} are taken into account  
 4144 only when @option{-pedantic} flag is used.

4146 @item -Wvariadic-macros  
 4147 @opindex Wvariadic-macros  
 4148 @opindex Wno-variadic-macros  
 4149 Warn if variadic macros are used in pedantic ISO C90 mode, or the GNU  
 4150 alternate syntax when in pedantic ISO C99 mode. This is default.  
 4151 To inhibit the warning messages, use @option{-Wno-variadic-macros}.

4153 @item -Wvla  
 4154 @opindex Wvla  
 4155 @opindex Wno-vla  
 4156 Warn if variable length array is used in the code.  
 4157 @option{-Wno-vla} will prevent the @option{-pedantic} warning of  
 4158 the variable length array.

4160 @item -Wvolatile-register-var  
 4161 @opindex Wvolatile-register-var  
 4162 @opindex Wno-volatile-register-var  
 4163 Warn if a register variable is declared volatile. The volatile  
 4164 modifier does not inhibit all optimizations that may eliminate reads  
 4165 and/or writes to register variables. This warning is enabled by  
 4166 @option{-Wall}.

4168 @item -Wdisabled-optimization  
 4169 @opindex Wdisabled-optimization  
 4170 @opindex Wno-disabled-optimization  
 4171 Warn if a requested optimization pass is disabled. This warning does  
 4172 not generally indicate that there is anything wrong with your code; it  
 4173 merely indicates that GCC's optimizers were unable to handle the code  
 4174 effectively. Often, the problem is that your code is too big or too  
 4175 complex; GCC will refuse to optimize programs when the optimization  
 4176 itself is likely to take inordinate amounts of time.

4178 @item -Wpointer-sign @r{(C and Objective-C only)}  
 4179 @opindex Wpointer-sign  
 4180 @opindex Wno-pointer-sign  
 4181 Warn for pointer argument passing or assignment with different signedness.  
 4182 This option is only supported for C and Objective-C. It is implied by

4183 @option{-Wall} and by @option{-pedantic}, which can be disabled with  
 4184 @option{-Wno-pointer-sign}.

4186 @item -Wstack-protector  
 4187 @opindex Wstack-protector  
 4188 @opindex Wno-stack-protector  
 4189 This option is only active when @option{-fstack-protector} is active. It  
 4190 warns about functions that will not be protected against stack smashing.

4192 @item -Wno-mudflap  
 4193 @opindex Wno-mudflap  
 4194 Suppress warnings about constructs that cannot be instrumented by  
 4195 @option{-fmudflap}.

4197 @item -Woverlength-strings  
 4198 @opindex Woverlength-strings  
 4199 @opindex Wno-overlength-strings  
 4200 Warn about string constants which are longer than the 'minimum  
 4201 maximum' length specified in the C standard. Modern compilers  
 4202 generally allow string constants which are much longer than the  
 4203 standard's minimum limit, but very portable programs should avoid  
 4204 using longer strings.

4206 The limit applies @emph{after} string constant concatenation, and does  
 4207 not count the trailing NUL. In C89, the limit was 509 characters; in  
 4208 C99, it was raised to 4095. C++98 does not specify a normative  
 4209 minimum maximum, so we do not diagnose overlength strings in C++.

4211 This option is implied by @option{-pedantic}, and can be disabled with  
 4212 @option{-Wno-overlength-strings}.

4213 @end table

4215 @node Debugging Options  
 4216 @section Options for Debugging Your Program or GCC  
 4217 @cindex options, debugging  
 4218 @cindex debugging information options

4220 GCC has various special options that are used for debugging  
 4221 either your program or GCC:

4223 @table @gcctabopt  
 4224 @item -g  
 4225 @opindex g  
 4226 Produce debugging information in the operating system's native format  
 4227 (stabs, COFF, XCOFF, or DWARF 2). GDB can work with this debugging  
 4228 information.

4230 On most systems that use stabs format, @option{-g} enables use of extra  
 4231 debugging information that only GDB can use; this extra information  
 4232 makes debugging work better in GDB but will probably make other debuggers  
 4233 crash or  
 4234 refuse to read the program. If you want to control for certain whether  
 4235 to generate the extra information, use @option{-gstabs+}, @option{-gstabs},  
 4236 @option{-gxcoff+}, @option{-gxcoff}, or @option{-gvms} (see below).

4238 GCC allows you to use @option{-g} with  
 4239 @option{-O}. The shortcuts taken by optimized code may occasionally  
 4240 produce surprising results: some variables you declared may not exist  
 4241 at all; flow of control may briefly move where you did not expect it;  
 4242 some statements may not be executed because they compute constant  
 4243 results or their values were already at hand; some statements may  
 4244 execute in different places because they were moved out of loops.

4246 Nevertheless it proves possible to debug optimized output. This makes  
 4247 it reasonable to use the optimizer for programs that might have bugs.

4249 The following options are useful when GCC is generated with the  
 4250 capability for more than one debugging format.

4252 @item -ggdb  
 4253 @opindex ggdb  
 4254 Produce debugging information for use by GDB@. This means to use the  
 4255 most expressive format available (DWARF 2, stabs, or the native format  
 4256 if neither of those are supported), including GDB extensions if at all  
 4257 possible.

4259 @item -gstabs  
 4260 @opindex gstabs  
 4261 Produce debugging information in stabs format (if that is supported),  
 4262 without GDB extensions. This is the format used by DBX on most BSD  
 4263 systems. On MIPS, Alpha and System V Release 4 systems this option  
 4264 produces stabs debugging output which is not understood by DBX or SDB@.  
 4265 On System V Release 4 systems this option requires the GNU assembler.

4267 @item -feliminate-unused-debug-symbols  
 4268 @opindex feliminate-unused-debug-symbols  
 4269 Produce debugging information in stabs format (if that is supported),  
 4270 for only symbols that are actually used.

4272 @item -femit-class-debug-always  
 4273 Instead of emitting debugging information for a C++ class in only one  
 4274 object file, emit it in all object files using the class. This option  
 4275 should be used only with debuggers that are unable to handle the way GCC  
 4276 normally emits debugging information for classes because using this  
 4277 option will increase the size of debugging information by as much as a  
 4278 factor of two.

4280 @item -gstabs+  
 4281 @opindex gstabs+  
 4282 Produce debugging information in stabs format (if that is supported),  
 4283 using GNU extensions understood only by the GNU debugger (GDB)@. The  
 4284 use of these extensions is likely to make other debuggers crash or  
 4285 refuse to read the program.

4287 @item -gcoff  
 4288 @opindex gcoff  
 4289 Produce debugging information in COFF format (if that is supported).  
 4290 This is the format used by SDB on most System V systems prior to  
 4291 System V Release 4.

4293 @item -gxcoff  
 4294 @opindex gxcoff  
 4295 Produce debugging information in XCOFF format (if that is supported).  
 4296 This is the format used by the DBX debugger on IBM RS/6000 systems.

4298 @item -gxcoff+  
 4299 @opindex gxcoff+  
 4300 Produce debugging information in XCOFF format (if that is supported),  
 4301 using GNU extensions understood only by the GNU debugger (GDB)@. The  
 4302 use of these extensions is likely to make other debuggers crash or  
 4303 refuse to read the program, and may cause assemblers other than the GNU  
 4304 assembler (GAS) to fail with an error.

4306 @item -gdwarf-2  
 4307 @opindex gdwarf-2  
 4308 Produce debugging information in DWARF version 2 format (if that is  
 4309 supported). This is the format used by DBX on IRIX 6. With this  
 4310 option, GCC uses features of DWARF version 3 when they are useful;  
 4311 version 3 is upward compatible with version 2, but may still cause  
 4312 problems for older debuggers.

4314 @item -gvms

4315 @opindex gvms  
 4316 Produce debugging information in VMS debug format (if that is  
 4317 supported). This is the format used by DEBUG on VMS systems.

4319 @item -g@var{level}  
 4320 @itemx -ggdb@var{level}  
 4321 @itemx -gstabs@var{level}  
 4322 @itemx -gcoff@var{level}  
 4323 @itemx -gxcoff@var{level}  
 4324 @itemx -gvms@var{level}  
 4325 Request debugging information and also use @var{level} to specify how  
 4326 much information. The default level is 2.

4328 Level 0 produces no debug information at all. Thus, @option{-g0} negates  
 4329 @option{-g}.

4331 Level 1 produces minimal information, enough for making backtraces in  
 4332 parts of the program that you don't plan to debug. This includes  
 4333 descriptions of functions and external variables, but no information  
 4334 about local variables and no line numbers.

4336 Level 3 includes extra information, such as all the macro definitions  
 4337 present in the program. Some debuggers support macro expansion when  
 4338 you use @option{-g3}.

4340 @option{-gdwarf-2} does not accept a concatenated debug level, because  
 4341 GCC used to support an option @option{-gdwarf} that meant to generate  
 4342 debug information in version 1 of the DWARF format (which is very  
 4343 different from version 2), and it would have been too confusing. That  
 4344 debug format is long obsolete, but the option cannot be changed now.  
 4345 Instead use an additional @option{-g@var{level}} option to change the  
 4346 debug level for DWARF2.

4348 @item -feliminate-dwarf2-dups  
 4349 @opindex feliminate-dwarf2-dups  
 4350 Compress DWARF2 debugging information by eliminating duplicated  
 4351 information about each symbol. This option only makes sense when  
 4352 generating DWARF2 debugging information with @option{-gdwarf-2}.

4354 @item -femit-struct-debug-baseonly  
 4355 Emit debug information for struct-like types  
 4356 only when the base name of the compilation source file  
 4357 matches the base name of file in which the struct was defined.

4359 This option substantially reduces the size of debugging information,  
 4360 but at significant potential loss in type information to the debugger.  
 4361 See @option{-femit-struct-debug-reduced} for a less aggressive option.  
 4362 See @option{-femit-struct-debug-detailed} for more detailed control.

4364 This option works only with DWARF 2.

4366 @item -femit-struct-debug-reduced  
 4367 Emit debug information for struct-like types  
 4368 only when the base name of the compilation source file  
 4369 matches the base name of file in which the type was defined,  
 4370 unless the struct is a template or defined in a system header.

4372 This option significantly reduces the size of debugging information,  
 4373 with some potential loss in type information to the debugger.  
 4374 See @option{-femit-struct-debug-baseonly} for a more aggressive option.  
 4375 See @option{-femit-struct-debug-detailed} for more detailed control.

4377 This option works only with DWARF 2.

4379 @item -femit-struct-debug-detailed@r{[ ]=@var{spec-list}@r{[ ]}  
 4380 Specify the struct-like types

4381 for which the compiler will generate debug information.  
 4382 The intent is to reduce duplicate struct debug information  
 4383 between different object files within the same program.

4385 This option is a detailed version of  
 4386 `@option{-femit-struct-debug-reduced}` and `@option{-femit-struct-debug-baseonly}`,  
 4387 which will serve for most needs.

4389 A specification has the syntax  
 4390 `[@samp{dir:}|@samp{ind:}][@samp{ord:}|@samp{gen:}](@samp{any}|@samp{sys}|@samp{b`

4392 The optional first word limits the specification to  
 4393 structs that are used directly (`@samp{dir:}`) or used indirectly (`@samp{ind:}`).  
 4394 A struct type is used directly when it is the type of a variable, member.  
 4395 Indirect uses arise through pointers to structs.  
 4396 That is, when use of an incomplete struct would be legal, the use is indirect.  
 4397 An example is  
 4398 `@samp{struct one direct; struct two * indirect;}`.

4400 The optional second word limits the specification to  
 4401 ordinary structs (`@samp{ord:}`) or generic structs (`@samp{gen:}`).  
 4402 Generic structs are a bit complicated to explain.  
 4403 For C++, these are non-explicit specializations of template classes,  
 4404 or non-template classes within the above.  
 4405 Other programming languages have generics,  
 4406 but `@samp{-femit-struct-debug-detailed}` does not yet implement them.

4408 The third word specifies the source files for those  
 4409 structs for which the compiler will emit debug information.  
 4410 The values `@samp{none}` and `@samp{any}` have the normal meaning.  
 4411 The value `@samp{base}` means that  
 4412 the base of name of the file in which the type declaration appears  
 4413 must match the base of the name of the main compilation file.  
 4414 In practice, this means that  
 4415 types declared in `@file{foo.c}` and `@file{foo.h}` will have debug information,  
 4416 but types declared in other header will not.  
 4417 The value `@samp{sys}` means those types satisfying `@samp{base}`  
 4418 or declared in system or compiler headers.

4420 You may need to experiment to determine the best settings for your application.

4422 The default is `@samp{-femit-struct-debug-detailed=all}`.

4424 This option works only with DWARF 2.

4426 `@item -fno-merge-debug-strings`  
 4427 `@opindex fmerge-debug-strings`  
 4428 `@opindex fno-merge-debug-strings`  
 4429 Direct the linker to not merge together strings in the debugging  
 4430 information which are identical in different object files. Merging is  
 4431 not supported by all assemblers or linkers. Merging decreases the size  
 4432 of the debug information in the output file at the cost of increasing  
 4433 link processing time. Merging is enabled by default.

4435 `@item -fdebug-prefix-map=@var{old}=@var{new}`  
 4436 `@opindex fdebug-prefix-map`  
 4437 When compiling files in directory `@file{@var{old}}`, record debugging  
 4438 information describing them as in `@file{@var{new}}` instead.

4440 `@item -fno-dwarf2-cfi-asm`  
 4441 `@opindex fdwarf2-cfi-asm`  
 4442 `@opindex fno-dwarf2-cfi-asm`  
 4443 Emit DWARF 2 unwind info as compiler generated `@code{.eh_frame}` section  
 4444 instead of using GAS `@code{.cfi_*}` directives.

4446 `@cindex @command{prof}`

4447 `@item -p`  
 4448 `@opindex p`  
 4449 Generate extra code to write profile information suitable for the  
 4450 analysis program `@command{prof}`. You must use this option when compiling  
 4451 the source files you want data about, and you must also use it when  
 4452 linking.

4454 `@cindex @command{gprof}`  
 4455 `@item -pg`  
 4456 `@opindex pg`  
 4457 Generate extra code to write profile information suitable for the  
 4458 analysis program `@command{gprof}`. You must use this option when compiling  
 4459 the source files you want data about, and you must also use it when  
 4460 linking.

4462 `@item -Q`  
 4463 `@opindex Q`  
 4464 Makes the compiler print out each function name as it is compiled, and  
 4465 print some statistics about each pass when it finishes.

4467 `@item -ftime-report`  
 4468 `@opindex ftime-report`  
 4469 Makes the compiler print some statistics about the time consumed by each  
 4470 pass when it finishes.

4472 `@item -fmem-report`  
 4473 `@opindex fmem-report`  
 4474 Makes the compiler print some statistics about permanent memory  
 4475 allocation when it finishes.

4477 `@item -fpre-ipa-mem-report`  
 4478 `@opindex fpre-ipa-mem-report`  
 4479 `@item -fpost-ipa-mem-report`  
 4480 `@opindex fpost-ipa-mem-report`  
 4481 Makes the compiler print some statistics about permanent memory  
 4482 allocation before or after interprocedural optimization.

4484 `@item -fprofile-arcs`  
 4485 `@opindex fprofile-arcs`  
 4486 Add code so that program flow `@dfn{arcs}` are instrumented. During  
 4487 execution the program records how many times each branch and call is  
 4488 executed and how many times it is taken or returns. When the compiled  
 4489 program exits it saves this data to a file called  
 4490 `@file{@var{auxname}.gcda}` for each source file. The data may be used for  
 4491 profile-directed optimizations (`@option{-fbranch-probabilities}`), or for  
 4492 test coverage analysis (`@option{-ftest-coverage}`). Each object file's  
 4493 `@var{auxname}` is generated from the name of the output file, if  
 4494 explicitly specified and it is not the final executable, otherwise it is  
 4495 the basename of the source file. In both cases any suffix is removed  
 4496 (e.g. `@file{foo.gcda}` for input file `@file{dir/foo.c}`, or  
 4497 `@file{dir/foo.gcda}` for output file specified as `@option{-o dir/foo.o}`).  
 4498 `@xref{Cross-profiling}`.

4500 `@cindex @command{gcov}`  
 4501 `@item --coverage`  
 4502 `@opindex coverage`

4504 This option is used to compile and link code instrumented for coverage  
 4505 analysis. The option is a synonym for `@option{-fprofile-arcs}`  
 4506 `@option{-ftest-coverage}` (when compiling) and `@option{-lgcov}` (when  
 4507 linking). See the documentation for those options for more details.

4509 `@itemize`

4511 `@item`  
 4512 Compile the source files with `@option{-fprofile-arcs}` plus optimization

4513 and code generation options. For test coverage analysis, use the  
 4514 additional `@option{-ftest-coverage}` option. You do not need to profile  
 4515 every source file in a program.

4517 @item  
 4518 Link your object files with `@option{-lgcov}` or `@option{-fprofile-arcs}`  
 4519 (the latter implies the former).

4521 @item  
 4522 Run the program on a representative workload to generate the arc profile  
 4523 information. This may be repeated any number of times. You can run  
 4524 concurrent instances of your program, and provided that the file system  
 4525 supports locking, the data files will be correctly updated. Also  
 4526 `@code{fork}` calls are detected and correctly handled (double counting  
 4527 will not happen).

4529 @item  
 4530 For profile-directed optimizations, compile the source files again with  
 4531 the same optimization and code generation options plus  
 4532 `@option{-fbranch-probabilities}` (`@pxref{Optimize Options,,Options that`  
 4533 `Control Optimization}`).

4535 @item  
 4536 For test coverage analysis, use `@command{gcov}` to produce human readable  
 4537 information from the `@file{.gcno}` and `@file{.gda}` files. Refer to the  
 4538 `@command{gcov}` documentation for further information.

4540 @end itemize

4542 With `@option{-fprofile-arcs}`, for each function of your program GCC  
 4543 creates a program flow graph, then finds a spanning tree for the graph.  
 4544 Only arcs that are not on the spanning tree have to be instrumented: the  
 4545 compiler adds code to count the number of times that these arcs are  
 4546 executed. When an arc is the only exit or only entrance to a block, the  
 4547 instrumentation code can be added to the block; otherwise, a new basic  
 4548 block must be created to hold the instrumentation code.

4550 @need 2000  
 4551 @item `-ftest-coverage`  
 4552 @opindex `ftest-coverage`  
 4553 Produce a notes file that the `@command{gcov}` code-coverage utility  
 4554 (`@pxref{Gcov,,@command{gcov}---a Test Coverage Program}`) can use to  
 4555 show program coverage. Each source file's note file is called  
 4556 `@file{@var{auxname}.gcno}`. Refer to the `@option{-fprofile-arcs}` option  
 4557 above for a description of `@var{auxname}` and instructions on how to  
 4558 generate test coverage data. Coverage data will match the source files  
 4559 more closely, if you do not optimize.

4561 @item `-fdbg-cnt-list`  
 4562 @opindex `fdbg-cnt-list`  
 4563 Print the name and the counter upperbound for all debug counters.

4565 @item `-fdbg-cnt=@var{counter-value-list}`  
 4566 @opindex `fdbg-cnt`  
 4567 Set the internal debug counter upperbound. `@var{counter-value-list}`  
 4568 is a comma-separated list of `@var{name}:@var{value}` pairs  
 4569 which sets the upperbound of each debug counter `@var{name}` to `@var{value}`.  
 4570 All debug counters have the initial upperbound of `@var{UINT_MAX}`,  
 4571 thus `dbg_cnt()` returns true always unless the upperbound is set by this option.  
 4572 e.g. With `-fdbg-cnt=dce:10,tail_call:0`  
 4573 `dbg_cnt(dce)` will return true only for first 10 invocations  
 4574 and `dbg_cnt(tail_call)` will return false always.

4576 @item `-d@var{letters}`  
 4577 @itemx `-fdump-rtl-@var{pass}`  
 4578 @opindex `d`

4579 Says to make debugging dumps during compilation at times specified by  
 4580 `@var{letters}`. This is used for debugging the RTL-based passes of the  
 4581 compiler. The file names for most of the dumps are made by appending a  
 4582 pass number and a word to the `@var{dumpname}`. `@var{dumpname}` is generated  
 4583 from the name of the output file, if explicitly specified and it is not  
 4584 an executable, otherwise it is the basename of the source file. These  
 4585 switches may have different effects when `@option{-E}` is used for  
 4586 preprocessing.

4588 Debug dumps can be enabled with a `@option{-fdump-rtl}` switch or some  
 4589 `@option{-d}` option `@var{letters}`. Here are the possible  
 4590 letters for use in `@var{pass}` and `@var{letters}`, and their meanings:

4592 @table @gcctabopt

4594 @item `-fdump-rtl-alignments`  
 4595 @opindex `fdump-rtl-alignments`  
 4596 Dump after branch alignments have been computed.

4598 @item `-fdump-rtl-asmcons`  
 4599 @opindex `fdump-rtl-asmcons`  
 4600 Dump after fixing rtl statements that have unsatisfied in/out constraints.

4602 @item `-fdump-rtl-auto_inc_dec`  
 4603 @opindex `fdump-rtl-auto_inc_dec`  
 4604 Dump after auto-inc-dec discovery. This pass is only run on  
 4605 architectures that have auto inc or auto dec instructions.

4607 @item `-fdump-rtl-barriers`  
 4608 @opindex `fdump-rtl-barriers`  
 4609 Dump after cleaning up the barrier instructions.

4611 @item `-fdump-rtl-bbpart`  
 4612 @opindex `fdump-rtl-bbpart`  
 4613 Dump after partitioning hot and cold basic blocks.

4615 @item `-fdump-rtl-bbro`  
 4616 @opindex `fdump-rtl-bbro`  
 4617 Dump after block reordering.

4619 @item `-fdump-rtl-btl1`  
 4620 @itemx `-fdump-rtl-btl2`  
 4621 @opindex `fdump-rtl-btl2`  
 4622 @opindex `fdump-rtl-btl1`  
 4623 `@option{-fdump-rtl-btl1}` and `@option{-fdump-rtl-btl2}` enable dumping  
 4624 after the two branch  
 4625 target load optimization passes.

4627 @item `-fdump-rtl-bypass`  
 4628 @opindex `fdump-rtl-bypass`  
 4629 Dump after jump bypassing and control flow optimizations.

4631 @item `-fdump-rtl-combine`  
 4632 @opindex `fdump-rtl-combine`  
 4633 Dump after the RTL instruction combination pass.

4635 @item `-fdump-rtl-compgotos`  
 4636 @opindex `fdump-rtl-compgotos`  
 4637 Dump after duplicating the computed gotos.

4639 @item `-fdump-rtl-ce1`  
 4640 @itemx `-fdump-rtl-ce2`  
 4641 @itemx `-fdump-rtl-ce3`  
 4642 @opindex `fdump-rtl-ce1`  
 4643 @opindex `fdump-rtl-ce2`  
 4644 @opindex `fdump-rtl-ce3`

4645 @option{-fdump-rtl-ce1}, @option{-fdump-rtl-ce2}, and  
 4646 @option{-fdump-rtl-ce3} enable dumping after the three  
 4647 if conversion passes.

4649 @itemx -fdump-rtl-cprop\_hardreg  
 4650 @opindex fdump-rtl-cprop\_hardreg  
 4651 Dump after hard register copy propagation.

4653 @itemx -fdump-rtl-csa  
 4654 @opindex fdump-rtl-csa  
 4655 Dump after combining stack adjustments.

4657 @item -fdump-rtl-csel  
 4658 @itemx -fdump-rtl-cse2  
 4659 @opindex fdump-rtl-csel  
 4660 @opindex fdump-rtl-cse2  
 4661 @option{-fdump-rtl-csel} and @option{-fdump-rtl-cse2} enable dumping after  
 4662 the two common sub-expression elimination passes.

4664 @itemx -fdump-rtl-dce  
 4665 @opindex fdump-rtl-dce  
 4666 Dump after the standalone dead code elimination passes.

4668 @itemx -fdump-rtl-dbr  
 4669 @opindex fdump-rtl-dbr  
 4670 Dump after delayed branch scheduling.

4672 @item -fdump-rtl-dcel  
 4673 @itemx -fdump-rtl-dce2  
 4674 @opindex fdump-rtl-dcel  
 4675 @opindex fdump-rtl-dce2  
 4676 @option{-fdump-rtl-dcel} and @option{-fdump-rtl-dce2} enable dumping after  
 4677 the two dead store elimination passes.

4679 @item -fdump-rtl-eh  
 4680 @opindex fdump-rtl-eh  
 4681 Dump after finalization of EH handling code.

4683 @item -fdump-rtl-eh\_ranges  
 4684 @opindex fdump-rtl-eh\_ranges  
 4685 Dump after conversion of EH handling range regions.

4687 @item -fdump-rtl-expand  
 4688 @opindex fdump-rtl-expand  
 4689 Dump after RTL generation.

4691 @item -fdump-rtl-fwprop1  
 4692 @itemx -fdump-rtl-fwprop2  
 4693 @opindex fdump-rtl-fwprop1  
 4694 @opindex fdump-rtl-fwprop2  
 4695 @option{-fdump-rtl-fwprop1} and @option{-fdump-rtl-fwprop2} enable  
 4696 dumping after the two forward propagation passes.

4698 @item -fdump-rtl-gcse1  
 4699 @itemx -fdump-rtl-gcse2  
 4700 @opindex fdump-rtl-gcse1  
 4701 @opindex fdump-rtl-gcse2  
 4702 @option{-fdump-rtl-gcse1} and @option{-fdump-rtl-gcse2} enable dumping  
 4703 after global common subexpression elimination.

4705 @item -fdump-rtl-init-regs  
 4706 @opindex fdump-rtl-init-regs  
 4707 Dump after the initialization of the registers.

4709 @item -fdump-rtl-initvals  
 4710 @opindex fdump-rtl-initvals

4711 Dump after the computation of the initial value sets.

4713 @itemx -fdump-rtl-into\_cfglayout  
 4714 @opindex fdump-rtl-into\_cfglayout  
 4715 Dump after converting to cfglayout mode.

4717 @item -fdump-rtl-ira  
 4718 @opindex fdump-rtl-ira  
 4719 Dump after iterated register allocation.

4721 @item -fdump-rtl-jump  
 4722 @opindex fdump-rtl-jump  
 4723 Dump after the second jump optimization.

4725 @item -fdump-rtl-loop2  
 4726 @opindex fdump-rtl-loop2  
 4727 @option{-fdump-rtl-loop2} enables dumping after the rtl  
 4728 loop optimization passes.

4730 @item -fdump-rtl-mach  
 4731 @opindex fdump-rtl-mach  
 4732 Dump after performing the machine dependent reorganization pass, if that  
 4733 pass exists.

4735 @item -fdump-rtl-mode\_sw  
 4736 @opindex fdump-rtl-mode\_sw  
 4737 Dump after removing redundant mode switches.

4739 @item -fdump-rtl-rnreg  
 4740 @opindex fdump-rtl-rnreg  
 4741 Dump after register renumbering.

4743 @itemx -fdump-rtl-outof\_cfglayout  
 4744 @opindex fdump-rtl-outof\_cfglayout  
 4745 Dump after converting from cfglayout mode.

4747 @item -fdump-rtl-peekhole2  
 4748 @opindex fdump-rtl-peekhole2  
 4749 Dump after the peekhole pass.

4751 @item -fdump-rtl-postreload  
 4752 @opindex fdump-rtl-postreload  
 4753 Dump after post-reload optimizations.

4755 @itemx -fdump-rtl-pro\_and\_epilogue  
 4756 @opindex fdump-rtl-pro\_and\_epilogue  
 4757 Dump after generating the function pro and epilogues.

4759 @item -fdump-rtl-regmove  
 4760 @opindex fdump-rtl-regmove  
 4761 Dump after the register move pass.

4763 @item -fdump-rtl-sched1  
 4764 @itemx -fdump-rtl-sched2  
 4765 @opindex fdump-rtl-sched1  
 4766 @opindex fdump-rtl-sched2  
 4767 @option{-fdump-rtl-sched1} and @option{-fdump-rtl-sched2} enable dumping  
 4768 after the basic block scheduling passes.

4770 @item -fdump-rtl-see  
 4771 @opindex fdump-rtl-see  
 4772 Dump after sign extension elimination.

4774 @item -fdump-rtl-segabstr  
 4775 @opindex fdump-rtl-segabstr  
 4776 Dump after common sequence discovery.

4778 @item -fdump-rtl-shorten  
 4779 @opindex fdump-rtl-shorten  
 4780 Dump after shortening branches.

4782 @item -fdump-rtl-sibling  
 4783 @opindex fdump-rtl-sibling  
 4784 Dump after sibling call optimizations.

4786 @item -fdump-rtl-split1  
 4787 @itemx -fdump-rtl-split2  
 4788 @itemx -fdump-rtl-split3  
 4789 @itemx -fdump-rtl-split4  
 4790 @itemx -fdump-rtl-split5  
 4791 @opindex fdump-rtl-split1  
 4792 @opindex fdump-rtl-split2  
 4793 @opindex fdump-rtl-split3  
 4794 @opindex fdump-rtl-split4  
 4795 @opindex fdump-rtl-split5  
 4796 @option{-fdump-rtl-split1}, @option{-fdump-rtl-split2},  
 4797 @option{-fdump-rtl-split3}, @option{-fdump-rtl-split4} and  
 4798 @option{-fdump-rtl-split5} enable dumping after five rounds of  
 4799 instruction splitting.

4801 @item -fdump-rtl-sms  
 4802 @opindex fdump-rtl-sms  
 4803 Dump after modulo scheduling. This pass is only run on some  
 4804 architectures.

4806 @item -fdump-rtl-stack  
 4807 @opindex fdump-rtl-stack  
 4808 Dump after conversion from GCC's "flat register file" registers to the  
 4809 x87's stack-like registers. This pass is only run on x86 variants.

4811 @item -fdump-rtl-subreg1  
 4812 @itemx -fdump-rtl-subreg2  
 4813 @opindex fdump-rtl-subreg1  
 4814 @opindex fdump-rtl-subreg2  
 4815 @option{-fdump-rtl-subreg1} and @option{-fdump-rtl-subreg2} enable dumping after  
 4816 the two subreg expansion passes.

4818 @item -fdump-rtl-unshare  
 4819 @opindex fdump-rtl-unshare  
 4820 Dump after all rtl has been unshared.

4822 @item -fdump-rtl-vartrack  
 4823 @opindex fdump-rtl-vartrack  
 4824 Dump after variable tracking.

4826 @item -fdump-rtl-vregs  
 4827 @opindex fdump-rtl-vregs  
 4828 Dump after converting virtual registers to hard registers.

4830 @item -fdump-rtl-web  
 4831 @opindex fdump-rtl-web  
 4832 Dump after live range splitting.

4834 @item -fdump-rtl-regclass  
 4835 @itemx -fdump-rtl-subregs\_of\_mode\_init  
 4836 @itemx -fdump-rtl-subregs\_of\_mode\_finish  
 4837 @itemx -fdump-rtl-dfinit  
 4838 @itemx -fdump-rtl-dfinish  
 4839 @opindex fdump-rtl-regclass  
 4840 @opindex fdump-rtl-subregs\_of\_mode\_init  
 4841 @opindex fdump-rtl-subregs\_of\_mode\_finish  
 4842 @opindex fdump-rtl-dfinit

4843 @opindex fdump-rtl-dfinish  
 4844 These dumps are defined but always produce empty files.

4846 @item -fdump-rtl-all  
 4847 @opindex fdump-rtl-all  
 4848 Produce all the dumps listed above.

4850 @item -dA  
 4851 @opindex dA  
 4852 Annotate the assembler output with miscellaneous debugging information.

4854 @item -dD  
 4855 @opindex dD  
 4856 Dump all macro definitions, at the end of preprocessing, in addition to  
 4857 normal output.

4859 @item -dH  
 4860 @opindex dH  
 4861 Produce a core dump whenever an error occurs.

4863 @item -dm  
 4864 @opindex dm  
 4865 Print statistics on memory usage, at the end of the run, to  
 4866 standard error.

4868 @item -dp  
 4869 @opindex dp  
 4870 Annotate the assembler output with a comment indicating which  
 4871 pattern and alternative was used. The length of each instruction is  
 4872 also printed.

4874 @item -dP  
 4875 @opindex dP  
 4876 Dump the RTL in the assembler output as a comment before each instruction.  
 4877 Also turns on @option{-dp} annotation.

4879 @item -dv  
 4880 @opindex dv  
 4881 For each of the other indicated dump files (@option{-fdump-rtl-@var{pass}}),  
 4882 dump a representation of the control flow graph suitable for viewing with VCG  
 4883 to @file{@var{file}.@var{pass}.vcg}.

4885 @item -dx  
 4886 @opindex dx  
 4887 Just generate RTL for a function instead of compiling it. Usually used  
 4888 with @option{-fdump-rtl-expand}.

4890 @item -dy  
 4891 @opindex dy  
 4892 Dump debugging information during parsing, to standard error.  
 4893 @end table

4895 @item -fdump-noaddr  
 4896 @opindex fdump-noaddr  
 4897 When doing debugging dumps, suppress address output. This makes it more  
 4898 feasible to use diff on debugging dumps for compiler invocations with  
 4899 different compiler binaries and/or different  
 4900 text / bss / data / heap / stack / dso start locations.

4902 @item -fdump-unnumbered  
 4903 @opindex fdump-unnumbered  
 4904 When doing debugging dumps, suppress instruction numbers and address output.  
 4905 This makes it more feasible to use diff on debugging dumps for compiler  
 4906 invocations with different options, in particular with and without  
 4907 @option{-g}.

4909 @item -fdump-translation-unit @r{(C++ only)}  
 4910 @itemx -fdump-translation-unit-@var{options} @r{(C++ only)}  
 4911 @opindex fdump-translation-unit  
 4912 Dump a representation of the tree structure for the entire translation  
 4913 unit to a file. The file name is made by appending @file{.tu} to the  
 4914 source file name. If the @samp{-@var{options}} form is used, @var{options}  
 4915 controls the details of the dump as described for the  
 4916 @option{-fdump-tree} options.

4918 @item -fdump-class-hierarchy @r{(C++ only)}  
 4919 @itemx -fdump-class-hierarchy-@var{options} @r{(C++ only)}  
 4920 @opindex fdump-class-hierarchy  
 4921 Dump a representation of each class's hierarchy and virtual function  
 4922 table layout to a file. The file name is made by appending @file{.class}  
 4923 to the source file name. If the @samp{-@var{options}} form is used,  
 4924 @var{options} controls the details of the dump as described for the  
 4925 @option{-fdump-tree} options.

4927 @item -fdump-ipa-@var{switch}  
 4928 @opindex fdump-ipa  
 4929 Control the dumping at various stages of inter-procedural analysis  
 4930 language tree to a file. The file name is generated by appending a switch  
 4931 specific suffix to the source file name. The following dumps are possible:

4933 @table @samp  
 4934 @item all  
 4935 Enables all inter-procedural analysis dumps.

4937 @item cgraph  
 4938 Dumps information about call-graph optimization, unused function removal,  
 4939 and inlining decisions.

4941 @item inline  
 4942 Dump after function inlining.

4944 @end table

4946 @item -fdump-statistics-@var{option}  
 4947 @opindex -fdump-statistics  
 4948 Enable and control dumping of pass statistics in a separate file. The  
 4949 file name is generated by appending a suffix ending in @samp{.statistics}  
 4950 to the source file name. If the @samp{-@var{option}} form is used,  
 4951 @samp{-stats} will cause counters to be summed over the whole compilation unit  
 4952 while @samp{-details} will dump every event as the passes generate them.  
 4953 The default with no option is to sum counters for each function compiled.

4955 @item -fdump-tree-@var{switch}  
 4956 @itemx -fdump-tree-@var{switch}-@var{options}  
 4957 @opindex fdump-tree  
 4958 Control the dumping at various stages of processing the intermediate  
 4959 language tree to a file. The file name is generated by appending a switch  
 4960 specific suffix to the source file name. If the @samp{-@var{options}}  
 4961 form is used, @var{options} is a list of @samp{-} separated options that  
 4962 control the details of the dump. Not all options are applicable to all  
 4963 dumps, those which are not meaningful will be ignored. The following  
 4964 options are available

4966 @table @samp  
 4967 @item address  
 4968 Print the address of each node. Usually this is not meaningful as it  
 4969 changes according to the environment and source file. Its primary use  
 4970 is for tying up a dump file with a debug environment.  
 4971 @item slim  
 4972 Inhibit dumping of members of a scope or body of a function merely  
 4973 because that scope has been reached. Only dump such items when they  
 4974 are directly reachable by some other path. When dumping pretty-printed

4975 trees, this option inhibits dumping the bodies of control structures.  
 4976 @item raw  
 4977 Print a raw representation of the tree. By default, trees are  
 4978 pretty-printed into a C-like representation.  
 4979 @item details  
 4980 Enable more detailed dumps (not honored by every dump option).  
 4981 @item stats  
 4982 Enable dumping various statistics about the pass (not honored by every dump  
 4983 option).  
 4984 @item blocks  
 4985 Enable showing basic block boundaries (disabled in raw dumps).  
 4986 @item vops  
 4987 Enable showing virtual operands for every statement.  
 4988 @item lineno  
 4989 Enable showing line numbers for statements.  
 4990 @item uid  
 4991 Enable showing the unique ID (@code{DECL\_UID}) for each variable.  
 4992 @item verbose  
 4993 Enable showing the tree dump for each statement.  
 4994 @item all  
 4995 Turn on all options, except @option{raw}, @option{slim}, @option{verbose}  
 4996 and @option{lineno}.

4997 @end table

4999 The following tree dumps are possible:  
 5000 @table @samp

5002 @item original  
 5003 Dump before any tree based optimization, to @file{@var{file}.original}.

5005 @item optimized  
 5006 Dump after all tree based optimization, to @file{@var{file}.optimized}.

5008 @item gimple  
 5009 @opindex fdump-tree-gimple  
 5010 Dump each function before and after the gimplification pass to a file. The  
 5011 file name is made by appending @file{.gimple} to the source file name.

5013 @item cfg  
 5014 @opindex fdump-tree-cfg  
 5015 Dump the control flow graph of each function to a file. The file name is  
 5016 made by appending @file{.cfg} to the source file name.

5018 @item vcg  
 5019 @opindex fdump-tree-vcg  
 5020 Dump the control flow graph of each function to a file in VCG format. The  
 5021 file name is made by appending @file{.vcg} to the source file name. Note  
 5022 that if the file contains more than one function, the generated file cannot  
 5023 be used directly by VCG@. You will need to cut and paste each function's  
 5024 graph into its own separate file first.

5026 @item ch  
 5027 @opindex fdump-tree-ch  
 5028 Dump each function after copying loop headers. The file name is made by  
 5029 appending @file{.ch} to the source file name.

5031 @item ssa  
 5032 @opindex fdump-tree-ssa  
 5033 Dump SSA related information to a file. The file name is made by appending  
 5034 @file{.ssa} to the source file name.

5036 @item alias  
 5037 @opindex fdump-tree-alias  
 5038 Dump aliasing information for each function. The file name is made by  
 5039 appending @file{.alias} to the source file name.

5041 @item ccp  
5042 @opindex fdump-tree-ccp  
5043 Dump each function after CCP@. The file name is made by appending  
5044 @file{.ccp} to the source file name.

5046 @item storeccp  
5047 @opindex fdump-tree-storeccp  
5048 Dump each function after STORE-CCP@. The file name is made by appending  
5049 @file{.storeccp} to the source file name.

5051 @item pre  
5052 @opindex fdump-tree-pre  
5053 Dump trees after partial redundancy elimination. The file name is made  
5054 by appending @file{.pre} to the source file name.

5056 @item fre  
5057 @opindex fdump-tree-fre  
5058 Dump trees after full redundancy elimination. The file name is made  
5059 by appending @file{.fre} to the source file name.

5061 @item copyprop  
5062 @opindex fdump-tree-copyprop  
5063 Dump trees after copy propagation. The file name is made  
5064 by appending @file{.copyprop} to the source file name.

5066 @item store\_copyprop  
5067 @opindex fdump-tree-store\_copyprop  
5068 Dump trees after store copy-propagation. The file name is made  
5069 by appending @file{.store\_copyprop} to the source file name.

5071 @item dce  
5072 @opindex fdump-tree-dce  
5073 Dump each function after dead code elimination. The file name is made by  
5074 appending @file{.dce} to the source file name.

5076 @item mudflap  
5077 @opindex fdump-tree-mudflap  
5078 Dump each function after adding mudflap instrumentation. The file name is  
5079 made by appending @file{.mudflap} to the source file name.

5081 @item sra  
5082 @opindex fdump-tree-sra  
5083 Dump each function after performing scalar replacement of aggregates. The  
5084 file name is made by appending @file{.sra} to the source file name.

5086 @item sink  
5087 @opindex fdump-tree-sink  
5088 Dump each function after performing code sinking. The file name is made  
5089 by appending @file{.sink} to the source file name.

5091 @item dom  
5092 @opindex fdump-tree-dom  
5093 Dump each function after applying dominator tree optimizations. The file  
5094 name is made by appending @file{.dom} to the source file name.

5096 @item dse  
5097 @opindex fdump-tree-dse  
5098 Dump each function after applying dead store elimination. The file  
5099 name is made by appending @file{.dse} to the source file name.

5101 @item phiopt  
5102 @opindex fdump-tree-phiopt  
5103 Dump each function after optimizing PHI nodes into straightline code. The file  
5104 name is made by appending @file{.phiopt} to the source file name.

5106 @item forwprop

5107 @opindex fdump-tree-forwprop  
5108 Dump each function after forward propagating single use variables. The file  
5109 name is made by appending @file{.forwprop} to the source file name.

5111 @item copyrename  
5112 @opindex fdump-tree-copyrename  
5113 Dump each function after applying the copy rename optimization. The file  
5114 name is made by appending @file{.copyrename} to the source file name.

5116 @item nrv  
5117 @opindex fdump-tree-nrv  
5118 Dump each function after applying the named return value optimization on  
5119 generic trees. The file name is made by appending @file{.nrv} to the source  
5120 file name.

5122 @item vect  
5123 @opindex fdump-tree-vect  
5124 Dump each function after applying vectorization of loops. The file name is  
5125 made by appending @file{.vect} to the source file name.

5127 @item vrp  
5128 @opindex fdump-tree-vrp  
5129 Dump each function after Value Range Propagation (VRP). The file name  
5130 is made by appending @file{.vrp} to the source file name.

5132 @item all  
5133 @opindex fdump-tree-all  
5134 Enable all the available tree dumps with the flags provided in this option.  
5135 @end table

5137 @item -ftree-vectorizer-verbose=@var{n}  
5138 @opindex ftree-vectorizer-verbose  
5139 This option controls the amount of debugging output the vectorizer prints.  
5140 This information is written to standard error, unless  
5141 @option{-fdump-tree-all} or @option{-fdump-tree-vect} is specified,  
5142 in which case it is output to the usual dump listing file, @file{.vect}.  
5143 For @var{n}=0 no diagnostic information is reported.  
5144 If @var{n}=1 the vectorizer reports each loop that got vectorized,  
5145 and the total number of loops that got vectorized.  
5146 If @var{n}=2 the vectorizer also reports non-vectorized loops that passed  
5147 the first analysis phase (vect\_analyze\_loop\_form) - i.e.: countable,  
5148 inner-most, single-bb, single-entry/exit loops. This is the same verbosity  
5149 level that @option{-fdump-tree-vect-stats} uses.  
5150 Higher verbosity levels mean either more information dumped for each  
5151 reported loop, or same amount of information reported for more loops:  
5152 If @var{n}=3, alignment related information is added to the reports.  
5153 If @var{n}=4, data-references related information (e.g.@: memory dependences,  
5154 memory access-patterns) is added to the reports.  
5155 If @var{n}=5, the vectorizer reports also non-vectorized inner-most loops  
5156 that did not pass the first analysis phase (i.e., may not be countable, or  
5157 may have complicated control-flow).  
5158 If @var{n}=6, the vectorizer reports also non-vectorized nested loops.  
5159 For @var{n}=7, all the information the vectorizer generates during its  
5160 analysis and transformation is reported. This is the same verbosity level  
5161 that @option{-fdump-tree-vect-details} uses.

5163 @item -frandom-seed=@var{string}  
5164 @opindex frandom-string  
5165 This option provides a seed that GCC uses when it would otherwise use  
5166 random numbers. It is used to generate certain symbol names  
5167 that have to be different in every compiled file. It is also used to  
5168 place unique stamps in coverage data files and the object files that  
5169 produce them. You can use the @option{-frandom-seed} option to produce  
5170 reproducibly identical object files.

5172 The @var{string} should be different for every file you compile.



5174 @item -fsched-verbose=@var{n}  
 5175 @opindex fsched-verbose  
 5176 On targets that use instruction scheduling, this option controls the  
 5177 amount of debugging output the scheduler prints. This information is  
 5178 written to standard error, unless @option{-fdump-rtl-sched1} or  
 5179 @option{-fdump-rtl-sched2} is specified, in which case it is output  
 5180 to the usual dump listing file, @file{.sched} or @file{.sched2}  
 5181 respectively. However for @var{n} greater than nine, the output is  
 5182 always printed to standard error.

5184 For @var{n} greater than zero, @option{-fsched-verbose} outputs the  
 5185 same information as @option{-fdump-rtl-sched1} and @option{-fdump-rtl-sched2}.  
 5186 For @var{n} greater than one, it also output basic block probabilities,  
 5187 detailed ready list information and unit/insn info. For @var{n} greater  
 5188 than two, it includes RTL at abort point, control-flow and regions info.  
 5189 And for @var{n} over four, @option{-fsched-verbose} also includes  
 5190 dependence info.

5192 @item -save-temps  
 5193 @opindex save-temps  
 5194 Store the usual ‘‘temporary’’ intermediate files permanently; place them  
 5195 in the current directory and name them based on the source file. Thus,  
 5196 compiling @file{foo.c} with @samp{-c -save-temps} would produce files  
 5197 @file{foo.i} and @file{foo.s}, as well as @file{foo.o}. This creates a  
 5198 preprocessed @file{foo.i} output file even though the compiler now  
 5199 normally uses an integrated preprocessor.

5201 When used in combination with the @option{-x} command line option,  
 5202 @option{-save-temps} is sensible enough to avoid over writing an  
 5203 input source file with the same extension as an intermediate file.  
 5204 The corresponding intermediate file may be obtained by renaming the  
 5205 source file before using @option{-save-temps}.

5207 @item -time  
 5208 @opindex time  
 5209 Report the CPU time taken by each subprocess in the compilation  
 5210 sequence. For C source files, this is the compiler proper and assembler  
 5211 (plus the linker if linking is done). The output looks like this:

```
5213 @smallexample
5214 # ccl 0.12 0.01
5215 # as 0.00 0.01
5216 @end smallexample
```

5218 The first number on each line is the ‘‘user time’’, that is time spent  
 5219 executing the program itself. The second number is ‘‘system time’’,  
 5220 time spent executing operating system routines on behalf of the program.  
 5221 Both numbers are in seconds.

5223 @item -fvar-tracking  
 5224 @opindex fvar-tracking  
 5225 Run variable tracking pass. It computes where variables are stored at each  
 5226 position in code. Better debugging information is then generated  
 5227 (if the debugging information format supports this information).

5229 It is enabled by default when compiling with optimization (@option{-Os},  
 5230 @option{-O}, @option{-O2}, @dots{}), debugging information (@option{-g}) and  
 5231 the debug info format supports it.

5233 @item -print-file-name=@var{library}  
 5234 @opindex print-file-name  
 5235 Print the full absolute name of the library file @var{library} that  
 5236 would be used when linking---and don't do anything else. With this  
 5237 option, GCC does not compile or link anything; it just prints the  
 5238 file name.

5240 @item -print-multi-directory  
 5241 @opindex print-multi-directory  
 5242 Print the directory name corresponding to the multilib selected by any  
 5243 other switches present in the command line. This directory is supposed  
 5244 to exist in @env{GCC\_EXEC\_PREFIX}.

5246 @item -print-multi-lib  
 5247 @opindex print-multi-lib  
 5248 Print the mapping from multilib directory names to compiler switches  
 5249 that enable them. The directory name is separated from the switches by  
 5250 @samp{;}, and each switch starts with an @samp{@@} instead of the  
 5251 @samp{-}, without spaces between multiple switches. This is supposed to  
 5252 ease shell-processing.

5254 @item -print-prog-name=@var{program}  
 5255 @opindex print-prog-name  
 5256 Like @option{-print-file-name}, but searches for a program such as @samp{cpp}.

5258 @item -print-libgcc-file-name  
 5259 @opindex print-libgcc-file-name  
 5260 Same as @option{-print-file-name=libgcc.a}.

5262 This is useful when you use @option{-nostdlib} or @option{-nodefaultlibs}  
 5263 but you do want to link with @file{libgcc.a}. You can do

```
5265 @smallexample
5266 gcc -nostdlib @var{files}@dots{} `gcc -print-libgcc-file-name`
5267 @end smallexample
```

5269 @item -print-search-dirs  
 5270 @opindex print-search-dirs  
 5271 Print the name of the configured installation directory and a list of  
 5272 program and library directories @command{gcc} will search---and don't do anything

5274 This is useful when @command{gcc} prints the error message  
 5275 @samp{installation problem, cannot exec cpp0: No such file or directory}.  
 5276 To resolve this you either need to put @file{cpp0} and the other compiler  
 5277 components where @command{gcc} expects to find them, or you can set the environm  
 5278 variable @env{GCC\_EXEC\_PREFIX} to the directory where you installed them.  
 5279 Don't forget the trailing @samp{//}.  
 5280 @xref{Environment Variables}.

5282 @item -print-sysroot  
 5283 @opindex print-sysroot  
 5284 Print the target sysroot directory that will be used during  
 5285 compilation. This is the target sysroot specified either at configure  
 5286 time or using the @option{--sysroot} option, possibly with an extra  
 5287 suffix that depends on compilation options. If no target sysroot is  
 5288 specified, the option prints nothing.

5290 @item -print-sysroot-headers-suffix  
 5291 @opindex print-sysroot-headers-suffix  
 5292 Print the suffix added to the target sysroot when searching for  
 5293 headers, or give an error if the compiler is not configured with such  
 5294 a suffix---and don't do anything else.

5296 @item -dumpmachine  
 5297 @opindex dumpmachine  
 5298 Print the compiler's target machine (for example,  
 5299 @samp{i686-pc-linux-gnu})---and don't do anything else.

5301 @item -dumpversion  
 5302 @opindex dumpversion  
 5303 Print the compiler version (for example, @samp{3.0})---and don't do  
 5304 anything else.

5306 @item -dumpspecs  
 5307 @opindex dumpspecs  
 5308 Print the compiler's built-in specs---and don't do anything else. (This  
 5309 is used when GCC itself is being built.) @xref{Spec Files}.

5311 @item -feliminate-unused-debug-types  
 5312 @opindex feliminate-unused-debug-types  
 5313 Normally, when producing DWARF2 output, GCC will emit debugging  
 5314 information for all types declared in a compilation  
 5315 unit, regardless of whether or not they are actually used  
 5316 in that compilation unit. Sometimes this is useful, such as  
 5317 if, in the debugger, you want to cast a value to a type that is  
 5318 not actually used in your program (but is declared). More often,  
 5319 however, this results in a significant amount of wasted space.  
 5320 With this option, GCC will avoid producing debug symbol output  
 5321 for types that are nowhere used in the source file being compiled.  
 5322 @end table

5324 @node Optimize Options  
 5325 @section Options That Control Optimization  
 5326 @cindex optimize options  
 5327 @cindex options, optimization

5329 These options control various sorts of optimizations.

5331 Without any optimization option, the compiler's goal is to reduce the  
 5332 cost of compilation and to make debugging produce the expected  
 5333 results. Statements are independent: if you stop the program with a  
 5334 breakpoint between statements, you can then assign a new value to any  
 5335 variable or change the program counter to any other statement in the  
 5336 function and get exactly the results you would expect from the source  
 5337 code.

5339 Turning on optimization flags makes the compiler attempt to improve  
 5340 the performance and/or code size at the expense of compilation time  
 5341 and possibly the ability to debug the program.

5343 The compiler performs optimization based on the knowledge it has of the  
 5344 program. Compiling multiple files at once to a single output file mode allows  
 5345 the compiler to use information gained from all of the files when compiling  
 5346 each of them.

5348 Not all optimizations are controlled directly by a flag. Only  
 5349 optimizations that have a flag are listed.

5351 @table @gcctabopt  
 5352 @item -O  
 5353 @itemx -O1  
 5354 @opindex O  
 5355 @opindex O1  
 5356 Optimize. Optimizing compilation takes somewhat more time, and a lot  
 5357 more memory for a large function.

5359 With @option{-O}, the compiler tries to reduce code size and execution  
 5360 time, without performing any optimizations that take a great deal of  
 5361 compilation time.

5363 @option{-O} turns on the following optimization flags:  
 5364 @gccoptlist{  
 5365 -fauto-inc-dec @gol  
 5366 -fprop-registers @gol  
 5367 -fdce @gol  
 5368 -fdefer-pop @gol  
 5369 -fdelayed-branch @gol  
 5370 -fdse @gol

5371 -fguess-branch-probability @gol  
 5372 -fif-conversion2 @gol  
 5373 -fif-conversion @gol  
 5374 -finline-small-functions @gol  
 5375 -fipa-pure-const @gol  
 5376 -fipa-reference @gol  
 5377 -fmerge-constants  
 5378 -fsplit-wide-types @gol  
 5379 -ftree-builtin-call-dce @gol  
 5380 -ftree-ccp @gol  
 5381 -ftree-ch @gol  
 5382 -ftree-copyrename @gol  
 5383 -ftree-dce @gol  
 5384 -ftree-dominator-opts @gol  
 5385 -ftree-dse @gol  
 5386 -ftree-fre @gol  
 5387 -ftree-sra @gol  
 5388 -ftree-ter @gol  
 5389 -funit-at-a-time}

5391 @option{-O} also turns on @option{-fomit-frame-pointer} on machines  
 5392 where doing so does not interfere with debugging.

5394 @item -O2  
 5395 @opindex O2  
 5396 Optimize even more. GCC performs nearly all supported optimizations  
 5397 that do not involve a space-speed tradeoff.  
 5398 As compared to @option{-O}, this option increases both compilation time  
 5399 and the performance of the generated code.

5401 @option{-O2} turns on all optimization flags specified by @option{-O}. It  
 5402 also turns on the following optimization flags:  
 5403 @gccoptlist{-fthread-jumps @gol  
 5404 -falign-functions -falign-jumps @gol  
 5405 -falign-loops -falign-labels @gol  
 5406 -fcaller-saves @gol  
 5407 -fcrossjumping @gol  
 5408 -fcse-follow-jumps -fcse-skip-blocks @gol  
 5409 -fdelete-null-pointer-checks @gol  
 5410 -fexpensive-optimizations @gol  
 5411 -fgcse -fgcse-lm @gol  
 5412 -findirect-inlining @gol  
 5413 -foptimize-sibling-calls @gol  
 5414 -fpeehole2 @gol  
 5415 -fregmove @gol  
 5416 -freorder-blocks -freorder-functions @gol  
 5417 -frerun-cse-after-loop @gol  
 5418 -fsched-interblock -fsched-spec @gol  
 5419 -fschedule-insns -fschedule-insns2 @gol  
 5420 -fstrict-aliasing -fstrict-overflow @gol  
 5421 -ftree-switch-conversion @gol  
 5422 -ftree-pre @gol  
 5423 -ftree-vrp}

5425 Please note the warning under @option{-fgcse} about  
 5426 invoking @option{-O2} on programs that use computed gotos.

5428 @item -O3  
 5429 @opindex O3  
 5430 Optimize yet more. @option{-O3} turns on all optimizations specified  
 5431 by @option{-O2} and also turns on the @option{-finline-functions},  
 5432 @option{-funswitch-loops}, @option{-fpredictive-commoning},  
 5433 @option{-fgcse-after-reload} and @option{-ftree-vectorize} options.

5435 @item -O0  
 5436 @opindex O0

5437 Reduce compilation time and make debugging produce the expected  
5438 results. This is the default.

5440 @item -Os  
5441 @opindex Os  
5442 Optimize for size. @option{-Os} enables all @option{-O2} optimizations that  
5443 do not typically increase code size. It also performs further  
5444 optimizations designed to reduce code size.

5446 @option{-Os} disables the following optimization flags:  
5447 @gcoptlist{-falign-functions -falign-jumps -falign-loops @gol  
5448 -falign-labels -freorder-blocks -freorder-blocks-and-partition @gol  
5449 -fprefetch-loop-arrays -ftree-vec-loop-version}

5451 If you use multiple @option{-O} options, with or without level numbers,  
5452 the last such option is the one that is effective.  
5453 @end table

5455 Options of the form @option{-f@var{flag}} specify machine-independent  
5456 flags. Most flags have both positive and negative forms; the negative  
5457 form of @option{-ffoo} would be @option{-fno-foo}. In the table  
5458 below, only one of the forms is listed--the one you typically will  
5459 use. You can figure out the other form by either removing @samp{no-}  
5460 or adding it.

5462 The following options control specific optimizations. They are either  
5463 activated by @option{-O} options or are related to ones that are. You  
5464 can use the following flags in the rare cases when ‘‘fine-tuning’’ of  
5465 optimizations to be performed is desired.

5467 @table @gcoptabopt  
5468 @item -fno-default-inline  
5469 @opindex fno-default-inline  
5470 Do not make member functions inline by default merely because they are  
5471 defined inside the class scope (C++ only). Otherwise, when you specify  
5472 @w{@option{-O}}, member functions defined inside class scope are compiled  
5473 inline by default; i.e., you don't need to add @samp{inline} in front of  
5474 the member function name.

5476 @item -fno-defer-pop  
5477 @opindex fno-defer-pop  
5478 Always pop the arguments to each function call as soon as that function  
5479 returns. For machines which must pop arguments after a function call,  
5480 the compiler normally lets arguments accumulate on the stack for several  
5481 function calls and pops them all at once.

5483 Disabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5485 @item -fforward-propagate  
5486 @opindex fforward-propagate  
5487 Perform a forward propagation pass on RTL@. The pass tries to combine two  
5488 instructions and checks if the result can be simplified. If loop unrolling  
5489 is active, two passes are performed and the second is scheduled after  
5490 loop unrolling.

5492 This option is enabled by default at optimization levels @option{-O2},  
5493 @option{-O3}, @option{-Os}.

5495 @item -fomit-frame-pointer  
5496 @opindex fomit-frame-pointer  
5497 Don't keep the frame pointer in a register for functions that  
5498 don't need one. This avoids the instructions to save, set up and  
5499 restore frame pointers; it also makes an extra register available  
5500 in many functions. @strong{It also makes debugging impossible on  
5501 some machines.}

5503 On some machines, such as the VAX, this flag has no effect, because  
5504 the standard calling sequence automatically handles the frame pointer  
5505 and nothing is saved by pretending it doesn't exist. The  
5506 machine-description macro @code{FRAME\_POINTER\_REQUIRED} controls  
5507 whether a target machine supports this flag. @xref{Registers,,Register  
5508 Usage, gccint, GNU Compiler Collection (GCC) Internals}.

5510 Enabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5512 @item -foptimize-sibling-calls  
5513 @opindex fo optimize-sibling-calls  
5514 Optimize sibling and tail recursive calls.

5516 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5518 @item -fno-inline  
5519 @opindex fno-inline  
5520 Don't pay attention to the @code{inline} keyword. Normally this option  
5521 is used to keep the compiler from expanding any functions inline.  
5522 Note that if you are not optimizing, no functions can be expanded inline.

5524 @item -finline-small-functions  
5525 @opindex finline-small-functions  
5526 Integrate functions into their callers when their body is smaller than expected  
5527 function call code (so overall size of program gets smaller). The compiler  
5528 heuristically decides which functions are simple enough to be worth integrating  
5529 in this way.

5531 Enabled at level @option{-O2}.

5533 @item -findirect-inlining  
5534 @opindex findirect-inlining  
5535 Inline also indirect calls that are discovered to be known at compile  
5536 time thanks to previous inlining. This option has any effect only  
5537 when inlining itself is turned on by the @option{-finline-functions}  
5538 or @option{-finline-small-functions} options.

5540 Enabled at level @option{-O2}.

5542 @item -finline-functions  
5543 @opindex finline-functions  
5544 Integrate all simple functions into their callers. The compiler  
5545 heuristically decides which functions are simple enough to be worth  
5546 integrating in this way.

5548 If all calls to a given function are integrated, and the function is  
5549 declared @code{static}, then the function is normally not output as  
5550 assembler code in its own right.

5552 Enabled at level @option{-O3}.

5554 @item -finline-functions-called-once  
5555 @opindex finline-functions-called-once  
5556 Consider all @code{static} functions called once for inlining into their  
5557 caller even if they are not marked @code{inline}. If a call to a given  
5558 function is integrated, then the function is not output as assembler code  
5559 in its own right.

5561 Enabled at levels @option{-O1}, @option{-O2}, @option{-O3} and @option{-Os}.

5563 @item -fearly-inlining  
5564 @opindex fearly-inlining  
5565 Inline functions marked by @code{always\_inline} and functions whose body seems  
5566 smaller than the function call overhead early before doing  
5567 @option{-fprofile-generate} instrumentation and real inlining pass. Doing so  
5568 makes profiling significantly cheaper and usually inlining faster on programs

5569 having large chains of nested wrapper functions.

5571 Enabled by default.

5573 @item -finline-limit=@var{n}  
 5574 @opindex finline-limit  
 5575 By default, GCC limits the size of functions that can be inlined. This flag  
 5576 allows coarse control of this limit. @var{n} is the size of functions that  
 5577 can be inlined in number of pseudo instructions.

5579 Inlining is actually controlled by a number of parameters, which may be  
 5580 specified individually by using @option{--param @var{name}=@var{value}}.  
 5581 The @option{-finline-limit=@var{n}} option sets some of these parameters  
 5582 as follows:

5584 @table @gcctabopt  
 5585 @item max-inline-insns-single  
 5586 is set to @var{n}/2.  
 5587 @item max-inline-insns-auto  
 5588 is set to @var{n}/2.  
 5589 @end table

5591 See below for a documentation of the individual  
 5592 parameters controlling inlining and for the defaults of these parameters.

5594 @emph{Note:} there may be no value to @option{-finline-limit} that results  
 5595 in default behavior.

5597 @emph{Note:} pseudo instruction represents, in this particular context, an  
 5598 abstract measurement of function's size. In no way does it represent a count  
 5599 of assembly instructions and as such its exact meaning might change from one  
 5600 release to another.

5602 @item -fkeep-inline-functions  
 5603 @opindex fkeep-inline-functions  
 5604 In C, emit @code{static} functions that are declared @code{inline}  
 5605 into the object file, even if the function has been inlined into all  
 5606 of its callers. This switch does not affect functions using the  
 5607 @code{extern inline} extension in GNU C89@. In C++, emit any and all  
 5608 inline functions into the object file.

5610 @item -fkeep-static-consts  
 5611 @opindex fkeep-static-consts  
 5612 Emit variables declared @code{static const} when optimization isn't turned  
 5613 on, even if the variables aren't referenced.

5615 GCC enables this option by default. If you want to force the compiler to  
 5616 check if the variable was referenced, regardless of whether or not  
 5617 optimization is turned on, use the @option{-fno-keep-static-consts} option.

5619 @item -fmerge-constants  
 5620 @opindex fmerge-constants  
 5621 Attempt to merge identical constants (string constants and floating point  
 5622 constants) across compilation units.

5624 This option is the default for optimized compilation if the assembler and  
 5625 linker support it. Use @option{-fno-merge-constants} to inhibit this  
 5626 behavior.

5628 Enabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5630 @item -fmerge-all-constants  
 5631 @opindex fmerge-all-constants  
 5632 Attempt to merge identical constants and identical variables.

5634 This option implies @option{-fmerge-constants}. In addition to

5635 @option{-fmerge-constants} this considers e.g.@: even constant initialized  
 5636 arrays or initialized constant variables with integral or floating point  
 5637 types. Languages like C or C++ require each variable, including multiple  
 5638 instances of the same variable in recursive calls, to have distinct locations,  
 5639 so using this option will result in non-conforming  
 5640 behavior.

5642 @item -fmodulo-sched  
 5643 @opindex fmodulo-sched  
 5644 Perform swing modulo scheduling immediately before the first scheduling  
 5645 pass. This pass looks at innermost loops and reorders their  
 5646 instructions by overlapping different iterations.

5648 @item -fmodulo-sched-allow-regmoves  
 5649 @opindex fmodulo-sched-allow-regmoves  
 5650 Perform more aggressive SMS based modulo scheduling with register moves  
 5651 allowed. By setting this flag certain anti-dependences edges will be  
 5652 deleted which will trigger the generation of reg-moves based on the  
 5653 life-range analysis. This option is effective only with  
 5654 @option{-fmodulo-sched} enabled.

5656 @item -fno-branch-count-reg  
 5657 @opindex fno-branch-count-reg  
 5658 Do not use "decrement and branch" instructions on a count register,  
 5659 but instead generate a sequence of instructions that decrement a  
 5660 register, compare it against zero, then branch based upon the result.  
 5661 This option is only meaningful on architectures that support such  
 5662 instructions, which include x86, PowerPC, IA-64 and S/390.

5664 The default is @option{-fbranch-count-reg}.

5666 @item -fno-function-cse  
 5667 @opindex fno-function-cse  
 5668 Do not put function addresses in registers; make each instruction that  
 5669 calls a constant function contain the function's address explicitly.

5671 This option results in less efficient code, but some strange hacks  
 5672 that alter the assembler output may be confused by the optimizations  
 5673 performed when this option is not used.

5675 The default is @option{-ffunction-cse}

5677 @item -fno-zero-initialized-in-bss  
 5678 @opindex fno-zero-initialized-in-bss  
 5679 If the target supports a BSS section, GCC by default puts variables that  
 5680 are initialized to zero into BSS@. This can save space in the resulting  
 5681 code.

5683 This option turns off this behavior because some programs explicitly  
 5684 rely on variables going to the data section. E.g., so that the  
 5685 resulting executable can find the beginning of that section and/or make  
 5686 assumptions based on that.

5688 The default is @option{-fzero-initialized-in-bss}.

5690 @item -fmudflap -fmudflapth -fmudflapir  
 5691 @opindex fmudflap  
 5692 @opindex fmudflapth  
 5693 @opindex fmudflapir  
 5694 @cindex bounds checking  
 5695 @cindex mudflap  
 5696 For front-ends that support it (C and C++), instrument all risky  
 5697 pointer/array dereferencing operations, some standard library  
 5698 string/heap functions, and some other associated constructs with  
 5699 range/validity tests. Modules so instrumented should be immune to  
 5700 buffer overflows, invalid heap use, and some other classes of C/C++

5701 programming errors. The instrumentation relies on a separate runtime  
5702 library (@file{libmudflap}), which will be linked into a program if  
5703 @option{-fmudflap} is given at link time. Run-time behavior of the  
5704 instrumented program is controlled by the @env{MUDFLAP\_OPTIONS}  
5705 environment variable. See @code{env MUDFLAP\_OPTIONS=-help a.out}  
5706 for its options.

5708 Use @option{-fmudflapth} instead of @option{-fmudflap} to compile and to  
5709 link if your program is multi-threaded. Use @option{-fmudflapir}, in  
5710 addition to @option{-fmudflap} or @option{-fmudflapth}, if  
5711 instrumentation should ignore pointer reads. This produces less  
5712 instrumentation (and therefore faster execution) and still provides  
5713 some protection against outright memory corrupting writes, but allows  
5714 erroneously read data to propagate within a program.

5716 @item -fthread-jumps  
5717 @opindex fthread-jumps  
5718 Perform optimizations where we check to see if a jump branches to a  
5719 location where another comparison subsumed by the first is found. If  
5720 so, the first branch is redirected to either the destination of the  
5721 second branch or a point immediately following it, depending on whether  
5722 the condition is known to be true or false.

5724 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5726 @item -fsplit-wide-types  
5727 @opindex fsplit-wide-types  
5728 When using a type that occupies multiple registers, such as @code{long}  
5729 @code{long} on a 32-bit system, split the registers apart and allocate them  
5730 independently. This normally generates better code for those types,  
5731 but may make debugging more difficult.

5733 Enabled at levels @option{-O}, @option{-O2}, @option{-O3},  
5734 @option{-Os}.

5736 @item -fcse-follow-jumps  
5737 @opindex fcse-follow-jumps  
5738 In common subexpression elimination (CSE), scan through jump instructions  
5739 when the target of the jump is not reached by any other path. For  
5740 example, when CSE encounters an @code{if} statement with an  
5741 @code{else} clause, CSE will follow the jump when the condition  
5742 tested is false.

5744 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5746 @item -fcse-skip-blocks  
5747 @opindex fcse-skip-blocks  
5748 This is similar to @option{-fcse-follow-jumps}, but causes CSE to  
5749 follow jumps which conditionally skip over blocks. When CSE  
5750 encounters a simple @code{if} statement with no else clause,  
5751 @option{-fcse-skip-blocks} causes CSE to follow the jump around the  
5752 body of the @code{if}.

5754 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5756 @item -frerun-cse-after-loop  
5757 @opindex frerun-cse-after-loop  
5758 Re-run common subexpression elimination after loop optimizations has been  
5759 performed.

5761 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5763 @item -fgcse  
5764 @opindex fgcse  
5765 Perform a global common subexpression elimination pass.  
5766 This pass also performs global constant and copy propagation.

5768 @emph{Note:} When compiling a program using computed gotos, a GCC  
5769 extension, you may get better runtime performance if you disable  
5770 the global common subexpression elimination pass by adding  
5771 @option{-fno-gcse} to the command line.

5773 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5775 @item -fgcse-lm  
5776 @opindex fgcse-lm  
5777 When @option{-fgcse-lm} is enabled, global common subexpression elimination will  
5778 attempt to move loads which are only killed by stores into themselves. This  
5779 allows a loop containing a load/store sequence to be changed to a load outside  
5780 the loop, and a copy/store within the loop.

5782 Enabled by default when gcse is enabled.

5784 @item -fgcse-sm  
5785 @opindex fgcse-sm  
5786 When @option{-fgcse-sm} is enabled, a store motion pass is run after  
5787 global common subexpression elimination. This pass will attempt to move  
5788 stores out of loops. When used in conjunction with @option{-fgcse-lm},  
5789 loops containing a load/store sequence can be changed to a load before  
5790 the loop and a store after the loop.

5792 Not enabled at any optimization level.

5794 @item -fgcse-las  
5795 @opindex fgcse-las  
5796 When @option{-fgcse-las} is enabled, the global common subexpression  
5797 elimination pass eliminates redundant loads that come after stores to the  
5798 same memory location (both partial and full redundancies).

5800 Not enabled at any optimization level.

5802 @item -fgcse-after-reload  
5803 @opindex fgcse-after-reload  
5804 When @option{-fgcse-after-reload} is enabled, a redundant load elimination  
5805 pass is performed after reload. The purpose of this pass is to cleanup  
5806 redundant spilling.

5808 @item -funsafe-loop-optimizations  
5809 @opindex unsafe-loop-optimizations  
5810 If given, the loop optimizer will assume that loop indices do not  
5811 overflow, and that the loops with nontrivial exit condition are not  
5812 infinite. This enables a wider range of loop optimizations even if  
5813 the loop optimizer itself cannot prove that these assumptions are valid.  
5814 Using @option{-Wunsafe-loop-optimizations}, the compiler will warn you  
5815 if it finds this kind of loop.

5817 @item -fcrossjumping  
5818 @opindex fcrossjumping  
5819 Perform cross-jumping transformation. This transformation unifies equivalent co  
5820 resulting code may or may not perform better than without cross-jumping.

5822 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5824 @item -fauto-inc-dec  
5825 @opindex auto-inc-dec  
5826 Combine increments or decrements of addresses with memory accesses.  
5827 This pass is always skipped on architectures that do not have  
5828 instructions to support this. Enabled by default at @option{-O} and  
5829 higher on architectures that support this.

5831 @item -fdce  
5832 @opindex fdce

5833 Perform dead code elimination (DCE) on RTL@.  
 5834 Enabled by default at @option{-O} and higher.

5836 @item -fdse  
 5837 @opindex fdse  
 5838 Perform dead store elimination (DSE) on RTL@.  
 5839 Enabled by default at @option{-O} and higher.

5841 @item -fif-conversion  
 5842 @opindex fif-conversion  
 5843 Attempt to transform conditional jumps into branch-less equivalents. This  
 5844 include use of conditional moves, min, max, set flags and abs instructions, and  
 5845 some tricks doable by standard arithmetics. The use of conditional execution  
 5846 on chips where it is available is controlled by @code{if-conversion2}.

5848 Enabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5850 @item -fif-conversion2  
 5851 @opindex fif-conversion2  
 5852 Use conditional execution (where available) to transform conditional jumps into  
 5853 branch-less equivalents.

5855 Enabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5857 @item -fdelete-null-pointer-checks  
 5858 @opindex fdelete-null-pointer-checks  
 5859 Use global dataflow analysis to identify and eliminate useless checks  
 5860 for null pointers. The compiler assumes that dereferencing a null  
 5861 pointer would have halted the program. If a pointer is checked after  
 5862 it has already been dereferenced, it cannot be null.

5864 In some environments, this assumption is not true, and programs can  
 5865 safely dereference null pointers. Use  
 5866 @option{-fno-delete-null-pointer-checks} to disable this optimization  
 5867 for programs which depend on that behavior.

5869 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5871 @item -fexpensive-optimizations  
 5872 @opindex fexpensive-optimizations  
 5873 Perform a number of minor optimizations that are relatively expensive.

5875 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5877 @item -foptimize-register-move  
 5878 @itemx -fregmove  
 5879 @opindex foptimize-register-move  
 5880 @opindex fregmove  
 5881 Attempt to reassign register numbers in move instructions and as  
 5882 operands of other simple instructions in order to maximize the amount of  
 5883 register tying. This is especially helpful on machines with two-operand  
 5884 instructions.

5886 Note @option{-fregmove} and @option{-foptimize-register-move} are the same  
 5887 optimization.

5889 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5891 @item -fira-algorithm=@var{algorithm}  
 5892 Use specified coloring algorithm for the integrated register  
 5893 allocator. The @var{algorithm} argument should be @code{priority} or  
 5894 @code{CB}. The first algorithm specifies Chow's priority coloring,  
 5895 the second one specifies Chaitin-Briggs coloring. The second  
 5896 algorithm can be unimplemented for some architectures. If it is  
 5897 implemented, it is the default because Chaitin-Briggs coloring as a  
 5898 rule generates a better code.

5900 @item -fira-region=@var{region}  
 5901 Use specified regions for the integrated register allocator. The  
 5902 @var{region} argument should be one of @code{all}, @code{mixed}, or  
 5903 @code{one}. The first value means using all loops as register  
 5904 allocation regions, the second value which is the default means using  
 5905 all loops except for loops with small register pressure as the  
 5906 regions, and third one means using all function as a single region.  
 5907 The first value can give best result for machines with small size and  
 5908 irregular register set, the third one results in faster and generates  
 5909 decent code and the smallest size code, and the default value usually  
 5910 give the best results in most cases and for most architectures.

5912 @item -fira-coalesce  
 5913 @opindex fira-coalesce  
 5914 Do optimistic register coalescing. This option might be profitable for  
 5915 architectures with big regular register files.

5917 @item -fno-ira-share-save-slots  
 5918 @opindex fno-ira-share-save-slots  
 5919 Switch off sharing stack slots used for saving call used hard  
 5920 registers living through a call. Each hard register will get a  
 5921 separate stack slot and as a result function stack frame will be  
 5922 bigger.

5924 @item -fno-ira-share-spill-slots  
 5925 @opindex fno-ira-share-spill-slots  
 5926 Switch off sharing stack slots allocated for pseudo-registers. Each  
 5927 pseudo-register which did not get a hard register will get a separate  
 5928 stack slot and as a result function stack frame will be bigger.

5930 @item -fira-verbose=@var{n}  
 5931 @opindex fira-verbose  
 5932 Set up how verbose dump file for the integrated register allocator  
 5933 will be. Default value is 5. If the value is greater or equal to 10,  
 5934 the dump file will be stderr as if the value were @var{n} minus 10.

5936 @item -fdelayed-branch  
 5937 @opindex fdelayed-branch  
 5938 If supported for the target machine, attempt to reorder instructions  
 5939 to exploit instruction slots available after delayed branch  
 5940 instructions.

5942 Enabled at levels @option{-O}, @option{-O2}, @option{-O3}, @option{-Os}.

5944 @item -fschedule-insns  
 5945 @opindex fschedule-insns  
 5946 If supported for the target machine, attempt to reorder instructions to  
 5947 eliminate execution stalls due to required data being unavailable. This  
 5948 helps machines that have slow floating point or memory load instructions  
 5949 by allowing other instructions to be issued until the result of the load  
 5950 or floating point instruction is required.

5952 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5954 @item -fschedule-insns2  
 5955 @opindex fschedule-insns2  
 5956 Similar to @option{-fschedule-insns}, but requests an additional pass of  
 5957 instruction scheduling after register allocation has been done. This is  
 5958 especially useful on machines with a relatively small number of  
 5959 registers and where memory load instructions take more than one cycle.

5961 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

5963 @item -fno-sched-interblock  
 5964 @opindex fno-sched-interblock

5965 Don't schedule instructions across basic blocks. This is normally  
5966 enabled by default when scheduling before register allocation, i.e.:@  
5967 with @option{-fschedule-insns} or at @option{-O2} or higher.

5969 @item -fno-sched-spec  
5970 @opindex fno-sched-spec  
5971 Don't allow speculative motion of non-load instructions. This is normally  
5972 enabled by default when scheduling before register allocation, i.e.:@  
5973 with @option{-fschedule-insns} or at @option{-O2} or higher.

5975 @item -fsched-spec-load  
5976 @opindex fsched-spec-load  
5977 Allow speculative motion of some load instructions. This only makes  
5978 sense when scheduling before register allocation, i.e.:@ with  
5979 @option{-fschedule-insns} or at @option{-O2} or higher.

5981 @item -fsched-spec-load-dangerous  
5982 @opindex fsched-spec-load-dangerous  
5983 Allow speculative motion of more load instructions. This only makes  
5984 sense when scheduling before register allocation, i.e.:@ with  
5985 @option{-fschedule-insns} or at @option{-O2} or higher.

5987 @item -fsched-stalled-insns  
5988 @itemx -fsched-stalled-insns=@var{n}  
5989 @opindex fsched-stalled-insns  
5990 Define how many insns (if any) can be moved prematurely from the queue  
5991 of stalled insns into the ready list, during the second scheduling pass.  
5992 @option{-fno-sched-stalled-insns} means that no insns will be moved  
5993 prematurely, @option{-fsched-stalled-insns=0} means there is no limit  
5994 on how many queued insns can be moved prematurely.  
5995 @option{-fsched-stalled-insns} without a value is equivalent to  
5996 @option{-fsched-stalled-insns=1}.

5998 @item -fsched-stalled-insns-dep  
5999 @itemx -fsched-stalled-insns-dep=@var{n}  
6000 @opindex fsched-stalled-insns-dep  
6001 Define how many insn groups (cycles) will be examined for a dependency  
6002 on a stalled insn that is candidate for premature removal from the queue  
6003 of stalled insns. This has an effect only during the second scheduling pass,  
6004 and only if @option{-fsched-stalled-insns} is used.  
6005 @option{-fno-sched-stalled-insns-dep} is equivalent to  
6006 @option{-fsched-stalled-insns-dep=0}.  
6007 @option{-fsched-stalled-insns-dep} without a value is equivalent to  
6008 @option{-fsched-stalled-insns-dep=1}.

6010 @item -fsched2-use-superblocks  
6011 @opindex fsched2-use-superblocks  
6012 When scheduling after register allocation, do use superblock scheduling  
6013 algorithm. Superblock scheduling allows motion across basic block boundaries  
6014 resulting on faster schedules. This option is experimental, as not all machine  
6015 descriptions used by GCC model the CPU closely enough to avoid unreliable  
6016 results from the algorithm.

6018 This only makes sense when scheduling after register allocation, i.e.:@ with  
6019 @option{-fschedule-insns2} or at @option{-O2} or higher.

6021 @item -fsched2-use-traces  
6022 @opindex fsched2-use-traces  
6023 Use @option{-fsched2-use-superblocks} algorithm when scheduling after register  
6024 allocation and additionally perform code duplication in order to increase the  
6025 size of superblocks using tracer pass. See @option{-ftracer} for details on  
6026 trace formation.

6028 This mode should produce faster but significantly longer programs. Also  
6029 without @option{-fbranch-probabilities} the traces constructed may not  
6030 match the reality and hurt the performance. This only makes

6031 sense when scheduling after register allocation, i.e.:@ with  
6032 @option{-fschedule-insns2} or at @option{-O2} or higher.

6034 @item -fsee  
6035 @opindex fsee  
6036 Eliminate redundant sign extension instructions and move the non-redundant  
6037 ones to optimal placement using lazy code motion (LCM).

6039 @item -freschedule-modulo-scheduled-loops  
6040 @opindex freschedule-modulo-scheduled-loops  
6041 The modulo scheduling comes before the traditional scheduling, if a loop  
6042 was modulo scheduled we may want to prevent the later scheduling passes  
6043 from changing its schedule, we use this option to control that.

6045 @item -fselective-scheduling  
6046 @opindex fselective-scheduling  
6047 Schedule instructions using selective scheduling algorithm. Selective  
6048 scheduling runs instead of the first scheduler pass.

6050 @item -fselective-scheduling2  
6051 @opindex fselective-scheduling2  
6052 Schedule instructions using selective scheduling algorithm. Selective  
6053 scheduling runs instead of the second scheduler pass.

6055 @item -fsel-sched-pipelining  
6056 @opindex fsel-sched-pipelining  
6057 Enable software pipelining of innermost loops during selective scheduling.  
6058 This option has no effect until one of @option{-fselective-scheduling} or  
6059 @option{-fselective-scheduling2} is turned on.

6061 @item -fsel-sched-pipelining-outer-loops  
6062 @opindex fsel-sched-pipelining-outer-loops  
6063 When pipelining loops during selective scheduling, also pipeline outer loops.  
6064 This option has no effect until @option{-fsel-sched-pipelining} is turned on.

6066 @item -fcaller-saves  
6067 @opindex fcallee-saves  
6068 Enable values to be allocated in registers that will be clobbered by  
6069 function calls, by emitting extra instructions to save and restore the  
6070 registers around such calls. Such allocation is done only when it  
6071 seems to result in better code than would otherwise be produced.

6073 This option is always enabled by default on certain machines, usually  
6074 those which have no call-preserved registers to use instead.

6076 Enabled at levels @option{-O2}, @option{-O3}, @option{-Os}.

6078 @item -fconserve-stack  
6079 @opindex fconserve-stack  
6080 Attempt to minimize stack usage. The compiler will attempt to use less  
6081 stack space, even if that makes the program slower. This option  
6082 implies setting the @option{large-stack-frame} parameter to 100  
6083 and the @option{large-stack-frame-growth} parameter to 400.

6085 @item -ftree-reassoc  
6086 @opindex ftree-reassoc  
6087 Perform reassociation on trees. This flag is enabled by default  
6088 at @option{-O} and higher.

6090 @item -ftree-pre  
6091 @opindex ftree-pre  
6092 Perform partial redundancy elimination (PRE) on trees. This flag is  
6093 enabled by default at @option{-O2} and @option{-O3}.

6095 @item -ftree-fre  
6096 @opindex ftree-fre

6097 Perform full redundancy elimination (FRE) on trees. The difference  
6098 between FRE and PRE is that FRE only considers expressions  
6099 that are computed on all paths leading to the redundant computation.  
6100 This analysis is faster than PRE, though it exposes fewer redundancies.  
6101 This flag is enabled by default at `@option{-O}` and higher.

6103 @item -ftree-copy-prop  
6104 @opindex ftree-copy-prop  
6105 Perform copy propagation on trees. This pass eliminates unnecessary  
6106 copy operations. This flag is enabled by default at `@option{-O}` and  
6107 higher.

6109 @item -fipa-pure-const  
6110 @opindex fipa-pure-const  
6111 Discover which functions are pure or constant.  
6112 Enabled by default at `@option{-O}` and higher.

6114 @item -fipa-reference  
6115 @opindex fipa-reference  
6116 Discover which static variables do not escape cannot escape the  
6117 compilation unit.  
6118 Enabled by default at `@option{-O}` and higher.

6120 @item -fipa-struct-reorg  
6121 @opindex fipa-struct-reorg  
6122 Perform structure reorganization optimization, that change C-like structures  
6123 layout in order to better utilize spatial locality. This transformation is  
6124 affective for programs containing arrays of structures. Available in two  
6125 compilation modes: profile-based (enabled with `@option{-fprofile-generate}`)  
6126 or static (which uses built-in heuristics). Require `@option{-fipa-type-escape}`  
6127 to provide the safety of this transformation. It works only in whole program  
6128 mode, so it requires `@option{-fwhole-program}` and `@option{-combine}` to be  
6129 enabled. Structures considered `@samp{cold}` by this transformation are not  
6130 affected (see `@option{--param struct-reorg-cold-struct-ratio=@var{value}}`).

6132 With this flag, the program debug info reflects a new structure layout.

6134 @item -fipa-pta  
6135 @opindex fipa-pta  
6136 Perform interprocedural pointer analysis. This option is experimental  
6137 and does not affect generated code.

6139 @item -fipa-cp  
6140 @opindex fipa-cp  
6141 Perform interprocedural constant propagation.  
6142 This optimization analyzes the program to determine when values passed  
6143 to functions are constants and then optimizes accordingly.  
6144 This optimization can substantially increase performance  
6145 if the application has constants passed to functions.  
6146 This flag is enabled by default at `@option{-O2}`, `@option{-Os}` and `@option{-O3}`.

6148 @item -fipa-cp-clone  
6149 @opindex fipa-cp-clone  
6150 Perform function cloning to make interprocedural constant propagation stronger.  
6151 When enabled, interprocedural constant propagation will perform function cloning  
6152 when externally visible function can be called with constant arguments.  
6153 Because this optimization can create multiple copies of functions,  
6154 it may significantly increase code size  
6155 (see `@option{--param ipcp-unit-growth=@var{value}}`).  
6156 This flag is enabled by default at `@option{-O3}`.

6158 @item -fipa-matrix-reorg  
6159 @opindex fipa-matrix-reorg  
6160 Perform matrix flattening and transposing.  
6161 Matrix flattening tries to replace a m-dimensional matrix  
6162 with its equivalent n-dimensional matrix, where  $n < m$ .

6163 This reduces the level of indirection needed for accessing the elements  
6164 of the matrix. The second optimization is matrix transposing that  
6165 attempts to change the order of the matrix's dimensions in order to  
6166 improve cache locality.  
6167 Both optimizations need the `@option{-fwhole-program}` flag.  
6168 Transposing is enabled only if profiling information is available.

6171 @item -ftree-sink  
6172 @opindex ftree-sink  
6173 Perform forward store motion on trees. This flag is  
6174 enabled by default at `@option{-O}` and higher.

6176 @item -ftree-ccp  
6177 @opindex ftree-ccp  
6178 Perform sparse conditional constant propagation (CCP) on trees. This  
6179 pass only operates on local scalar variables and is enabled by default  
6180 at `@option{-O}` and higher.

6182 @item -ftree-switch-conversion  
6183 Perform conversion of simple initializations in a switch to  
6184 initializations from a scalar array. This flag is enabled by default  
6185 at `@option{-O2}` and higher.

6187 @item -ftree-dce  
6188 @opindex ftree-dce  
6189 Perform dead code elimination (DCE) on trees. This flag is enabled by  
6190 default at `@option{-O}` and higher.

6192 @item -ftree-builtin-call-dce  
6193 @opindex ftree-builtin-call-dce  
6194 Perform conditional dead code elimination (DCE) for calls to builtin functions  
6195 that may set `@code{errno}` but are otherwise side-effect free. This flag is  
6196 enabled by default at `@option{-O2}` and higher if `@option{-Os}` is not also  
6197 specified.

6199 @item -ftree-dominator-opts  
6200 @opindex ftree-dominator-opts  
6201 Perform a variety of simple scalar cleanups (constant/copy  
6202 propagation, redundancy elimination, range propagation and expression  
6203 simplification) based on a dominator tree traversal. This also  
6204 performs jump threading (to reduce jumps to jumps). This flag is  
6205 enabled by default at `@option{-O}` and higher.

6207 @item -ftree-dse  
6208 @opindex ftree-dse  
6209 Perform dead store elimination (DSE) on trees. A dead store is a store into  
6210 a memory location which will later be overwritten by another store without  
6211 any intervening loads. In this case the earlier store can be deleted. This  
6212 flag is enabled by default at `@option{-O}` and higher.

6214 @item -ftree-ch  
6215 @opindex ftree-ch  
6216 Perform loop header copying on trees. This is beneficial since it increases  
6217 effectiveness of code motion optimizations. It also saves one jump. This flag  
6218 is enabled by default at `@option{-O}` and higher. It is not enabled  
6219 for `@option{-Os}`, since it usually increases code size.

6221 @item -ftree-loop-optimize  
6222 @opindex ftree-loop-optimize  
6223 Perform loop optimizations on trees. This flag is enabled by default  
6224 at `@option{-O}` and higher.

6226 @item -ftree-loop-linear  
6227 @opindex ftree-loop-linear  
6228 Perform linear loop transformations on tree. This flag can improve cache



6229 performance and allow further loop optimizations to take place.

6231 @item -floop-interchange  
 6232 Perform loop interchange transformations on loops. Interchanging two  
 6233 nested loops switches the inner and outer loops. For example, given a  
 6234 loop like:  
 6235 @smallexample  
 6236 DO J = 1, M  
 6237 DO I = 1, N  
 6238 A(J, I) = A(J, I) \* C  
 6239 ENDDO  
 6240 ENDDO  
 6241 @end smallexample  
 6242 loop interchange will transform the loop as if the user had written:  
 6243 @smallexample  
 6244 DO I = 1, N  
 6245 DO J = 1, M  
 6246 A(J, I) = A(J, I) \* C  
 6247 ENDDO  
 6248 ENDDO  
 6249 @end smallexample  
 6250 which can be beneficial when @code{N} is larger than the caches,  
 6251 because in Fortran, the elements of an array are stored in memory  
 6252 contiguously by column, and the original loop iterates over rows,  
 6253 potentially creating at each access a cache miss. This optimization  
 6254 applies to all the languages supported by GCC and is not limited to  
 6255 Fortran. To use this code transformation, GCC has to be configured  
 6256 with @option{--with-ppl} and @option{--with-cloog} to enable the  
 6257 Graphite loop transformation infrastructure.

6259 @item -floop-strip-mine  
 6260 Perform loop strip mining transformations on loops. Strip mining  
 6261 splits a loop into two nested loops. The outer loop has strides  
 6262 equal to the strip size and the inner loop has strides of the  
 6263 original loop within a strip. For example, given a loop like:  
 6264 @smallexample  
 6265 DO I = 1, N  
 6266 A(I) = A(I) + C  
 6267 ENDDO  
 6268 @end smallexample  
 6269 loop strip mining will transform the loop as if the user had written:  
 6270 @smallexample  
 6271 DO II = 1, N, 4  
 6272 DO I = II, min (II + 3, N)  
 6273 A(I) = A(I) + C  
 6274 ENDDO  
 6275 ENDDO  
 6276 @end smallexample  
 6277 This optimization applies to all the languages supported by GCC and is  
 6278 not limited to Fortran. To use this code transformation, GCC has to  
 6279 be configured with @option{--with-ppl} and @option{--with-cloog} to  
 6280 enable the Graphite loop transformation infrastructure.

6282 @item -floop-block  
 6283 Perform loop blocking transformations on loops. Blocking strip mines  
 6284 each loop in the loop nest such that the memory accesses of the  
 6285 element loops fit inside caches. For example, given a loop like:  
 6286 @smallexample  
 6287 DO I = 1, N  
 6288 DO J = 1, M  
 6289 A(J, I) = B(I) + C(J)  
 6290 ENDDO  
 6291 ENDDO  
 6292 @end smallexample  
 6293 loop blocking will transform the loop as if the user had written:  
 6294 @smallexample

6295 DO II = 1, N, 64  
 6296 DO JJ = 1, M, 64  
 6297 DO I = II, min (II + 63, N)  
 6298 DO J = JJ, min (JJ + 63, M)  
 6299 A(J, I) = B(I) + C(J)  
 6300 ENDDO  
 6301 ENDDO  
 6302 ENDDO  
 6303 ENDDO  
 6304 @end smallexample  
 6305 which can be beneficial when @code{M} is larger than the caches,  
 6306 because the innermost loop will iterate over a smaller amount of data  
 6307 that can be kept in the caches. This optimization applies to all the  
 6308 languages supported by GCC and is not limited to Fortran. To use this  
 6309 code transformation, GCC has to be configured with @option{--with-ppl}  
 6310 and @option{--with-cloog} to enable the Graphite loop transformation  
 6311 infrastructure.

6313 @item -fcheck-data-deps  
 6314 @opindex fcheck-data-deps  
 6315 Compare the results of several data dependence analyzers. This option  
 6316 is used for debugging the data dependence analyzers.

6318 @item -ftree-loop-distribution  
 6319 Perform loop distribution. This flag can improve cache performance on  
 6320 big loop bodies and allow further loop optimizations, like  
 6321 parallelization or vectorization, to take place. For example, the loop  
 6322 @smallexample  
 6323 DO I = 1, N  
 6324 A(I) = B(I) + C  
 6325 D(I) = E(I) \* F  
 6326 ENDDO  
 6327 @end smallexample  
 6328 is transformed to  
 6329 @smallexample  
 6330 DO I = 1, N  
 6331 A(I) = B(I) + C  
 6332 ENDDO  
 6333 DO I = 1, N  
 6334 D(I) = E(I) \* F  
 6335 ENDDO  
 6336 @end smallexample

6338 @item -ftree-loop-im  
 6339 @opindex ftree-loop-im  
 6340 Perform loop invariant motion on trees. This pass moves only invariants that  
 6341 would be hard to handle at RTL level (function calls, operations that expand to  
 6342 nontrivial sequences of insns). With @option{--funswitch-loops} it also moves  
 6343 operands of conditions that are invariant out of the loop, so that we can use  
 6344 just trivial invariance analysis in loop unswitching. The pass also includes  
 6345 store motion.

6347 @item -ftree-loop-ivcanon  
 6348 @opindex ftree-loop-ivcanon  
 6349 Create a canonical counter for number of iterations in the loop for that  
 6350 determining number of iterations requires complicated analysis. Later  
 6351 optimizations then may determine the number easily. Useful especially  
 6352 in connection with unrolling.

6354 @item -fivopts  
 6355 @opindex fivopts  
 6356 Perform induction variable optimizations (strength reduction, induction  
 6357 variable merging and induction variable elimination) on trees.

6359 @item -ftree-parallelize-loops=n  
 6360 @opindex ftree-parallelize-loops

6361 Parallelize loops, i.e., split their iteration space to run in  $n$  threads.  
 6362 This is only possible for loops whose iterations are independent  
 6363 and can be arbitrarily reordered. The optimization is only  
 6364 profitable on multiprocessor machines, for loops that are CPU-intensive,  
 6365 rather than constrained e.g. by memory bandwidth. This option  
 6366 implies `@option{-pthread}`, and thus is only supported on targets  
 6367 that have support for `@option{-pthread}`.

6369 `@item -ftree-sra`  
 6370 `@opindex ftree-sra`  
 6371 Perform scalar replacement of aggregates. This pass replaces structure  
 6372 references with scalars to prevent committing structures to memory too  
 6373 early. This flag is enabled by default at `@option{-O}` and higher.

6375 `@item -ftree-copyrename`  
 6376 `@opindex ftree-copyrename`  
 6377 Perform copy renaming on trees. This pass attempts to rename compiler  
 6378 temporaries to other variables at copy locations, usually resulting in  
 6379 variable names which more closely resemble the original variables. This flag  
 6380 is enabled by default at `@option{-O}` and higher.

6382 `@item -ftree-ter`  
 6383 `@opindex ftree-ter`  
 6384 Perform temporary expression replacement during the SSA->normal phase. Single  
 6385 use/single def temporaries are replaced at their use location with their  
 6386 defining expression. This results in non-GIMPLE code, but gives the expanders  
 6387 much more complex trees to work on resulting in better RTL generation. This is  
 6388 enabled by default at `@option{-O}` and higher.

6390 `@item -ftree-vectorize`  
 6391 `@opindex ftree-vectorize`  
 6392 Perform loop vectorization on trees. This flag is enabled by default at  
 6393 `@option{-O3}`.

6395 `@item -ftree-vect-loop-version`  
 6396 `@opindex ftree-vect-loop-version`  
 6397 Perform loop versioning when doing loop vectorization on trees. When a loop  
 6398 appears to be vectorizable except that data alignment or data dependence cannot  
 6399 be determined at compile time then vectorized and non-vectorized versions of  
 6400 the loop are generated along with runtime checks for alignment or dependence  
 6401 to control which version is executed. This option is enabled by default  
 6402 except at level `@option{-Os}` where it is disabled.

6404 `@item -fvect-cost-model`  
 6405 `@opindex fvect-cost-model`  
 6406 Enable cost model for vectorization.

6408 `@item -ftree-vrp`  
 6409 `@opindex ftree-vrp`  
 6410 Perform Value Range Propagation on trees. This is similar to the  
 6411 constant propagation pass, but instead of values, ranges of values are  
 6412 propagated. This allows the optimizers to remove unnecessary range  
 6413 checks like array bound checks and null pointer checks. This is  
 6414 enabled by default at `@option{-O2}` and higher. Null pointer check  
 6415 elimination is only done if `@option{-fdelete-null-pointer-checks}` is  
 6416 enabled.

6418 `@item -ftracer`  
 6419 `@opindex ftracer`  
 6420 Perform tail duplication to enlarge superblock size. This transformation  
 6421 simplifies the control flow of the function allowing other optimizations to do  
 6422 better job.

6424 `@item -funroll-loops`  
 6425 `@opindex funroll-loops`  
 6426 Unroll loops whose number of iterations can be determined at compile

6427 time or upon entry to the loop. `@option{-funroll-loops}` implies  
 6428 `@option{-frerun-cse-after-loop}`. This option makes code larger,  
 6429 and may or may not make it run faster.

6431 `@item -funroll-all-loops`  
 6432 `@opindex funroll-all-loops`  
 6433 Unroll all loops, even if their number of iterations is uncertain when  
 6434 the loop is entered. This usually makes programs run more slowly.  
 6435 `@option{-funroll-all-loops}` implies the same options as  
 6436 `@option{-funroll-loops}`,

6438 `@item -fsplit-ivs-in-unroller`  
 6439 `@opindex fsplit-ivs-in-unroller`  
 6440 Enables expressing of values of induction variables in later iterations  
 6441 of the unrolled loop using the value in the first iteration. This breaks  
 6442 long dependency chains, thus improving efficiency of the scheduling passes.

6444 Combination of `@option{-fweb}` and CSE is often sufficient to obtain the  
 6445 same effect. However in cases the loop body is more complicated than  
 6446 a single basic block, this is not reliable. It also does not work at all  
 6447 on some of the architectures due to restrictions in the CSE pass.

6449 This optimization is enabled by default.

6451 `@item -fvariable-expansion-in-unroller`  
 6452 `@opindex fvariable-expansion-in-unroller`  
 6453 With this option, the compiler will create multiple copies of some  
 6454 local variables when unrolling a loop which can result in superior code.

6456 `@item -fpredictive-commoning`  
 6457 `@opindex fpredictive-commoning`  
 6458 Perform predictive commoning optimization, i.e., reusing computations  
 6459 (especially memory loads and stores) performed in previous  
 6460 iterations of loops.

6462 This option is enabled at level `@option{-O3}`.

6464 `@item -fprefetch-loop-arrays`  
 6465 `@opindex fprefetch-loop-arrays`  
 6466 If supported by the target machine, generate instructions to prefetch  
 6467 memory to improve the performance of loops that access large arrays.

6469 This option may generate better or worse code; results are highly  
 6470 dependent on the structure of loops within the source code.

6472 Disabled at level `@option{-Os}`.

6474 `@item -fno-peekhole`  
 6475 `@itemx -fno-peekhole2`  
 6476 `@opindex fno-peekhole`  
 6477 `@opindex fno-peekhole2`  
 6478 Disable any machine-specific peekhole optimizations. The difference  
 6479 between `@option{-fno-peekhole}` and `@option{-fno-peekhole2}` is in how they  
 6480 are implemented in the compiler; some targets use one, some use the  
 6481 other, a few use both.

6483 `@option{-fpeekhole}` is enabled by default.  
 6484 `@option{-fpeekhole2}` enabled at levels `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6486 `@item -fno-guess-branch-probability`  
 6487 `@opindex fno-guess-branch-probability`  
 6488 Do not guess branch probabilities using heuristics.

6490 GCC will use heuristics to guess branch probabilities if they are  
 6491 not provided by profiling feedback (`@option{-fprofile-arcs}`). These  
 6492 heuristics are based on the control flow graph. If some branch probabilities

6493 are specified by `@samp{__builtin_expect}`, then the heuristics will be  
 6494 used to guess branch probabilities for the rest of the control flow graph,  
 6495 taking the `@samp{__builtin_expect}` info into account. The interactions  
 6496 between the heuristics and `@samp{__builtin_expect}` can be complex, and in  
 6497 some cases, it may be useful to disable the heuristics so that the effects  
 6498 of `@samp{__builtin_expect}` are easier to understand.

6500 The default is `@option{-fguess-branch-probability}` at levels  
 6501 `@option{-O}`, `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6503 `@item -freorder-blocks`  
 6504 `@opindex freorder-blocks`  
 6505 Reorder basic blocks in the compiled function in order to reduce number of  
 6506 taken branches and improve code locality.

6508 Enabled at levels `@option{-O2}`, `@option{-O3}`.

6510 `@item -freorder-blocks-and-partition`  
 6511 `@opindex freorder-blocks-and-partition`  
 6512 In addition to reordering basic blocks in the compiled function, in order  
 6513 to reduce number of taken branches, partitions hot and cold basic blocks  
 6514 into separate sections of the assembly and `.o` files, to improve  
 6515 paging and cache locality performance.

6517 This optimization is automatically turned off in the presence of  
 6518 exception handling, for linkonce sections, for functions with a user-defined  
 6519 section attribute and on any architecture that does not support named  
 6520 sections.

6522 `@item -freorder-functions`  
 6523 `@opindex freorder-functions`  
 6524 Reorder functions in the object file in order to  
 6525 improve code locality. This is implemented by using special  
 6526 subsections `@code{.text.hot}` for most frequently executed functions and  
 6527 `@code{.text.unlikely}` for unlikely executed functions. Reordering is done by  
 6528 the linker so object file format must support named sections and linker must  
 6529 place them in a reasonable way.

6531 Also profile feedback must be available in to make this option effective. See  
 6532 `@option{-fprofile-arcs}` for details.

6534 Enabled at levels `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6536 `@item -fstrict-aliasing`  
 6537 `@opindex fstrict-aliasing`  
 6538 Allow the compiler to assume the strictest aliasing rules applicable to  
 6539 the language being compiled. For C (and C++), this activates  
 6540 optimizations based on the type of expressions. In particular, an  
 6541 object of one type is assumed never to reside at the same address as an  
 6542 object of a different type, unless the types are almost the same. For  
 6543 example, an `@code{unsigned int}` can alias an `@code{int}`, but not a  
 6544 `@code{void*}` or a `@code{double}`. A character type may alias any other  
 6545 type.

6547 `@anchor{Type-punning}` Pay special attention to code like this:

```
6548 @smallexample
6549 union a_union @
6550   int i;
6551   double d;
6552 @};
```

```
6554 int f() @
6555   union a_union t;
6556   t.d = 3.0;
6557   return t.i;
6558 @}
```

```
6559 @end smallexample
6560 The practice of reading from a different union member than the one most
6561 recently written to (called "type-punning") is common. Even with
6562 @option{-fstrict-aliasing}, type-punning is allowed, provided the memory
6563 is accessed through the union type. So, the code above will work as
6564 expected. @xref{Structures unions enumerations and bit-fields}
6565 implementation}. However, this code might not:
6566 @smallexample
6567 int f() @
6568   union a_union t;
6569   int* ip;
6570   t.d = 3.0;
6571   ip = &t.i;
6572   return *ip;
6573 @}
6574 @end smallexample
```

6576 Similarly, access by taking the address, casting the resulting pointer  
 6577 and dereferencing the result has undefined behavior, even if the cast  
 6578 uses a union type, e.g.:

```
6579 @smallexample
6580 int f() @
6581   double d = 3.0;
6582   return ((union a_union *) &d)->i;
6583 @}
6584 @end smallexample
```

6586 The `@option{-fstrict-aliasing}` option is enabled at levels  
 6587 `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6589 `@item -fstrict-calling-conventions`  
 6590 `@opindex mstrict-calling-conventions`  
 6591 Use strict ABI calling conventions even with local functions.  
 6592 This disable certain optimizations that may cause GCC to call local  
 6593 functions in a manner other than that described by the ABI.

```
6595 #endif /* ! codereview */
6596 @item -fstrict-overflow
6597 @opindex fstrict-overflow
6598 Allow the compiler to assume strict signed overflow rules, depending
6599 on the language being compiled. For C (and C++) this means that
6600 overflow when doing arithmetic with signed numbers is undefined, which
6601 means that the compiler may assume that it will not happen. This
6602 permits various optimizations. For example, the compiler will assume
6603 that an expression like @code{i + 10 > i} will always be true for
6604 signed @code{i}. This assumption is only valid if signed overflow is
6605 undefined, as the expression is false if @code{i + 10} overflows when
6606 using twos complement arithmetic. When this option is in effect any
6607 attempt to determine whether an operation on signed numbers will
6608 overflow must be written carefully to not actually involve overflow.
```

6610 This option also allows the compiler to assume strict pointer  
 6611 semantics: given a pointer to an object, if adding an offset to that  
 6612 pointer does not produce a pointer to the same object, the addition is  
 6613 undefined. This permits the compiler to conclude that `@code{p + u >`  
 6614 `p}` is always true for a pointer `@code{p}` and unsigned integer  
 6615 `@code{u}`. This assumption is only valid because pointer wraparound is  
 6616 undefined, as the expression is false if `@code{p + u}` overflows using  
 6617 twos complement arithmetic.

6619 See also the `@option{-fwrapv}` option. Using `@option{-fwrapv}` means  
 6620 that integer signed overflow is fully defined: it wraps. When  
 6621 `@option{-fwrapv}` is used, there is no difference between  
 6622 `@option{-fstrict-overflow}` and `@option{-fno-strict-overflow}` for  
 6623 integers. With `@option{-fwrapv}` certain types of overflow are  
 6624 permitted. For example, if the compiler gets an overflow when doing

6625 arithmetic on constants, the overflowed value can still be used with  
6626 `@option{-fwrapv}`, but not otherwise.

6628 The `@option{-fstrict-overflow}` option is enabled at levels  
6629 `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6631 `@item -falign-functions`  
6632 `@itemx -falign-functions=@var{n}`  
6633 `@opindex falign-functions`  
6634 Align the start of functions to the next power-of-two greater than  
6635 `@var{n}`, skipping up to `@var{n}` bytes. For instance,  
6636 `@option{-falign-functions=32}` aligns functions to the next 32-byte  
6637 boundary, but `@option{-falign-functions=24}` would align to the next  
6638 32-byte boundary only if this can be done by skipping 23 bytes or less.

6640 `@option{-fno-align-functions}` and `@option{-falign-functions=1}` are  
6641 equivalent and mean that functions will not be aligned.

6643 Some assemblers only support this flag when `@var{n}` is a power of two;  
6644 in that case, it is rounded up.

6646 If `@var{n}` is not specified or is zero, use a machine-dependent default.

6648 Enabled at levels `@option{-O2}`, `@option{-O3}`.

6650 `@item -falign-labels`  
6651 `@itemx -falign-labels=@var{n}`  
6652 `@opindex falign-labels`  
6653 Align all branch targets to a power-of-two boundary, skipping up to  
6654 `@var{n}` bytes like `@option{-falign-functions}`. This option can easily  
6655 make code slower, because it must insert dummy operations for when the  
6656 branch target is reached in the usual flow of the code.

6658 `@option{-fno-align-labels}` and `@option{-falign-labels=1}` are  
6659 equivalent and mean that labels will not be aligned.

6661 If `@option{-falign-loops}` or `@option{-falign-jumps}` are applicable and  
6662 are greater than this value, then their values are used instead.

6664 If `@var{n}` is not specified or is zero, use a machine-dependent default  
6665 which is very likely to be `@samp{1}`, meaning no alignment.

6667 Enabled at levels `@option{-O2}`, `@option{-O3}`.

6669 `@item -falign-loops`  
6670 `@itemx -falign-loops=@var{n}`  
6671 `@opindex falign-loops`  
6672 Align loops to a power-of-two boundary, skipping up to `@var{n}` bytes  
6673 like `@option{-falign-functions}`. The hope is that the loop will be  
6674 executed many times, which will make up for any execution of the dummy  
6675 operations.

6677 `@option{-fno-align-loops}` and `@option{-falign-loops=1}` are  
6678 equivalent and mean that loops will not be aligned.

6680 If `@var{n}` is not specified or is zero, use a machine-dependent default.

6682 Enabled at levels `@option{-O2}`, `@option{-O3}`.

6684 `@item -falign-jumps`  
6685 `@itemx -falign-jumps=@var{n}`  
6686 `@opindex falign-jumps`  
6687 Align branch targets to a power-of-two boundary, for branch targets  
6688 where the targets can only be reached by jumping, skipping up to `@var{n}`  
6689 bytes like `@option{-falign-functions}`. In this case, no dummy operations  
6690 need be executed.

6692 `@option{-fno-align-jumps}` and `@option{-falign-jumps=1}` are  
6693 equivalent and mean that loops will not be aligned.

6695 If `@var{n}` is not specified or is zero, use a machine-dependent default.

6697 Enabled at levels `@option{-O2}`, `@option{-O3}`.

6699 `@item -funit-at-a-time`  
6700 `@opindex funit-at-a-time`  
6701 This option is left for compatibility reasons. `@option{-funit-at-a-time}`  
6702 has no effect, while `@option{-fno-unit-at-a-time}` implies  
6703 `@option{-fno-toplevel-reorder}` and `@option{-fno-section-anchors}`.

6705 Enabled by default.

6707 `@item -fno-toplevel-reorder`  
6708 `@opindex fno-toplevel-reorder`  
6709 Do not reorder top-level functions, variables, and `@code{asm}`  
6710 statements. Output them in the same order that they appear in the  
6711 input file. When this option is used, unreferenced static variables  
6712 will not be removed. This option is intended to support existing code  
6713 which relies on a particular ordering. For new code, it is better to  
6714 use attributes.

6716 Enabled at level `@option{-O0}`. When disabled explicitly, it also imply  
6717 `@option{-fno-section-anchors}` that is otherwise enabled at `@option{-O0}` on some  
6718 targets.

6720 `@item -fweb`  
6721 `@opindex fweb`  
6722 Constructs webs as commonly used for register allocation purposes and assign  
6723 each web individual pseudo register. This allows the register allocation pass  
6724 to operate on pseudos directly, but also strengthens several other optimization  
6725 passes, such as CSE, loop optimizer and trivial dead code remover. It can,  
6726 however, make debugging impossible, since variables will no longer stay in a  
6727 `''home register''`.

6729 Enabled by default with `@option{-funroll-loops}`.

6731 `@item -fwhole-program`  
6732 `@opindex fwhole-program`  
6733 Assume that the current compilation unit represents whole program being  
6734 compiled. All public functions and variables with the exception of `@code{main}`  
6735 and those merged by attribute `@code{externally_visible}` become static functions  
6736 and in a affect gets more aggressively optimized by interprocedural optimizers.  
6737 While this option is equivalent to proper use of `@code{static}` keyword for  
6738 programs consisting of single file, in combination with option  
6739 `@option{--combine}` this flag can be used to compile most of smaller scale C  
6740 programs since the functions and variables become local for the whole combined  
6741 compilation unit, not for the single source file itself.

6743 This option is not supported for Fortran programs.

6745 `@item -fcprop-registers`  
6746 `@opindex fcprop-registers`  
6747 After register allocation and post-register allocation instruction splitting,  
6748 we perform a copy-propagation pass to try to reduce scheduling dependencies  
6749 and occasionally eliminate the copy.

6751 Enabled at levels `@option{-O}`, `@option{-O2}`, `@option{-O3}`, `@option{-Os}`.

6753 `@item -fprofile-correction`  
6754 `@opindex fprofile-correction`  
6755 Profiles collected using an instrumented binary for multi-threaded programs may  
6756 be inconsistent due to missed counter updates. When this option is specified,

6757 GCC will use heuristics to correct or smooth out such inconsistencies. By  
6758 default, GCC will emit an error message when an inconsistent profile is detected

6760 @item -fprofile-dir=@var{path}  
6761 @opindex fprofile-dir

6763 Set the directory to search the profile data files in to @var{path}.  
6764 This option affects only the profile data generated by  
6765 @option{-fprofile-generate}, @option{-ftest-coverage}, @option{-fprofile-arcs}  
6766 and used by @option{-fprofile-use} and @option{-fbranch-probabilities}  
6767 and its related options.  
6768 By default, GCC will use the current directory as @var{path}  
6769 thus the profile data file will appear in the same directory as the object file.

6771 @item -fprofile-generate  
6772 @itemx -fprofile-generate=@var{path}  
6773 @opindex fprofile-generate

6775 Enable options usually used for instrumenting application to produce  
6776 profile useful for later recompilation with profile feedback based  
6777 optimization. You must use @option{-fprofile-generate} both when  
6778 compiling and when linking your program.

6780 The following options are enabled: @code{-fprofile-arcs}, @code{-fprofile-values}

6782 If @var{path} is specified, GCC will look at the @var{path} to find  
6783 the profile feedback data files. See @option{-fprofile-dir}.

6785 @item -fprofile-use  
6786 @itemx -fprofile-use=@var{path}  
6787 @opindex fprofile-use  
6788 Enable profile feedback directed optimizations, and optimizations  
6789 generally profitable only with profile feedback available.

6791 The following options are enabled: @code{-fbranch-probabilities}, @code{-fvpt},  
6792 @code{-funroll-loops}, @code{-fpeel-loops}, @code{-ftracer}

6794 By default, GCC emits an error message if the feedback profiles do not  
6795 match the source code. This error can be turned into a warning by using  
6796 @option{-Wcoverage-mismatch}. Note this may result in poorly optimized  
6797 code.

6799 If @var{path} is specified, GCC will look at the @var{path} to find  
6800 the profile feedback data files. See @option{-fprofile-dir}.  
6801 @end table

6803 The following options control compiler behavior regarding floating  
6804 point arithmetic. These options trade off between speed and  
6805 correctness. All must be specifically enabled.

6807 @table @gcctabopt  
6808 @item -ffloat-store  
6809 @opindex ffloat-store  
6810 Do not store floating point variables in registers, and inhibit other  
6811 options that might change whether a floating point value is taken from a  
6812 register or memory.

6814 @cindex floating point precision  
6815 This option prevents undesirable excess precision on machines such as  
6816 the 68000 where the floating registers (of the 68881) keep more  
6817 precision than a @code{double} is supposed to have. Similarly for the  
6818 x86 architecture. For most programs, the excess precision does only  
6819 good, but a few programs rely on the precise definition of IEEE floating  
6820 point. Use @option{-ffloat-store} for such programs, after modifying  
6821 them to store all pertinent intermediate computations into variables.

6823 @item -ffast-math  
6824 @opindex ffast-math  
6825 Sets @option{-fno-math-errno}, @option{-funsafe-math-optimizations},  
6826 @option{-ffinite-math-only}, @option{-fno-rounding-math},  
6827 @option{-fno-signaling-nans} and @option{-fcx-limited-range}.

6829 This option causes the preprocessor macro @code{\_\_FAST\_MATH\_\_} to be defined.

6831 This option is not turned on by any @option{-O} option since  
6832 it can result in incorrect output for programs which depend on  
6833 an exact implementation of IEEE or ISO rules/specifications for  
6834 math functions. It may, however, yield faster code for programs  
6835 that do not require the guarantees of these specifications.

6837 @item -fno-math-errno  
6838 @opindex fno-math-errno  
6839 Do not set ERRNO after calling math functions that are executed  
6840 with a single instruction, e.g., sqrt. A program that relies on  
6841 IEEE exceptions for math error handling may want to use this flag  
6842 for speed while maintaining IEEE arithmetic compatibility.

6844 This option is not turned on by any @option{-O} option since  
6845 it can result in incorrect output for programs which depend on  
6846 an exact implementation of IEEE or ISO rules/specifications for  
6847 math functions. It may, however, yield faster code for programs  
6848 that do not require the guarantees of these specifications.

6850 The default is @option{-fmath-errno}.

6852 On Darwin systems, the math library never sets @code{errno}. There is  
6853 therefore no reason for the compiler to consider the possibility that  
6854 it might, and @option{-fno-math-errno} is the default.

6856 @item -funsafe-math-optimizations  
6857 @opindex funsafe-math-optimizations

6859 Allow optimizations for floating-point arithmetic that (a) assume  
6860 that arguments and results are valid and (b) may violate IEEE or  
6861 ANSI standards. When used at link-time, it may include libraries  
6862 or startup files that change the default FPU control word or other  
6863 similar optimizations.

6865 This option is not turned on by any @option{-O} option since  
6866 it can result in incorrect output for programs which depend on  
6867 an exact implementation of IEEE or ISO rules/specifications for  
6868 math functions. It may, however, yield faster code for programs  
6869 that do not require the guarantees of these specifications.  
6870 Enables @option{-fno-signed-zeros}, @option{-fno-trapping-math},  
6871 @option{-fassociative-math} and @option{-freciprocal-math}.

6873 The default is @option{-fno-unsafe-math-optimizations}.

6875 @item -fassociative-math  
6876 @opindex fassociative-math

6878 Allow re-association of operands in series of floating-point operations.  
6879 This violates the ISO C and C++ language standard by possibly changing  
6880 computation result. NOTE: re-ordering may change the sign of zero as  
6881 well as ignore NaNs and inhibit or create underflow or overflow (and  
6882 thus cannot be used on a code which relies on rounding behavior like  
6883 @code{(x + 2\*\*52) - 2\*\*52}). May also reorder floating-point comparisons  
6884 and thus may not be used when ordered comparisons are required.  
6885 This option requires that both @option{-fno-signed-zeros} and  
6886 @option{-fno-trapping-math} be in effect. Moreover, it doesn't make  
6887 much sense with @option{-frounding-math}.

6889 The default is `@option{-fno-associative-math}`.

6891 @item -freciprocal-math  
6892 @opindex freciprocal-math

6894 Allow the reciprocal of a value to be used instead of dividing by  
6895 the value if this enables optimizations. For example `@code{x / y}`  
6896 can be replaced with `@code{x * (1/y)}` which is useful if `@code{(1/y)}`  
6897 is subject to common subexpression elimination. Note that this loses  
6898 precision and increases the number of flops operating on the value.

6900 The default is `@option{-fno-reciprocal-math}`.

6902 @item -ffinite-math-only  
6903 @opindex ffinite-math-only  
6904 Allow optimizations for floating-point arithmetic that assume  
6905 that arguments and results are not NaNs or +-Infs.

6907 This option is not turned on by any `@option{-O}` option since  
6908 it can result in incorrect output for programs which depend on  
6909 an exact implementation of IEEE or ISO rules/specifications for  
6910 math functions. It may, however, yield faster code for programs  
6911 that do not require the guarantees of these specifications.

6913 The default is `@option{-fno-finite-math-only}`.

6915 @item -fno-signed-zeros  
6916 @opindex fno-signed-zeros  
6917 Allow optimizations for floating point arithmetic that ignore the  
6918 signedness of zero. IEEE arithmetic specifies the behavior of  
6919 distinct `+0.0` and `@minus{}0.0` values, which then prohibits simplification  
6920 of expressions such as `x+0.0` or `0.0*x` (even with `@option{-ffinite-math-only}`).  
6921 This option implies that the sign of a zero result isn't significant.

6923 The default is `@option{-fsigned-zeros}`.

6925 @item -fno-trapping-math  
6926 @opindex fno-trapping-math  
6927 Compile code assuming that floating-point operations cannot generate  
6928 user-visible traps. These traps include division by zero, overflow,  
6929 underflow, inexact result and invalid operation. This option requires  
6930 that `@option{-fno-signaling-nans}` be in effect. Setting this option may  
6931 allow faster code if one relies on "non-stop" IEEE arithmetic, for example.

6933 This option should never be turned on by any `@option{-O}` option since  
6934 it can result in incorrect output for programs which depend on  
6935 an exact implementation of IEEE or ISO rules/specifications for  
6936 math functions.

6938 The default is `@option{-ftrapping-math}`.

6940 @item -frounding-math  
6941 @opindex frounding-math  
6942 Disable transformations and optimizations that assume default floating  
6943 point rounding behavior. This is round-to-zero for all floating point  
6944 to integer conversions, and round-to-nearest for all other arithmetic  
6945 truncations. This option should be specified for programs that change  
6946 the FP rounding mode dynamically, or that may be executed with a  
6947 non-default rounding mode. This option disables constant folding of  
6948 floating point expressions at compile-time (which may be affected by  
6949 rounding mode) and arithmetic transformations that are unsafe in the  
6950 presence of sign-dependent rounding modes.

6952 The default is `@option{-fno-rounding-math}`.

6954 This option is experimental and does not currently guarantee to

6955 disable all GCC optimizations that are affected by rounding mode.  
6956 Future versions of GCC may provide finer control of this setting  
6957 using C99's `@code{FENV_ACCESS}` pragma. This command line option  
6958 will be used to specify the default state for `@code{FENV_ACCESS}`.

6960 @item -frtl-abstract-sequences  
6961 @opindex frtl-abstract-sequences  
6962 It is a size optimization method. This option is to find identical  
6963 sequences of code, which can be turned into pseudo-procedures and  
6964 then replace all occurrences with calls to the newly created  
6965 subroutine. It is kind of an opposite of `@option{-finline-functions}`.  
6966 This optimization runs at RTL level.

6968 @item -fsignaling-nans  
6969 @opindex fsignaling-nans  
6970 Compile code assuming that IEEE signaling NaNs may generate user-visible  
6971 traps during floating-point operations. Setting this option disables  
6972 optimizations that may change the number of exceptions visible with  
6973 signaling NaNs. This option implies `@option{-ftrapping-math}`.

6975 This option causes the preprocessor macro `@code{__SUPPORT_SNAN__}` to  
6976 be defined.

6978 The default is `@option{-fno-signaling-nans}`.

6980 This option is experimental and does not currently guarantee to  
6981 disable all GCC optimizations that affect signaling NaN behavior.

6983 @item -fsingle-precision-constant  
6984 @opindex fsingle-precision-constant  
6985 Treat floating point constant as single precision constant instead of  
6986 implicitly converting it to double precision constant.

6988 @item -fcx-limited-range  
6989 @opindex fcx-limited-range  
6990 When enabled, this option states that a range reduction step is not  
6991 needed when performing complex division. Also, there is no checking  
6992 whether the result of a complex multiplication or division is `@code{NaN}`  
6993 + `I*NaN`, with an attempt to rescue the situation in that case. The  
6994 default is `@option{-fno-cx-limited-range}`, but is enabled by  
6995 `@option{-ffast-math}`.

6997 This option controls the default setting of the ISO C99  
6998 `@code{CX_LIMITED_RANGE}` pragma. Nevertheless, the option applies to  
6999 all languages.

7001 @item -fcx-fortran-rules  
7002 @opindex fcx-fortran-rules  
7003 Complex multiplication and division follow Fortran rules. Range  
7004 reduction is done as part of complex division, but there is no checking  
7005 whether the result of a complex multiplication or division is `@code{NaN}`  
7006 + `I*NaN`, with an attempt to rescue the situation in that case.

7008 The default is `@option{-fno-cx-fortran-rules}`.

7010 @end table

7012 The following options control optimizations that may improve  
7013 performance, but are not enabled by any `@option{-O}` options. This  
7014 section includes experimental options that may produce broken code.

7016 @table @gcctabopt  
7017 @item -fbranch-probabilities  
7018 @opindex fbranch-probabilities  
7019 After running a program compiled with `@option{-fprofile-arcs}`  
7020 ([@pxref{Debugging Options,, Options for Debugging Your Program or](#)

7021 @command{gcc}}, you can compile it a second time using  
 7022 @option{-fbranch-probabilities}, to improve optimizations based on  
 7023 the number of times each branch was taken. When the program  
 7024 compiled with @option{-fprofile-arcs} exits it saves arc execution  
 7025 counts to a file called @file{@var{sourcename}.gcda} for each source  
 7026 file. The information in this data file is very dependent on the  
 7027 structure of the generated code, so you must use the same source code  
 7028 and the same optimization options for both compilations.

7030 With @option{-fbranch-probabilities}, GCC puts a  
 7031 @samp{REG\_BR\_PROB} note on each @samp{JUMP\_INSN} and @samp{CALL\_INSN}.  
 7032 These can be used to improve optimization. Currently, they are only  
 7033 used in one place: in @file{reorg.c}, instead of guessing which path a  
 7034 branch is mostly to take, the @samp{REG\_BR\_PROB} values are used to  
 7035 exactly determine which path is taken more often.

7037 @item -fprofile-values  
 7038 @opindex fprofile-values  
 7039 If combined with @option{-fprofile-arcs}, it adds code so that some  
 7040 data about values of expressions in the program is gathered.

7042 With @option{-fbranch-probabilities}, it reads back the data gathered  
 7043 from profiling values of expressions and adds @samp{REG\_VALUE\_PROFILE}  
 7044 notes to instructions for their later usage in optimizations.

7046 Enabled with @option{-fprofile-generate} and @option{-fprofile-use}.

7048 @item -fvpt  
 7049 @opindex fvpt  
 7050 If combined with @option{-fprofile-arcs}, it instructs the compiler to add  
 7051 a code to gather information about values of expressions.

7053 With @option{-fbranch-probabilities}, it reads back the data gathered  
 7054 and actually performs the optimizations based on them.  
 7055 Currently the optimizations include specialization of division operation  
 7056 using the knowledge about the value of the denominator.

7058 @item -frename-registers  
 7059 @opindex frename-registers  
 7060 Attempt to avoid false dependencies in scheduled code by making use  
 7061 of registers left over after register allocation. This optimization  
 7062 will most benefit processors with lots of registers. Depending on the  
 7063 debug information format adopted by the target, however, it can  
 7064 make debugging impossible, since variables will no longer stay in  
 7065 a ``home register''.

7067 Enabled by default with @option{-funroll-loops}.

7069 @item -ftracer  
 7070 @opindex ftracer  
 7071 Perform tail duplication to enlarge superblock size. This transformation  
 7072 simplifies the control flow of the function allowing other optimizations to do  
 7073 better job.

7075 Enabled with @option{-fprofile-use}.

7077 @item -funroll-loops  
 7078 @opindex funroll-loops  
 7079 Unroll loops whose number of iterations can be determined at compile time or  
 7080 upon entry to the loop. @option{-funroll-loops} implies  
 7081 @option{-frerun-cse-after-loop}, @option{-fweb} and @option{-frename-registers}.  
 7082 It also turns on complete loop peeling (i.e.@: complete removal of loops with  
 7083 small constant number of iterations). This option makes code larger, and may  
 7084 or may not make it run faster.

7086 Enabled with @option{-fprofile-use}.

7088 @item -funroll-all-loops  
 7089 @opindex funroll-all-loops  
 7090 Unroll all loops, even if their number of iterations is uncertain when  
 7091 the loop is entered. This usually makes programs run more slowly.  
 7092 @option{-funroll-all-loops} implies the same options as  
 7093 @option{-funroll-loops}.

7095 @item -fpeel-loops  
 7096 @opindex fpeel-loops  
 7097 Peels the loops for that there is enough information that they do not  
 7098 roll much (from profile feedback). It also turns on complete loop peeling  
 7099 (i.e.@: complete removal of loops with small constant number of iterations).

7101 Enabled with @option{-fprofile-use}.

7103 @item -fmove-loop-invariants  
 7104 @opindex fmove-loop-invariants  
 7105 Enables the loop invariant motion pass in the RTL loop optimizer. Enabled  
 7106 at level @option{-O1}

7108 @item -funswitch-loops  
 7109 @opindex funswitch-loops  
 7110 Move branches with loop invariant conditions out of the loop, with duplicates  
 7111 of the loop on both branches (modified according to result of the condition).

7113 @item -ffunction-sections  
 7114 @itemx -fdata-sections  
 7115 @opindex ffunction-sections  
 7116 @opindex fdata-sections  
 7117 Place each function or data item into its own section in the output  
 7118 file if the target supports arbitrary sections. The name of the  
 7119 function or the name of the data item determines the section's name  
 7120 in the output file.

7122 Use these options on systems where the linker can perform optimizations  
 7123 to improve locality of reference in the instruction space. Most systems  
 7124 using the ELF object format and SPARC processors running Solaris 2 have  
 7125 linkers with such optimizations. AIX may have these optimizations in  
 7126 the future.

7128 Only use these options when there are significant benefits from doing  
 7129 so. When you specify these options, the assembler and linker will  
 7130 create larger object and executable files and will also be slower.  
 7131 You will not be able to use @code{gprof} on all systems if you  
 7132 specify this option and you may have problems with debugging if  
 7133 you specify both this option and @option{-g}.

7135 @item -fbranch-target-load-optimize  
 7136 @opindex fbranch-target-load-optimize  
 7137 Perform branch target register load optimization before prologue / epilogue  
 7138 threading.  
 7139 The use of target registers can typically be exposed only during reload,  
 7140 thus hoisting loads out of loops and doing inter-block scheduling needs  
 7141 a separate optimization pass.

7143 @item -fbranch-target-load-optimize2  
 7144 @opindex fbranch-target-load-optimize2  
 7145 Perform branch target register load optimization after prologue / epilogue  
 7146 threading.

7148 @item -fbtr-bb-exclusive  
 7149 @opindex fbtr-bb-exclusive  
 7150 When performing branch target register load optimization, don't reuse  
 7151 branch target registers in within any basic block.

7153 @item -fstack-protector  
 7154 @opindex fstack-protector  
 7155 Emit extra code to check for buffer overflows, such as stack smashing  
 7156 attacks. This is done by adding a guard variable to functions with  
 7157 vulnerable objects. This includes functions that call `alloca`, and  
 7158 functions with buffers larger than 8 bytes. The guards are initialized  
 7159 when a function is entered and then checked when the function exits.  
 7160 If a guard check fails, an error message is printed and the program exits.

7162 @item -fstack-protector-all  
 7163 @opindex fstack-protector-all  
 7164 Like @option{-fstack-protector} except that all functions are protected.

7166 @item -fsection-anchors  
 7167 @opindex fsection-anchors  
 7168 Try to reduce the number of symbolic address calculations by using  
 7169 shared ‘‘anchor’’ symbols to address nearby objects. This transformation  
 7170 can help to reduce the number of GOT entries and GOT accesses on some  
 7171 targets.

7173 For example, the implementation of the following function @code{foo}:

```
7175 @smallexample
7176 static int a, b, c;
7177 int foo (void) @{ return a + b + c; @}
7178 @end smallexample
```

7180 would usually calculate the addresses of all three variables, but if you  
 7181 compile it with @option{-fsection-anchors}, it will access the variables  
 7182 from a common anchor point instead. The effect is similar to the  
 7183 following pseudocode (which isn't valid C):

```
7185 @smallexample
7186 int foo (void)
7187 @{
7188     register int *xr = &x;
7189     return xr[&a - &x] + xr[&b - &x] + xr[&c - &x];
7190 @}
7191 @end smallexample
```

7193 Not all targets support this option.

7195 @item --param @var{name}=@var{value}  
 7196 @opindex param  
 7197 In some places, GCC uses various constants to control the amount of  
 7198 optimization that is done. For example, GCC will not inline functions  
 7199 that contain more than a certain number of instructions. You can  
 7200 control some of these constants on the command-line using the  
 7201 @option{--param} option.

7203 The names of specific parameters, and the meaning of the values, are  
 7204 tied to the internals of the compiler, and are subject to change  
 7205 without notice in future releases.

7207 In each case, the @var{value} is an integer. The allowable choices for  
 7208 @var{name} are given in the following table:

7210 @table @gcctabopt  
 7211 @item sra-max-structure-size  
 7212 The maximum structure size, in bytes, at which the scalar replacement  
 7213 of aggregates (SRA) optimization will perform block copies. The  
 7214 default value, 0, implies that GCC will select the most appropriate  
 7215 size itself.

7217 @item sra-field-structure-ratio  
 7218 The threshold ratio (as a percentage) between instantiated fields and

7219 the complete structure size. We say that if the ratio of the number  
 7220 of bytes in instantiated fields to the number of bytes in the complete  
 7221 structure exceeds this parameter, then block copies are not used. The  
 7222 default is 75.

7224 @item struct-reorg-cold-struct-ratio  
 7225 The threshold ratio (as a percentage) between a structure frequency  
 7226 and the frequency of the hottest structure in the program. This parameter  
 7227 is used by struct-reorg optimization enabled by @option{-fipa-struct-reorg}.  
 7228 We say that if the ratio of a structure frequency, calculated by profiling,  
 7229 to the hottest structure frequency in the program is less than this  
 7230 parameter, then structure reorganization is not applied to this structure.  
 7231 The default is 10.

7233 @item predictable-branch-cost-outcome  
 7234 When branch is predicted to be taken with probability lower than this threshold  
 7235 (in percent), then it is considered well predictable. The default is 10.

7237 @item max-crossjump-edges  
 7238 The maximum number of incoming edges to consider for crossjumping.  
 7239 The algorithm used by @option{-fcrossjumping} is  $O(N^2)$  in  
 7240 the number of edges incoming to each block. Increasing values mean  
 7241 more aggressive optimization, making the compile time increase with  
 7242 probably small improvement in executable size.

7244 @item min-crossjump-insns  
 7245 The minimum number of instructions which must be matched at the end  
 7246 of two blocks before crossjumping will be performed on them. This  
 7247 value is ignored in the case where all instructions in the block being  
 7248 crossjumped from are matched. The default value is 5.

7250 @item max-grow-copy-bb-insns  
 7251 The maximum code size expansion factor when copying basic blocks  
 7252 instead of jumping. The expansion is relative to a jump instruction.  
 7253 The default value is 8.

7255 @item max-goto-duplication-insns  
 7256 The maximum number of instructions to duplicate to a block that jumps  
 7257 to a computed goto. To avoid  $O(N^2)$  behavior in a number of  
 7258 passes, GCC factors computed gotos early in the compilation process,  
 7259 and un-factors them as late as possible. Only computed jumps at the  
 7260 end of a basic blocks with no more than max-goto-duplication-insns are  
 7261 un-factored. The default value is 8.

7263 @item max-delay-slot-insn-search  
 7264 The maximum number of instructions to consider when looking for an  
 7265 instruction to fill a delay slot. If more than this arbitrary number of  
 7266 instructions is searched, the time savings from filling the delay slot  
 7267 will be minimal so stop searching. Increasing values mean more  
 7268 aggressive optimization, making the compile time increase with probably  
 7269 small improvement in executable run time.

7271 @item max-delay-slot-live-search  
 7272 When trying to fill delay slots, the maximum number of instructions to  
 7273 consider when searching for a block with valid live register  
 7274 information. Increasing this arbitrarily chosen value means more  
 7275 aggressive optimization, increasing the compile time. This parameter  
 7276 should be removed when the delay slot code is rewritten to maintain the  
 7277 control-flow graph.

7279 @item max-gcse-memory  
 7280 The approximate maximum amount of memory that will be allocated in  
 7281 order to perform the global common subexpression elimination  
 7282 optimization. If more memory than specified is required, the  
 7283 optimization will not be done.



7285 @item max-gcse-passes  
7286 The maximum number of passes of GCSE to run. The default is 1.

7288 @item max-pending-list-length  
7289 The maximum number of pending dependencies scheduling will allow  
7290 before flushing the current state and starting over. Large functions  
7291 with few branches or calls can create excessively large lists which  
7292 needlessly consume memory and resources.

7294 @item max-inline-insns-single  
7295 Several parameters control the tree inliner used in gcc.  
7296 This number sets the maximum number of instructions (counted in GCC's  
7297 internal representation) in a single function that the tree inliner  
7298 will consider for inlining. This only affects functions declared  
7299 inline and methods implemented in a class declaration (C++).  
7300 The default value is 450.

7302 @item max-inline-insns-auto  
7303 When you use @option{-finline-functions} (included in @option{-O3}),  
7304 a lot of functions that would otherwise not be considered for inlining  
7305 by the compiler will be investigated. To those functions, a different  
7306 (more restrictive) limit compared to functions declared inline can  
7307 be applied.  
7308 The default value is 90.

7310 @item large-function-insns  
7311 The limit specifying really large functions. For functions larger than this  
7312 limit after inlining, inlining is constrained by  
7313 @option{--param large-function-growth}. This parameter is useful primarily  
7314 to avoid extreme compilation time caused by non-linear algorithms used by the  
7315 backend.  
7316 The default value is 2700.

7318 @item large-function-growth  
7319 Specifies maximal growth of large function caused by inlining in percents.  
7320 The default value is 100 which limits large function growth to 2.0 times  
7321 the original size.

7323 @item large-unit-insns  
7324 The limit specifying large translation unit. Growth caused by inlining of  
7325 units larger than this limit is limited by @option{--param inline-unit-growth}.  
7326 For small units this might be too tight (consider unit consisting of function A  
7327 that is inline and B that just calls A three time. If B is small relative to  
7328 A, the growth of unit is 300% and yet such inlining is very sane. For very  
7329 large units consisting of small inlineable functions however the overall unit  
7330 growth limit is needed to avoid exponential explosion of code size. Thus for  
7331 smaller units, the size is increased to @option{--param large-unit-insns}  
7332 before applying @option{--param inline-unit-growth}. The default is 10000

7334 @item inline-unit-growth  
7335 Specifies maximal overall growth of the compilation unit caused by inlining.  
7336 The default value is 30 which limits unit growth to 1.3 times the original  
7337 size.

7339 @item icpc-unit-growth  
7340 Specifies maximal overall growth of the compilation unit caused by  
7341 interprocedural constant propagation. The default value is 10 which limits  
7342 unit growth to 1.1 times the original size.

7344 @item large-stack-frame  
7345 The limit specifying large stack frames. While inlining the algorithm is trying  
7346 to not grow past this limit too much. Default value is 256 bytes.

7348 @item large-stack-frame-growth  
7349 Specifies maximal growth of large stack frames caused by inlining in percents.  
7350 The default value is 1000 which limits large stack frame growth to 11 times

7351 the original size.

7353 @item max-inline-insns-recursive  
7354 @itemx max-inline-insns-recursive-auto  
7355 Specifies maximum number of instructions out-of-line copy of self recursive inli  
7356 function can grow into by performing recursive inlining.

7358 For functions declared inline @option{--param max-inline-insns-recursive} is  
7359 taken into account. For function not declared inline, recursive inlining  
7360 happens only when @option{-finline-functions} (included in @option{-O3}) is  
7361 enabled and @option{--param max-inline-insns-recursive-auto} is used. The  
7362 default value is 450.

7364 @item max-inline-recursive-depth  
7365 @itemx max-inline-recursive-depth-auto  
7366 Specifies maximum recursion depth used by the recursive inlining.

7368 For functions declared inline @option{--param max-inline-recursive-depth} is  
7369 taken into account. For function not declared inline, recursive inlining  
7370 happens only when @option{-finline-functions} (included in @option{-O3}) is  
7371 enabled and @option{--param max-inline-recursive-depth-auto} is used. The  
7372 default value is 8.

7374 @item min-inline-recursive-probability  
7375 Recursive inlining is profitable only for function having deep recursion  
7376 in average and can hurt for function having little recursion depth by  
7377 increasing the prologue size or complexity of function body to other  
7378 optimizers.

7380 When profile feedback is available (see @option{-fprofile-generate}) the actual  
7381 recursion depth can be guessed from probability that function will recurse via  
7382 given call expression. This parameter limits inlining only to call expression  
7383 whose probability exceeds given threshold (in percents). The default value is  
7384 10.

7386 @item inline-call-cost  
7387 Specify cost of call instruction relative to simple arithmetics operations  
7388 (having cost of 1). Increasing this cost disqualifies inlining of non-leaf  
7389 functions and at the same time increases size of leaf function that is believed  
7390 reduce function size by being inlined. In effect it increases amount of  
7391 inlining for code having large abstraction penalty (many functions that just  
7392 pass the arguments to other functions) and decrease inlining for code with low  
7393 abstraction penalty. The default value is 12.

7395 @item min-vect-loop-bound  
7396 The minimum number of iterations under which a loop will not get vectorized  
7397 when @option{-ftree-vectorize} is used. The number of iterations after  
7398 vectorization needs to be greater than the value specified by this option  
7399 to allow vectorization. The default value is 0.

7401 @item max-unrolled-insns  
7402 The maximum number of instructions that a loop should have if that loop  
7403 is unrolled, and if the loop is unrolled, it determines how many times  
7404 the loop code is unrolled.

7406 @item max-average-unrolled-insns  
7407 The maximum number of instructions biased by probabilities of their execution  
7408 that a loop should have if that loop is unrolled, and if the loop is unrolled,  
7409 it determines how many times the loop code is unrolled.

7411 @item max-unroll-times  
7412 The maximum number of unrollings of a single loop.

7414 @item max-peeled-insns  
7415 The maximum number of instructions that a loop should have if that loop  
7416 is peeled, and if the loop is peeled, it determines how many times

7417 the loop code is peeled.

7419 @item max-peel-times  
7420 The maximum number of peelings of a single loop.

7422 @item max-completely-peeled-insns  
7423 The maximum number of insns of a completely peeled loop.

7425 @item max-completely-peel-times  
7426 The maximum number of iterations of a loop to be suitable for complete peeling.

7428 @item max-completely-peel-loop-nest-depth  
7429 The maximum depth of a loop nest suitable for complete peeling.

7431 @item max-unswitch-insns  
7432 The maximum number of insns of an unswitched loop.

7434 @item max-unswitch-level  
7435 The maximum number of branches unswitched in a single loop.

7437 @item lim-expensive  
7438 The minimum cost of an expensive expression in the loop invariant motion.

7440 @item iv-consider-all-candidates-bound  
7441 Bound on number of candidates for induction variables below that  
7442 all candidates are considered for each use in induction variable  
7443 optimizations. Only the most relevant candidates are considered  
7444 if there are more candidates, to avoid quadratic time complexity.

7446 @item iv-max-considered-uses  
7447 The induction variable optimizations give up on loops that contain more  
7448 induction variable uses.

7450 @item iv-always-prune-cand-set-bound  
7451 If number of candidates in the set is smaller than this value,  
7452 we always try to remove unnecessary ivs from the set during its  
7453 optimization when a new iv is added to the set.

7455 @item scev-max-expr-size  
7456 Bound on size of expressions used in the scalar evolutions analyzer.  
7457 Large expressions slow the analyzer.

7459 @item omega-max-vars  
7460 The maximum number of variables in an Omega constraint system.  
7461 The default value is 128.

7463 @item omega-max-geqs  
7464 The maximum number of inequalities in an Omega constraint system.  
7465 The default value is 256.

7467 @item omega-max-egs  
7468 The maximum number of equalities in an Omega constraint system.  
7469 The default value is 128.

7471 @item omega-max-wild-cards  
7472 The maximum number of wildcard variables that the Omega solver will  
7473 be able to insert. The default value is 18.

7475 @item omega-hash-table-size  
7476 The size of the hash table in the Omega solver. The default value is  
7477 550.

7479 @item omega-max-keys  
7480 The maximal number of keys used by the Omega solver. The default  
7481 value is 500.

7483 @item omega-eliminate-redundant-constraints  
7484 When set to 1, use expensive methods to eliminate all redundant  
7485 constraints. The default value is 0.

7487 @item vect-max-version-for-alignment-checks  
7488 The maximum number of runtime checks that can be performed when  
7489 doing loop versioning for alignment in the vectorizer. See option  
7490 `ftree-vect-loop-version` for more information.

7492 @item vect-max-version-for-alias-checks  
7493 The maximum number of runtime checks that can be performed when  
7494 doing loop versioning for alias in the vectorizer. See option  
7495 `ftree-vect-loop-version` for more information.

7497 @item max-iterations-to-track  
7499 The maximum number of iterations of a loop the brute force algorithm  
7500 for analysis of # of iterations of the loop tries to evaluate.

7502 @item hot-bb-count-fraction  
7503 Select fraction of the maximal count of repetitions of basic block in program  
7504 given basic block needs to have to be considered hot.

7506 @item hot-bb-frequency-fraction  
7507 Select fraction of the maximal frequency of executions of basic block in  
7508 function given basic block needs to have to be considered hot

7510 @item max-predicted-iterations  
7511 The maximum number of loop iterations we predict statically. This is useful  
7512 in cases where function contain single loop with known bound and other loop  
7513 with unknown. We predict the known number of iterations correctly, while  
7514 the unknown number of iterations average to roughly 10. This means that the  
7515 loop without bounds would appear artificially cold relative to the other one.

7517 @item align-threshold  
7519 Select fraction of the maximal frequency of executions of basic block in  
7520 function given basic block will get aligned.

7522 @item align-loop-iterations  
7524 A loop expected to iterate at least the selected number of iterations will get  
7525 aligned.

7527 @item tracer-dynamic-coverage  
7528 @itemx tracer-dynamic-coverage-feedback  
7530 This value is used to limit superblock formation once the given percentage of  
7531 executed instructions is covered. This limits unnecessary code size  
7532 expansion.

7534 The @option{tracer-dynamic-coverage-feedback} is used only when profile  
7535 feedback is available. The real profiles (as opposed to statically estimated  
7536 ones) are much less balanced allowing the threshold to be larger value.

7538 @item tracer-max-code-growth  
7539 Stop tail duplication once code growth has reached given percentage. This is  
7540 rather hokey argument, as most of the duplicates will be eliminated later in  
7541 cross jumping, so it may be set to much higher values than is the desired code  
7542 growth.

7544 @item tracer-min-branch-ratio  
7546 Stop reverse growth when the reverse probability of best edge is less than this  
7547 threshold (in percent).

7549 @item tracer-min-branch-ratio  
7550 @itemx tracer-min-branch-ratio-feedback

7552 Stop forward growth if the best edge do have probability lower than this  
7553 threshold.

7555 Similarly to @option{tracer-dynamic-coverage} two values are present, one for  
7556 compilation for profile feedback and one for compilation without. The value  
7557 for compilation with profile feedback needs to be more conservative (higher) in  
7558 order to make tracer effective.

7560 @item max-cse-path-length

7562 Maximum number of basic blocks on path that cse considers. The default is 10.

7564 @item max-cse-insns  
7565 The maximum instructions CSE process before flushing. The default is 1000.

7567 @item max-aliased-vops

7569 Maximum number of virtual operands per function allowed to represent  
7570 aliases before triggering the alias partitioning heuristic. Alias  
7571 partitioning reduces compile times and memory consumption needed for  
7572 aliasing at the expense of precision loss in alias information. The  
7573 default value for this parameter is 100 for -O1, 500 for -O2 and 1000  
7574 for -O3.

7576 Notice that if a function contains more memory statements than the  
7577 value of this parameter, it is not really possible to achieve this  
7578 reduction. In this case, the compiler will use the number of memory  
7579 statements as the value for @option{max-aliased-vops}.

7581 @item avg-aliased-vops

7583 Average number of virtual operands per statement allowed to represent  
7584 aliases before triggering the alias partitioning heuristic. This  
7585 works in conjunction with @option{max-aliased-vops}. If a function  
7586 contains more than @option{max-aliased-vops} virtual operators, then  
7587 memory symbols will be grouped into memory partitions until either the  
7588 total number of virtual operators is below @option{max-aliased-vops}  
7589 or the average number of virtual operators per memory statement is  
7590 below @option{avg-aliased-vops}. The default value for this parameter  
7591 is 1 for -O1 and -O2, and 3 for -O3.

7593 @item gcc-min-expand

7595 GCC uses a garbage collector to manage its own memory allocation. This  
7596 parameter specifies the minimum percentage by which the garbage  
7597 collector's heap should be allowed to expand between collections.  
7598 Tuning this may improve compilation speed; it has no effect on code  
7599 generation.

7601 The default is  $30\% + 70\% * (\text{RAM}/1\text{GB})$  with an upper bound of 100% when  
7602  $\text{RAM} \geq 1\text{GB}$ . If @code{getrlimit} is available, the notion of "RAM" is  
7603 the smallest of actual RAM and @code{RLIMIT\_DATA} or @code{RLIMIT\_AS}. If  
7604 GCC is not able to calculate RAM on a particular platform, the lower  
7605 bound of 30% is used. Setting this parameter and  
7606 @option{gcc-min-heapsize} to zero causes a full collection to occur at  
7607 every opportunity. This is extremely slow, but can be useful for  
7608 debugging.

7610 @item gcc-min-heapsize

7612 Minimum size of the garbage collector's heap before it begins bothering  
7613 to collect garbage. The first collection occurs after the heap expands  
7614 by @option{gcc-min-expand}% beyond @option{gcc-min-heapsize}. Again,

7615 tuning this may improve compilation speed, and has no effect on code  
7616 generation.

7618 The default is the smaller of  $\text{RAM}/8$ , `RLIMIT_RSS`, or a limit which  
7619 tries to ensure that `RLIMIT_DATA` or `RLIMIT_AS` are not exceeded, but  
7620 with a lower bound of 4096 (four megabytes) and an upper bound of  
7621 131072 (128 megabytes). If GCC is not able to calculate RAM on a  
7622 particular platform, the lower bound is used. Setting this parameter  
7623 very large effectively disables garbage collection. Setting this  
7624 parameter and @option{gcc-min-expand} to zero causes a full collection  
7625 to occur at every opportunity.

7627 @item max-reload-search-insns  
7628 The maximum number of instruction reload should look backward for equivalent  
7629 register. Increasing values mean more aggressive optimization, making the  
7630 compile time increase with probably slightly better performance. The default  
7631 value is 100.

7633 @item max-cselib-memory-locations  
7634 The maximum number of memory locations cselib should take into account.  
7635 Increasing values mean more aggressive optimization, making the compile time  
7636 increase with probably slightly better performance. The default value is 500.

7638 @item reorder-blocks-duplicate  
7639 @itemx reorder-blocks-duplicate-feedback

7641 Used by basic block reordering pass to decide whether to use unconditional  
7642 branch or duplicate the code on its destination. Code is duplicated when its  
7643 estimated size is smaller than this value multiplied by the estimated size of  
7644 unconditional jump in the hot spots of the program.

7646 The @option{reorder-block-duplicate-feedback} is used only when profile  
7647 feedback is available and may be set to higher values than  
7648 @option{reorder-block-duplicate} since information about the hot spots is more  
7649 accurate.

7651 @item max-sched-ready-insns  
7652 The maximum number of instructions ready to be issued the scheduler should  
7653 consider at any given time during the first scheduling pass. Increasing  
7654 values mean more thorough searches, making the compilation time increase  
7655 with probably little benefit. The default value is 100.

7657 @item max-sched-region-blocks  
7658 The maximum number of blocks in a region to be considered for  
7659 interblock scheduling. The default value is 10.

7661 @item max-pipeline-region-blocks  
7662 The maximum number of blocks in a region to be considered for  
7663 pipelining in the selective scheduler. The default value is 15.

7665 @item max-sched-region-insns  
7666 The maximum number of insns in a region to be considered for  
7667 interblock scheduling. The default value is 100.

7669 @item max-pipeline-region-insns  
7670 The maximum number of insns in a region to be considered for  
7671 pipelining in the selective scheduler. The default value is 200.

7673 @item min-spec-prob  
7674 The minimum probability (in percents) of reaching a source block  
7675 for interblock speculative scheduling. The default value is 40.

7677 @item max-sched-extend-regions-iters  
7678 The maximum number of iterations through CFG to extend regions.  
7679 0 - disable region extension,  
7680 N - do at most N iterations.

7681 The default value is 0.

7683 @item max-sched-insn-conflict-delay  
7684 The maximum conflict delay for an insn to be considered for speculative motion.  
7685 The default value is 3.

7687 @item sched-spec-prob-cuttoff  
7688 The minimal probability of speculation success (in percents), so that  
7689 speculative insn will be scheduled.  
7690 The default value is 40.

7692 @item sched-mem-true-dep-cost  
7693 Minimal distance (in CPU cycles) between store and load targeting same  
7694 memory locations. The default value is 1.

7696 @item selsched-max-lookahead  
7697 The maximum size of the lookahead window of selective scheduling. It is a  
7698 depth of search for available instructions.  
7699 The default value is 50.

7701 @item selsched-max-sched-times  
7702 The maximum number of times that an instruction will be scheduled during  
7703 selective scheduling. This is the limit on the number of iterations  
7704 through which the instruction may be pipelined. The default value is 2.

7706 @item selsched-max-insns-to-rename  
7707 The maximum number of best instructions in the ready list that are considered  
7708 for renaming in the selective scheduler. The default value is 2.

7710 @item max-last-value-rtl  
7711 The maximum size measured as number of RTLs that can be recorded in an expressio  
7712 in combiner for a pseudo register as last known value of that register. The def  
7713 is 10000.

7715 @item integer-share-limit  
7716 Small integer constants can use a shared data structure, reducing the  
7717 compiler's memory usage and increasing its speed. This sets the maximum  
7718 value of a shared integer constant. The default value is 256.

7720 @item min-virtual-mappings  
7721 Specifies the minimum number of virtual mappings in the incremental  
7722 SSA updater that should be registered to trigger the virtual mappings  
7723 heuristic defined by virtual-mappings-ratio. The default value is  
7724 100.

7726 @item virtual-mappings-ratio  
7727 If the number of virtual mappings is virtual-mappings-ratio bigger  
7728 than the number of virtual symbols to be updated, then the incremental  
7729 SSA updater switches to a full update for those symbols. The default  
7730 ratio is 3.

7732 @item ssp-buffer-size  
7733 The minimum size of buffers (i.e.@: arrays) that will receive stack smashing  
7734 protection when @option{-fstack-protection} is used.

7736 @item max-jump-thread-duplication-stmts  
7737 Maximum number of statements allowed in a block that needs to be  
7738 duplicated when threading jumps.

7740 @item max-fields-for-field-sensitive  
7741 Maximum number of fields in a structure we will treat in  
7742 a field sensitive manner during pointer analysis. The default is zero  
7743 for -O0, and -O1 and 100 for -Os, -O2, and -O3.

7745 @item prefetch-latency  
7746 Estimate on average number of instructions that are executed before

7747 prefetch finishes. The distance we prefetch ahead is proportional  
7748 to this constant. Increasing this number may also lead to less  
7749 streams being prefetched (see @option{simultaneous-prefetches}).

7751 @item simultaneous-prefetches  
7752 Maximum number of prefetches that can run at the same time.

7754 @item l1-cache-line-size  
7755 The size of cache line in L1 cache, in bytes.

7757 @item l1-cache-size  
7758 The size of L1 cache, in kilobytes.

7760 @item l2-cache-size  
7761 The size of L2 cache, in kilobytes.

7763 @item use-canonical-types  
7764 Whether the compiler should use the "canonical" type system. By  
7765 default, this should always be 1, which uses a more efficient internal  
7766 mechanism for comparing types in C++ and Objective-C++. However, if  
7767 bugs in the canonical type system are causing compilation failures,  
7768 set this value to 0 to disable canonical types.

7770 @item switch-conversion-max-branch-ratio  
7771 Switch initialization conversion will refuse to create arrays that are  
7772 bigger than @option{switch-conversion-max-branch-ratio} times the number of  
7773 branches in the switch.

7775 @item max-partial-antic-length  
7776 Maximum length of the partial antic set computed during the tree  
7777 partial redundancy elimination optimization (@option{-ftree-pre}) when  
7778 optimizing at @option{-O3} and above. For some sorts of source code  
7779 the enhanced partial redundancy elimination optimization can run away,  
7780 consuming all of the memory available on the host machine. This  
7781 parameter sets a limit on the length of the sets that are computed,  
7782 which prevents the runaway behavior. Setting a value of 0 for  
7783 this parameter will allow an unlimited set length.

7785 @item sccvn-max-scc-size  
7786 Maximum size of a strongly connected component (SCC) during SCCVN  
7787 processing. If this limit is hit, SCCVN processing for the whole  
7788 function will not be done and optimizations depending on it will  
7789 be disabled. The default maximum SCC size is 10000.

7791 @item ira-max-loops-num  
7792 IRA uses a regional register allocation by default. If a function  
7793 contains loops more than number given by the parameter, only at most  
7794 given number of the most frequently executed loops will form regions  
7795 for the regional register allocation. The default value of the  
7796 parameter is 100.

7798 @item ira-max-conflict-table-size  
7799 Although IRA uses a sophisticated algorithm of compression conflict  
7800 table, the table can be still big for huge functions. If the conflict  
7801 table for a function could be more than size in MB given by the  
7802 parameter, the conflict table is not built and faster, simpler, and  
7803 lower quality register allocation algorithm will be used. The  
7804 algorithm do not use pseudo-register conflicts. The default value of  
7805 the parameter is 2000.

7807 @item loop-invariant-max-bbs-in-loop  
7808 Loop invariant motion can be very expensive, both in compile time and  
7809 in amount of needed compile time memory, with very large loops. Loops  
7810 with more basic blocks than this parameter won't have loop invariant  
7811 motion optimization performed on them. The default value of the  
7812 parameter is 1000 for -O1 and 10000 for -O2 and above.

```

7814 @end table
7815 @end table

7817 @node Preprocessor Options
7818 @section Options Controlling the Preprocessor
7819 @cindex preprocessor options
7820 @cindex options, preprocessor

7822 These options control the C preprocessor, which is run on each C source
7823 file before actual compilation.

7825 If you use the @option{-E} option, nothing is done except preprocessing.
7826 Some of these options make sense only together with @option{-E} because
7827 they cause the preprocessor output to be unsuitable for actual
7828 compilation.

7830 @table @gcctabopt
7831 @item -Wp,@var{option}
7832 @opindex Wp
7833 You can use @option{-Wp,@var{option}} to bypass the compiler driver
7834 and pass @var{option} directly through to the preprocessor. If
7835 @var{option} contains commas, it is split into multiple options at the
7836 commas. However, many options are modified, translated or interpreted
7837 by the compiler driver before being passed to the preprocessor, and
7838 @option{-Wp} forcibly bypasses this phase. The preprocessor's direct
7839 interface is undocumented and subject to change, so whenever possible
7840 you should avoid using @option{-Wp} and let the driver handle the
7841 options instead.

7843 @item -Xpreprocessor @var{option}
7844 @opindex Xpreprocessor
7845 Pass @var{option} as an option to the preprocessor. You can use this to
7846 supply system-specific preprocessor options which GCC does not know how to
7847 recognize.

7849 If you want to pass an option that takes an argument, you must use
7850 @option{-Xpreprocessor} twice, once for the option and once for the argument.
7851 @end table

7853 @include cppopts.texi

7855 @node Assembler Options
7856 @section Passing Options to the Assembler

7858 @c prevent bad page break with this line
7859 You can pass options to the assembler.

7861 @table @gcctabopt
7862 @item -Wa,@var{option}
7863 @opindex Wa
7864 Pass @var{option} as an option to the assembler. If @var{option}
7865 contains commas, it is split into multiple options at the commas.

7867 @item -Xassembler @var{option}
7868 @opindex Xassembler
7869 Pass @var{option} as an option to the assembler. You can use this to
7870 supply system-specific assembler options which GCC does not know how to
7871 recognize.

7873 If you want to pass an option that takes an argument, you must use
7874 @option{-Xassembler} twice, once for the option and once for the argument.

7876 @end table

7878 @node Link Options

```

```

7879 @section Options for Linking
7880 @cindex link options
7881 @cindex options, linking

7883 These options come into play when the compiler links object files into
7884 an executable output file. They are meaningless if the compiler is
7885 not doing a link step.

7887 @table @gcctabopt
7888 @cindex file names
7889 @item @var{object-file-name}
7890 A file name that does not end in a special recognized suffix is
7891 considered to name an object file or library. (Object files are
7892 distinguished from libraries by the linker according to the file
7893 contents.) If linking is done, these object files are used as input
7894 to the linker.

7896 @item -c
7897 @itemx -S
7898 @itemx -E
7899 @opindex c
7900 @opindex S
7901 @opindex E
7902 If any of these options is used, then the linker is not run, and
7903 object file names should not be used as arguments. @xref{Overall
7904 Options}.

7906 @cindex Libraries
7907 @item -l@var{library}
7908 @itemx -l @var{library}
7909 @opindex l
7910 Search the library named @var{library} when linking. (The second
7911 alternative with the library as a separate argument is only for
7912 POSIX compliance and is not recommended.)

7914 It makes a difference where in the command you write this option; the
7915 linker searches and processes libraries and object files in the order they
7916 are specified. Thus, @samp{foo.o -lz bar.o} searches library @samp{z}
7917 after file @file{foo.o} but before @file{bar.o}. If @file{bar.o} refers
7918 to functions in @samp{z}, those functions may not be loaded.

7920 The linker searches a standard list of directories for the library,
7921 which is actually a file named @file{lib@var{library}.a}. The linker
7922 then uses this file as if it had been specified precisely by name.

7924 The directories searched include several standard system directories
7925 plus any that you specify with @option{-L}.

7927 Normally the files found this way are library files---archive files
7928 whose members are object files. The linker handles an archive file by
7929 scanning through it for members which define symbols that have so far
7930 been referenced but not defined. But if the file that is found is an
7931 ordinary object file, it is linked in the usual fashion. The only
7932 difference between using an @option{-l} option and specifying a file name
7933 is that @option{-l} surrounds @var{library} with @samp{lib} and @samp{.a}
7934 and searches several directories.

7936 @item -lobjc
7937 @opindex lobjc
7938 You need this special case of the @option{-l} option in order to
7939 link an Objective-C or Objective-C++ program.

7941 @item -nostartfiles
7942 @opindex nostartfiles
7943 Do not use the standard system startup files when linking.
7944 The standard system libraries are used normally, unless @option{-nostdlib}

```

7945 or `@option{-nodefaultlibs}` is used.

7947 `@item -nodefaultlibs`  
 7948 `@opindex nodefaultlibs`  
 7949 Do not use the standard system libraries when linking.  
 7950 Only the libraries you specify will be passed to the linker.  
 7951 The standard startup files are used normally, unless `@option{-nostartfiles}`  
 7952 is used. The compiler may generate calls to `@code{memcpy}`,  
 7953 `@code{memset}`, `@code{memcpy}` and `@code{memmove}`.  
 7954 These entries are usually resolved by entries in  
 7955 `libc`. These entry points should be supplied through some other  
 7956 mechanism when this option is specified.

7958 `@item -nostdlib`  
 7959 `@opindex nostdlib`  
 7960 Do not use the standard system startup files or libraries when linking.  
 7961 No startup files and only the libraries you specify will be passed to  
 7962 the linker. The compiler may generate calls to `@code{memcpy}`, `@code{memset}`,  
 7963 `@code{memcpy}` and `@code{memmove}`.  
 7964 These entries are usually resolved by entries in  
 7965 `libc`. These entry points should be supplied through some other  
 7966 mechanism when this option is specified.

7968 `@cindex @option{-lgcc}`, use with `@option{-nostdlib}`  
 7969 `@cindex @option{-nostdlib}` and unresolved references  
 7970 `@cindex unresolved references` and `@option{-nostdlib}`  
 7971 `@cindex @option{-lgcc}`, use with `@option{-nodefaultlibs}`  
 7972 `@cindex @option{-nodefaultlibs}` and unresolved references  
 7973 `@cindex unresolved references` and `@option{-nodefaultlibs}`  
 7974 One of the standard libraries bypassed by `@option{-nostdlib}` and  
 7975 `@option{-nodefaultlibs}` is `@file{libgcc.a}`, a library of internal subroutines  
 7976 that GCC uses to overcome shortcomings of particular machines, or special  
 7977 needs for some languages.  
 7978 ([@xref{Interface,,Interfacing to GCC Output,gccint,GNU Compiler](#)  
 7979 [Collection \(GCC\) Internals}](#)),  
 7980 for more discussion of `@file{libgcc.a}`.)  
 7981 In most cases, you need `@file{libgcc.a}` even when you want to avoid  
 7982 other standard libraries. In other words, when you specify `@option{-nostdlib}`  
 7983 or `@option{-nodefaultlibs}` you should usually specify `@option{-lgcc}` as well.  
 7984 This ensures that you have no unresolved references to internal GCC  
 7985 library subroutines. (For example, `@samp{__main}`, used to ensure C++  
 7986 constructors will be called; `@pxref{Collect2,,@code{collect2}, gccint,`  
 7987 `GNU Compiler Collection (GCC) Internals}`.)

7989 `@item -pie`  
 7990 `@opindex pie`  
 7991 Produce a position independent executable on targets which support it.  
 7992 For predictable results, you must also specify the same set of options  
 7993 that were used to generate code (`@option{-fpie}`, `@option{-fPIE}`,  
 7994 or model suboptions) when you specify this option.

7996 `@item -rdynamic`  
 7997 `@opindex rdynamic`  
 7998 Pass the flag `@option{-export-dynamic}` to the ELF linker, on targets  
 7999 that support it. This instructs the linker to add all symbols, not  
 8000 only used ones, to the dynamic symbol table. This option is needed  
 8001 for some uses of `@code{dlopen}` or to allow obtaining backtraces  
 8002 from within a program.

8004 `@item -s`  
 8005 `@opindex s`  
 8006 Remove all symbol table and relocation information from the executable.

8008 `@item -static`  
 8009 `@opindex static`  
 8010 On systems that support dynamic linking, this prevents linking with the shared

8011 libraries. On other systems, this option has no effect.

8013 `@item -shared`  
 8014 `@opindex shared`  
 8015 Produce a shared object which can then be linked with other objects to  
 8016 form an executable. Not all systems support this option. For predictable  
 8017 results, you must also specify the same set of options that were used to  
 8018 generate code (`@option{-fpic}`, `@option{-fPIC}`, or model suboptions)  
 8019 when you specify this option.[@footnote{On some systems, @samp{gcc -shared}](#)  
 8020 [needs to build supplementary stub code for constructors to work. On](#)  
 8021 [multi-libbed systems, @samp{gcc -shared} must select the correct support](#)  
 8022 [libraries to link against. Failing to supply the correct flags may lead](#)  
 8023 [to subtle defects. Supplying them in cases where they are not necessary](#)  
 8024 [is innocuous.}](#)

8026 `@item -shared-libgcc`  
 8027 `@itemx -static-libgcc`  
 8028 `@opindex shared-libgcc`  
 8029 `@opindex static-libgcc`  
 8030 On systems that provide `@file{libgcc}` as a shared library, these options  
 8031 force the use of either the shared or static version respectively.  
 8032 If no shared version of `@file{libgcc}` was built when the compiler was  
 8033 configured, these options have no effect.

8035 There are several situations in which an application should use the  
 8036 shared `@file{libgcc}` instead of the static version. The most common  
 8037 of these is when the application wishes to throw and catch exceptions  
 8038 across different shared libraries. In that case, each of the libraries  
 8039 as well as the application itself should use the shared `@file{libgcc}`.

8041 Therefore, the G++ and GCJ drivers automatically add  
 8042 `@option{-shared-libgcc}` whenever you build a shared library or a main  
 8043 executable, because C++ and Java programs typically use exceptions, so  
 8044 this is the right thing to do.

8046 If, instead, you use the GCC driver to create shared libraries, you may  
 8047 find that they will not always be linked with the shared `@file{libgcc}`.  
 8048 If GCC finds, at its configuration time, that you have a non-GNU linker  
 8049 or a GNU linker that does not support option `@option{--eh-frame-hdr}`,  
 8050 it will link the shared version of `@file{libgcc}` into shared libraries  
 8051 by default. Otherwise, it will take advantage of the linker and optimize  
 8052 away the linking with the shared version of `@file{libgcc}`, linking with  
 8053 the static version of `libgcc` by default. This allows exceptions to  
 8054 propagate through such shared libraries, without incurring relocation  
 8055 costs at library load time.

8057 However, if a library or main executable is supposed to throw or catch  
 8058 exceptions, you must link it using the G++ or GCJ driver, as appropriate  
 8059 for the languages used in the program, or using the option  
 8060 `@option{-shared-libgcc}`, such that it is linked with the shared  
 8061 `@file{libgcc}`.

8063 `@item -symbolic`  
 8064 `@opindex symbolic`  
 8065 Bind references to global symbols when building a shared object. Warn  
 8066 about any unresolved references (unless overridden by the link editor  
 8067 option `@samp{-Xlinker -z -Xlinker defs}`). Only a few systems support  
 8068 this option.

8070 `@item -T @var{script}`  
 8071 `@opindex T`  
 8072 `@cindex linker script`  
 8073 Use `@var{script}` as the linker script. This option is supported by most  
 8074 systems using the GNU linker. On some targets, such as bare-board  
 8075 targets without an operating system, the `@option{-T}` option may be required  
 8076 when linking to avoid references to undefined symbols.

```

8078 @item -Xlinker @var{option}
8079 @opindex Xlinker
8080 Pass @var{option} as an option to the linker. You can use this to
8081 supply system-specific linker options which GCC does not know how to
8082 recognize.

8084 If you want to pass an option that takes a separate argument, you must use
8085 @option{-Xlinker} twice, once for the option and once for the argument.
8086 For example, to pass @option{-assert definitions}, you must write
8087 @samp{-Xlinker -assert -Xlinker definitions}. It does not work to write
8088 @option{-Xlinker "-assert definitions"}, because this passes the entire
8089 string as a single argument, which is not what the linker expects.

8091 When using the GNU linker, it is usually more convenient to pass
8092 arguments to linker options using the @option{@var{option}=@var{value}}
8093 syntax than as separate arguments. For example, you can specify
8094 @samp{-Xlinker -Map=output.map} rather than
8095 @samp{-Xlinker -Map -Xlinker output.map}. Other linkers may not support
8096 this syntax for command-line options.

8098 @item -Wl,@var{option}
8099 @opindex Wl
8100 Pass @var{option} as an option to the linker. If @var{option} contains
8101 commas, it is split into multiple options at the commas. You can use this
8102 syntax to pass an argument to the option.
8103 For example, @samp{-Wl,-Map,output.map} passes @samp{-Map output.map} to the
8104 linker. When using the GNU linker, you can also get the same effect with
8105 @samp{-Wl,-Map=output.map}.

8107 @item -u @var{symbol}
8108 @opindex u
8109 Pretend the symbol @var{symbol} is undefined, to force linking of
8110 library modules to define it. You can use @option{-u} multiple times with
8111 different symbols to force loading of additional library modules.
8112 @end table

8114 @node Directory Options
8115 @section Options for Directory Search
8116 @cindex directory options
8117 @cindex options, directory search
8118 @cindex search path

8120 These options specify directories to search for header files, for
8121 libraries and for parts of the compiler:

8123 @table @gcctabopt
8124 @item -I@var{dir}
8125 @opindex I
8126 Add the directory @var{dir} to the head of the list of directories to be
8127 searched for header files. This can be used to override a system header
8128 file, substituting your own version, since these directories are
8129 searched before the system header file directories. However, you should
8130 not use this option to add directories that contain vendor-supplied
8131 system header files (use @option{-isystem} for that). If you use more than
8132 one @option{-I} option, the directories are scanned in left-to-right
8133 order; the standard system directories come after.

8135 If a standard system include directory, or a directory specified with
8136 @option{-isystem}, is also specified with @option{-I}, the @option{-I}
8137 option will be ignored. The directory will still be searched but as a
8138 system directory at its normal position in the system include chain.
8139 This is to ensure that GCC's procedure to fix buggy system headers and
8140 the ordering for the include_next directive are not inadvertently changed.
8141 If you really need to change the search order for system directories,
8142 use the @option{-nostdinc} and/or @option{-isystem} options.

```

```

8144 @item -iquote@var{dir}
8145 @opindex iquote
8146 Add the directory @var{dir} to the head of the list of directories to
8147 be searched for header files only for the case of @samp{#include
8148 "@var{file}"}; they are not searched for @samp{#include <@var{file}>},
8149 otherwise just like @option{-I}.

8151 @item -L@var{dir}
8152 @opindex L
8153 Add directory @var{dir} to the list of directories to be searched
8154 for @option{-l}.

8156 @item -B@var{prefix}
8157 @opindex B
8158 This option specifies where to find the executables, libraries,
8159 include files, and data files of the compiler itself.

8161 The compiler driver program runs one or more of the subprograms
8162 @file{cpp}, @file{ccl}, @file{as} and @file{ld}. It tries
8163 @var{prefix} as a prefix for each program it tries to run, both with and
8164 without @samp{@var{machine}/@var{version}/} (@pxref{Target Options}).

8166 For each subprogram to be run, the compiler driver first tries the
8167 @option{-B} prefix, if any. If that name is not found, or if @option{-B}
8168 was not specified, the driver tries two standard prefixes, which are
8169 @file{/usr/lib/gcc/} and @file{/usr/local/lib/gcc/}. If neither of
8170 those results in a file name that is found, the unmodified program
8171 name is searched for using the directories specified in your
8172 @env{PATH} environment variable.

8174 The compiler will check to see if the path provided by the @option{-B}
8175 refers to a directory, and if necessary it will add a directory
8176 separator character at the end of the path.

8178 @option{-B} prefixes that effectively specify directory names also apply
8179 to libraries in the linker, because the compiler translates these
8180 options into @option{-L} options for the linker. They also apply to
8181 includes files in the preprocessor, because the compiler translates these
8182 options into @option{-isystem} options for the preprocessor. In this case,
8183 the compiler appends @samp{include} to the prefix.

8185 The run-time support file @file{libgcc.a} can also be searched for using
8186 the @option{-B} prefix, if needed. If it is not found there, the two
8187 standard prefixes above are tried, and that is all. The file is left
8188 out of the link if it is not found by those means.

8190 Another way to specify a prefix much like the @option{-B} prefix is to use
8191 the environment variable @env{GCC_EXEC_PREFIX}. @xref{Environment
8192 Variables}.

8194 As a special kludge, if the path provided by @option{-B} is
8195 @file{[dir/]stage@var{N}/}, where @var{N} is a number in the range 0 to
8196 9, then it will be replaced by @file{[dir/]include}. This is to help
8197 with boot-strapping the compiler.

8199 @item -specs=@var{file}
8200 @opindex specs
8201 Process @var{file} after the compiler reads in the standard @file{specs}
8202 file, in order to override the defaults that the @file{gcc} driver
8203 program uses when determining what switches to pass to @file{ccl},
8204 @file{cclplus}, @file{as}, @file{ld}, etc. More than one
8205 @option{-specs=@var{file}} can be specified on the command line, and they
8206 are processed in order, from left to right.

8208 @item --sysroot=@var{dir}

```

```

8209 @opindex sysroot
8210 Use @var{dir} as the logical root directory for headers and libraries.
8211 For example, if the compiler would normally search for headers in
8212 @file{/usr/include} and libraries in @file{/usr/lib}, it will instead
8213 search @file{@var{dir}/usr/include} and @file{@var{dir}/usr/lib}.

8215 If you use both this option and the @option{-isysroot} option, then
8216 the @option{--sysroot} option will apply to libraries, but the
8217 @option{-isysroot} option will apply to header files.

8219 The GNU linker (beginning with version 2.16) has the necessary support
8220 for this option. If your linker does not support this option, the
8221 header file aspect of @option{--sysroot} will still work, but the
8222 library aspect will not.

8224 @item -I-
8225 @opindex I-
8226 This option has been deprecated. Please use @option{-iquote} instead for
8227 @option{-I} directories before the @option{-I-} and remove the @option{-I-}.
8228 Any directories you specify with @option{-I} options before the @option{-I-}
8229 option are searched only for the case of @samp{#include "var{file}"};
8230 they are not searched for @samp{#include <var{file}>}.

8232 If additional directories are specified with @option{-I} options after
8233 the @option{-I-}, these directories are searched for all @samp{#include}
8234 directives. (Ordinarily @emph{all} @option{-I} directories are used
8235 this way.)

8237 In addition, the @option{-I-} option inhibits the use of the current
8238 directory (where the current input file came from) as the first search
8239 directory for @samp{#include "var{file}"}. There is no way to
8240 override this effect of @option{-I-}. With @option{-I-} you can specify
8241 searching the directory which was current when the compiler was
8242 invoked. That is not exactly the same as what the preprocessor does
8243 by default, but it is often satisfactory.

8245 @option{-I-} does not inhibit the use of the standard system directories
8246 for header files. Thus, @option{-I-} and @option{-nostdinc} are
8247 independent.
8248 @end table

8250 @c man end

8252 @node Spec Files
8253 @section Specifying subprocesses and the switches to pass to them
8254 @cindex Spec Files

8256 @command{gcc} is a driver program. It performs its job by invoking a
8257 sequence of other programs to do the work of compiling, assembling and
8258 linking. GCC interprets its command-line parameters and uses these to
8259 deduce which programs it should invoke, and which command-line options
8260 it ought to place on their command lines. This behavior is controlled
8261 by @dfn{spec strings}. In most cases there is one spec string for each
8262 program that GCC can invoke, but a few programs have multiple spec
8263 strings to control their behavior. The spec strings built into GCC can
8264 be overridden by using the @option{-specs=} command-line switch to specify
8265 a spec file.

8267 @dfn{Spec files} are plaintext files that are used to construct spec
8268 strings. They consist of a sequence of directives separated by blank
8269 lines. The type of directive is determined by the first non-whitespace
8270 character on the line and it can be one of the following:

8272 @table @code
8273 @item %@var{command}
8274 Issues a @var{command} to the spec file processor. The commands that can

```

```

8275 appear here are:

8277 @table @code
8278 @item %include <@var{file}>
8279 @cindex %include
8280 Search for @var{file} and insert its text at the current point in the
8281 specs file.

8283 @item %include_noerr <@var{file}>
8284 @cindex %include_noerr
8285 Just like @samp{#include}, but do not generate an error message if the include
8286 file cannot be found.

8288 @item %rename @var{old_name} @var{new_name}
8289 @cindex %rename
8290 Rename the spec string @var{old_name} to @var{new_name}.

8292 @end table

8294 @item *[@var{spec_name}]:
8295 This tells the compiler to create, override or delete the named spec
8296 string. All lines after this directive up to the next directive or
8297 blank line are considered to be the text for the spec string. If this
8298 results in an empty string then the spec will be deleted. (Or, if the
8299 spec did not exist, then nothing will happen.) Otherwise, if the spec
8300 does not currently exist a new spec will be created. If the spec does
8301 exist then its contents will be overridden by the text of this
8302 directive, unless the first character of that text is the @samp{+}
8303 character, in which case the text will be appended to the spec.

8305 @item [:@var{suffix}]:
8306 Creates a new @samp{[@var{suffix}] spec} pair. All lines after this directive
8307 and up to the next directive or blank line are considered to make up the
8308 spec string for the indicated suffix. When the compiler encounters an
8309 input file with the named suffix, it will process the spec string in
8310 order to work out how to compile that file. For example:

8312 @smallexample
8313 .ZZ:
8314 z-compile -input %i
8315 @end smallexample

8317 This says that any input file whose name ends in @samp{.ZZ} should be
8318 passed to the program @samp{z-compile}, which should be invoked with the
8319 command-line switch @option{-input} and with the result of performing the
8320 @samp{%i} substitution. (See below.)

8322 As an alternative to providing a spec string, the text that follows a
8323 suffix directive can be one of the following:

8325 @table @code
8326 @item @@@var{language}
8327 This says that the suffix is an alias for a known @var{language}. This is
8328 similar to using the @option{-x} command-line switch to GCC to specify a
8329 language explicitly. For example:

8331 @smallexample
8332 .ZZ:
8333 @c++
8334 @end smallexample

8336 Says that .ZZ files are, in fact, C++ source files.

8338 @item #@var{name}
8339 This causes an error messages saying:

```



```

8341 @smallexample
8342 @var{name} compiler not installed on this system.
8343 @end smallexample
8344 @end table

8346 GCC already has an extensive list of suffixes built into it.
8347 This directive will add an entry to the end of the list of suffixes, but
8348 since the list is searched from the end backwards, it is effectively
8349 possible to override earlier entries using this technique.

```

```
8351 @end table
```

```

8353 GCC has the following spec strings built into it. Spec files can
8354 override these strings or create their own. Note that individual
8355 targets can also add their own spec strings to this list.

```

```

8357 @smallexample
8358 asm      Options to pass to the assembler
8359 asm_final Options to pass to the assembler post-processor
8360 cpp      Options to pass to the C preprocessor
8361 ccl      Options to pass to the C compiler
8362 cclplus  Options to pass to the C++ compiler
8363 endfile  Object files to include at the end of the link
8364 link     Options to pass to the linker
8365 lib      Libraries to include on the command line to the linker
8366 libgcc  Decides which GCC support library to pass to the linker
8367 linker   Sets the name of the linker
8368 predefines Defines to be passed to the C preprocessor
8369 signed_char Defines to pass to CPP to say whether @code{char} is signed
8370          by default
8371 startfile Object files to include at the start of the link
8372 @end smallexample

```

```
8374 Here is a small example of a spec file:
```

```

8376 @smallexample
8377 %rename lib      old_lib

8379 *lib:
8380 --start-group -lgcc -lc -levall --end-group %(old_lib)
8381 @end smallexample

```

```

8383 This example renames the spec called @samp{lib} to @samp{old_lib} and
8384 then overrides the previous definition of @samp{lib} with a new one.
8385 The new definition adds in some extra command-line options before
8386 including the text of the old definition.

```

```

8388 @dfn{Spec strings} are a list of command-line options to be passed to their
8389 corresponding program. In addition, the spec strings can contain
8390 @samp{%}-prefixed sequences to substitute variable text or to
8391 conditionally insert text into the command line. Using these constructs
8392 it is possible to generate quite complex command lines.

```

```

8394 Here is a table of all defined @samp{%}-sequences for spec
8395 strings. Note that spaces are not generated automatically around the
8396 results of expanding these sequences. Therefore you can concatenate them
8397 together or combine them with constant text in a single argument.

```

```

8399 @table @code
8400 @item %%
8401 Substitute one @samp{%} into the program name or argument.

```

```

8403 @item %i
8404 Substitute the name of the input file being processed.

```

```
8406 @item %b
```

```

8407 Substitute the basename of the input file being processed.
8408 This is the substring up to (and not including) the last period
8409 and not including the directory.

```

```

8411 @item %B
8412 This is the same as @samp{%b}, but include the file suffix (text after
8413 the last period).

```

```

8415 @item %d
8416 Marks the argument containing or following the @samp{%d} as a
8417 temporary file name, so that that file will be deleted if GCC exits
8418 successfully. Unlike @samp{%g}, this contributes no text to the
8419 argument.

```

```

8421 @item %g@var{suffix}
8422 Substitute a file name that has suffix @var{suffix} and is chosen
8423 once per compilation, and mark the argument in the same way as
8424 @samp{%d}. To reduce exposure to denial-of-service attacks, the file
8425 name is now chosen in a way that is hard to predict even when previously
8426 chosen file names are known. For example, @samp{%g.s @dots{} %g.o @dots{} %g.s}
8427 might turn into @samp{ccUVUUUAU.s ccXYAXZ12.o ccUVUUUAU.s}. @var{suffix} matches
8428 the regexp @samp{[.A-Za-z]*} or the special string @samp{%0}, which is
8429 treated exactly as if @samp{%0} had been preprocessed. Previously, @samp{%g}
8430 was simply substituted with a file name chosen once per compilation,
8431 without regard to any appended suffix (which was therefore treated
8432 just like ordinary text), making such attacks more likely to succeed.

```

```

8434 @item %u@var{suffix}
8435 Like @samp{%g}, but generates a new temporary file name even if
8436 @samp{%u@var{suffix}} was already seen.

```

```

8438 @item %U@var{suffix}
8439 Substitutes the last file name generated with @samp{%u@var{suffix}}, generating
8440 new one if there is no such last file name. In the absence of any
8441 @samp{%u@var{suffix}}, this is just like @samp{%g@var{suffix}}, except they don'
8442 the same suffix @emph{space}, so @samp{%g.s @dots{} %U.s @dots{} %g.s @dots{} %U
8443 would involve the generation of two distinct file names, one
8444 for each @samp{%g.s} and another for each @samp{%U.s}. Previously, @samp{%U} wa
8445 simply substituted with a file name chosen for the previous @samp{%u},
8446 without regard to any appended suffix.

```

```

8448 @item %j@var{suffix}
8449 Substitutes the name of the @code{HOST_BIT_BUCKET}, if any, and if it is
8450 writable, and if save-temps is off; otherwise, substitute the name
8451 of a temporary file, just like @samp{%u}. This temporary file is not
8452 meant for communication between processes, but rather as a junk
8453 disposal mechanism.

```

```

8455 @item %|@var{suffix}
8456 @itemx %m@var{suffix}
8457 Like @samp{%g}, except if @option{-pipe} is in effect. In that case
8458 @samp{%|} substitutes a single dash and @samp{%m} substitutes nothing at
8459 all. These are the two most common ways to instruct a program that it
8460 should read from standard input or write to standard output. If you
8461 need something more elaborate you can use an @samp{%@{pipe:@code{X}@}}
8462 construct: see for example @file{/lang-specs.h}.

```

```

8464 @item %.@var{SUFFIX}
8465 Substitutes @var{.SUFFIX} for the suffixes of a matched switch's args
8466 when it is subsequently output with @samp{%*}. @var{SUFFIX} is
8467 terminated by the next space or %.

```

```

8469 @item %w
8470 Marks the argument containing or following the @samp{%w} as the
8471 designated output file of this compilation. This puts the argument
8472 into the sequence of arguments that @samp{%o} will substitute later.

```

8474 @item %o  
8475 Substitutes the names of all the output files, with spaces  
8476 automatically placed around them. You should write spaces  
8477 around the @samp{%o} as well or the results are undefined.  
8478 @samp{%o} is for use in the specs for running the linker.  
8479 Input files whose names have no recognized suffix are not compiled  
8480 at all, but they are included among the output files, so they will  
8481 be linked.

8483 @item %O  
8484 Substitutes the suffix for object files. Note that this is  
8485 handled specially when it immediately follows @samp{%g, %u, or %U},  
8486 because of the need for those to form complete file names. The  
8487 handling is such that @samp{%O} is treated exactly as if it had already  
8488 been substituted, except that @samp{%g, %u, and %U} do not currently  
8489 support additional @var{suffix} characters following @samp{%O} as they would  
8490 following, for example, @samp{.o}.

8492 @item %p  
8493 Substitutes the standard macro predefinitions for the  
8494 current target machine. Use this when running @code{cpp}.

8496 @item %P  
8497 Like @samp{p}, but puts @samp{\_\_} before and after the name of each  
8498 predefined macro, except for macros that start with @samp{\_\_} or with  
8499 @samp{\_\_var{L}}, where @var{L} is an uppercase letter. This is for ISO  
8500 C@.

8502 @item %I  
8503 Substitute any of @option{-iprefix} (made from @env{GCC\_EXEC\_PREFIX}),  
8504 @option{-isysroot} (made from @env{TARGET\_SYSTEM\_ROOT}),  
8505 @option{-isystem} (made from @env{COMPILER\_PATH} and @option{-B} options)  
8506 and @option{-imultilib} as necessary.

8508 @item %s  
8509 Current argument is the name of a library or startup file of some sort.  
8510 Search for that file in a standard list of directories and substitute  
8511 the full name found.

8513 @item %e@var{str}  
8514 Print @var{str} as an error message. @var{str} is terminated by a newline.  
8515 Use this when inconsistent options are detected.

8517 @item %(@var{name})  
8518 Substitute the contents of spec string @var{name} at this point.

8520 @item %[@var{name}]  
8521 Like @samp{%(dots{)}} but put @samp{\_\_} around @option{-D} arguments.

8523 @item %x@{var{option}@}  
8524 Accumulate an option for @samp{%X}.

8526 @item %X  
8527 Output the accumulated linker options specified by @option{-Wl} or a @samp{%x}  
8528 spec string.

8530 @item %Y  
8531 Output the accumulated assembler options specified by @option{-Wa}.

8533 @item %Z  
8534 Output the accumulated preprocessor options specified by @option{-Wp}.

8536 @item %a  
8537 Process the @code{asm} spec. This is used to compute the  
8538 switches to be passed to the assembler.

8540 @item %A  
8541 Process the @code{asm\_final} spec. This is a spec string for  
8542 passing switches to an assembler post-processor, if such a program is  
8543 needed.

8545 @item %l  
8546 Process the @code{link} spec. This is the spec for computing the  
8547 command line passed to the linker. Typically it will make use of the  
8548 @samp{%L %G %S %D and %E} sequences.

8550 @item %D  
8551 Dump out a @option{-L} option for each directory that GCC believes might  
8552 contain startup files. If the target supports multilibs then the  
8553 current multilib directory will be prepended to each of these paths.

8555 @item %L  
8556 Process the @code{lib} spec. This is a spec string for deciding which  
8557 libraries should be included on the command line to the linker.

8559 @item %G  
8560 Process the @code{libgcc} spec. This is a spec string for deciding  
8561 which GCC support library should be included on the command line to the linker.

8563 @item %S  
8564 Process the @code{startfile} spec. This is a spec for deciding which  
8565 object files should be the first ones passed to the linker. Typically  
8566 this might be a file named @file{crt0.o}.

8568 @item %E  
8569 Process the @code{endfile} spec. This is a spec string that specifies  
8570 the last object files that will be passed to the linker.

8572 @item %C  
8573 Process the @code{cpp} spec. This is used to construct the arguments  
8574 to be passed to the C preprocessor.

8576 @item %l  
8577 Process the @code{cc1} spec. This is used to construct the options to be  
8578 passed to the actual C compiler (@samp{cc1}).

8580 @item %2  
8581 Process the @code{cc1plus} spec. This is used to construct the options to be  
8582 passed to the actual C++ compiler (@samp{cc1plus}).

8584 @item %\*  
8585 Substitute the variable part of a matched option. See below.  
8586 Note that each comma in the substituted string is replaced by  
8587 a single space.

8589 @item %<@code{S}  
8590 Remove all occurrences of @code{-S} from the command line. Note---this  
8591 command is position dependent. @samp{%} commands in the spec string  
8592 before this one will see @code{-S}, @samp{%} commands in the spec string  
8593 after this one will not.

8595 @item %:@var{function}(@var{args})  
8596 Call the named function @var{function}, passing it @var{args}.  
8597 @var{args} is first processed as a nested spec string, then split  
8598 into an argument vector in the usual fashion. The function returns  
8599 a string which is processed as if it had appeared literally as part  
8600 of the current spec.

8602 The following built-in spec functions are provided:

8604 @table @code

```

8605 @item @code{getenv}
8606 The @code{getenv} spec function takes two arguments: an environment
8607 variable name and a string. If the environment variable is not
8608 defined, a fatal error is issued. Otherwise, the return value is the
8609 value of the environment variable concatenated with the string. For
8610 example, if @code{env{TOPDIR}} is defined as @code{file{/path/to/top}}, then:

8612 @smallexample
8613 %:getenv(TOPDIR /include)
8614 @end smallexample

8616 expands to @code{file{/path/to/top/include}}.

8618 @item @code{if-exists}
8619 The @code{if-exists} spec function takes one argument, an absolute
8620 pathname to a file. If the file exists, @code{if-exists} returns the
8621 pathname. Here is a small example of its usage:

8623 @smallexample
8624 *startfile:
8625 crt0%O% %:if-exists(crti%O%) crtbegin%O%
8626 @end smallexample

8628 @item @code{if-exists-else}
8629 The @code{if-exists-else} spec function is similar to the @code{if-exists}
8630 spec function, except that it takes two arguments. The first argument is
8631 an absolute pathname to a file. If the file exists, @code{if-exists-else}
8632 returns the pathname. If it does not exist, it returns the second argument.
8633 This way, @code{if-exists-else} can be used to select one file or another,
8634 based on the existence of the first. Here is a small example of its usage:

8636 @smallexample
8637 *startfile:
8638 crt0%O% %:if-exists(crti%O%) \
8639 %:if-exists-else(crtbegin%O% crtbegin%O%)
8640 @end smallexample

8642 @item @code{replace-outfile}
8643 The @code{replace-outfile} spec function takes two arguments. It looks for the
8644 first argument in the outfiles array and replaces it with the second argument.
8645 is a small example of its usage:

8647 @smallexample
8648 %@{fgnu-runtime%:replace-outfile(-lobjc -lobjc-gnu)}
8649 @end smallexample

8651 @item @code{print-asm-header}
8652 The @code{print-asm-header} function takes no arguments and simply
8653 prints a banner like:

8655 @smallexample
8656 Assembler options
8657 =====

8659 Use "--Wa,OPTION" to pass "OPTION" to the assembler.
8660 @end smallexample

8662 It is used to separate compiler options from assembler options
8663 in the @code{option{--target-help}} output.
8664 @end table

8666 @item %@{@code{S}@}
8667 Substitutes the @code{-S} switch, if that switch was given to GCC@.
8668 If that switch was not specified, this substitutes nothing. Note that
8669 the leading dash is omitted when specifying this option, and it is
8670 automatically inserted if the substitution is performed. Thus the spec

```

```

8671 string @samp{%@{foo@}} would match the command-line option @code{option{-foo}}
8672 and would output the command line option @code{option{-foo}}.

8674 @item %W@{@code{S}@}
8675 Like %@{@code{S}@} but mark last argument supplied within as a file to be
8676 deleted on failure.

8678 @item %@{@code{S}*@}
8679 Substitutes all the switches specified to GCC whose names start
8680 with @code{-S}, but which also take an argument. This is used for
8681 switches like @code{option{-o}}, @code{option{-D}}, @code{option{-I}}, etc.
8682 GCC considers @code{option{-o foo}} as being
8683 one switch whose names starts with @code{o*}. %@{o*@} would substitute this
8684 text, including the space. Thus two arguments would be generated.

8686 @item %@{@code{S}*@&@code{T}*@}
8687 Like %@{@code{S}*@}, but preserve order of @code{S} and @code{T} options
8688 (the order of @code{S} and @code{T} in the spec is not significant).
8689 There can be any number of ampersand-separated variables; for each the
8690 wild card is optional. Useful for CPP as @code{samp{%@{D*&U*&A*}}}.

8692 @item %@{@code{S}:@code{X}@}
8693 Substitutes @code{X}, if the @code{samp{-S}} switch was given to GCC@.

8695 @item %@{!@code{S}:@code{X}@}
8696 Substitutes @code{X}, if the @code{samp{-S}} switch was @code{emph{not}} given to GCC@.

8698 @item %@{@code{S}*:@code{X}@}
8699 Substitutes @code{X} if one or more switches whose names start with
8700 @code{-S} are specified to GCC@. Normally @code{X} is substituted only
8701 once, no matter how many such switches appeared. However, if @code{*}
8702 appears somewhere in @code{X}, then @code{X} will be substituted once
8703 for each matching switch, with the @code{*} replaced by the part of
8704 that switch that matched the @code{*}.

8706 @item %@{.@code{S}:@code{X}@}
8707 Substitutes @code{X}, if processing a file with suffix @code{S}.

8709 @item %@{!.@code{S}:@code{X}@}
8710 Substitutes @code{X}, if @code{emph{not}} processing a file with suffix @code{S}.

8712 @item %@{,@code{S}:@code{X}@}
8713 Substitutes @code{X}, if processing a file for language @code{S}.

8715 @item %@{!.@code{S}:@code{X}@}
8716 Substitutes @code{X}, if not processing a file for language @code{S}.

8718 @item %@{@code{S}|@code{P}:@code{X}@}
8719 Substitutes @code{X} if either @code{-S} or @code{-P} was given to
8720 GCC@. This may be combined with @code{samp{!}}, @code{samp{.}}, @code{samp{,}}, and
8721 @code{*} sequences as well, although they have a stronger binding than
8722 the @code{samp{||}}. If @code{*} appears in @code{X}, all of the
8723 alternatives must be starred, and only the first matching alternative
8724 is substituted.

8726 For example, a spec string like this:

8728 @smallexample
8729 %@{.c:-foo@} %@{!.c:-bar@} %@{.c|d:-baz@} %@{!.c|d:-boggle@}
8730 @end smallexample

8732 will output the following command-line options from the following input
8733 command-line options:

8735 @smallexample
8736 fred.c -foo -baz

```

```

8737 jim.d      -bar -boggle
8738 -d fred.c   -foo -baz -boggle
8739 -d jim.d    -bar -baz -boggle
8740 @end smallexample

8742 @item %@{S:X; T:Y; :D@}

8744 If @code{S} was given to GCC, substitutes @code{X}; else if @code{T} was
8745 given to GCC, substitutes @code{Y}; else substitutes @code{D}. There can
8746 be as many clauses as you need. This may be combined with @code{.},
8747 @code{,}, @code{!}, @code{[]}, and @code{*} as needed.

8750 @end table

8752 The conditional text @code{X} in a %@{code{S}:code{X}@} or similar
8753 construct may contain other nested @samp{%} constructs or spaces, or
8754 even newlines. They are processed as usual, as described above.
8755 Trailing white space in @code{X} is ignored. White space may also
8756 appear anywhere on the left side of the colon in these constructs,
8757 except between @code{.} or @code{*} and the corresponding word.

8759 The @option{-O}, @option{-f}, @option{-m}, and @option{-W} switches are
8760 handled specifically in these constructs. If another value of
8761 @option{-O} or the negated form of a @option{-f}, @option{-m}, or
8762 @option{-W} switch is found later in the command line, the earlier
8763 switch value is ignored, except with @code{S}*@ where @code{S} is
8764 just one letter, which passes all matching options.

8766 The character @samp{[]} at the beginning of the predicate text is used to
8767 indicate that a command should be piped to the following command, but
8768 only if @option{-pipe} is specified.

8770 It is built into GCC which switches take arguments and which do not.
8771 (You might think it would be useful to generalize this to allow each
8772 compiler's spec to say which switches take arguments. But this cannot
8773 be done in a consistent fashion. GCC cannot even decide which input
8774 files have been specified without knowing which switches take arguments,
8775 and it must know which input files to compile in order to tell which
8776 compilers to run).

8778 GCC also knows implicitly that arguments starting in @option{-l} are to be
8779 treated as compiler output files, and passed to the linker in their
8780 proper position among the other output files.

8782 @c man begin OPTIONS

8784 @node Target Options
8785 @section Specifying Target Machine and Compiler Version
8786 @cindex target options
8787 @cindex cross compiling
8788 @cindex specifying machine version
8789 @cindex specifying compiler version and target machine
8790 @cindex compiler version, specifying
8791 @cindex target machine, specifying

8793 The usual way to run GCC is to run the executable called @file{gcc}, or
8794 @file{<machine>-gcc} when cross-compiling, or
8795 @file{<machine>-gcc-<version>} to run a version other than the one that
8796 was installed last. Sometimes this is inconvenient, so GCC provides
8797 options that will switch to another cross-compiler or version.

8799 @table @gcctabopt
8800 @item -b @var{machine}
8801 @opindex b
8802 The argument @var{machine} specifies the target machine for compilation.

```

```

8804 The value to use for @var{machine} is the same as was specified as the
8805 machine type when configuring GCC as a cross-compiler. For
8806 example, if a cross-compiler was configured with @samp{configure
8807 arm-elf}, meaning to compile for an arm processor with elf binaries,
8808 then you would specify @option{-b arm-elf} to run that cross compiler.
8809 Because there are other options beginning with @option{-b}, the
8810 configuration must contain a hyphen, or @option{-b} alone should be one
8811 argument followed by the configuration in the next argument.

8813 @item -V @var{version}
8814 @opindex V
8815 The argument @var{version} specifies which version of GCC to run.
8816 This is useful when multiple versions are installed. For example,
8817 @var{version} might be @samp{4.0}, meaning to run GCC version 4.0.
8818 @end table

8820 The @option{-V} and @option{-b} options work by running the
8821 @file{<machine>-gcc-<version>} executable, so there's no real reason to
8822 use them if you can just run that directly.

8824 @node Submodel Options
8825 @section Hardware Models and Configurations
8826 @cindex submodel options
8827 @cindex specifying hardware config
8828 @cindex hardware models and configurations, specifying
8829 @cindex machine dependent options

8831 Earlier we discussed the standard option @option{-b} which chooses among
8832 different installed compilers for completely different target
8833 machines, such as VAX vs.@: 68000 vs.@: 80386.

8835 In addition, each of these target machine types can have its own
8836 special options, starting with @samp{-m}, to choose among various
8837 hardware models or configurations---for example, 68010 vs 68020,
8838 floating coprocessor or none. A single installed version of the
8839 compiler can compile for any model or configuration, according to the
8840 options specified.

8842 Some configurations of the compiler also support additional special
8843 options, usually for compatibility with other compilers on the same
8844 platform.

8846 @c This list is ordered alphanumerically by subsection name.
8847 @c It should be the same order and spelling as these options are listed
8848 @c in Machine Dependent Options

8850 @menu
8851 * ARC Options::
8852 * ARM Options::
8853 * AVR Options::
8854 * Blackfin Options::
8855 * CRIS Options::
8856 * CRX Options::
8857 * Darwin Options::
8858 * DEC Alpha Options::
8859 * DEC Alpha/VMS Options::
8860 * FR30 Options::
8861 * FRV Options::
8862 * GNU/Linux Options::
8863 * H8/300 Options::
8864 * HPPA Options::
8865 * i386 and x86-64 Options::
8866 * i386 and x86-64 Windows Options::
8867 * IA-64 Options::
8868 * M32C Options::

```

```

8869 * M32R/D Options::
8870 * M680x0 Options::
8871 * M68hc1x Options::
8872 * MCore Options::
8873 * MIPS Options::
8874 * MMIX Options::
8875 * MN10300 Options::
8876 * PDP-11 Options::
8877 * picoChip Options::
8878 * PowerPC Options::
8879 * RS/6000 and PowerPC Options::
8880 * S/390 and zSeries Options::
8881 * Score Options::
8882 * SH Options::
8883 * SPARC Options::
8884 * SPU Options::
8885 * System V Options::
8886 * V850 Options::
8887 * VAX Options::
8888 * VxWorks Options::
8889 * x86-64 Options::
8890 * Xstormy16 Options::
8891 * Xtensa Options::
8892 * zSeries Options::
8893 @end menu

8895 @node ARC Options
8896 @subsection ARC Options
8897 @cindex ARC Options

8899 These options are defined for ARC implementations:

8901 @table @gcctabopt
8902 @item -EL
8903 @opindex EL
8904 Compile code for little endian mode. This is the default.

8906 @item -EB
8907 @opindex EB
8908 Compile code for big endian mode.

8910 @item -mmangle-cpu
8911 @opindex mmangle-cpu
8912 Prepend the name of the cpu to all public symbol names.
8913 In multiple-processor systems, there are many ARC variants with different
8914 instruction and register set characteristics. This flag prevents code
8915 compiled for one cpu to be linked with code compiled for another.
8916 No facility exists for handling variants that are "almost identical".
8917 This is an all or nothing option.

8919 @item -mcpu=@var{cpu}
8920 @opindex mcpu
8921 Compile code for ARC variant @var{cpu}.
8922 Which variants are supported depend on the configuration.
8923 All variants support @option{-mcpu=base}, this is the default.

8925 @item -mtext=@var{text-section}
8926 @itemx -mdata=@var{data-section}
8927 @itemx -mrodata=@var{readonly-data-section}
8928 @opindex mtext
8929 @opindex mdata
8930 @opindex mrodata
8931 Put functions, data, and readonly data in @var{text-section},
8932 @var{data-section}, and @var{readonly-data-section} respectively
8933 by default. This can be overridden with the @code{section} attribute.
8934 @xref{Variable Attributes}.

```

```

8936 @item -mfix-cortex-m3-ldrd
8937 @opindex mfix-cortex-m3-ldrd
8938 Some Cortex-M3 cores can cause data corruption when @code{ldrd} instructions
8939 with overlapping destination and base registers are used. This option avoids
8940 generating these instructions. This option is enabled by default when
8941 @option{-mcpu=cortex-m3} is specified.

8943 @end table

8945 @node ARM Options
8946 @subsection ARM Options
8947 @cindex ARM options

8949 These @samp{-m} options are defined for Advanced RISC Machines (ARM)
8950 architectures:

8952 @table @gcctabopt
8953 @item -mabi=@var{name}
8954 @opindex mabi
8955 Generate code for the specified ABI@. Permissible values are: @samp{apcs-gnu},
8956 @samp{atpcs}, @samp{apcs}, @samp{apcs-linux} and @samp{iwmmxt}.

8958 @item -mapcs-frame
8959 @opindex mapcs-frame
8960 Generate a stack frame that is compliant with the ARM Procedure Call
8961 Standard for all functions, even if this is not strictly necessary for
8962 correct execution of the code. Specifying @option{-fomit-frame-pointer}
8963 with this option will cause the stack frames not to be generated for
8964 leaf functions. The default is @option{-mno-apcs-frame}.

8966 @item -mapcs
8967 @opindex mapcs
8968 This is a synonym for @option{-mapcs-frame}.

8970 @ignore
8971 @c not currently implemented
8972 @item -mapcs-stack-check
8973 @opindex mapcs-stack-check
8974 Generate code to check the amount of stack space available upon entry to
8975 every function (that actually uses some stack space). If there is
8976 insufficient space available then either the function
8977 @samp{__rt_stkovf_split_small} or @samp{__rt_stkovf_split_big} will be
8978 called, depending upon the amount of stack space required. The run time
8979 system is required to provide these functions. The default is
8980 @option{-mno-apcs-stack-check}, since this produces smaller code.

8982 @c not currently implemented
8983 @item -mapcs-float
8984 @opindex mapcs-float
8985 Pass floating point arguments using the float point registers. This is
8986 one of the variants of the APCS@. This option is recommended if the
8987 target hardware has a floating point unit or if a lot of floating point
8988 arithmetic is going to be performed by the code. The default is
8989 @option{-mno-apcs-float}, since integer only code is slightly increased in
8990 size if @option{-mapcs-float} is used.

8992 @c not currently implemented
8993 @item -mapcs-reentrant
8994 @opindex mapcs-reentrant
8995 Generate reentrant, position independent code. The default is
8996 @option{-mno-apcs-reentrant}.
8997 @end ignore

8999 @item -mthumb-interwork
9000 @opindex mthumb-interwork

```

9001 Generate code which supports calling between the ARM and Thumb  
 9002 instruction sets. Without this option the two instruction sets cannot  
 9003 be reliably used inside one program. The default is  
 9004 `@option{-mno-thumb-interwork}`, since slightly larger code is generated  
 9005 when `@option{-mthumb-interwork}` is specified.

9007 `@item -mno-sched-prolog`  
 9008 `@opindex mno-sched-prolog`  
 9009 Prevent the reordering of instructions in the function prolog, or the  
 9010 merging of those instruction with the instructions in the function's  
 9011 body. This means that all functions will start with a recognizable set  
 9012 of instructions (or in fact one of a choice from a small set of  
 9013 different function prologues), and this information can be used to  
 9014 locate the start if functions inside an executable piece of code. The  
 9015 default is `@option{-msched-prolog}`.

9017 `@item -mfloat-abi=@var{name}`  
 9018 `@opindex mfloat-abi`  
 9019 Specifies which floating-point ABI to use. Permissible values  
 9020 are: `@samp{soft}`, `@samp{softfp}` and `@samp{hard}`.

9022 Specifying `@samp{soft}` causes GCC to generate output containing  
 9023 library calls for floating-point operations.  
 9024 `@samp{softfp}` allows the generation of code using hardware floating-point  
 9025 instructions, but still uses the soft-float calling conventions.  
 9026 `@samp{hard}` allows generation of floating-point instructions  
 9027 and uses FPU-specific calling conventions.

9029 Using `@option{-mfloat-abi=hard}` with VFP coprocessors is not supported.  
 9030 Use `@option{-mfloat-abi=softfp}` with the appropriate `@option{-mfpu}` option  
 9031 to allow the compiler to generate code that makes use of the hardware  
 9032 floating-point capabilities for these CPUs.

9034 The default depends on the specific target configuration. Note that  
 9035 the hard-float and soft-float ABIs are not link-compatible; you must  
 9036 compile your entire program with the same ABI, and link with a  
 9037 compatible set of libraries.

9039 `@item -mhard-float`  
 9040 `@opindex mhard-float`  
 9041 Equivalent to `@option{-mfloat-abi=hard}`.

9043 `@item -msoft-float`  
 9044 `@opindex msoft-float`  
 9045 Equivalent to `@option{-mfloat-abi=soft}`.

9047 `@item -mlittle-endian`  
 9048 `@opindex mlittle-endian`  
 9049 Generate code for a processor running in little-endian mode. This is  
 9050 the default for all standard configurations.

9052 `@item -mbig-endian`  
 9053 `@opindex mbig-endian`  
 9054 Generate code for a processor running in big-endian mode; the default is  
 9055 to compile code for a little-endian processor.

9057 `@item -mwords-little-endian`  
 9058 `@opindex mwords-little-endian`  
 9059 This option only applies when generating code for big-endian processors.  
 9060 Generate code for a little-endian word order but a big-endian byte  
 9061 order. That is, a byte order of the form `@samp{32107654}`. Note: this  
 9062 option should only be used if you require compatibility with code for  
 9063 big-endian ARM processors generated by versions of the compiler prior to  
 9064 2.8.

9066 `@item -mcpu=@var{name}`

9067 `@opindex mcpu`  
 9068 This specifies the name of the target ARM processor. GCC uses this name  
 9069 to determine what kind of instructions it can emit when generating  
 9070 assembly code. Permissible names are: `@samp{arm2}`, `@samp{arm250}`,  
 9071 `@samp{arm3}`, `@samp{arm6}`, `@samp{arm60}`, `@samp{arm600}`, `@samp{arm610}`,  
 9072 `@samp{arm620}`, `@samp{arm7}`, `@samp{arm7m}`, `@samp{arm7d}`, `@samp{arm7dm}`,  
 9073 `@samp{arm7di}`, `@samp{arm7dmi}`, `@samp{arm70}`, `@samp{arm700}`,  
 9074 `@samp{arm700i}`, `@samp{arm710}`, `@samp{arm710c}`, `@samp{arm7100}`,  
 9075 `@samp{arm720}`,  
 9076 `@samp{arm7500}`, `@samp{arm7500fe}`, `@samp{arm7tdmi}`, `@samp{arm7tdmi-s}`,  
 9077 `@samp{arm710t}`, `@samp{arm720t}`, `@samp{arm740t}`,  
 9078 `@samp{strongarm}`, `@samp{strongarm110}`, `@samp{strongarm1100}`,  
 9079 `@samp{strongarm1110}`,  
 9080 `@samp{arm8}`, `@samp{arm810}`, `@samp{arm9}`, `@samp{arm9e}`, `@samp{arm920}`,  
 9081 `@samp{arm920t}`, `@samp{arm922t}`, `@samp{arm946e-s}`, `@samp{arm966e-s}`,  
 9082 `@samp{arm968e-s}`, `@samp{arm926ej-s}`, `@samp{arm940t}`, `@samp{arm9tdmi}`,  
 9083 `@samp{arm10tdmi}`, `@samp{arm1020t}`, `@samp{arm1026ej-s}`,  
 9084 `@samp{arm10e}`, `@samp{arm1020e}`, `@samp{arm1022e}`,  
 9085 `@samp{arm1136j-s}`, `@samp{arm1136jf-s}`, `@samp{mpcore}`, `@samp{mpcorenovfp}`,  
 9086 `@samp{arm1156t2-s}`, `@samp{arm1176jz-s}`, `@samp{arm1176jzf-s}`,  
 9087 `@samp{cortex-a8}`, `@samp{cortex-a9}`,  
 9088 `@samp{cortex-r4}`, `@samp{cortex-r4f}`, `@samp{cortex-m3}`,  
 9089 `@samp{cortex-m1}`,  
 9090 `@samp{xscale}`, `@samp{iwmmxt}`, `@samp{iwmmxt2}`, `@samp{ep9312}`.

9092 `@item -mtune=@var{name}`  
 9093 `@opindex mtune`  
 9094 This option is very similar to the `@option{-mcpu=}` option, except that  
 9095 instead of specifying the actual target processor type, and hence  
 9096 restricting which instructions can be used, it specifies that GCC should  
 9097 tune the performance of the code as if the target were of the type  
 9098 specified in this option, but still choosing the instructions that it  
 9099 will generate based on the cpu specified by a `@option{-mcpu=}` option.  
 9100 For some ARM implementations better performance can be obtained by using  
 9101 this option.

9103 `@item -march=@var{name}`  
 9104 `@opindex march`  
 9105 This specifies the name of the target ARM architecture. GCC uses this  
 9106 name to determine what kind of instructions it can emit when generating  
 9107 assembly code. This option can be used in conjunction with or instead  
 9108 of the `@option{-mcpu=}` option. Permissible names are: `@samp{armv2}`,  
 9109 `@samp{armv2a}`, `@samp{armv3}`, `@samp{armv3m}`, `@samp{armv4}`, `@samp{armv4t}`,  
 9110 `@samp{armv5}`, `@samp{armv5t}`, `@samp{armv5e}`, `@samp{armv5te}`,  
 9111 `@samp{armv6}`, `@samp{armv6j}`,  
 9112 `@samp{armv6t2}`, `@samp{armv6z}`, `@samp{armv6zk}`, `@samp{armv6-m}`,  
 9113 `@samp{armv7}`, `@samp{armv7-a}`, `@samp{armv7-r}`, `@samp{armv7-m}`,  
 9114 `@samp{iwmmxt}`, `@samp{iwmmxt2}`, `@samp{ep9312}`.

9116 `@item -mfpu=@var{name}`  
 9117 `@itemx -mfpe=@var{number}`  
 9118 `@itemx -mfp=@var{number}`  
 9119 `@opindex mfpu`  
 9120 `@opindex mfpe`  
 9121 `@opindex mfp`  
 9122 This specifies what floating point hardware (or hardware emulation) is  
 9123 available on the target. Permissible names are: `@samp{fpa}`, `@samp{fpe2}`,  
 9124 `@samp{fpe3}`, `@samp{maverick}`, `@samp{vfp}`, `@samp{vfpv3}`, `@samp{vfpv3-d16}` and  
 9125 `@samp{neon}`. `@option{-mfp}` and `@option{-mfpe}`  
 9126 are synonyms for `@option{-mfpu=@var{number}}`, for compatibility  
 9127 with older versions of GCC.

9129 If `@option{-msoft-float}` is specified this specifies the format of  
 9130 floating point values.

9132 `@item -mstructure-size-boundary=@var{n}`

9133 `@opindex mstructure-size-boundary`  
 9134 The size of all structures and unions will be rounded up to a multiple  
 9135 of the number of bits set by this option. Permissible values are 8, 32  
 9136 and 64. The default value varies for different toolchains. For the COFF  
 9137 targeted toolchain the default value is 8. A value of 64 is only allowed  
 9138 if the underlying ABI supports it.

9140 Specifying the larger number can produce faster, more efficient code, but  
 9141 can also increase the size of the program. Different values are potentially  
 9142 incompatible. Code compiled with one value cannot necessarily expect to  
 9143 work with code or libraries compiled with another value, if they exchange  
 9144 information using structures or unions.

9146 `@item -mabort-on-noreturn`  
 9147 `@opindex mabort-on-noreturn`  
 9148 Generate a call to the function `@code{abort}` at the end of a  
 9149 `@code{noreturn}` function. It will be executed if the function tries to  
 9150 return.

9152 `@item -mlong-calls`  
 9153 `@itemx -mno-long-calls`  
 9154 `@opindex mlong-calls`  
 9155 `@opindex mno-long-calls`  
 9156 Tells the compiler to perform function calls by first loading the  
 9157 address of the function into a register and then performing a subroutine  
 9158 call on this register. This switch is needed if the target function  
 9159 will lie outside of the 64 megabyte addressing range of the offset based  
 9160 version of subroutine call instruction.

9162 Even if this switch is enabled, not all function calls will be turned  
 9163 into long calls. The heuristic is that static functions, functions  
 9164 which have the `@samp{short-call}` attribute, functions that are inside  
 9165 the scope of a `@samp{#pragma no_long_calls}` directive and functions whose  
 9166 definitions have already been compiled within the current compilation  
 9167 unit, will not be turned into long calls. The exception to this rule is  
 9168 that weak function definitions, functions with the `@samp{long-call}`  
 9169 attribute or the `@samp{section}` attribute, and functions that are within  
 9170 the scope of a `@samp{#pragma long_calls}` directive, will always be  
 9171 turned into long calls.

9173 This feature is not enabled by default. Specifying  
 9174 `@option{-mno-long-calls}` will restore the default behavior, as will  
 9175 placing the function calls within the scope of a `@samp{#pragma`  
 9176 `long_calls_off}` directive. Note these switches have no effect on how  
 9177 the compiler generates code to handle function calls via function  
 9178 pointers.

9180 `@item -msingle-pic-base`  
 9181 `@opindex msingle-pic-base`  
 9182 Treat the register used for PIC addressing as read-only, rather than  
 9183 loading it in the prologue for each function. The run-time system is  
 9184 responsible for initializing this register with an appropriate value  
 9185 before execution begins.

9187 `@item -mpic-register=@var{reg}`  
 9188 `@opindex mpic-register`  
 9189 Specify the register to be used for PIC addressing. The default is R10  
 9190 unless stack-checking is enabled, when R9 is used.

9192 `@item -mcirrus-fix-invalid-insns`  
 9193 `@opindex mcirrus-fix-invalid-insns`  
 9194 `@opindex mno-cirrus-fix-invalid-insns`  
 9195 Insert NOPs into the instruction stream to in order to work around  
 9196 problems with invalid Maverick instruction combinations. This option  
 9197 is only valid if the `@option{-mcpu=ep9312}` option has been used to  
 9198 enable generation of instructions for the Cirrus Maverick floating

9199 `point` co-processor. This option is not enabled by default, since the  
 9200 problem is only present in older Maverick implementations. The default  
 9201 can be re-enabled by use of the `@option{-mno-cirrus-fix-invalid-insns}`  
 9202 switch.

9204 `@item -mpoke-function-name`  
 9205 `@opindex mpoke-function-name`  
 9206 Write the name of each function into the text section, directly  
 9207 preceding the function prologue. The generated code is similar to this:

```
9209 @smallexample
9210     t0
9211     .ascii "arm_poke_function_name", 0
9212     .align
9213     t1
9214     .word 0xff000000 + (t1 - t0)
9215     arm_poke_function_name
9216     mov     ip, sp
9217     stmfd  sp!, @{fp, ip, lr, pc@}
9218     sub     fp, ip, #4
9219 @end smallexample
```

9221 When performing a stack backtrace, code can inspect the value of  
 9222 `@code{pc}` stored at `@code{fp + 0}`. If the trace function then looks at  
 9223 location `@code{pc - 12}` and the top 8 bits are set, then we know that  
 9224 there is a function name embedded immediately preceding this location  
 9225 and has length `@code{((pc[-3]) & 0xff000000)}`.

9227 `@item -mthumb`  
 9228 `@opindex mthumb`  
 9229 Generate code for the Thumb instruction set. The default is to  
 9230 use the 32-bit ARM instruction set.  
 9231 This option automatically enables either 16-bit Thumb-1 or  
 9232 mixed 16/32-bit Thumb-2 instructions based on the `@option{-mcpu=@var{name}}`  
 9233 and `@option{-march=@var{name}}` options.

9235 `@item -mtpcs-frame`  
 9236 `@opindex mtpcs-frame`  
 9237 Generate a stack frame that is compliant with the Thumb Procedure Call  
 9238 Standard for all non-leaf functions. (A leaf function is one that does  
 9239 not call any other functions.) The default is `@option{-mno-tpcs-frame}`.

9241 `@item -mtpcs-leaf-frame`  
 9242 `@opindex mtpcs-leaf-frame`  
 9243 Generate a stack frame that is compliant with the Thumb Procedure Call  
 9244 Standard for all leaf functions. (A leaf function is one that does  
 9245 not call any other functions.) The default is `@option{-mno-apcs-leaf-frame}`.

9247 `@item -mcallee-super-interworking`  
 9248 `@opindex mcallee-super-interworking`  
 9249 Gives all externally visible functions in the file being compiled an ARM  
 9250 instruction set header which switches to Thumb mode before executing the  
 9251 rest of the function. This allows these functions to be called from  
 9252 non-interworking code.

9254 `@item -mcaller-super-interworking`  
 9255 `@opindex mcaller-super-interworking`  
 9256 Allows calls via function pointers (including virtual functions) to  
 9257 execute correctly regardless of whether the target code has been  
 9258 compiled for interworking or not. There is a small overhead in the cost  
 9259 of executing a function pointer if this option is enabled.

9261 `@item -mtp=@var{name}`  
 9262 `@opindex mtp`  
 9263 Specify the access model for the thread local storage pointer. The valid  
 9264 models are `@option{soft}`, which generates calls to `@code{__aeabi_read_tp}`,

9265 @option{cp15}, which fetches the thread pointer from @code{cp15} directly  
 9266 (supported in the arm6k architecture), and @option{auto}, which uses the  
 9267 best available method for the selected processor. The default setting is  
 9268 @option{auto}.

9270 @item -mword-relocations  
 9271 @opindex mword-relocations  
 9272 Only generate absolute relocations on word sized values (i.e. R\_ARM\_ABS32).  
 9273 This is enabled by default on targets (uClinux, SymbianOS) where the runtime  
 9274 loader imposes this restriction, and when @option{-fpic} or @option{-fPIC}  
 9275 is specified.

9277 @end table

9279 @node AVR Options  
 9280 @subsection AVR Options  
 9281 @cindex AVR Options

9283 These options are defined for AVR implementations:

9285 @table @gcctabopt  
 9286 @item -mmcu=@var{mcu}  
 9287 @opindex mmcu  
 9288 Specify ATMEL AVR instruction set or MCU type.

9290 Instruction set avr1 is for the minimal AVR core, not supported by the C  
 9291 compiler, only for assembler programs (MCU types: at90s1200, attiny10,  
 9292 attiny11, attiny12, attiny15, attiny28).

9294 Instruction set avr2 (default) is for the classic AVR core with up to  
 9295 8K program memory space (MCU types: at90s2313, at90s2323, attiny22,  
 9296 at90s2333, at90s2343, at90s4414, at90s4433, at90s4434, at90s8515,  
 9297 at90c8534, at90s8535).

9299 Instruction set avr3 is for the classic AVR core with up to 128K program  
 9300 memory space (MCU types: atmega103, atmega603, at43usb320, at76c711).

9302 Instruction set avr4 is for the enhanced AVR core with up to 8K program  
 9303 memory space (MCU types: atmega8, atmega83, atmega85).

9305 Instruction set avr5 is for the enhanced AVR core with up to 128K program  
 9306 memory space (MCU types: atmega16, atmega161, atmega163, atmega32, atmega323,  
 9307 atmega64, atmega128, at43usb355, at94k).

9309 @item -msize  
 9310 @opindex msize  
 9311 Output instruction sizes to the asm file.

9313 @item -mno-interrupts  
 9314 @opindex mno-interrupts  
 9315 Generated code is not compatible with hardware interrupts.  
 9316 Code size will be smaller.

9318 @item -mcall-prologues  
 9319 @opindex mcall-prologues  
 9320 Functions prologues/epilogues expanded as call to appropriate  
 9321 subroutines. Code size will be smaller.

9323 @item -mno-tablejump  
 9324 @opindex mno-tablejump  
 9325 Do not generate tablejump insns which sometimes increase code size.  
 9326 The option is now deprecated in favor of the equivalent  
 9327 @option{-fno-jump-tables}

9329 @item -mtiny-stack  
 9330 @opindex mtiny-stack

9331 Change only the low 8 bits of the stack pointer.

9333 @item -mint8  
 9334 @opindex mint8  
 9335 Assume int to be 8 bit integer. This affects the sizes of all types: A  
 9336 char will be 1 byte, an int will be 1 byte, a long will be 2 bytes  
 9337 and long long will be 4 bytes. Please note that this option does not  
 9338 comply to the C standards, but it will provide you with smaller code  
 9339 size.  
 9340 @end table

9342 @node Blackfin Options  
 9343 @subsection Blackfin Options  
 9344 @cindex Blackfin Options

9346 @table @gcctabopt  
 9347 @item -mcpu=@var{cpu}@r{[!]-@var{sirevision}@r{]}  
 9348 @opindex mcpu=  
 9349 Specifies the name of the target Blackfin processor. Currently, @var{cpu}  
 9350 can be one of @samp{bf512}, @samp{bf514}, @samp{bf516}, @samp{bf518},  
 9351 @samp{bf522}, @samp{bf523}, @samp{bf524}, @samp{bf525}, @samp{bf526},  
 9352 @samp{bf527}, @samp{bf531}, @samp{bf532}, @samp{bf533},  
 9353 @samp{bf534}, @samp{bf536}, @samp{bf537}, @samp{bf538}, @samp{bf539},  
 9354 @samp{bf542}, @samp{bf544}, @samp{bf547}, @samp{bf548}, @samp{bf549},  
 9355 @samp{bf561}.

9356 The optional @var{sirevision} specifies the silicon revision of the target  
 9357 Blackfin processor. Any workarounds available for the targeted silicon revision  
 9358 will be enabled. If @var{sirevision} is @samp{none}, no workarounds are enabled  
 9359 If @var{sirevision} is @samp{any}, all workarounds for the targeted processor  
 9360 will be enabled. The @code{\_\_SILICON\_REVISION\_\_} macro is defined to two  
 9361 hexadecimal digits representing the major and minor numbers in the silicon  
 9362 revision. If @var{sirevision} is @samp{none}, the @code{\_\_SILICON\_REVISION\_\_}  
 9363 is not defined. If @var{sirevision} is @samp{any}, the  
 9364 @code{\_\_SILICON\_REVISION\_\_} is defined to be @code{0xffff}.

9365 If this optional @var{sirevision} is not used, GCC assumes the latest known  
 9366 silicon revision of the targeted Blackfin processor.

9368 Support for @samp{bf561} is incomplete. For @samp{bf561},  
 9369 Only the processor macro is defined.

9370 Without this option, @samp{bf532} is used as the processor by default.  
 9371 The corresponding predefined processor macros for @var{cpu} is to  
 9372 be defined. And for @samp{bfin-elf} toolchain, this causes the hardware BSP  
 9373 provided by libgloss to be linked in if @option{-msim} is not given.

9375 @item -msim  
 9376 @opindex msim  
 9377 Specifies that the program will be run on the simulator. This causes  
 9378 the simulator BSP provided by libgloss to be linked in. This option  
 9379 has effect only for @samp{bfin-elf} toolchain.  
 9380 Certain other options, such as @option{-mid-shared-library} and  
 9381 @option{-mfdpic}, imply @option{-msim}.

9383 @item -momit-leaf-frame-pointer  
 9384 @opindex momit-leaf-frame-pointer  
 9385 Don't keep the frame pointer in a register for leaf functions. This  
 9386 avoids the instructions to save, set up and restore frame pointers and  
 9387 makes an extra register available in leaf functions. The option  
 9388 @option{-fomit-frame-pointer} removes the frame pointer for all functions  
 9389 which might make debugging harder.

9391 @item -mspecld-anomaly  
 9392 @opindex mspecld-anomaly  
 9393 When enabled, the compiler will ensure that the generated code does not  
 9394 contain speculative loads after jump instructions. If this option is used,  
 9395 @code{\_\_WORKAROUND\_SPECULATIVE\_LOADS} is defined.



9397 @item -mno-specld-anomaly  
 9398 @opindex mno-specld-anomaly  
 9399 Don't generate extra code to prevent speculative loads from occurring.

9401 @item -mcsync-anomaly  
 9402 @opindex mcsync-anomaly  
 9403 When enabled, the compiler will ensure that the generated code does not  
 9404 contain CSYNC or SSYNC instructions too soon after conditional branches.  
 9405 If this option is used, @code{\_\_WORKAROUND\_SPECULATIVE\_SYNCS} is defined.

9407 @item -mno-csync-anomaly  
 9408 @opindex mno-csync-anomaly  
 9409 Don't generate extra code to prevent CSYNC or SSYNC instructions from  
 9410 occurring too soon after a conditional branch.

9412 @item -mlow-64k  
 9413 @opindex mlow-64k  
 9414 When enabled, the compiler is free to take advantage of the knowledge that  
 9415 the entire program fits into the low 64k of memory.

9417 @item -mno-low-64k  
 9418 @opindex mno-low-64k  
 9419 Assume that the program is arbitrarily large. This is the default.

9421 @item -mstack-check-ll  
 9422 @opindex mstack-check-ll  
 9423 Do stack checking using information placed into L1 scratchpad memory by the  
 9424 uClinux kernel.

9426 @item -mid-shared-library  
 9427 @opindex mid-shared-library  
 9428 Generate code that supports shared libraries via the library ID method.  
 9429 This allows for execute in place and shared libraries in an environment  
 9430 without virtual memory management. This option implies @option{-fPIC}.  
 9431 With a @samp{bfin-elf} target, this option implies @option{-msim}.

9433 @item -mno-id-shared-library  
 9434 @opindex mno-id-shared-library  
 9435 Generate code that doesn't assume ID based shared libraries are being used.  
 9436 This is the default.

9438 @item -mleaf-id-shared-library  
 9439 @opindex mleaf-id-shared-library  
 9440 Generate code that supports shared libraries via the library ID method,  
 9441 but assumes that this library or executable won't link against any other  
 9442 ID shared libraries. That allows the compiler to use faster code for jumps  
 9443 and calls.

9445 @item -mno-leaf-id-shared-library  
 9446 @opindex mno-leaf-id-shared-library  
 9447 Do not assume that the code being compiled won't link against any ID shared  
 9448 libraries. Slower code will be generated for jump and call insns.

9450 @item -mshared-library-id=n  
 9451 @opindex mshared-library-id  
 9452 Specified the identification number of the ID based shared library being  
 9453 compiled. Specifying a value of 0 will generate more compact code, specifying  
 9454 other values will force the allocation of that number to the current  
 9455 library but is no more space or time efficient than omitting this option.

9457 @item -msep-data  
 9458 @opindex msep-data  
 9459 Generate code that allows the data segment to be located in a different  
 9460 area of memory from the text segment. This allows for execute in place in  
 9461 an environment without virtual memory management by eliminating relocations  
 9462 against the text section.

9464 @item -mno-sep-data  
 9465 @opindex mno-sep-data  
 9466 Generate code that assumes that the data segment follows the text segment.  
 9467 This is the default.

9469 @item -mlong-calls  
 9470 @itemx -mno-long-calls  
 9471 @opindex mlong-calls  
 9472 @opindex mno-long-calls  
 9473 Tells the compiler to perform function calls by first loading the  
 9474 address of the function into a register and then performing a subroutine  
 9475 call on this register. This switch is needed if the target function  
 9476 will lie outside of the 24 bit addressing range of the offset based  
 9477 version of subroutine call instruction.

9479 This feature is not enabled by default. Specifying  
 9480 @option{-mno-long-calls} will restore the default behavior. Note these  
 9481 switches have no effect on how the compiler generates code to handle  
 9482 function calls via function pointers.

9484 @item -mfast-fp  
 9485 @opindex mfast-fp  
 9486 Link with the fast floating-point library. This library relaxes some of  
 9487 the IEEE floating-point standard's rules for checking inputs against  
 9488 Not-a-Number (NaN), in the interest of performance.

9490 @item -minline-plt  
 9491 @opindex minline-plt  
 9492 Enable inlining of PLT entries in function calls to functions that are  
 9493 not known to bind locally. It has no effect without @option{-mfdpic}.

9495 @item -mmulticore  
 9496 @opindex mmulticore  
 9497 Build standalone application for multicore Blackfin processor. Proper  
 9498 start files and link scripts will be used to support multicore.  
 9499 This option defines @code{\_\_BFIN\_MULTICORE}. It can only be used with  
 9500 @option{-mcpu=bf561@r{[]}-var{sirevision}@r{[]}}. It can be used with  
 9501 @option{-mcorea} or @option{-mcoreb}. If it's used without  
 9502 @option{-mcorea} or @option{-mcoreb}, single application/dual core  
 9503 programming model is used. In this model, the main function of Core B  
 9504 should be named as coreb\_main. If it's used with @option{-mcorea} or  
 9505 @option{-mcoreb}, one application per core programming model is used.  
 9506 If this option is not used, single core application programming  
 9507 model is used.

9509 @item -mcorea  
 9510 @opindex mcorea  
 9511 Build standalone application for Core A of BF561 when using  
 9512 one application per core programming model. Proper start files  
 9513 and link scripts will be used to support Core A. This option  
 9514 defines @code{\_\_BFIN\_COREA}. It must be used with @option{-mmulticore}.

9516 @item -mcoreb  
 9517 @opindex mcoreb  
 9518 Build standalone application for Core B of BF561 when using  
 9519 one application per core programming model. Proper start files  
 9520 and link scripts will be used to support Core B. This option  
 9521 defines @code{\_\_BFIN\_COREB}. When this option is used, coreb\_main  
 9522 should be used instead of main. It must be used with  
 9523 @option{-mmulticore}.

9525 @item -msdram  
 9526 @opindex msdram  
 9527 Build standalone application for SDRAM. Proper start files and  
 9528 link scripts will be used to put the application into SDRAM.

9529 Loader should initialize SDRAM before loading the application  
 9530 into SDRAM. This option defines @code{\_\_BFIN\_SDRAM}.

9532 @item -micplb  
 9533 @opindex micplb  
 9534 Assume that ICPLBs are enabled at runtime. This has an effect on certain  
 9535 anomaly workarounds. For Linux targets, the default is to assume ICPLBs  
 9536 are enabled; for standalone applications the default is off.  
 9537 @end table

9539 @node CRIS Options  
 9540 @subsection CRIS Options  
 9541 @cindex CRIS Options

9543 These options are defined specifically for the CRIS ports.

9545 @table @gcctabopt  
 9546 @item -march=@var{architecture-type}  
 9547 @itemx -mcpu=@var{architecture-type}  
 9548 @opindex march  
 9549 @opindex mcpu  
 9550 Generate code for the specified architecture. The choices for  
 9551 @var{architecture-type} are @samp{v3}, @samp{v8} and @samp{v10} for  
 9552 respectively ETRAX@w{ }4, ETRAX@w{ }100, and ETRAX@w{ }100@w{ }LX@.  
 9553 Default is @samp{v0} except for cris-axis-linux-gnu, where the default is  
 9554 @samp{v10}.

9556 @item -mtune=@var{architecture-type}  
 9557 @opindex mtune  
 9558 Tune to @var{architecture-type} everything applicable about the generated  
 9559 code, except for the ABI and the set of available instructions. The  
 9560 choices for @var{architecture-type} are the same as for  
 9561 @option{-march=@var{architecture-type}}.

9563 @item -mmax-stack-frame=@var{n}  
 9564 @opindex mmax-stack-frame  
 9565 Warn when the stack frame of a function exceeds @var{n} bytes.

9567 @item -metrax4  
 9568 @itemx -metrax100  
 9569 @opindex metrax4  
 9570 @opindex metrax100  
 9571 The options @option{-metrax4} and @option{-metrax100} are synonyms for  
 9572 @option{-march=v3} and @option{-march=v8} respectively.

9574 @item -mmul-bug-workaround  
 9575 @itemx -mno-mul-bug-workaround  
 9576 @opindex mmul-bug-workaround  
 9577 @opindex mno-mul-bug-workaround  
 9578 Work around a bug in the @code{muls} and @code{mulu} instructions for CPU  
 9579 models where it applies. This option is active by default.

9581 @item -mpdebug  
 9582 @opindex mpdebug  
 9583 Enable CRIS-specific verbose debug-related information in the assembly  
 9584 code. This option also has the effect to turn off the @samp{#NO\_APP}  
 9585 formatted-code indicator to the assembler at the beginning of the  
 9586 assembly file.

9588 @item -mcc-init  
 9589 @opindex mcc-init  
 9590 Do not use condition-code results from previous instruction; always emit  
 9591 compare and test instructions before use of condition codes.

9593 @item -mno-side-effects  
 9594 @opindex mno-side-effects

9595 Do not emit instructions with side-effects in addressing modes other than  
 9596 post-increment.

9598 @item -mstack-align  
 9599 @itemx -mno-stack-align  
 9600 @itemx -mdata-align  
 9601 @itemx -mno-data-align  
 9602 @itemx -mconst-align  
 9603 @itemx -mno-const-align  
 9604 @opindex mstack-align  
 9605 @opindex mno-stack-align  
 9606 @opindex mdata-align  
 9607 @opindex mno-data-align  
 9608 @opindex mconst-align  
 9609 @opindex mno-const-align  
 9610 These options (no-options) arranges (eliminate arrangements) for the  
 9611 stack-frame, individual data and constants to be aligned for the maximum  
 9612 single data access size for the chosen CPU model. The default is to  
 9613 arrange for 32-bit alignment. ABI details such as structure layout are  
 9614 not affected by these options.

9616 @item -m32-bit  
 9617 @itemx -m16-bit  
 9618 @itemx -m8-bit  
 9619 @opindex m32-bit  
 9620 @opindex m16-bit  
 9621 @opindex m8-bit  
 9622 Similar to the stack- data- and const-align options above, these options  
 9623 arrange for stack-frame, writable data and constants to all be 32-bit,  
 9624 16-bit or 8-bit aligned. The default is 32-bit alignment.

9626 @item -mno-prologue-epilogue  
 9627 @itemx -mprologue-epilogue  
 9628 @opindex mno-prologue-epilogue  
 9629 @opindex mprologue-epilogue  
 9630 With @option{-mno-prologue-epilogue}, the normal function prologue and  
 9631 epilogue that sets up the stack-frame are omitted and no return  
 9632 instructions or return sequences are generated in the code. Use this  
 9633 option only together with visual inspection of the compiled code: no  
 9634 warnings or errors are generated when call-saved registers must be saved,  
 9635 or storage for local variable needs to be allocated.

9637 @item -mno-gotplt  
 9638 @itemx -mgotplt  
 9639 @opindex mno-gotplt  
 9640 @opindex mgotplt  
 9641 With @option{-fpic} and @option{-fPIC}, don't generate (do generate)  
 9642 instruction sequences that load addresses for functions from the PLT part  
 9643 of the GOT rather than (traditional on other architectures) calls to the  
 9644 PLT@. The default is @option{-mgotplt}.

9646 @item -melf  
 9647 @opindex melf  
 9648 Legacy no-op option only recognized with the cris-axis-elf and  
 9649 cris-axis-linux-gnu targets.

9651 @item -mlinux  
 9652 @opindex mlinux  
 9653 Legacy no-op option only recognized with the cris-axis-linux-gnu target.

9655 @item -sim  
 9656 @opindex sim  
 9657 This option, recognized for the cris-axis-elf arranges  
 9658 to link with input-output functions from a simulator library. Code,  
 9659 initialized data and zero-initialized data are allocated consecutively.

9661 @item -sim2  
 9662 @opindex sim2  
 9663 Like @option{-sim}, but pass linker options to locate initialized data at  
 9664 0x40000000 and zero-initialized data at 0x80000000.  
 9665 @end table

9667 @node CRX Options  
 9668 @subsection CRX Options  
 9669 @cindex CRX Options

9671 These options are defined specifically for the CRX ports.

9673 @table @gcctabopt

9675 @item -mmac  
 9676 @opindex mmac  
 9677 Enable the use of multiply-accumulate instructions. Disabled by default.

9679 @item -mpush-args  
 9680 @opindex mpush-args  
 9681 Push instructions will be used to pass outgoing arguments when functions  
 9682 are called. Enabled by default.  
 9683 @end table

9685 @node Darwin Options  
 9686 @subsection Darwin Options  
 9687 @cindex Darwin options

9689 These options are defined for all architectures running the Darwin operating  
 9690 system.

9692 FSF GCC on Darwin does not create "fat" object files; it will create  
 9693 an object file for the single architecture that it was built to  
 9694 target. Apple's GCC on Darwin does create "fat" files if multiple  
 9695 @option{-arch} options are used; it does so by running the compiler or  
 9696 linker multiple times and joining the results together with  
 9697 @file{lipo}.

9699 The subtype of the file created (like @samp{ppc7400} or @samp{ppc970} or  
 9700 @samp{i686}) is determined by the flags that specify the ISA  
 9701 that GCC is targeting, like @option{-mcpu} or @option{-march}. The  
 9702 @option{-force\_cpusubtype\_ALL} option can be used to override this.

9704 The Darwin tools vary in their behavior when presented with an ISA  
 9705 mismatch. The assembler, @file{as}, will only permit instructions to  
 9706 be used that are valid for the subtype of the file it is generating,  
 9707 so you cannot put 64-bit instructions in an @samp{ppc750} object file.  
 9708 The linker for shared libraries, @file{/usr/bin/libtool}, will fail  
 9709 and print an error if asked to create a shared library with a less  
 9710 restrictive subtype than its input files (for instance, trying to put  
 9711 a @samp{ppc970} object file in a @samp{ppc7400} library). The linker  
 9712 for executables, @file{ld}, will quietly give the executable the most  
 9713 restrictive subtype of any of its input files.

9715 @table @gcctabopt  
 9716 @item -F@var{dir}  
 9717 @opindex F  
 9718 Add the framework directory @var{dir} to the head of the list of  
 9719 directories to be searched for header files. These directories are  
 9720 interleaved with those specified by @option{-I} options and are  
 9721 scanned in a left-to-right order.

9723 A framework directory is a directory with frameworks in it. A  
 9724 framework is a directory with a @samp{"Headers"} and/or  
 9725 @samp{"PrivateHeaders"} directory contained directly in it that ends  
 9726 in @samp{.framework"}. The name of a framework is the name of this

9727 directory excluding the @samp{.framework"}. Headers associated with  
 9728 the framework are found in one of those two directories, with  
 9729 @samp{"Headers"} being searched first. A subframework is a framework  
 9730 directory that is in a framework's @samp{"Frameworks"} directory.  
 9731 Includes of subframework headers can only appear in a header of a  
 9732 framework that contains the subframework, or in a sibling subframework  
 9733 header. Two subframeworks are siblings if they occur in the same  
 9734 framework. A subframework should not have the same name as a  
 9735 framework, a warning will be issued if this is violated. Currently a  
 9736 subframework cannot have subframeworks, in the future, the mechanism  
 9737 may be extended to support this. The standard frameworks can be found  
 9738 in @samp{"/System/Library/Frameworks"} and  
 9739 @samp{"/Library/Frameworks"}. An example include looks like  
 9740 @code{#include <Framework/header.h>}, where @samp{Framework} denotes  
 9741 the name of the framework and header.h is found in the  
 9742 @samp{"PrivateHeaders"} or @samp{"Headers"} directory.

9744 @item -iframework@var{dir}  
 9745 @opindex iframework  
 9746 Like @option{-F} except the directory is treated as a system  
 9747 directory. The main difference between this @option{-iframework} and  
 9748 @option{-F} is that with @option{-iframework} the compiler does not  
 9749 warn about constructs contained within header files found via  
 9750 @var{dir}. This option is valid only for the C family of languages.

9752 @item -gused  
 9753 @opindex gused  
 9754 Emit debugging information for symbols that are used. For STABS  
 9755 debugging format, this enables @option{-feliminate-unused-debug-symbols}.  
 9756 This is by default ON@.

9758 @item -gfull  
 9759 @opindex gfull  
 9760 Emit debugging information for all symbols and types.

9762 @item -mmacosx-version-min=@var{version}  
 9763 The earliest version of MacOS X that this executable will run on  
 9764 is @var{version}. Typical values of @var{version} include @code{10.1},  
 9765 @code{10.2}, and @code{10.3.9}.

9767 If the compiler was built to use the system's headers by default,  
 9768 then the default for this option is the system version on which the  
 9769 compiler is running, otherwise the default is to make choices which  
 9770 are compatible with as many systems and code bases as possible.

9772 @item -mkernel  
 9773 @opindex mkernel  
 9774 Enable kernel development mode. The @option{-mkernel} option sets  
 9775 @option{-static}, @option{-fno-common}, @option{-fno-cxa-atexit},  
 9776 @option{-fno-exceptions}, @option{-fno-non-call-exceptions},  
 9777 @option{-fapple-kext}, @option{-fno-weak} and @option{-fno-rtti} where  
 9778 applicable. This mode also sets @option{-mno-altivec},  
 9779 @option{-msoft-float}, @option{-fno-builtin} and  
 9780 @option{-mlong-branch} for PowerPC targets.

9782 @item -mone-byte-bool  
 9783 @opindex mone-byte-bool  
 9784 Override the defaults for @samp{bool} so that @samp{sizeof(bool)=1}.  
 9785 By default @samp{sizeof(bool)} is @samp{4} when compiling for  
 9786 Darwin/PowerPC and @samp{1} when compiling for Darwin/x86, so this  
 9787 option has no effect on x86.

9789 @strong{Warning:} The @option{-mone-byte-bool} switch causes GCC  
 9790 to generate code that is not binary compatible with code generated  
 9791 without that switch. Using this switch may require recompiling all  
 9792 other modules in a program, including system libraries. Use this

```

9793 switch to conform to a non-default data model.

9795 @item -mfix-and-continue
9796 @itemx -ffix-and-continue
9797 @itemx -findirect-data
9798 @opindex mfix-and-continue
9799 @opindex ffix-and-continue
9800 @opindex findirect-data
9801 Generate code suitable for fast turn around development. Needed to
9802 enable gdb to dynamically load @code{.o} files into already running
9803 programs. @option{-findirect-data} and @option{-ffix-and-continue}
9804 are provided for backwards compatibility.

9806 @item -all_load
9807 @opindex all_load
9808 Loads all members of static archive libraries.
9809 See man ld(1) for more information.

9811 @item -arch_errors_fatal
9812 @opindex arch_errors_fatal
9813 Cause the errors having to do with files that have the wrong architecture
9814 to be fatal.

9816 @item -bind_at_load
9817 @opindex bind_at_load
9818 Causes the output file to be marked such that the dynamic linker will
9819 bind all undefined references when the file is loaded or launched.

9821 @item -bundle
9822 @opindex bundle
9823 Produce a Mach-o bundle format file.
9824 See man ld(1) for more information.

9826 @item -bundle_loader @var{executable}
9827 @opindex bundle_loader
9828 This option specifies the @var{executable} that will be loading the build
9829 output file being linked. See man ld(1) for more information.

9831 @item -dynamiclib
9832 @opindex dynamiclib
9833 When passed this option, GCC will produce a dynamic library instead of
9834 an executable when linking, using the Darwin @file{libtool} command.

9836 @item -force_cpusubtype_ALL
9837 @opindex force_cpusubtype_ALL
9838 This causes GCC's output file to have the @var{ALL} subtype, instead of
9839 one controlled by the @option{-mcpu} or @option{-march} option.

9841 @item -allowable_client @var{client_name}
9842 @itemx -client_name
9843 @itemx -compatibility_version
9844 @itemx -current_version
9845 @itemx -dead_strip
9846 @itemx -dependency-file
9847 @itemx -dylib_file
9848 @itemx -dylinker_install_name
9849 @itemx -dynamic
9850 @itemx -exported_symbols_list
9851 @itemx -filelist
9852 @itemx -flat_namespace
9853 @itemx -force_flat_namespace
9854 @itemx -headerpad_max_install_names
9855 @itemx -image_base
9856 @itemx -init
9857 @itemx -install_name
9858 @itemx -keep_private_externs

```

```

9859 @itemx -multi_module
9860 @itemx -multiply_defined
9861 @itemx -multiply_defined_unused
9862 @itemx -noall_load
9863 @itemx -no_dead_strip_inits_and_terms
9864 @itemx -nofixprebinding
9865 @itemx -nomultidefs
9866 @itemx -noprebind
9867 @itemx -noseglinkedit
9868 @itemx -pagezero_size
9869 @itemx -prebind
9870 @itemx -prebind_all_twolevel_modules
9871 @itemx -private_bundle
9872 @itemx -read_only_relocs
9873 @itemx -sectalign
9874 @itemx -sectobjectsymbols
9875 @itemx -whyload
9876 @itemx -segladdr
9877 @itemx -sectcreate
9878 @itemx -sectobjectsymbols
9879 @itemx -sectororder
9880 @itemx -segaddr
9881 @itemx -segs_read_only_addr
9882 @itemx -segs_read_write_addr
9883 @itemx -seg_addr_table
9884 @itemx -seg_addr_table_filename
9885 @itemx -seglinkedit
9886 @itemx -segprot
9887 @itemx -segs_read_only_addr
9888 @itemx -segs_read_write_addr
9889 @itemx -single_module
9890 @itemx -static
9891 @itemx -sub_library
9892 @itemx -sub_umbrella
9893 @itemx -twolevel_namespace
9894 @itemx -umbrella
9895 @itemx -undefined
9896 @itemx -unexported_symbols_list
9897 @itemx -weak_reference_mismatches
9898 @itemx -whatsloaded
9899 @opindex allowable_client
9900 @opindex client_name
9901 @opindex compatibility_version
9902 @opindex current_version
9903 @opindex dead_strip
9904 @opindex dependency-file
9905 @opindex dylib_file
9906 @opindex dylinker_install_name
9907 @opindex dynamic
9908 @opindex exported_symbols_list
9909 @opindex filelist
9910 @opindex flat_namespace
9911 @opindex force_flat_namespace
9912 @opindex headerpad_max_install_names
9913 @opindex image_base
9914 @opindex init
9915 @opindex install_name
9916 @opindex keep_private_externs
9917 @opindex multi_module
9918 @opindex multiply_defined
9919 @opindex multiply_defined_unused
9920 @opindex noall_load
9921 @opindex no_dead_strip_inits_and_terms
9922 @opindex nofixprebinding
9923 @opindex nomultidefs
9924 @opindex noprebind

```

```

9925 @opindex noseblinkedit
9926 @opindex pagezero_size
9927 @opindex prebind
9928 @opindex prebind_all_twolevel_modules
9929 @opindex private_bundle
9930 @opindex read_only_relocs
9931 @opindex sectalign
9932 @opindex sectobjectsymbols
9933 @opindex whyload
9934 @opindex segladdr
9935 @opindex sectcreate
9936 @opindex sectobjectsymbols
9937 @opindex sectorder
9938 @opindex segaddr
9939 @opindex segs_read_only_addr
9940 @opindex segs_read_write_addr
9941 @opindex seg_addr_table
9942 @opindex seg_addr_table_filename
9943 @opindex seglinkedit
9944 @opindex segprot
9945 @opindex segs_read_only_addr
9946 @opindex segs_read_write_addr
9947 @opindex single_module
9948 @opindex static
9949 @opindex sub_library
9950 @opindex sub_umbrella
9951 @opindex twolevel_namespace
9952 @opindex umbrella
9953 @opindex undefined
9954 @opindex unexported_symbols_list
9955 @opindex weak_reference_mismatches
9956 @opindex whatloaded
9957 These options are passed to the Darwin linker. The Darwin linker man page
9958 describes them in detail.
9959 @end table

```

```

9961 @node DEC Alpha Options
9962 @subsection DEC Alpha Options

```

9964 These @samp{-m} options are defined for the DEC Alpha implementations:

```

9966 @table @gcctabopt
9967 @item -mno-soft-float
9968 @itemx -msoft-float
9969 @opindex mno-soft-float
9970 @opindex msoft-float
9971 Use (do not use) the hardware floating-point instructions for
9972 floating-point operations. When @option{-msoft-float} is specified,
9973 functions in @file{libgcc.a} will be used to perform floating-point
9974 operations. Unless they are replaced by routines that emulate the
9975 floating-point operations, or compiled in such a way as to call such
9976 emulations routines, these routines will issue floating-point
9977 operations. If you are compiling for an Alpha without floating-point
9978 operations, you must ensure that the library is built so as not to call
9979 them.

```

9981 Note that Alpha implementations without floating-point operations are  
9982 required to have floating-point registers.

```

9984 @item -mfp-reg
9985 @itemx -mno-fp-regs
9986 @opindex mfp-reg
9987 @opindex mno-fp-regs
9988 Generate code that uses (does not use) the floating-point register set.
9989 @option{-mno-fp-regs} implies @option{-msoft-float}. If the floating-point
9990 register set is not used, floating point operands are passed in integer

```

```

9991 registers as if they were integers and floating-point results are passed
9992 in @code{$0} instead of @code{$f0}. This is a non-standard calling sequence,
9993 so any function with a floating-point argument or return value called by code
9994 compiled with @option{-mno-fp-regs} must also be compiled with that
9995 option.

```

9997 A typical use of this option is building a kernel that does not use,  
9998 and hence need not save and restore, any floating-point registers.

```

10000 @item -mieee
10001 @opindex mieee
10002 The Alpha architecture implements floating-point hardware optimized for
10003 maximum performance. It is mostly compliant with the IEEE floating
10004 point standard. However, for full compliance, software assistance is
10005 required. This option generates code fully IEEE compliant code
10006 @emph{except} that the @var{inexact-flag} is not maintained (see below).
10007 If this option is turned on, the preprocessor macro @code{IEEE_FP} is
10008 defined during compilation. The resulting code is less efficient but is
10009 able to correctly support denormalized numbers and exceptional IEEE
10010 values such as not-a-number and plus/minus infinity. Other Alpha
10011 compilers call this option @option{-ieee_with_no_inexact}.

```

```

10013 @item -mieee-with-inexact
10014 @opindex mieee-with-inexact
10015 This is like @option{-mieee} except the generated code also maintains
10016 the IEEE @var{inexact-flag}. Turning on this option causes the
10017 generated code to implement fully-compliant IEEE math. In addition to
10018 @code{IEEE_FP}, @code{IEEE_FP_EXACT} is defined as a preprocessor
10019 macro. On some Alpha implementations the resulting code may execute
10020 significantly slower than the code generated by default. Since there is
10021 very little code that depends on the @var{inexact-flag}, you should
10022 normally not specify this option. Other Alpha compilers call this
10023 option @option{-ieee_with_inexact}.

```

```

10025 @item -mfp-trap-mode=@var{trap-mode}
10026 @opindex mfp-trap-mode
10027 This option controls what floating-point related traps are enabled.
10028 Other Alpha compilers call this option @option{-fptm @var{trap-mode}}.
10029 The trap mode can be set to one of four values:

```

```

10031 @table @samp
10032 @item n
10033 This is the default (normal) setting. The only traps that are enabled
10034 are the ones that cannot be disabled in software (e.g., division by zero
10035 trap).

```

```

10037 @item u
10038 In addition to the traps enabled by @samp{n}, underflow traps are enabled
10039 as well.

```

```

10041 @item su
10042 Like @samp{u}, but the instructions are marked to be safe for software
10043 completion (see Alpha architecture manual for details).

```

```

10045 @item sui
10046 Like @samp{su}, but inexact traps are enabled as well.
10047 @end table

```

```

10049 @item -mfp-rounding-mode=@var{rounding-mode}
10050 @opindex mfp-rounding-mode
10051 Selects the IEEE rounding mode. Other Alpha compilers call this option
10052 @option{-fprm @var{rounding-mode}}. The @var{rounding-mode} can be one
10053 of:

```

```

10055 @table @samp
10056 @item n

```

10057 Normal IEEE rounding mode. Floating point numbers are rounded towards  
 10058 the nearest machine number or towards the even machine number in case  
 10059 of a tie.

10061 @item m  
 10062 Round towards minus infinity.

10064 @item c  
 10065 Chopped rounding mode. Floating point numbers are rounded towards zero.

10067 @item d  
 10068 Dynamic rounding mode. A field in the floating point control register  
 10069 (@var{fpcr}, see Alpha architecture reference manual) controls the  
 10070 rounding mode in effect. The C library initializes this register for  
 10071 rounding towards plus infinity. Thus, unless your program modifies the  
 10072 @var{fpcr}, @samp{d} corresponds to round towards plus infinity.  
 10073 @end table

10075 @item -mtrap-precision=@var{trap-precision}  
 10076 @opindex mtrap-precision  
 10077 In the Alpha architecture, floating point traps are imprecise. This  
 10078 means without software assistance it is impossible to recover from a  
 10079 floating trap and program execution normally needs to be terminated.  
 10080 GCC can generate code that can assist operating system trap handlers  
 10081 in determining the exact location that caused a floating point trap.  
 10082 Depending on the requirements of an application, different levels of  
 10083 precisions can be selected:

10085 @table @samp  
 10086 @item p  
 10087 Program precision. This option is the default and means a trap handler  
 10088 can only identify which program caused a floating point exception.

10090 @item f  
 10091 Function precision. The trap handler can determine the function that  
 10092 caused a floating point exception.

10094 @item i  
 10095 Instruction precision. The trap handler can determine the exact  
 10096 instruction that caused a floating point exception.  
 10097 @end table

10099 Other Alpha compilers provide the equivalent options called  
 10100 @option{-scope\_safe} and @option{-resumption\_safe}.

10102 @item -mieee-conformant  
 10103 @opindex mieee-conformant  
 10104 This option marks the generated code as IEEE conformant. You must not  
 10105 use this option unless you also specify @option{-mtrap-precision=i} and either  
 10106 @option{-mfp-trap-mode=su} or @option{-mfp-trap-mode=sui}. Its only effect  
 10107 is to emit the line @samp{.eflag 48} in the function prologue of the  
 10108 generated assembly file. Under DEC Unix, this has the effect that  
 10109 IEEE-conformant math library routines will be linked in.

10111 @item -mbuild-constants  
 10112 @opindex mbuild-constants  
 10113 Normally GCC examines a 32- or 64-bit integer constant to  
 10114 see if it can construct it from smaller constants in two or three  
 10115 instructions. If it cannot, it will output the constant as a literal and  
 10116 generate code to load it from the data segment at runtime.

10118 Use this option to require GCC to construct @emph{all} integer constants  
 10119 using code, even if it takes more instructions (the maximum is six).

10121 You would typically use this option to build a shared library dynamic  
 10122 loader. Itself a shared library, it must relocate itself in memory

10123 before it can find the variables and constants in its own data segment.

10125 @item -malpha-as  
 10126 @itemx -mgas  
 10127 @opindex malpha-as  
 10128 @opindex mgas  
 10129 Select whether to generate code to be assembled by the vendor-supplied  
 10130 assembler (@option{-malpha-as}) or by the GNU assembler @option{-mgas}.

10132 @item -mbwx  
 10133 @itemx -mno-bwx  
 10134 @itemx -mcix  
 10135 @itemx -mno-cix  
 10136 @itemx -mfix  
 10137 @itemx -mno-fix  
 10138 @itemx -mmax  
 10139 @itemx -mno-max  
 10140 @opindex mbwx  
 10141 @opindex mno-bwx  
 10142 @opindex mcix  
 10143 @opindex mno-cix  
 10144 @opindex mfix  
 10145 @opindex mno-fix  
 10146 @opindex mmax  
 10147 @opindex mno-max  
 10148 Indicate whether GCC should generate code to use the optional BWX,  
 10149 CIX, FIX and MAX instruction sets. The default is to use the instruction  
 10150 sets supported by the CPU type specified via @option{-mcpu=} option or that  
 10151 of the CPU on which GCC was built if none was specified.

10153 @item -mfloat-vax  
 10154 @itemx -mfloat-ieee  
 10155 @opindex mfloat-vax  
 10156 @opindex mfloat-ieee  
 10157 Generate code that uses (does not use) VAX F and G floating point  
 10158 arithmetic instead of IEEE single and double precision.

10160 @item -mexplicit-relocs  
 10161 @itemx -mno-explicit-relocs  
 10162 @opindex mexplicit-relocs  
 10163 @opindex mno-explicit-relocs  
 10164 Older Alpha assemblers provided no way to generate symbol relocations  
 10165 except via assembler macros. Use of these macros does not allow  
 10166 optimal instruction scheduling. GNU binutils as of version 2.12  
 10167 supports a new syntax that allows the compiler to explicitly mark  
 10168 which relocations should apply to which instructions. This option  
 10169 is mostly useful for debugging, as GCC detects the capabilities of  
 10170 the assembler when it is built and sets the default accordingly.

10172 @item -msmall-data  
 10173 @itemx -mlarge-data  
 10174 @opindex msmall-data  
 10175 @opindex mlarge-data  
 10176 When @option{-mexplicit-relocs} is in effect, static data is  
 10177 accessed via @dfn{gp-relative} relocations. When @option{-msmall-data}  
 10178 is used, objects 8 bytes long or smaller are placed in a @dfn{small data area}  
 10179 (the @code{.sdata} and @code{.sbss} sections) and are accessed via  
 10180 16-bit relocations off of the @code{\$gp} register. This limits the  
 10181 size of the small data area to 64KB, but allows the variables to be  
 10182 directly accessed via a single instruction.

10184 The default is @option{-mlarge-data}. With this option the data area  
 10185 is limited to just below 2GB@. Programs that require more than 2GB of  
 10186 data must use @code{malloc} or @code{mmap} to allocate the data in the  
 10187 heap instead of in the program's data segment.

10189 When generating code for shared libraries, `@option{-fpic}` implies  
 10190 `@option{-msmall-data}` and `@option{-fPIC}` implies `@option{-mlarge-data}`.

10192 `@item -msmall-text`  
 10193 `@itemx -mlarge-text`  
 10194 `@opindex msmall-text`  
 10195 `@opindex mlarge-text`  
 10196 When `@option{-msmall-text}` is used, the compiler assumes that the  
 10197 code of the entire program (or shared library) fits in 4MB, and is  
 10198 thus reachable with a branch instruction. When `@option{-msmall-data}`  
 10199 is used, the compiler can assume that all local symbols share the  
 10200 same `@code{$gp}` value, and thus reduce the number of instructions  
 10201 required for a function call from 4 to 1.

10203 The default is `@option{-mlarge-text}`.

10205 `@item -mcpu=@var{cpu_type}`  
 10206 `@opindex mcpu`  
 10207 Set the instruction set and instruction scheduling parameters for  
 10208 machine type `@var{cpu_type}`. You can specify either the `@samp{EV}`  
 10209 style name or the corresponding chip number. GCC supports scheduling  
 10210 parameters for the EV4, EV5 and EV6 family of processors and will  
 10211 choose the default values for the instruction set from the processor  
 10212 you specify. If you do not specify a processor type, GCC will default  
 10213 to the processor on which the compiler was built.

10215 Supported values for `@var{cpu_type}` are

10217 `@table @samp`  
 10218 `@item ev4`  
 10219 `@itemx ev45`  
 10220 `@itemx 21064`  
 10221 Schedules as an EV4 and has no instruction set extensions.

10223 `@item ev5`  
 10224 `@itemx 21164`  
 10225 Schedules as an EV5 and has no instruction set extensions.

10227 `@item ev56`  
 10228 `@itemx 21164a`  
 10229 Schedules as an EV5 and supports the BWX extension.

10231 `@item pca56`  
 10232 `@itemx 21164pc`  
 10233 `@itemx 21164PC`  
 10234 Schedules as an EV5 and supports the BWX and MAX extensions.

10236 `@item ev6`  
 10237 `@itemx 21264`  
 10238 Schedules as an EV6 and supports the BWX, FIX, and MAX extensions.

10240 `@item ev67`  
 10241 `@itemx 21264a`  
 10242 Schedules as an EV6 and supports the BWX, CIX, FIX, and MAX extensions.  
 10243 `@end table`

10245 Native Linux/GNU toolchains also support the value `@samp{native}`,  
 10246 which selects the best architecture option for the host processor.  
 10247 `@option{-mcpu=native}` has no effect if GCC does not recognize  
 10248 the processor.

10250 `@item -mtune=@var{cpu_type}`  
 10251 `@opindex mtune`  
 10252 Set only the instruction scheduling parameters for machine type  
 10253 `@var{cpu_type}`. The instruction set is not changed.

10255 Native Linux/GNU toolchains also support the value `@samp{native}`,  
 10256 which selects the best architecture option for the host processor.  
 10257 `@option{-mtune=native}` has no effect if GCC does not recognize  
 10258 the processor.

10260 `@item -mmemory-latency=@var{time}`  
 10261 `@opindex mmemory-latency`  
 10262 Sets the latency the scheduler should assume for typical memory  
 10263 references as seen by the application. This number is highly  
 10264 dependent on the memory access patterns used by the application  
 10265 and the size of the external cache on the machine.

10267 Valid options for `@var{time}` are

10269 `@table @samp`  
 10270 `@item @var{number}`  
 10271 A decimal number representing clock cycles.

10273 `@item L1`  
 10274 `@itemx L2`  
 10275 `@itemx L3`  
 10276 `@itemx main`  
 10277 The compiler contains estimates of the number of clock cycles for  
 10278 ``typical'' EV4 & EV5 hardware for the Level 1, 2 & 3 caches  
 10279 (also called Dcache, Scache, and Bcache), as well as to main memory.  
 10280 Note that L3 is only valid for EV5.

10282 `@end table`  
 10283 `@end table`

10285 `@node DEC Alpha/VMS Options`  
 10286 `@subsection DEC Alpha/VMS Options`

10288 These `@samp{-m}` options are defined for the DEC Alpha/VMS implementations:

10290 `@table @gcctabopt`  
 10291 `@item -mvms-return-codes`  
 10292 `@opindex mvms-return-codes`  
 10293 Return VMS condition codes from main. The default is to return POSIX  
 10294 style condition (e.g. @: error) codes.  
 10295 `@end table`

10297 `@node FR30 Options`  
 10298 `@subsection FR30 Options`  
 10299 `@cindex FR30 Options`

10301 These options are defined specifically for the FR30 port.

10303 `@table @gcctabopt`

10305 `@item -msmall-model`  
 10306 `@opindex msmall-model`  
 10307 Use the small address space model. This can produce smaller code, but  
 10308 it does assume that all symbolic values and addresses will fit into a  
 10309 20-bit range.

10311 `@item -mno-lsim`  
 10312 `@opindex mno-lsim`  
 10313 Assume that run-time support has been provided and so there is no need  
 10314 to include the simulator library (`@file{libsim.a}`) on the linker  
 10315 command line.

10317 `@end table`

10319 `@node FRV Options`  
 10320 `@subsection FRV Options`

10321 @cindex FRV Options

10323 @table @gcctabopt  
 10324 @item -mgpr-32  
 10325 @opindex mgpr-32

10327 Only use the first 32 general purpose registers.

10329 @item -mgpr-64  
 10330 @opindex mgpr-64

10332 Use all 64 general purpose registers.

10334 @item -mfpr-32  
 10335 @opindex mfpr-32

10337 Use only the first 32 floating point registers.

10339 @item -mfpr-64  
 10340 @opindex mfpr-64

10342 Use all 64 floating point registers

10344 @item -mhard-float  
 10345 @opindex mhard-float

10347 Use hardware instructions for floating point operations.

10349 @item -msoft-float  
 10350 @opindex msoft-float

10352 Use library routines for floating point operations.

10354 @item -malloca-cc  
 10355 @opindex malloc-cc

10357 Dynamically allocate condition code registers.

10359 @item -mfixed-cc  
 10360 @opindex mfixed-cc

10362 Do not try to dynamically allocate condition code registers, only  
 10363 use @code{icc0} and @code{fcc0}.

10365 @item -mdword  
 10366 @opindex mdword

10368 Change ABI to use double word insns.

10370 @item -mno-dword  
 10371 @opindex mno-dword

10373 Do not use double word instructions.

10375 @item -mdouble  
 10376 @opindex mdouble

10378 Use floating point double instructions.

10380 @item -mno-double  
 10381 @opindex mno-double

10383 Do not use floating point double instructions.

10385 @item -mmedia  
 10386 @opindex mmedia

10388 Use media instructions.

10390 @item -mno-media  
 10391 @opindex mno-media

10393 Do not use media instructions.

10395 @item -mmuladd  
 10396 @opindex mmuladd

10398 Use multiply and add/subtract instructions.

10400 @item -mno-muladd  
 10401 @opindex mno-muladd

10403 Do not use multiply and add/subtract instructions.

10405 @item -mfdpic  
 10406 @opindex mfdpic

10408 Select the FDPIC ABI, that uses function descriptors to represent  
 10409 pointers to functions. Without any PIC/PIE-related options, it  
 10410 implies @option{-fPIE}. With @option{-fpic} or @option{-fpie}, it  
 10411 assumes GOT entries and small data are within a 12-bit range from the  
 10412 GOT base address; with @option{-fPIC} or @option{-fPIE}, GOT offsets  
 10413 are computed with 32 bits.  
 10414 With a @samp{bfin-elf} target, this option implies @option{-msim}.

10416 @item -minline-plt  
 10417 @opindex minline-plt

10419 Enable inlining of PLT entries in function calls to functions that are  
 10420 not known to bind locally. It has no effect without @option{-mfdpic}.

10421 It's enabled by default if optimizing for speed and compiling for  
 10422 shared libraries (i.e., @option{-fPIC} or @option{-fpie}), or when an  
 10423 optimization option such as @option{-O3} or above is present in the  
 10424 command line.

10426 @item -mTLS  
 10427 @opindex TLS

10429 Assume a large TLS segment when generating thread-local code.

10431 @item -mtls  
 10432 @opindex tls

10434 Do not assume a large TLS segment when generating thread-local code.

10436 @item -mgprel-ro  
 10437 @opindex mgprel-ro

10439 Enable the use of @code{GPREL} relocations in the FDPIC ABI for data  
 10440 that is known to be in read-only sections. It's enabled by default,  
 10441 except for @option{-fpic} or @option{-fpie}: even though it may help  
 10442 make the global offset table smaller, it trades 1 instruction for 4.  
 10443 With @option{-fPIC} or @option{-fPIE}, it trades 3 instructions for 4,  
 10444 one of which may be shared by multiple symbols, and it avoids the need  
 10445 for a GOT entry for the referenced symbol, so it's more likely to be a  
 10446 win. If it is not, @option{-mno-gprel-ro} can be used to disable it.

10448 @item -multilib-library-pic  
 10449 @opindex multilib-library-pic

10451 Link with the (library, not FD) pic libraries. It's implied by  
 10452 @option{-mlibrary-pic}, as well as by @option{-fPIC} and



10453 @option{-fpic} without @option{-mfdpic}. You should never have to use  
 10454 it explicitly.

10456 @item -mlinked-fp  
 10457 @opindexmlinked-fp

10459 Follow the EABI requirement of always creating a frame pointer whenever  
 10460 a stack frame is allocated. This option is enabled by default and can  
 10461 be disabled with @option{-mno-linked-fp}.

10463 @item -mlong-calls  
 10464 @opindexmlinked-calls

10466 Use indirect addressing to call functions outside the current  
 10467 compilation unit. This allows the functions to be placed anywhere  
 10468 within the 32-bit address space.

10470 @item -malign-labels  
 10471 @opindexmlinked-labels

10473 Try to align labels to an 8-byte boundary by inserting nops into the  
 10474 previous packet. This option only has an effect when VLIW packing  
 10475 is enabled. It doesn't create new packets; it merely adds nops to  
 10476 existing ones.

10478 @item -mlibrary-pic  
 10479 @opindexmlinked-pic

10481 Generate position-independent EABI code.

10483 @item -macc-4  
 10484 @opindexmlinked-4

10486 Use only the first four media accumulator registers.

10488 @item -macc-8  
 10489 @opindexmlinked-8

10491 Use all eight media accumulator registers.

10493 @item -mpack  
 10494 @opindexmlinked-pack

10496 Pack VLIW instructions.

10498 @item -mno-pack  
 10499 @opindexmlinked-no-pack

10501 Do not pack VLIW instructions.

10503 @item -mno-eflags  
 10504 @opindexmlinked-no-eflags

10506 Do not mark ABI switches in `e_flags`.

10508 @item -mcond-move  
 10509 @opindexmlinked-cond-move

10511 Enable the use of conditional-move instructions (default).

10513 This switch is mainly for debugging the compiler and will likely be removed  
 10514 in a future version.

10516 @item -mno-cond-move  
 10517 @opindexmlinked-no-cond-move

10519 Disable the use of conditional-move instructions.

10521 This switch is mainly for debugging the compiler and will likely be removed  
 10522 in a future version.

10524 @item -mscc  
 10525 @opindexmlinked-mscc

10527 Enable the use of conditional set instructions (default).

10529 This switch is mainly for debugging the compiler and will likely be removed  
 10530 in a future version.

10532 @item -mno-scc  
 10533 @opindexmlinked-no-scc

10535 Disable the use of conditional set instructions.

10537 This switch is mainly for debugging the compiler and will likely be removed  
 10538 in a future version.

10540 @item -mcond-exec  
 10541 @opindexmlinked-cond-exec

10543 Enable the use of conditional execution (default).

10545 This switch is mainly for debugging the compiler and will likely be removed  
 10546 in a future version.

10548 @item -mno-cond-exec  
 10549 @opindexmlinked-no-cond-exec

10551 Disable the use of conditional execution.

10553 This switch is mainly for debugging the compiler and will likely be removed  
 10554 in a future version.

10556 @item -mvliw-branch  
 10557 @opindexmlinked-vliw-branch

10559 Run a pass to pack branches into VLIW instructions (default).

10561 This switch is mainly for debugging the compiler and will likely be removed  
 10562 in a future version.

10564 @item -mno-vliw-branch  
 10565 @opindexmlinked-no-vliw-branch

10567 Do not run a pass to pack branches into VLIW instructions.

10569 This switch is mainly for debugging the compiler and will likely be removed  
 10570 in a future version.

10572 @item -mmulti-cond-exec  
 10573 @opindexmlinked-multi-cond-exec

10575 Enable optimization of @code{&&} and @code{||} in conditional execution  
 10576 (default).

10578 This switch is mainly for debugging the compiler and will likely be removed  
 10579 in a future version.

10581 @item -mno-multi-cond-exec  
 10582 @opindexmlinked-no-multi-cond-exec

10584 Disable optimization of @code{&&} and @code{||} in conditional execution.

```

10586 This switch is mainly for debugging the compiler and will likely be removed
10587 in a future version.
10589 @item -mnested-cond-exec
10590 @opindex mnested-cond-exec
10592 Enable nested conditional execution optimizations (default).
10594 This switch is mainly for debugging the compiler and will likely be removed
10595 in a future version.
10597 @item -mno-nested-cond-exec
10598 @opindex mno-nested-cond-exec
10600 Disable nested conditional execution optimizations.
10602 This switch is mainly for debugging the compiler and will likely be removed
10603 in a future version.
10605 @item -moptimize-membar
10606 @opindex moptimize-membar
10608 This switch removes redundant @code{membar} instructions from the
10609 compiler generated code. It is enabled by default.
10611 @item -mno-optimize-membar
10612 @opindex mno-optimize-membar
10614 This switch disables the automatic removal of redundant @code{membar}
10615 instructions from the generated code.
10617 @item -mtomcat-stats
10618 @opindex mtomcat-stats
10620 Cause gas to print out tomcat statistics.
10622 @item -mcpu=@var{cpu}
10623 @opindex mcpu
10625 Select the processor type for which to generate code. Possible values are
10626 @samp{frv}, @samp{fr550}, @samp{tomcat}, @samp{fr500}, @samp{fr450},
10627 @samp{fr405}, @samp{fr400}, @samp{fr300} and @samp{simple}.
10629 @end table
10631 @node GNU/Linux Options
10632 @subsection GNU/Linux Options
10634 These @samp{-m} options are defined for GNU/Linux targets:
10636 @table @gcctabopt
10637 @item -mglibc
10638 @opindex mglibc
10639 Use the GNU C library instead of uclibc. This is the default except
10640 on @samp{*-linux-*uclibc*} targets.
10642 @item -muclibc
10643 @opindex muclibc
10644 Use uclibc instead of the GNU C library. This is the default on
10645 @samp{*-linux-*uclibc*} targets.
10646 @end table
10648 @node H8/300 Options
10649 @subsection H8/300 Options

```

```

10651 These @samp{-m} options are defined for the H8/300 implementations:
10653 @table @gcctabopt
10654 @item -mrelax
10655 @opindex mrelax
10656 Shorten some address references at link time, when possible; uses the
10657 linker option @option{-relax}. @xref{H8/300,, @code{ld}} and the H8/300,
10658 ld, Using ld}, for a fuller description.
10660 @item -mh
10661 @opindex mh
10662 Generate code for the H8/300H@.
10664 @item -ms
10665 @opindex ms
10666 Generate code for the H8S@.
10668 @item -mn
10669 @opindex mn
10670 Generate code for the H8S and H8/300H in the normal mode. This switch
10671 must be used either with @option{-mh} or @option{-ms}.
10673 @item -ms2600
10674 @opindex ms2600
10675 Generate code for the H8S/2600. This switch must be used with @option{-ms}.
10677 @item -mint32
10678 @opindex mint32
10679 Make @code{int} data 32 bits by default.
10681 @item -malign-300
10682 @opindex malign-300
10683 On the H8/300H and H8S, use the same alignment rules as for the H8/300.
10684 The default for the H8/300H and H8S is to align longs and floats on 4
10685 byte boundaries.
10686 @option{-malign-300} causes them to be aligned on 2 byte boundaries.
10687 This option has no effect on the H8/300.
10688 @end table
10690 @node HPPA Options
10691 @subsection HPPA Options
10692 @cindex HPPA Options
10694 These @samp{-m} options are defined for the HPPA family of computers:
10696 @table @gcctabopt
10697 @item -march=@var{architecture-type}
10698 @opindex march
10699 Generate code for the specified architecture. The choices for
10700 @var{architecture-type} are @samp{1.0} for PA 1.0, @samp{1.1} for PA
10701 1.1, and @samp{2.0} for PA 2.0 processors. Refer to
10702 @file{/usr/lib/sched.models} on an HP-UX system to determine the proper
10703 architecture option for your machine. Code compiled for lower numbered
10704 architectures will run on higher numbered architectures, but not the
10705 other way around.
10707 @item -mpa-risc-1-0
10708 @itemx -mpa-risc-1-1
10709 @itemx -mpa-risc-2-0
10710 @opindex mpa-risc-1-0
10711 @opindex mpa-risc-1-1
10712 @opindex mpa-risc-2-0
10713 Synonyms for @option{-march=1.0}, @option{-march=1.1}, and @option{-march=2.0} r
10715 @item -mbig-switch
10716 @opindex mbig-switch

```

10717 Generate code suitable for big switch tables. Use this option only if  
 10718 the assembler/linker complain about out of range branches within a switch  
 10719 table.

10721 @item -mjump-in-delay  
 10722 @opindex mjump-in-delay  
 10723 Fill delay slots of function calls with unconditional jump instructions  
 10724 by modifying the return pointer for the function call to be the target  
 10725 of the conditional jump.

10727 @item -mdisable-fpregs  
 10728 @opindex mdisable-fpregs  
 10729 Prevent floating point registers from being used in any manner. This is  
 10730 necessary for compiling kernels which perform lazy context switching of  
 10731 floating point registers. If you use this option and attempt to perform  
 10732 floating point operations, the compiler will abort.

10734 @item -mdisable-indexing  
 10735 @opindex mdisable-indexing  
 10736 Prevent the compiler from using indexing address modes. This avoids some  
 10737 rather obscure problems when compiling MIG generated code under MACH@.

10739 @item -mno-space-regs  
 10740 @opindex mno-space-regs  
 10741 Generate code that assumes the target has no space registers. This allows  
 10742 GCC to generate faster indirect calls and use unscalled index address modes.

10744 Such code is suitable for level 0 PA systems and kernels.

10746 @item -mfast-indirect-calls  
 10747 @opindex mfast-indirect-calls  
 10748 Generate code that assumes calls never cross space boundaries. This  
 10749 allows GCC to emit code which performs faster indirect calls.

10751 This option will not work in the presence of shared libraries or nested  
 10752 functions.

10754 @item -mfixed-range=@var{register-range}  
 10755 @opindex mfixed-range  
 10756 Generate code treating the given register range as fixed registers.  
 10757 A fixed register is one that the register allocator can not use. This is  
 10758 useful when compiling kernel code. A register range is specified as  
 10759 two registers separated by a dash. Multiple register ranges can be  
 10760 specified separated by a comma.

10762 @item -mlong-load-store  
 10763 @opindex mlong-load-store  
 10764 Generate 3-instruction load and store sequences as sometimes required by  
 10765 the HP-UX 10 linker. This is equivalent to the @samp{+k} option to  
 10766 the HP compilers.

10768 @item -mportable-runtime  
 10769 @opindex mportable-runtime  
 10770 Use the portable calling conventions proposed by HP for ELF systems.

10772 @item -mgas  
 10773 @opindex mgas  
 10774 Enable the use of assembler directives only GAS understands.

10776 @item -mschedule=@var{cpu-type}  
 10777 @opindex mschedule  
 10778 Schedule code according to the constraints for the machine type  
 10779 @var{cpu-type}. The choices for @var{cpu-type} are @samp{700},  
 10780 @samp{7100}, @samp{7100LC}, @samp{7200}, @samp{7300} and @samp{8000}. Refer  
 10781 to @file{/usr/lib/sched.models} on an HP-UX system to determine the  
 10782 proper scheduling option for your machine. The default scheduling is

10783 @samp{8000}.

10785 @item -mlinker-opt  
 10786 @opindex mlinker-opt  
 10787 Enable the optimization pass in the HP-UX linker. Note this makes symbolic  
 10788 debugging impossible. It also triggers a bug in the HP-UX 8 and HP-UX 9  
 10789 linkers in which they give bogus error messages when linking some programs.

10791 @item -msoft-float  
 10792 @opindex msoft-float  
 10793 Generate output containing library calls for floating point.  
 10794 @strong{Warning:} the requisite libraries are not available for all HPPA  
 10795 targets. Normally the facilities of the machine's usual C compiler are  
 10796 used, but this cannot be done directly in cross-compilation. You must make  
 10797 your own arrangements to provide suitable library functions for  
 10798 cross-compilation.

10800 @option{-msoft-float} changes the calling convention in the output file;  
 10801 therefore, it is only useful if you compile @emph{all} of a program with  
 10802 this option. In particular, you need to compile @file{libgcc.a}, the  
 10803 library that comes with GCC, with @option{-msoft-float} in order for  
 10804 this to work.

10806 @item -msio  
 10807 @opindex msio  
 10808 Generate the predefine, @code{\_SIO}, for server IO@. The default is  
 10809 @option{-mwsio}. This generates the predefines, @code{\_\_hp9000s700},  
 10810 @code{\_\_hp9000s700\_\_} and @code{\_WSIO}, for workstation IO@. These  
 10811 options are available under HP-UX and HI-UX@.

10813 @item -mgnu-ld  
 10814 @opindex gnu-ld  
 10815 Use GNU ld specific options. This passes @option{-shared} to ld when  
 10816 building a shared library. It is the default when GCC is configured,  
 10817 explicitly or implicitly, with the GNU linker. This option does not  
 10818 have any affect on which ld is called, it only changes what parameters  
 10819 are passed to that ld. The ld that is called is determined by the  
 10820 @option{--with-ld} configure option, GCC's program search path, and  
 10821 finally by the user's @env{PATH}. The linker used by GCC can be printed  
 10822 using @samp{which 'gcc -print-prog-name=ld'}. This option is only available  
 10823 on the 64 bit HP-UX GCC, i.e.@: configured with @samp{hppa\*64\*-\*hpux\*}.

10825 @item -mhp-ld  
 10826 @opindex hp-ld  
 10827 Use HP ld specific options. This passes @option{-b} to ld when building  
 10828 a shared library and passes @option{+AcceptTypeMismatch} to ld on all  
 10829 links. It is the default when GCC is configured, explicitly or  
 10830 implicitly, with the HP linker. This option does not have any affect on  
 10831 which ld is called, it only changes what parameters are passed to that  
 10832 ld. The ld that is called is determined by the @option{--with-ld}  
 10833 configure option, GCC's program search path, and finally by the user's  
 10834 @env{PATH}. The linker used by GCC can be printed using @samp{which  
 10835 'gcc -print-prog-name=ld'}. This option is only available on the 64 bit  
 10836 HP-UX GCC, i.e.@: configured with @samp{hppa\*64\*-\*hpux\*}.

10838 @item -mlong-calls  
 10839 @opindex mno-long-calls  
 10840 Generate code that uses long call sequences. This ensures that a call  
 10841 is always able to reach linker generated stubs. The default is to generate  
 10842 long calls only when the distance from the call site to the beginning  
 10843 of the function or translation unit, as the case may be, exceeds a  
 10844 predefined limit set by the branch type being used. The limits for  
 10845 normal calls are 7,600,000 and 240,000 bytes, respectively for the  
 10846 PA 2.0 and PA 1.X architectures. Sibcalls are always limited at  
 10847 240,000 bytes.

10849 Distances are measured from the beginning of functions when using the  
 10850 `@option{-function-sections}` option, or when using the `@option{-mgas}`  
 10851 and `@option{-mno-portable-runtime}` options together under HP-UX with  
 10852 the SOM linker.

10854 It is normally not desirable to use this option as it will degrade  
 10855 performance. However, it may be useful in large applications,  
 10856 particularly when partial linking is used to build the application.

10858 The types of long calls used depends on the capabilities of the  
 10859 assembler and linker, and the type of code being generated. The  
 10860 impact on systems that support long absolute calls, and long pic  
 10861 symbol-difference or pc-relative calls should be relatively small.  
 10862 However, an indirect call is used on 32-bit ELF systems in pic code  
 10863 and it is quite long.

10865 `@item -munix=@var{unix-std}`  
 10866 `@opindex march`  
 10867 Generate compiler predefines and select a startfile for the specified  
 10868 UNIX standard. The choices for `@var{unix-std}` are `@samp{93}`, `@samp{95}`  
 10869 and `@samp{98}`. `@samp{93}` is supported on all HP-UX versions. `@samp{95}`  
 10870 is available on HP-UX 10.10 and later. `@samp{98}` is available on HP-UX  
 10871 11.11 and later. The default values are `@samp{93}` for HP-UX 10.00,  
 10872 `@samp{95}` for HP-UX 10.10 though to 11.00, and `@samp{98}` for HP-UX 11.11  
 10873 and later.

10875 `@option{-munix=93}` provides the same predefines as GCC 3.3 and 3.4.  
 10876 `@option{-munix=95}` provides additional predefines for `@code{XOPEN_UNIX}`  
 10877 and `@code{XOPEN_SOURCE_EXTENDED}`, and the startfile `@file{unix95.o}`.  
 10878 `@option{-munix=98}` provides additional predefines for `@code{XOPEN_UNIX}`,  
 10879 `@code{XOPEN_SOURCE_EXTENDED}`, `@code{INCLUDE_STDC_A1_SOURCE}` and  
 10880 `@code{INCLUDE_XOPEN_SOURCE_500}`, and the startfile `@file{unix98.o}`.

10882 It is `@emph{important}` to note that this option changes the interfaces  
 10883 for various library routines. It also affects the operational behavior  
 10884 of the C library. Thus, `@emph{extreme}` care is needed in using this  
 10885 option.

10887 Library code that is intended to operate with more than one UNIX  
 10888 standard must test, set and restore the variable `@var{__xpg4_extended_mask}`  
 10889 as appropriate. Most GNU software doesn't provide this capability.

10891 `@item -nolibld`  
 10892 `@opindex nolibld`  
 10893 Suppress the generation of link options to search `libld.sl` when the  
 10894 `@option{-static}` option is specified on HP-UX 10 and later.

10896 `@item -static`  
 10897 `@opindex static`  
 10898 The HP-UX implementation of `setlocale` in `libc` has a dependency on  
 10899 `libld.sl`. There isn't an archive version of `libld.sl`. Thus,  
 10900 when the `@option{-static}` option is specified, special link options  
 10901 are needed to resolve this dependency.

10903 On HP-UX 10 and later, the GCC driver adds the necessary options to  
 10904 link with `libld.sl` when the `@option{-static}` option is specified.  
 10905 This causes the resulting binary to be dynamic. On the 64-bit port,  
 10906 the linkers generate dynamic binaries by default in any case. The  
 10907 `@option{-nolibld}` option can be used to prevent the GCC driver from  
 10908 adding these link options.

10910 `@item -threads`  
 10911 `@opindex threads`  
 10912 Add support for multithreading with the `@dfn{dce thread}` library  
 10913 under HP-UX. This option sets flags for both the preprocessor and  
 10914 linker.

10915 `@end table`

10917 `@node i386 and x86-64 Options`  
 10918 `@subsection Intel 386 and AMD x86-64 Options`  
 10919 `@cindex i386 Options`  
 10920 `@cindex x86-64 Options`  
 10921 `@cindex Intel 386 Options`  
 10922 `@cindex AMD x86-64 Options`

10924 These `@samp{-m}` options are defined for the i386 and x86-64 family of  
 10925 computers:

10927 `@table @gcctabopt`  
 10928 `@item -mtune=@var{cpu-type}`  
 10929 `@opindex mtune`  
 10930 Tune to `@var{cpu-type}` everything applicable about the generated code, except  
 10931 for the ABI and the set of available instructions. The choices for  
 10932 `@var{cpu-type}` are:  
 10933 `@table @emph`  
 10934 `@item generic`  
 10935 Produce code optimized for the most common IA32/AMD64/EM64T processors.  
 10936 If you know the CPU on which your code will run, then you should use  
 10937 the corresponding `@option{-mtune}` option instead of  
 10938 `@option{-mtune=generic}`. But, if you do not know exactly what CPU users  
 10939 of your application will have, then you should use this option.

10941 As new processors are deployed in the marketplace, the behavior of this  
 10942 option will change. Therefore, if you upgrade to a newer version of  
 10943 GCC, the code generated option will change to reflect the processors  
 10944 that were most common when that version of GCC was released.

10946 There is no `@option{-march=generic}` option because `@option{-march}`  
 10947 indicates the instruction set the compiler can use, and there is no  
 10948 generic instruction set applicable to all processors. In contrast,  
 10949 `@option{-mtune}` indicates the processor (or, in this case, collection of  
 10950 processors) for which the code is optimized.

10951 `@item native`  
 10952 This selects the CPU to tune for at compilation time by determining  
 10953 the processor type of the compiling machine. Using `@option{-mtune=native}`  
 10954 will produce code optimized for the local machine under the constraints  
 10955 of the selected instruction set. Using `@option{-march=native}` will  
 10956 enable all instruction subsets supported by the local machine (hence  
 10957 the result might not run on different machines).

10958 `@item i386`  
 10959 Original Intel's i386 CPU@.  
 10960 `@item i486`  
 10961 Intel's i486 CPU@. (No scheduling is implemented for this chip.)  
 10962 `@item i586, pentium`  
 10963 Intel Pentium CPU with no MMX support.  
 10964 `@item pentium-mmx`  
 10965 Intel PentiumMMX CPU based on Pentium core with MMX instruction set support.  
 10966 `@item pentiumpro`  
 10967 Intel PentiumPro CPU@.  
 10968 `@item i686`  
 10969 Same as `@code{generic}`, but when used as `@code{march}` option, PentiumPro  
 10970 instruction set will be used, so the code will run on all i686 family chips.  
 10971 `@item pentium2`  
 10972 Intel Pentium2 CPU based on PentiumPro core with MMX instruction set support.  
 10973 `@item pentium3, pentium3m`  
 10974 Intel Pentium3 CPU based on PentiumPro core with MMX and SSE instruction set  
 10975 support.  
 10976 `@item pentium-m`  
 10977 Low power version of Intel Pentium3 CPU with MMX, SSE and SSE2 instruction set  
 10978 support. Used by Centrino notebooks.  
 10979 `@item pentium4, pentium4m`  
 10980 Intel Pentium4 CPU with MMX, SSE and SSE2 instruction set support.

```

10981 @item prescott
10982 Improved version of Intel Pentium4 CPU with MMX, SSE, SSE2 and SSE3 instruction
10983 set support.
10984 @item nocona
10985 Improved version of Intel Pentium4 CPU with 64-bit extensions, MMX, SSE,
10986 SSE2 and SSE3 instruction set support.
10987 @item core2
10988 Intel Core2 CPU with 64-bit extensions, MMX, SSE, SSE2, SSE3 and SSSE3
10989 instruction set support.
10990 @item k6
10991 AMD K6 CPU with MMX instruction set support.
10992 @item k6-2, k6-3
10993 Improved versions of AMD K6 CPU with MMX and 3dNOW!@: instruction set support.
10994 @item athlon, athlon-tbird
10995 AMD Athlon CPU with MMX, 3dNOW!, enhanced 3dNOW!@: and SSE prefetch instructions
10996 support.
10997 @item athlon-4, athlon-xp, athlon-mp
10998 Improved AMD Athlon CPU with MMX, 3dNOW!, enhanced 3dNOW!@: and full SSE
10999 instruction set support.
11000 @item k8, opteron, athlon64, athlon-fx
11001 AMD K8 core based CPUs with x86-64 instruction set support. (This superset
11002 MMX, SSE, SSE2, 3dNOW!, enhanced 3dNOW!@: and 64-bit instruction set extensions.
11003 @item k8-sse3, opteron-sse3, athlon64-sse3
11004 Improved versions of k8, opteron and athlon64 with SSE3 instruction set support.
11005 @item amd64, barcelona
11006 AMD Family 10h core based CPUs with x86-64 instruction set support. (This
11007 superset MMX, SSE, SSE2, SSE3, SSE4A, 3dNOW!, enhanced 3dNOW!, ABM and 64-bit
11008 instruction set extensions.)
11009 @item winchip-c6
11010 IDT Winchip C6 CPU, dealt in same way as i486 with additional MMX instruction
11011 set support.
11012 @item winchip2
11013 IDT Winchip2 CPU, dealt in same way as i486 with additional MMX and 3dNOW!@:
11014 instruction set support.
11015 @item c3
11016 Via C3 CPU with MMX and 3dNOW!@: instruction set support. (No scheduling is
11017 implemented for this chip.)
11018 @item c3-2
11019 Via C3-2 CPU with MMX and SSE instruction set support. (No scheduling is
11020 implemented for this chip.)
11021 @item geode
11022 Embedded AMD CPU with MMX and 3dNOW! instruction set support.
11023 @end table

11025 While picking a specific @var{cpu-type} will schedule things appropriately
11026 for that particular chip, the compiler will not generate any code that
11027 does not run on the i386 without the @option{-march=@var{cpu-type}} option
11028 being used.

11030 @item -march=@var{cpu-type}
11031 @opindex march
11032 Generate instructions for the machine type @var{cpu-type}. The choices
11033 for @var{cpu-type} are the same as for @option{-mtune}. Moreover,
11034 specifying @option{-march=@var{cpu-type}} implies @option{-mtune=@var{cpu-type}}

11036 @item -mcpu=@var{cpu-type}
11037 @opindex mcpu
11038 A deprecated synonym for @option{-mtune}.

11040 @item -mfpmath=@var{unit}
11041 @opindex march
11042 Generate floating point arithmetics for selected unit @var{unit}. The choices
11043 for @var{unit} are:

11045 @table @samp
11046 @item 387

```

```

11047 Use the standard 387 floating point coprocessor present majority of chips and
11048 emulated otherwise. Code compiled with this option will run almost everywhere.
11049 The temporary results are computed in 80bit precision instead of precision
11050 specified by the type resulting in slightly different results compared to most
11051 of other chips. See @option{-ffloat-store} for more detailed description.

11053 This is the default choice for i386 compiler.

11055 @item sse
11056 Use scalar floating point instructions present in the SSE instruction set.
11057 This instruction set is supported by Pentium3 and newer chips, in the AMD line
11058 by Athlon-4, Athlon-xp and Athlon-mp chips. The earlier version of SSE
11059 instruction set supports only single precision arithmetics, thus the double and
11060 extended precision arithmetics is still done using 387. Later version, present
11061 only in Pentium4 and the future AMD x86-64 chips supports double precision
11062 arithmetics too.

11064 For the i386 compiler, you need to use @option{-march=@var{cpu-type}}, @option{-
11065 or @option{-msse2} switches to enable SSE extensions and make this option
11066 effective. For the x86-64 compiler, these extensions are enabled by default.

11068 The resulting code should be considerably faster in the majority of cases and av
11069 the numerical instability problems of 387 code, but may break some existing
11070 code that expects temporaries to be 80bit.

11072 This is the default choice for the x86-64 compiler.

11074 @item sse,387
11075 @itemx sse+387
11076 @itemx both
11077 Attempt to utilize both instruction sets at once. This effectively double the
11078 amount of available registers and on chips with separate execution units for
11079 387 and SSE the execution resources too. Use this option with care, as it is
11080 still experimental, because the GCC register allocator does not model separate
11081 functional units well resulting in instable performance.
11082 @end table

11084 @item -masm=@var{dialect}
11085 @opindex masm=@var{dialect}
11086 Output asm instructions using selected @var{dialect}. Supported
11087 choices are @samp{intel} or @samp{att} (the default one). Darwin does
11088 not support @samp{intel}.

11090 @item -mieee-fp
11091 @itemx -mno-ieee-fp
11092 @opindex mieee-fp
11093 @opindex mno-ieee-fp
11094 Control whether or not the compiler uses IEEE floating point
11095 comparisons. These handle correctly the case where the result of a
11096 comparison is unordered.

11098 @item -msoft-float
11099 @opindex msoft-float
11100 Generate output containing library calls for floating point.
11101 @strong{Warning:} the requisite libraries are not part of GCC@.
11102 Normally the facilities of the machine's usual C compiler are used, but
11103 this can't be done directly in cross-compilation. You must make your
11104 own arrangements to provide suitable library functions for
11105 cross-compilation.

11107 On machines where a function returns floating point results in the 80387
11108 register stack, some floating point opcodes may be emitted even if
11109 @option{-msoft-float} is used.

11111 @item -mno-fp-ret-in-387
11112 @opindex mno-fp-ret-in-387

```

11113 Do not use the FPU registers for return values of functions.

11115 The usual calling convention has functions return values of types  
 11116 `@code{float}` and `@code{double}` in an FPU register, even if there  
 11117 is no FPU@. The idea is that the operating system should emulate  
 11118 an FPU@.

11120 The option `@option{-mno-fp-ret-in-387}` causes such values to be returned  
 11121 in ordinary CPU registers instead.

11123 `@item -mno-fancy-math-387`  
 11124 `@opindex mno-fancy-math-387`  
 11125 Some 387 emulators do not support the `@code{sin}`, `@code{cos}` and  
 11126 `@code{sqrt}` instructions for the 387. Specify this option to avoid  
 11127 generating those instructions. This option is the default on FreeBSD,  
 11128 OpenBSD and NetBSD@. This option is overridden when `@option{-march}`  
 11129 indicates that the target cpu will always have an FPU and so the  
 11130 instruction will not need emulation. As of revision 2.6.1, these  
 11131 instructions are not generated unless you also use the  
 11132 `@option{-funsafe-math-optimizations}` switch.

11134 `@item -malign-double`  
 11135 `@itemx -mno-align-double`  
 11136 `@opindex mallow-align-double`  
 11137 `@opindex mno-align-double`  
 11138 Control whether GCC aligns `@code{double}`, `@code{long double}`, and  
 11139 `@code{long long}` variables on a two word boundary or a one word  
 11140 boundary. Aligning `@code{double}` variables on a two word boundary will  
 11141 produce code that runs somewhat faster on a `@samp{Pentium}` at the  
 11142 expense of more memory.

11144 On x86-64, `@option{-malign-double}` is enabled by default.

11146 `@strong{Warning:}` if you use the `@option{-malign-double}` switch,  
 11147 structures containing the above types will be aligned differently than  
 11148 the published application binary interface specifications for the 386  
 11149 and will not be binary compatible with structures in code compiled  
 11150 without that switch.

11152 `@item -m96bit-long-double`  
 11153 `@itemx -m128bit-long-double`  
 11154 `@opindex m96bit-long-double`  
 11155 `@opindex m128bit-long-double`  
 11156 These switches control the size of `@code{long double}` type. The i386  
 11157 application binary interface specifies the size to be 96 bits,  
 11158 so `@option{-m96bit-long-double}` is the default in 32 bit mode.

11160 Modern architectures (Pentium and newer) would prefer `@code{long double}`  
 11161 to be aligned to an 8 or 16 byte boundary. In arrays or structures  
 11162 conforming to the ABI, this would not be possible. So specifying a  
 11163 `@option{-m128bit-long-double}` will align `@code{long double}`  
 11164 to a 16 byte boundary by padding the `@code{long double}` with an additional  
 11165 32 bit zero.

11167 In the x86-64 compiler, `@option{-m128bit-long-double}` is the default choice as  
 11168 its ABI specifies that `@code{long double}` is to be aligned on 16 byte boundary.

11170 Notice that neither of these options enable any extra precision over the x87  
 11171 standard of 80 bits for a `@code{long double}`.

11173 `@strong{Warning:}` if you override the default value for your target ABI, the  
 11174 structures and arrays containing `@code{long double}` variables will change  
 11175 their size as well as function calling convention for function taking  
 11176 `@code{long double}` will be modified. Hence they will not be binary  
 11177 compatible with arrays or structures in code compiled without that switch.

11179 `@item -mlarge-data-threshold=@var{number}`  
 11180 `@opindex mlarge-data-threshold=@var{number}`  
 11181 When `@option{-mcmmodel=medium}` is specified, the data greater than  
 11182 `@var{threshold}` are placed in large data section. This value must be the  
 11183 same across all object linked into the binary and defaults to 65535.

11185 `@item -mrtld`  
 11186 `@opindex mrtld`  
 11187 Use a different function-calling convention, in which functions that  
 11188 take a fixed number of arguments return with the `@code{ret}` `@var{num}`  
 11189 instruction, which pops their arguments while returning. This saves one  
 11190 instruction in the caller since there is no need to pop the arguments  
 11191 there.

11193 You can specify that an individual function is called with this calling  
 11194 sequence with the function attribute `@samp{stdcall}`. You can also  
 11195 override the `@option{-mrtld}` option by using the function attribute  
 11196 `@samp{cdecl}`. `@xref{Function Attributes}`.

11198 `@strong{Warning:}` this calling convention is incompatible with the one  
 11199 normally used on Unix, so you cannot use it if you need to call  
 11200 libraries compiled with the Unix compiler.

11202 Also, you must provide function prototypes for all functions that  
 11203 take variable numbers of arguments (including `@code{printf}`);  
 11204 otherwise incorrect code will be generated for calls to those  
 11205 functions.

11207 In addition, seriously incorrect code will result if you call a  
 11208 function with too many arguments. (Normally, extra arguments are  
 11209 harmlessly ignored.)

11211 `@item -mregparm=@var{num}`  
 11212 `@opindex mregparm`  
 11213 Control how many registers are used to pass integer arguments. By  
 11214 default, no registers are used to pass arguments, and at most 3  
 11215 registers can be used. You can control this behavior for a specific  
 11216 function by using the function attribute `@samp{regparm}`.  
 11217 `@xref{Function Attributes}`.

11219 `@strong{Warning:}` if you use this switch, and  
 11220 `@var{num}` is nonzero, then you must build all modules with the same  
 11221 value, including any libraries. This includes the system libraries and  
 11222 startup modules.

11224 `@item -msseregparm`  
 11225 `@opindex msseregparm`  
 11226 Use SSE register passing conventions for float and double arguments  
 11227 and return values. You can control this behavior for a specific  
 11228 function by using the function attribute `@samp{sseregparm}`.  
 11229 `@xref{Function Attributes}`.

11231 `@strong{Warning:}` if you use this switch then you must build all  
 11232 modules with the same value, including any libraries. This includes  
 11233 the system libraries and startup modules.

11235 `@item -mpc32`  
 11236 `@itemx -mpc64`  
 11237 `@itemx -mpc80`  
 11238 `@opindex mpc32`  
 11239 `@opindex mpc64`  
 11240 `@opindex mpc80`

11242 Set 80387 floating-point precision to 32, 64 or 80 bits. When `@option{-mpc32}`  
 11243 is specified, the significands of results of floating-point operations are  
 11244 rounded to 24 bits (single precision); `@option{-mpc64}` rounds the

11245 significands of results of floating-point operations to 53 bits (double  
 11246 precision) and `@option{-mpc80}` rounds the significands of results of  
 11247 floating-point operations to 64 bits (extended double precision), which is  
 11248 the default. When this option is used, floating-point operations in higher  
 11249 precisions are not available to the programmer without setting the FPU  
 11250 control word explicitly.

11252 Setting the rounding of floating-point operations to less than the default  
 11253 80 bits can speed some programs by 2% or more. Note that some mathematical  
 11254 libraries assume that extended precision (80 bit) floating-point operations  
 11255 are enabled by default; routines in such libraries could suffer significant  
 11256 loss of accuracy, typically through so-called "catastrophic cancellation",  
 11257 when this option is used to set the precision to less than extended precision.

11259 `@item -mstackrealign`  
 11260 `@opindex mstackrealign`  
 11261 Realign the stack at entry. On the Intel x86, the `@option{-mstackrealign}`  
 11262 option will generate an alternate prologue and epilogue that realigns the  
 11263 runtime stack if necessary. This supports mixing legacy codes that keep a  
 11264 4-byte aligned stack with modern codes that keep a 16-byte stack for  
 11265 SSE compatibility. See also the attribute `@code{force_align_arg_pointer}`,  
 11266 applicable to individual functions.

11268 `@item -mpreferred-stack-boundary=@var{num}`  
 11269 `@opindex mpreferred-stack-boundary`  
 11270 Attempt to keep the stack boundary aligned to a 2 raised to `@var{num}`  
 11271 byte boundary. If `@option{-mpreferred-stack-boundary}` is not specified,  
 11272 the default is 4 (16 bytes or 128 bits).

11274 `@item -mincoming-stack-boundary=@var{num}`  
 11275 `@opindex mincoming-stack-boundary`  
 11276 Assume the incoming stack is aligned to a 2 raised to `@var{num}` byte  
 11277 boundary. If `@option{-mincoming-stack-boundary}` is not specified,  
 11278 the one specified by `@option{-mpreferred-stack-boundary}` will be used.

11280 On Pentium and PentiumPro, `@code{double}` and `@code{long double}` values  
 11281 should be aligned to an 8 byte boundary (see `@option{-malign-double}`) or  
 11282 suffer significant run time performance penalties. On Pentium III, the  
 11283 Streaming SIMD Extension (SSE) data type `@code{__m128}` may not work  
 11284 properly if it is not 16 byte aligned.

11286 To ensure proper alignment of this values on the stack, the stack boundary  
 11287 must be as aligned as that required by any value stored on the stack.  
 11288 Further, every function must be generated such that it keeps the stack  
 11289 aligned. Thus calling a function compiled with a higher preferred  
 11290 stack boundary from a function compiled with a lower preferred stack  
 11291 boundary will most likely misalign the stack. It is recommended that  
 11292 libraries that use callbacks always use the default setting.

11294 This extra alignment does consume extra stack space, and generally  
 11295 increases code size. Code that is sensitive to stack space usage, such  
 11296 as embedded systems and operating system kernels, may want to reduce the  
 11297 preferred alignment to `@option{-mpreferred-stack-boundary=2}`.

11299 `@item -mmmx`  
 11300 `@itemx -mno-mmx`  
 11301 `@itemx -msse`  
 11302 `@itemx -mno-sse`  
 11303 `@itemx -msse2`  
 11304 `@itemx -mno-sse2`  
 11305 `@itemx -msse3`  
 11306 `@itemx -mno-sse3`  
 11307 `@itemx -msse3`  
 11308 `@itemx -mno-ssse3`  
 11309 `@itemx -msse4.1`  
 11310 `@itemx -mno-sse4.1`

11311 `@itemx -msse4.2`  
 11312 `@itemx -mno-sse4.2`  
 11313 `@itemx -msse4`  
 11314 `@itemx -mno-sse4`  
 11315 `@itemx -mavx`  
 11316 `@itemx -mno-avx`  
 11317 `@itemx -maes`  
 11318 `@itemx -mno-aes`  
 11319 `@itemx -mpclmul`  
 11320 `@itemx -mno-pclmul`  
 11321 `@itemx -msse4a`  
 11322 `@itemx -mno-sse4a`  
 11323 `@itemx -msse5`  
 11324 `@itemx -mno-sse5`  
 11325 `@itemx -m3dnow`  
 11326 `@itemx -mno-3dnow`  
 11327 `@itemx -mpopcnt`  
 11328 `@itemx -mno-popcnt`  
 11329 `@itemx -mabm`  
 11330 `@itemx -mno-abm`  
 11331 `@opindex mmx`  
 11332 `@opindex mno-mmx`  
 11333 `@opindex msse`  
 11334 `@opindex mno-sse`  
 11335 `@opindex m3dnow`  
 11336 `@opindex mno-3dnow`  
 11337 These switches enable or disable the use of instructions in the MMX,  
 11338 SSE, SSE2, SSE3, SSSE3, SSE4.1, AVX, AES, PCLMUL, SSE4A, SSE5, ABM or  
 11339 3DNow!@: extended instruction sets.  
 11340 These extensions are also available as built-in functions: see  
 11341 `@ref{X86 Built-in Functions}`, for details of the functions enabled and  
 11342 disabled by these switches.

11344 To have SSE/SSE2 instructions generated automatically from floating-point  
 11345 code (as opposed to 387 instructions), see `@option{-mfpmath=sse}`.

11347 GCC depresses SSEx instructions when `@option{-mavx}` is used. Instead, it  
 11348 generates new AVX instructions or AVX equivalence for all SSEx instructions  
 11349 when needed.

11351 These options will enable GCC to use these extended instructions in  
 11352 generated code, even without `@option{-mfpmath=sse}`. Applications which  
 11353 perform runtime CPU detection must compile separate files for each  
 11354 supported architecture, using the appropriate flags. In particular,  
 11355 the file containing the CPU detection code should be compiled without  
 11356 these options.

11358 `@item -mclld`  
 11359 `@opindex mclld`  
 11360 This option instructs GCC to emit a `@code{cld}` instruction in the prologue  
 11361 of functions that use string instructions. String instructions depend on  
 11362 the DF flag to select between autoincrement or autodecrement mode. While the  
 11363 ABI specifies the DF flag to be cleared on function entry, some operating  
 11364 systems violate this specification by not clearing the DF flag in their  
 11365 exception dispatchers. The exception handler can be invoked with the DF flag  
 11366 set which leads to wrong direction mode, when string instructions are used.  
 11367 This option can be enabled by default on 32-bit x86 targets by configuring  
 11368 GCC with the `@option{--enable-cld}` configure option. Generation of `@code{cld}`  
 11369 instructions can be suppressed with the `@option{-mno-cld}` compiler option  
 11370 in this case.

11372 `@item -mcx16`  
 11373 `@opindex mcx16`  
 11374 This option will enable GCC to use CMPXCHG16B instruction in generated code.  
 11375 CMPXCHG16B allows for atomic operations on 128-bit double quadword (or oword)  
 11376 data types. This is useful for high resolution counters that could be updated

11377 by multiple processors (or cores). This instruction is generated as part of  
 11378 atomic built-in functions: see @ref{Atomic Builtins} for details.

11380 @item -msahf  
 11381 @opindex msahf  
 11382 This option will enable GCC to use SAHF instruction in generated 64-bit code.  
 11383 Early Intel CPUs with Intel 64 lacked LAHF and SAHF instructions supported  
 11384 by AMD64 until introduction of Pentium 4 G1 step in December 2005. LAHF and  
 11385 SAHF are load and store instructions, respectively, for certain status flags.  
 11386 In 64-bit mode, SAHF instruction is used to optimize @code{fmod}, @code{drem}  
 11387 or @code{remainder} built-in functions: see @ref{Other Builtins} for details.

11389 @item -mrecip  
 11390 @opindex mrecip  
 11391 This option will enable GCC to use RCPSS and RSQRTSS instructions (and their  
 11392 vectorized variants RCPPS and RSQRTPS) with an additional Newton-Raphson step  
 11393 to increase precision instead of DIVSS and SQRTSS (and their vectorized  
 11394 variants) for single precision floating point arguments. These instructions  
 11395 are generated only when @option{-funsafe-math-optimizations} is enabled  
 11396 together with @option{-finite-math-only} and @option{-fno-trapping-math}.  
 11397 Note that while the throughput of the sequence is higher than the throughput  
 11398 of the non-reciprocal instruction, the precision of the sequence can be  
 11399 decreased by up to 2 ulp (i.e. the inverse of 1.0 equals 0.99999994).

11401 @item -mveclibabi=@var{type}  
 11402 @opindex mveclibabi  
 11403 Specifies the ABI type to use for vectorizing intrinsics using an  
 11404 external library. Supported types are @code{svml} for the Intel short  
 11405 vector math library and @code{acml} for the AMD math core library style  
 11406 of interfacing. GCC will currently emit calls to @code{vmlExp2},  
 11407 @code{vmlLn2}, @code{vmlLog102}, @code{vmlPow2}, @code{vmlTanh2},  
 11408 @code{vmlTan2}, @code{vmlAtan2}, @code{vmlAtanh2},  
 11409 @code{vmlCbrt2}, @code{vmlSinh2}, @code{vmlSin2}, @code{vmlAsinh2},  
 11410 @code{vmlAsin2}, @code{vmlCosh2}, @code{vmlCos2}, @code{vmlAcosh2},  
 11411 @code{vmlAcos2}, @code{vmlExp4}, @code{vmlLn4}, @code{vmlLog104},  
 11412 @code{vmlPow4}, @code{vmlTanh4}, @code{vmlTan4},  
 11413 @code{vmlAtan4}, @code{vmlAtanh4}, @code{vmlCbrt4}, @code{vmlSinh4},  
 11414 @code{vmlSin4}, @code{vmlAsinh4}, @code{vmlAsin4}, @code{vmlCosh4},  
 11415 @code{vmlCos4}, @code{vmlAcosh4} and @code{vmlAcos4} for corresponding  
 11416 function type when @option{-mveclibabi=svml} is used and @code{\_\_vrd2\_sin},  
 11417 @code{\_\_vrd2\_cos}, @code{\_\_vrd2\_exp}, @code{\_\_vrd2\_log}, @code{\_\_vrd2\_log2},  
 11418 @code{\_\_vrd2\_log10}, @code{\_\_vrs4\_sinf}, @code{\_\_vrs4\_cosf},  
 11419 @code{\_\_vrs4\_expf}, @code{\_\_vrs4\_logf}, @code{\_\_vrs4\_log2f},  
 11420 @code{\_\_vrs4\_log10f} and @code{\_\_vrs4\_powf} for corresponding function type  
 11421 when @option{-mveclibabi=acml} is used. Both @option{-ftree-vectorize} and  
 11422 @option{-funsafe-math-optimizations} have to be enabled. A SVML or ACML ABI  
 11423 compatible library will have to be specified at link time.

11425 @item -mpush-args  
 11426 @itemx -mno-push-args  
 11427 @opindex mpush-args  
 11428 @opindex mno-push-args  
 11429 Use PUSH operations to store outgoing parameters. This method is shorter  
 11430 and usually equally fast as method using SUB/MOV operations and is enabled  
 11431 by default. In some cases disabling it may improve performance because of  
 11432 improved scheduling and reduced dependencies.

11434 @item -maccumulate-outgoing-args  
 11435 @opindex maccumulate-outgoing-args  
 11436 If enabled, the maximum amount of space required for outgoing arguments will be  
 11437 computed in the function prologue. This is faster on most modern CPUs  
 11438 because of reduced dependencies, improved scheduling and reduced stack usage  
 11439 when preferred stack boundary is not equal to 2. The drawback is a notable  
 11440 increase in code size. This switch implies @option{-mno-push-args}.

11442 @item -mthreads

11443 @opindex mthreads  
 11444 Support thread-safe exception handling on @samp{Mingw32}. Code that relies  
 11445 on thread-safe exception handling must compile and link all code with the  
 11446 @option{-mthreads} option. When compiling, @option{-mthreads} defines  
 11447 @option{-D\_MT}; when linking, it links in a special thread helper library  
 11448 @option{-lmingwthrd} which cleans up per thread exception handling data.

11450 @item -mno-align-stringops  
 11451 @opindex mno-align-stringops  
 11452 Do not align destination of inlined string operations. This switch reduces  
 11453 code size and improves performance in case the destination is already aligned,  
 11454 but GCC doesn't know about it.

11456 @item -minline-all-stringops  
 11457 @opindex minline-all-stringops  
 11458 By default GCC inlines string operations only when destination is known to be  
 11459 aligned at least to 4 byte boundary. This enables more inlining, increase code  
 11460 size, but may improve performance of code that depends on fast memcopy, strlen  
 11461 and memset for short lengths.

11463 @item -minline-stringops-dynamically  
 11464 @opindex minline-stringops-dynamically  
 11465 For string operation of unknown size, inline runtime checks so for small  
 11466 blocks inline code is used, while for large blocks library call is used.

11468 @item -mstringop-strategy=@var{alg}  
 11469 @opindex mstringop-strategy=@var{alg}  
 11470 Overwrite internal decision heuristic about particular algorithm to inline  
 11471 string operation with. The allowed values are @code{rep\_byte},  
 11472 @code{rep\_4byte}, @code{rep\_8byte} for expanding using i386 @code{rep} prefix  
 11473 of specified size, @code{byte\_loop}, @code{loop}, @code{unrolled\_loop} for  
 11474 expanding inline loop, @code{libcall} for always expanding library call.

11476 @item -momit-leaf-frame-pointer  
 11477 @opindex momit-leaf-frame-pointer  
 11478 Don't keep the frame pointer in a register for leaf functions. This  
 11479 avoids the instructions to save, set up and restore frame pointers and  
 11480 makes an extra register available in leaf functions. The option  
 11481 @option{-fomit-frame-pointer} removes the frame pointer for all functions  
 11482 which might make debugging harder.

11484 @item -mtls-direct-seg-refs  
 11485 @itemx -mno-tls-direct-seg-refs  
 11486 @opindex mtls-direct-seg-refs  
 11487 Controls whether TLS variables may be accessed with offsets from the  
 11488 TLS segment register (@code{%gs} for 32-bit, @code{%fs} for 64-bit),  
 11489 or whether the thread base pointer must be added. Whether or not this  
 11490 is legal depends on the operating system, and whether it maps the  
 11491 segment to cover the entire TLS area.

11493 For systems that use GNU libc, the default is on.

11495 @item -mfused-madd  
 11496 @itemx -mno-fused-madd  
 11497 @opindex mfused-madd  
 11498 Enable automatic generation of fused floating point multiply-add instructions  
 11499 if the ISA supports such instructions. The -mfused-madd option is on by  
 11500 default. The fused multiply-add instructions have a different  
 11501 rounding behavior compared to executing a multiply followed by an add.

11503 @item -msse2avx  
 11504 @itemx -mno-sse2avx  
 11505 @opindex msse2avx  
 11506 Specify that the assembler should encode SSE instructions with VEX  
 11507 prefix. The option @option{-mavx} turns this on by default.  
 11508 @end table



```

11510 These @samp{-m} switches are supported in addition to the above
11511 on AMD x86-64 processors in 64-bit environments.

11513 @table @gcctabopt
11514 @item -m32
11515 @itemx -m64
11516 @opindex m32
11517 @opindex m64
11518 Generate code for a 32-bit or 64-bit environment.
11519 The 32-bit environment sets int, long and pointer to 32 bits and
11520 generates code that runs on any i386 system.
11521 The 64-bit environment sets int to 32 bits and long and pointer
11522 to 64 bits and generates code for AMD's x86-64 architecture. For
11523 darwin only the -m64 option turns off the @option{-fno-pic} and
11524 @option{-mdynamic-no-pic} options.

11526 @item -mno-red-zone
11527 @opindex no-red-zone
11528 Do not use a so called red zone for x86-64 code. The red zone is mandated
11529 by the x86-64 ABI, it is a 128-byte area beyond the location of the
11530 stack pointer that will not be modified by signal or interrupt handlers
11531 and therefore can be used for temporary data without adjusting the stack
11532 pointer. The flag @option{-mno-red-zone} disables this red zone.

11534 @item -mcmmodel=small
11535 @opindex mcmmodel=small
11536 Generate code for the small code model: the program and its symbols must
11537 be linked in the lower 2 GB of the address space. Pointers are 64 bits.
11538 Programs can be statically or dynamically linked. This is the default
11539 code model.

11541 @item -mcmmodel=kernel
11542 @opindex mcmmodel=kernel
11543 Generate code for the kernel code model. The kernel runs in the
11544 negative 2 GB of the address space.
11545 This model has to be used for Linux kernel code.

11547 @item -mcmmodel=medium
11548 @opindex mcmmodel=medium
11549 Generate code for the medium model: The program is linked in the lower 2
11550 GB of the address space. Small symbols are also placed there. Symbols
11551 with sizes larger than @option{-mlarge-data-threshold} are put into
11552 large data or bss sections and can be located above 2GB. Programs can
11553 be statically or dynamically linked.

11555 @item -mcmmodel=large
11556 @opindex mcmmodel=large
11557 Generate code for the large model: This model makes no assumptions
11558 about addresses and sizes of sections.

11560 @item -msave-args
11561 @opindex msave-args
11562 Save integer arguments on the stack at function entry.
11563 @end table

11565 @node IA-64 Options
11566 @subsection IA-64 Options
11567 @cindex IA-64 Options

11569 These are the @samp{-m} options defined for the Intel IA-64 architecture.

11571 @table @gcctabopt
11572 @item -mbig-endian
11573 @opindex mbig-endian
11574 Generate code for a big endian target. This is the default for HP-UX@.

```

```

11576 @item -mlittle-endian
11577 @opindex mlittle-endian
11578 Generate code for a little endian target. This is the default for AIX5
11579 and GNU/Linux.

11581 @item -mgnu-as
11582 @itemx -mno-gnu-as
11583 @opindex mgnu-as
11584 @opindex mno-gnu-as
11585 Generate (or don't) code for the GNU assembler. This is the default.
11586 @c Also, this is the default if the configure option @option{--with-gnu-as}
11587 @c is used.

11589 @item -mgnu-ld
11590 @itemx -mno-gnu-ld
11591 @opindex mgnu-ld
11592 @opindex mno-gnu-ld
11593 Generate (or don't) code for the GNU linker. This is the default.
11594 @c Also, this is the default if the configure option @option{--with-gnu-ld}
11595 @c is used.

11597 @item -mno-pic
11598 @opindex mno-pic
11599 Generate code that does not use a global pointer register. The result
11600 is not position independent code, and violates the IA-64 ABI@.

11602 @item -mvolatile-asm-stop
11603 @itemx -mno-volatile-asm-stop
11604 @opindex mvolatile-asm-stop
11605 @opindex mno-volatile-asm-stop
11606 Generate (or don't) a stop bit immediately before and after volatile asm
11607 statements.

11609 @item -mregister-names
11610 @itemx -mno-register-names
11611 @opindex mregister-names
11612 @opindex mno-register-names
11613 Generate (or don't) @samp{in}, @samp{loc}, and @samp{out} register names for
11614 the stacked registers. This may make assembler output more readable.

11616 @item -mno-sdata
11617 @itemx -msdata
11618 @opindex mno-sdata
11619 @opindex msdata
11620 Disable (or enable) optimizations that use the small data section. This may
11621 be useful for working around optimizer bugs.

11623 @item -mconstant-gp
11624 @opindex mconstant-gp
11625 Generate code that uses a single constant global pointer value. This is
11626 useful when compiling kernel code.

11628 @item -mauto-pic
11629 @opindex mauto-pic
11630 Generate code that is self-relocatable. This implies @option{-mconstant-gp}.
11631 This is useful when compiling firmware code.

11633 @item -minline-float-divide-min-latency
11634 @opindex minline-float-divide-min-latency
11635 Generate code for inline divides of floating point values
11636 using the minimum latency algorithm.

11638 @item -minline-float-divide-max-throughput
11639 @opindex minline-float-divide-max-throughput
11640 Generate code for inline divides of floating point values

```

```

11641 using the maximum throughput algorithm.

11643 @item -minline-int-divide-min-latency
11644 @opindex minline-int-divide-min-latency
11645 Generate code for inline divides of integer values
11646 using the minimum latency algorithm.

11648 @item -minline-int-divide-max-throughput
11649 @opindex minline-int-divide-max-throughput
11650 Generate code for inline divides of integer values
11651 using the maximum throughput algorithm.

11653 @item -minline-sqrt-min-latency
11654 @opindex minline-sqrt-min-latency
11655 Generate code for inline square roots
11656 using the minimum latency algorithm.

11658 @item -minline-sqrt-max-throughput
11659 @opindex minline-sqrt-max-throughput
11660 Generate code for inline square roots
11661 using the maximum throughput algorithm.

11663 @item -mno-dwarf2-asm
11664 @itemx -mdwarf2-asm
11665 @opindex mno-dwarf2-asm
11666 @opindex mdwarf2-asm
11667 Don't (or do) generate assembler code for the DWARF2 line number debugging
11668 info. This may be useful when not using the GNU assembler.

11670 @item -mearly-stop-bits
11671 @itemx -mno-early-stop-bits
11672 @opindex mearly-stop-bits
11673 @opindex mno-early-stop-bits
11674 Allow stop bits to be placed earlier than immediately preceding the
11675 instruction that triggered the stop bit. This can improve instruction
11676 scheduling, but does not always do so.

11678 @item -mfixed-range=@var{register-range}
11679 @opindex mfixed-range
11680 Generate code treating the given register range as fixed registers.
11681 A fixed register is one that the register allocator can not use. This is
11682 useful when compiling kernel code. A register range is specified as
11683 two registers separated by a dash. Multiple register ranges can be
11684 specified separated by a comma.

11686 @item -mtls-size=@var{tls-size}
11687 @opindex mtls-size
11688 Specify bit size of immediate TLS offsets. Valid values are 14, 22, and
11689 64.

11691 @item -mtune=@var{cpu-type}
11692 @opindex mtune
11693 Tune the instruction scheduling for a particular CPU, Valid values are
11694 itanium, itanium1, merced, itanium2, and mckinley.

11696 @item -mt
11697 @itemx -pthread
11698 @opindex mt
11699 @opindex pthread
11700 Add support for multithreading using the POSIX threads library. This
11701 option sets flags for both the preprocessor and linker. It does
11702 not affect the thread safety of object code produced by the compiler or
11703 that of libraries supplied with it. These are HP-UX specific flags.

11705 @item -milp32
11706 @itemx -mlp64

```

```

11707 @opindex milp32
11708 @opindex mlp64
11709 Generate code for a 32-bit or 64-bit environment.
11710 The 32-bit environment sets int, long and pointer to 32 bits.
11711 The 64-bit environment sets int to 32 bits and long and pointer
11712 to 64 bits. These are HP-UX specific flags.

11714 @item -mno-sched-br-data-spec
11715 @itemx -msched-br-data-spec
11716 @opindex mno-sched-br-data-spec
11717 @opindex msched-br-data-spec
11718 (Dis/En)able data speculative scheduling before reload.
11719 This will result in generation of the ld.a instructions and
11720 the corresponding check instructions (ld.c / chk.a).
11721 The default is 'disable'.

11723 @item -msched-ar-data-spec
11724 @itemx -mno-sched-ar-data-spec
11725 @opindex msched-ar-data-spec
11726 @opindex mno-sched-ar-data-spec
11727 (En/Dis)able data speculative scheduling after reload.
11728 This will result in generation of the ld.a instructions and
11729 the corresponding check instructions (ld.c / chk.a).
11730 The default is 'enable'.

11732 @item -mno-sched-control-spec
11733 @itemx -msched-control-spec
11734 @opindex mno-sched-control-spec
11735 @opindex msched-control-spec
11736 (Dis/En)able control speculative scheduling. This feature is
11737 available only during region scheduling (i.e.@: before reload).
11738 This will result in generation of the ld.s instructions and
11739 the corresponding check instructions chk.s .
11740 The default is 'disable'.

11742 @item -msched-br-in-data-spec
11743 @itemx -mno-sched-br-in-data-spec
11744 @opindex msched-br-in-data-spec
11745 @opindex mno-sched-br-in-data-spec
11746 (En/Dis)able speculative scheduling of the instructions that
11747 are dependent on the data speculative loads before reload.
11748 This is effective only with @option{-msched-br-data-spec} enabled.
11749 The default is 'enable'.

11751 @item -msched-ar-in-data-spec
11752 @itemx -mno-sched-ar-in-data-spec
11753 @opindex msched-ar-in-data-spec
11754 @opindex mno-sched-ar-in-data-spec
11755 (En/Dis)able speculative scheduling of the instructions that
11756 are dependent on the data speculative loads after reload.
11757 This is effective only with @option{-msched-ar-data-spec} enabled.
11758 The default is 'enable'.

11760 @item -msched-in-control-spec
11761 @itemx -mno-sched-in-control-spec
11762 @opindex msched-in-control-spec
11763 @opindex mno-sched-in-control-spec
11764 (En/Dis)able speculative scheduling of the instructions that
11765 are dependent on the control speculative loads.
11766 This is effective only with @option{-msched-control-spec} enabled.
11767 The default is 'enable'.

11769 @item -msched-ldc
11770 @itemx -mno-sched-ldc
11771 @opindex msched-ldc
11772 @opindex mno-sched-ldc

```

```

11773 (En/Dis)able use of simple data speculation checks ld.c .
11774 If disabled, only chk.a instructions will be emitted to check
11775 data speculative loads.
11776 The default is 'enable'.

11778 @item -mno-sched-control-ldc
11779 @itemx -msched-control-ldc
11780 @opindex mno-sched-control-ldc
11781 @opindex msched-control-ldc
11782 (Dis/En)able use of ld.c instructions to check control speculative loads.
11783 If enabled, in case of control speculative load with no speculatively
11784 scheduled dependent instructions this load will be emitted as ld.sa and
11785 ld.c will be used to check it.
11786 The default is 'disable'.

11788 @item -mno-sched-spec-verbose
11789 @itemx -msched-spec-verbose
11790 @opindex mno-sched-spec-verbose
11791 @opindex msched-spec-verbose
11792 (Dis/En)able printing of the information about speculative motions.

11794 @item -mno-sched-prefer-non-data-spec-insns
11795 @itemx -msched-prefer-non-data-spec-insns
11796 @opindex mno-sched-prefer-non-data-spec-insns
11797 @opindex msched-prefer-non-data-spec-insns
11798 If enabled, data speculative instructions will be chosen for schedule
11799 only if there are no other choices at the moment. This will make
11800 the use of the data speculation much more conservative.
11801 The default is 'disable'.

11803 @item -mno-sched-prefer-non-control-spec-insns
11804 @itemx -msched-prefer-non-control-spec-insns
11805 @opindex mno-sched-prefer-non-control-spec-insns
11806 @opindex msched-prefer-non-control-spec-insns
11807 If enabled, control speculative instructions will be chosen for schedule
11808 only if there are no other choices at the moment. This will make
11809 the use of the control speculation much more conservative.
11810 The default is 'disable'.

11812 @item -mno-sched-count-spec-in-critical-path
11813 @itemx -msched-count-spec-in-critical-path
11814 @opindex mno-sched-count-spec-in-critical-path
11815 @opindex msched-count-spec-in-critical-path
11816 If enabled, speculative dependencies will be considered during
11817 computation of the instructions priorities. This will make the use of the
11818 speculation a bit more conservative.
11819 The default is 'disable'.

11821 @end table

11823 @node M32C Options
11824 @subsection M32C Options
11825 @cindex M32C options

11827 @table @gcctabopt
11828 @item -mcpu=@var{name}
11829 @opindex mcpu=
11830 Select the CPU for which code is generated. @var{name} may be one of
11831 @samp{r8c} for the R8C/Tiny series, @samp{m16c} for the M16C (up to
11832 /60) series, @samp{m32cm} for the M16C/80 series, or @samp{m32c} for
11833 the M32C/80 series.

11835 @item -msim
11836 @opindex msim
11837 Specifies that the program will be run on the simulator. This causes
11838 an alternate runtime library to be linked in which supports, for

```

```

11839 example, file I/O@. You must not use this option when generating
11840 programs that will run on real hardware; you must provide your own
11841 runtime library for whatever I/O functions are needed.

11843 @item -memregs=@var{number}
11844 @opindex memregs=
11845 Specifies the number of memory-based pseudo-registers GCC will use
11846 during code generation. These pseudo-registers will be used like real
11847 registers, so there is a tradeoff between GCC's ability to fit the
11848 code into available registers, and the performance penalty of using
11849 memory instead of registers. Note that all modules in a program must
11850 be compiled with the same value for this option. Because of that, you
11851 must not use this option with the default runtime libraries gcc
11852 builds.

11854 @end table

11856 @node M32R/D Options
11857 @subsection M32R/D Options
11858 @cindex M32R/D options

11860 These @option{-m} options are defined for Renesas M32R/D architectures:

11862 @table @gcctabopt
11863 @item -m32r2
11864 @opindex m32r2
11865 Generate code for the M32R/2@.

11867 @item -m32rx
11868 @opindex m32rx
11869 Generate code for the M32R/X@.

11871 @item -m32r
11872 @opindex m32r
11873 Generate code for the M32R@. This is the default.

11875 @item -mmodel=small
11876 @opindex mmodel=small
11877 Assume all objects live in the lower 16MB of memory (so that their addresses
11878 can be loaded with the @code{ld24} instruction), and assume all subroutines
11879 are reachable with the @code{bl} instruction.
11880 This is the default.

11882 The addressability of a particular object can be set with the
11883 @code{model} attribute.

11885 @item -mmodel=medium
11886 @opindex mmodel=medium
11887 Assume objects may be anywhere in the 32-bit address space (the compiler
11888 will generate @code{seth/add3} instructions to load their addresses), and
11889 assume all subroutines are reachable with the @code{bl} instruction.

11891 @item -mmodel=large
11892 @opindex mmodel=large
11893 Assume objects may be anywhere in the 32-bit address space (the compiler
11894 will generate @code{seth/add3} instructions to load their addresses), and
11895 assume subroutines may not be reachable with the @code{bl} instruction
11896 (the compiler will generate the much slower @code{seth/add3/jl}
11897 instruction sequence).

11899 @item -msdata=none
11900 @opindex msdata=none
11901 Disable use of the small data area. Variables will be put into
11902 one of @samp{.data}, @samp{.bss}, or @samp{.rodata} (unless the
11903 @code{section} attribute has been specified).
11904 This is the default.

```

```

11906 The small data area consists of sections @samp{.sdata} and @samp{.sbss}.
11907 Objects may be explicitly put in the small data area with the
11908 @code{section} attribute using one of these sections.

11910 @item -msdata=sdata
11911 @opindex msdata=sdata
11912 Put small global and static data in the small data area, but do not
11913 generate special code to reference them.

11915 @item -msdata=use
11916 @opindex msdata=use
11917 Put small global and static data in the small data area, and generate
11918 special instructions to reference them.

11920 @item -G @var{num}
11921 @opindex G
11922 @cindex smaller data references
11923 Put global and static objects less than or equal to @var{num} bytes
11924 into the small data or bss sections instead of the normal data or bss
11925 sections. The default value of @var{num} is 8.
11926 The @option{-msdata} option must be set to one of @samp{sdata} or @samp{use}
11927 for this option to have any effect.

11929 All modules should be compiled with the same @option{-G @var{num}} value.
11930 Compiling with different values of @var{num} may or may not work; if it
11931 doesn't the linker will give an error message---incorrect code will not be
11932 generated.

11934 @item -mdebug
11935 @opindex mdebug
11936 Makes the M32R specific code in the compiler display some statistics
11937 that might help in debugging programs.

11939 @item -malign-loops
11940 @opindex mallow-align-loops
11941 Align all loops to a 32-byte boundary.

11943 @item -mno-align-loops
11944 @opindex mno-align-loops
11945 Do not enforce a 32-byte alignment for loops. This is the default.

11947 @item -missue-rate=@var{number}
11948 @opindex missue-rate=@var{number}
11949 Issue @var{number} instructions per cycle. @var{number} can only be 1
11950 or 2.

11952 @item -mbranch-cost=@var{number}
11953 @opindex mbranch-cost=@var{number}
11954 @var{number} can only be 1 or 2. If it is 1 then branches will be
11955 preferred over conditional code, if it is 2, then the opposite will
11956 apply.

11958 @item -mflush-trap=@var{number}
11959 @opindex mflush-trap=@var{number}
11960 Specifies the trap number to use to flush the cache. The default is
11961 12. Valid numbers are between 0 and 15 inclusive.

11963 @item -mno-flush-trap
11964 @opindex mno-flush-trap
11965 Specifies that the cache cannot be flushed by using a trap.

11967 @item -mflush-func=@var{name}
11968 @opindex mflush-func=@var{name}
11969 Specifies the name of the operating system function to call to flush
11970 the cache. The default is @code{flush_cache}, but a function call

```

```

11971 will only be used if a trap is not available.

11973 @item -mno-flush-func
11974 @opindex mno-flush-func
11975 Indicates that there is no OS function for flushing the cache.

11977 @end table

11979 @node M680x0 Options
11980 @subsection M680x0 Options
11981 @cindex M680x0 options

11983 These are the @samp{-m} options defined for M680x0 and ColdFire processors.
11984 The default settings depend on which architecture was selected when
11985 the compiler was configured; the defaults for the most common choices
11986 are given below.

11988 @table @gcctabopt
11989 @item -march=@var{arch}
11990 @opindex march
11991 Generate code for a specific M680x0 or ColdFire instruction set
11992 architecture. Permissible values of @var{arch} for M680x0
11993 architectures are: @samp{68000}, @samp{68010}, @samp{68020},
11994 @samp{68030}, @samp{68040}, @samp{68060} and @samp{cpu32}. ColdFire
11995 architectures are selected according to Freescale's ISA classification
11996 and the permissible values are: @samp{isaa}, @samp{isaaplus},
11997 @samp{isab} and @samp{isac}.

11999 gcc defines a macro @samp{__mcf@var{arch}} whenever it is generating
12000 code for a ColdFire target. The @var{arch} in this macro is one of the
12001 @option{-march} arguments given above.

12003 When used together, @option{-march} and @option{-mtune} select code
12004 that runs on a family of similar processors but that is optimized
12005 for a particular microarchitecture.

12007 @item -mcpu=@var{cpu}
12008 @opindex mcpu
12009 Generate code for a specific M680x0 or ColdFire processor.
12010 The M680x0 @var{cpu}s are: @samp{68000}, @samp{68010}, @samp{68020},
12011 @samp{68030}, @samp{68040}, @samp{68060}, @samp{68302}, @samp{68332}
12012 and @samp{cpu32}. The ColdFire @var{cpu}s are given by the table
12013 below, which also classifies the CPUs into families:

12015 @multitable @columnfractions 0.20 0.80
12016 @item @strong{Family} @tab @strong{@samp{-mcpu} arguments}
12017 @item @samp{51qe} @tab @samp{51qe}
12018 @item @samp{5206} @tab @samp{5202} @samp{5204} @samp{5206}
12019 @item @samp{5206e} @tab @samp{5206e}
12020 @item @samp{5208} @tab @samp{5207} @samp{5208}
12021 @item @samp{5211a} @tab @samp{5210a} @samp{5211a}
12022 @item @samp{5213} @tab @samp{5211} @samp{5212} @samp{5213}
12023 @item @samp{5216} @tab @samp{5214} @samp{5216}
12024 @item @samp{52235} @tab @samp{52230} @samp{52231} @samp{52232} @samp{52233} @sam
12025 @item @samp{5225} @tab @samp{5224} @samp{5225}
12026 @item @samp{5235} @tab @samp{5232} @samp{5233} @samp{5234} @samp{5235} @samp{523
12027 @item @samp{5249} @tab @samp{5249}
12028 @item @samp{5250} @tab @samp{5250}
12029 @item @samp{5271} @tab @samp{5270} @samp{5271}
12030 @item @samp{5272} @tab @samp{5272}
12031 @item @samp{5275} @tab @samp{5274} @samp{5275}
12032 @item @samp{5282} @tab @samp{5280} @samp{5281} @samp{5282} @samp{528x}
12033 @item @samp{5307} @tab @samp{5307}
12034 @item @samp{5329} @tab @samp{5327} @samp{5328} @samp{5329} @samp{532x}
12035 @item @samp{5373} @tab @samp{5372} @samp{5373} @samp{537x}
12036 @item @samp{5407} @tab @samp{5407}

```

```

12037 @item @samp{5475} @tab @samp{5470} @samp{5471} @samp{5472} @samp{5473} @samp{547}
12038 @end multitable

12040 @option{-mcpu=@var{cpu}} overrides @option{-march=@var{arch}} if
12041 @var{arch} is compatible with @var{cpu}. Other combinations of
12042 @option{-mcpu} and @option{-march} are rejected.

12044 gcc defines the macro @samp{__mcf_cpu_@var{cpu}} when ColdFire target
12045 @var{cpu} is selected. It also defines @samp{__mcf_family_@var{family}},
12046 where the value of @var{family} is given by the table above.

12048 @item -mtune=@var{tune}
12049 @opindex mtune
12050 Tune the code for a particular microarchitecture, within the
12051 constraints set by @option{-march} and @option{-mcpu}.
12052 The M680x0 microarchitectures are: @samp{68000}, @samp{68010},
12053 @samp{68020}, @samp{68030}, @samp{68040}, @samp{68060}
12054 and @samp{cpu32}. The ColdFire microarchitectures
12055 are: @samp{cfv1}, @samp{cfv2}, @samp{cfv3}, @samp{cfv4} and @samp{cfv4e}.

12057 You can also use @option{-mtune=68020-40} for code that needs
12058 to run relatively well on 68020, 68030 and 68040 targets.
12059 @option{-mtune=68020-60} is similar but includes 68060 targets
12060 as well. These two options select the same tuning decisions as
12061 @option{-m68020-40} and @option{-m68020-60} respectively.

12063 gcc defines the macros @samp{__mc@var{arch}} and @samp{__mc@var{arch}__}
12064 when tuning for 680x0 architecture @var{arch}. It also defines
12065 @samp{mc@var{arch}} unless either @option{-ansi} or a non-GNU @option{-std}
12066 option is used. If gcc is tuning for a range of architectures,
12067 as selected by @option{-mtune=68020-40} or @option{-mtune=68020-60},
12068 it defines the macros for every architecture in the range.

12070 gcc also defines the macro @samp{__m@var{uarch}__} when tuning for
12071 ColdFire microarchitecture @var{uarch}, where @var{uarch} is one
12072 of the arguments given above.

12074 @item -m68000
12075 @itemx -mc68000
12076 @opindex m68000
12077 @opindex mc68000
12078 Generate output for a 68000. This is the default
12079 when the compiler is configured for 68000-based systems.
12080 It is equivalent to @option{-march=68000}.

12082 Use this option for microcontrollers with a 68000 or EC000 core,
12083 including the 68008, 68302, 68306, 68307, 68322, 68328 and 68356.

12085 @item -m68010
12086 @opindex m68010
12087 Generate output for a 68010. This is the default
12088 when the compiler is configured for 68010-based systems.
12089 It is equivalent to @option{-march=68010}.

12091 @item -m68020
12092 @itemx -mc68020
12093 @opindex m68020
12094 @opindex mc68020
12095 Generate output for a 68020. This is the default
12096 when the compiler is configured for 68020-based systems.
12097 It is equivalent to @option{-march=68020}.

12099 @item -m68030
12100 @opindex m68030
12101 Generate output for a 68030. This is the default when the compiler is
12102 configured for 68030-based systems. It is equivalent to

```

```

12103 @option{-march=68030}.

12105 @item -m68040
12106 @opindex m68040
12107 Generate output for a 68040. This is the default when the compiler is
12108 configured for 68040-based systems. It is equivalent to
12109 @option{-march=68040}.

12111 This option inhibits the use of 68881/68882 instructions that have to be
12112 emulated by software on the 68040. Use this option if your 68040 does not
12113 have code to emulate those instructions.

12115 @item -m68060
12116 @opindex m68060
12117 Generate output for a 68060. This is the default when the compiler is
12118 configured for 68060-based systems. It is equivalent to
12119 @option{-march=68060}.

12121 This option inhibits the use of 68020 and 68881/68882 instructions that
12122 have to be emulated by software on the 68060. Use this option if your 68060
12123 does not have code to emulate those instructions.

12125 @item -mcpu32
12126 @opindex mcpu32
12127 Generate output for a CPU32. This is the default
12128 when the compiler is configured for CPU32-based systems.
12129 It is equivalent to @option{-march=cpu32}.

12131 Use this option for microcontrollers with a
12132 CPU32 or CPU32+ core, including the 68330, 68331, 68332, 68333, 68334,
12133 68336, 68340, 68341, 68349 and 68360.

12135 @item -m5200
12136 @opindex m5200
12137 Generate output for a 520X ColdFire CPU@. This is the default
12138 when the compiler is configured for 520X-based systems.
12139 It is equivalent to @option{-mcpu=5206}, and is now deprecated
12140 in favor of that option.

12142 Use this option for microcontroller with a 5200 core, including
12143 the MCF5202, MCF5203, MCF5204 and MCF5206.

12145 @item -m5206e
12146 @opindex m5206e
12147 Generate output for a 5206e ColdFire CPU@. The option is now
12148 deprecated in favor of the equivalent @option{-mcpu=5206e}.

12150 @item -m528x
12151 @opindex m528x
12152 Generate output for a member of the ColdFire 528X family.
12153 The option is now deprecated in favor of the equivalent
12154 @option{-mcpu=528x}.

12156 @item -m5307
12157 @opindex m5307
12158 Generate output for a ColdFire 5307 CPU@. The option is now deprecated
12159 in favor of the equivalent @option{-mcpu=5307}.

12161 @item -m5407
12162 @opindex m5407
12163 Generate output for a ColdFire 5407 CPU@. The option is now deprecated
12164 in favor of the equivalent @option{-mcpu=5407}.

12166 @item -mcfv4e
12167 @opindex mcfv4e
12168 Generate output for a ColdFire V4e family CPU (e.g.: 547x/548x).

```

12169 This includes use of hardware floating point instructions.  
 12170 The option is equivalent to `@option{-mcpu=547x}`, and is now  
 12171 deprecated in favor of that option.

12173 `@item -m68020-40`  
 12174 `@opindex m68020-40`  
 12175 Generate output for a 68040, without using any of the new instructions.  
 12176 This results in code which can run relatively efficiently on either a  
 12177 68020/68881 or a 68030 or a 68040. The generated code does use the  
 12178 68881 instructions that are emulated on the 68040.

12180 The option is equivalent to `@option{-march=68020} @option{-mtune=68020-40}`.

12182 `@item -m68020-60`  
 12183 `@opindex m68020-60`  
 12184 Generate output for a 68060, without using any of the new instructions.  
 12185 This results in code which can run relatively efficiently on either a  
 12186 68020/68881 or a 68030 or a 68040. The generated code does use the  
 12187 68881 instructions that are emulated on the 68060.

12189 The option is equivalent to `@option{-march=68020} @option{-mtune=68020-60}`.

12191 `@item -mhard-float`  
 12192 `@itemx -m68881`  
 12193 `@opindex mhard-float`  
 12194 `@opindex m68881`  
 12195 Generate floating-point instructions. This is the default for 68020  
 12196 and above, and for ColdFire devices that have an FPU@. It defines the  
 12197 macro `@samp{__HAVE_68881__}` on M680x0 targets and `@samp{__mcfcpu__}`  
 12198 on ColdFire targets.

12200 `@item -msoft-float`  
 12201 `@opindex msoft-float`  
 12202 Do not generate floating-point instructions; use library calls instead.  
 12203 This is the default for 68000, 68010, and 68832 targets. It is also  
 12204 the default for ColdFire devices that have no FPU.

12206 `@item -mdiv`  
 12207 `@itemx -mno-div`  
 12208 `@opindex mdiv`  
 12209 `@opindex mno-div`  
 12210 Generate (do not generate) ColdFire hardware divide and remainder  
 12211 instructions. If `@option{-march}` is used without `@option{-mcpu}`,  
 12212 the default is `'on'` for ColdFire architectures and `'off'` for M680x0  
 12213 architectures. Otherwise, the default is taken from the target CPU  
 12214 (either the default CPU, or the one specified by `@option{-mcpu}`). For  
 12215 example, the default is `'off'` for `@option{-mcpu=5206}` and `'on'` for  
 12216 `@option{-mcpu=5206e}`.

12218 gcc defines the macro `@samp{__mcfhwdiv__}` when this option is enabled.

12220 `@item -mshort`  
 12221 `@opindex mshort`  
 12222 Consider type `@code{int}` to be 16 bits wide, like `@code{short int}`.  
 12223 Additionally, parameters passed on the stack are also aligned to a  
 12224 16-bit boundary even on targets whose API mandates promotion to 32-bit.

12226 `@item -mno-short`  
 12227 `@opindex mno-short`  
 12228 Do not consider type `@code{int}` to be 16 bits wide. This is the default.

12230 `@item -mnobitfield`  
 12231 `@itemx -mno-bitfield`  
 12232 `@opindex mnobitfield`  
 12233 `@opindex mno-bitfield`  
 12234 Do not use the bit-field instructions. The `@option{-m68000}`, `@option{-mcpu32}`

12235 and `@option{-m5200}` options imply `@w{@option{-mnobitfield}}`.

12237 `@item -mbitfield`  
 12238 `@opindex mbitfield`  
 12239 Do use the bit-field instructions. The `@option{-m68020}` option implies  
 12240 `@option{-mbitfield}`. This is the default if you use a configuration  
 12241 designed for a 68020.

12243 `@item -mrtd`  
 12244 `@opindex mrtd`  
 12245 Use a different function-calling convention, in which functions  
 12246 that take a fixed number of arguments return with the `@code{rtd}`  
 12247 instruction, which pops their arguments while returning. This  
 12248 saves one instruction in the caller since there is no need to pop  
 12249 the arguments there.

12251 This calling convention is incompatible with the one normally  
 12252 used on Unix, so you cannot use it if you need to call libraries  
 12253 compiled with the Unix compiler.

12255 Also, you must provide function prototypes for all functions that  
 12256 take variable numbers of arguments (including `@code{printf}`);  
 12257 otherwise incorrect code will be generated for calls to those  
 12258 functions.

12260 In addition, seriously incorrect code will result if you call a  
 12261 function with too many arguments. (Normally, extra arguments are  
 12262 harmlessly ignored.)

12264 The `@code{rtd}` instruction is supported by the 68010, 68020, 68030,  
 12265 68040, 68060 and CPU32 processors, but not by the 68000 or 5200.

12267 `@item -mno-rtd`  
 12268 `@opindex mno-rtd`  
 12269 Do not use the calling conventions selected by `@option{-mrtd}`.  
 12270 This is the default.

12272 `@item -malign-int`  
 12273 `@itemx -mno-align-int`  
 12274 `@opindex malign-int`  
 12275 `@opindex mno-align-int`  
 12276 Control whether GCC aligns `@code{int}`, `@code{long}`, `@code{long long}`,  
 12277 `@code{float}`, `@code{double}`, and `@code{long double}` variables on a 32-bit  
 12278 boundary (`@option{-malign-int}`) or a 16-bit boundary (`@option{-mno-align-int}`).  
 12279 Aligning variables on 32-bit boundaries produces code that runs somewhat  
 12280 faster on processors with 32-bit busses at the expense of more memory.

12282 `@strong{Warning:}` if you use the `@option{-malign-int}` switch, GCC will  
 12283 align structures containing the above types differently than  
 12284 most published application binary interface specifications for the m68k.

12286 `@item -mpcrel`  
 12287 `@opindex mprel`  
 12288 Use the pc-relative addressing mode of the 68000 directly, instead of  
 12289 using a global offset table. At present, this option implies `@option{-fpic}`,  
 12290 allowing at most a 16-bit offset for pc-relative addressing. `@option{-fPIC}` is  
 12291 not presently supported with `@option{-mpcrel}`, though this could be supported fo  
 12292 68020 and higher processors.

12294 `@item -mno-strict-align`  
 12295 `@itemx -mstrict-align`  
 12296 `@opindex mno-strict-align`  
 12297 `@opindex mstrict-align`  
 12298 Do not (do) assume that unaligned memory references will be handled by  
 12299 the system.

```

12301 @item -msep-data
12302 Generate code that allows the data segment to be located in a different
12303 area of memory from the text segment. This allows for execute in place in
12304 an environment without virtual memory management. This option implies
12305 @option{-fPIC}.

12307 @item -mno-sep-data
12308 Generate code that assumes that the data segment follows the text segment.
12309 This is the default.

12311 @item -mid-shared-library
12312 Generate code that supports shared libraries via the library ID method.
12313 This allows for execute in place and shared libraries in an environment
12314 without virtual memory management. This option implies @option{-fPIC}.

12316 @item -mno-id-shared-library
12317 Generate code that doesn't assume ID based shared libraries are being used.
12318 This is the default.

12320 @item -mshared-library-id=n
12321 Specified the identification number of the ID based shared library being
12322 compiled. Specifying a value of 0 will generate more compact code, specifying
12323 other values will force the allocation of that number to the current
12324 library but is no more space or time efficient than omitting this option.

12326 @item -mxgot
12327 @itemx -mno-xgot
12328 @opindex mxgot
12329 @opindex mno-xgot
12330 When generating position-independent code for ColdFire, generate code
12331 that works if the GOT has more than 8192 entries. This code is
12332 larger and slower than code generated without this option. On M680x0
12333 processors, this option is not needed; @option{-fPIC} suffices.

12335 GCC normally uses a single instruction to load values from the GOT@.
12336 While this is relatively efficient, it only works if the GOT
12337 is smaller than about 64k. Anything larger causes the linker
12338 to report an error such as:

12340 @cindex relocation truncated to fit (ColdFire)
12341 @smallexample
12342 relocation truncated to fit: R_68K_GOT160 foobar
12343 @end smallexample

12345 If this happens, you should recompile your code with @option{-mxgot}.
12346 It should then work with very large GOTs. However, code generated with
12347 @option{-mxgot} is less efficient, since it takes 4 instructions to fetch
12348 the value of a global symbol.

12350 Note that some linkers, including newer versions of the GNU linker,
12351 can create multiple GOTs and sort GOT entries. If you have such a linker,
12352 you should only need to use @option{-mxgot} when compiling a single
12353 object file that accesses more than 8192 GOT entries. Very few do.

12355 These options have no effect unless GCC is generating
12356 position-independent code.

12358 @end table

12360 @node M68hc1x Options
12361 @subsection M68hc1x Options
12362 @cindex M68hc1x options

12364 These are the @samp{-m} options defined for the 68hc11 and 68hc12
12365 microcontrollers. The default values for these options depends on
12366 which style of microcontroller was selected when the compiler was configured;

```

```

12367 the defaults for the most common choices are given below.

12369 @table @gcctabopt
12370 @item -m6811
12371 @itemx -m68hc11
12372 @opindex m6811
12373 @opindex m68hc11
12374 Generate output for a 68HC11. This is the default
12375 when the compiler is configured for 68HC11-based systems.

12377 @item -m6812
12378 @itemx -m68hc12
12379 @opindex m6812
12380 @opindex m68hc12
12381 Generate output for a 68HC12. This is the default
12382 when the compiler is configured for 68HC12-based systems.

12384 @item -m68S12
12385 @itemx -m68hcs12
12386 @opindex m68S12
12387 @opindex m68hcs12
12388 Generate output for a 68HCS12.

12390 @item -mauto-incdec
12391 @opindex mauto-incdec
12392 Enable the use of 68HC12 pre and post auto-increment and auto-decrement
12393 addressing modes.

12395 @item -minmax
12396 @itemx -nominmax
12397 @opindex minmax
12398 @opindex mnominmax
12399 Enable the use of 68HC12 min and max instructions.

12401 @item -mlong-calls
12402 @itemx -mno-long-calls
12403 @opindex mlong-calls
12404 @opindex mno-long-calls
12405 Treat all calls as being far away (near). If calls are assumed to be
12406 far away, the compiler will use the @code{call} instruction to
12407 call a function and the @code{rtc} instruction for returning.

12409 @item -mshort
12410 @opindex mshort
12411 Consider type @code{int} to be 16 bits wide, like @code{short int}.

12413 @item -msoft-reg-count=@var{count}
12414 @opindex msoft-reg-count
12415 Specify the number of pseudo-soft registers which are used for the
12416 code generation. The maximum number is 32. Using more pseudo-soft
12417 register may or may not result in better code depending on the program.
12418 The default is 4 for 68HC11 and 2 for 68HC12.

12420 @end table

12422 @node MCore Options
12423 @subsection MCore Options
12424 @cindex MCore options

12426 These are the @samp{-m} options defined for the Motorola M*Core
12427 processors.

12429 @table @gcctabopt

12431 @item -mhardlit
12432 @itemx -mno-hardlit

```

```

12433 @opindex mhardlit
12434 @opindex mno-hardlit
12435 Inline constants into the code stream if it can be done in two
12436 instructions or less.

12438 @item -mdiv
12439 @itemx -mno-div
12440 @opindex mdiv
12441 @opindex mno-div
12442 Use the divide instruction. (Enabled by default).

12444 @item -mrelax-immediate
12445 @itemx -mno-relax-immediate
12446 @opindex mrelax-immediate
12447 @opindex mno-relax-immediate
12448 Allow arbitrary sized immediates in bit operations.

12450 @item -mwide-bitfields
12451 @itemx -mno-wide-bitfields
12452 @opindex mwide-bitfields
12453 @opindex mno-wide-bitfields
12454 Always treat bit-fields as int-sized.

12456 @item -m4byte-functions
12457 @itemx -mno-4byte-functions
12458 @opindex m4byte-functions
12459 @opindex mno-4byte-functions
12460 Force all functions to be aligned to a four byte boundary.

12462 @item -mcallgraph-data
12463 @itemx -mno-callgraph-data
12464 @opindex mcallgraph-data
12465 @opindex mno-callgraph-data
12466 Emit callgraph information.

12468 @item -mslow-bytes
12469 @itemx -mno-slow-bytes
12470 @opindex mslow-bytes
12471 @opindex mno-slow-bytes
12472 Prefer word access when reading byte quantities.

12474 @item -mlittle-endian
12475 @itemx -mbig-endian
12476 @opindex mlittle-endian
12477 @opindex mbig-endian
12478 Generate code for a little endian target.

12480 @item -m210
12481 @itemx -m340
12482 @opindex m210
12483 @opindex m340
12484 Generate code for the 210 processor.

12486 @item -mno-lsim
12487 @opindex no-lsim
12488 Assume that run-time support has been provided and so omit the
12489 simulator library (@file{libsim.a}) from the linker command line.

12491 @item -mstack-increment=@var{size}
12492 @opindex mstack-increment
12493 Set the maximum amount for a single stack increment operation. Large
12494 values can increase the speed of programs which contain functions
12495 that need a large amount of stack space, but they can also trigger a
12496 segmentation fault if the stack is extended too much. The default
12497 value is 0x1000.

```

```

12499 @end table

12501 @node MIPS Options
12502 @subsection MIPS Options
12503 @cindex MIPS options

12505 @table @gcctabopt

12507 @item -EB
12508 @opindex EB
12509 Generate big-endian code.

12511 @item -EL
12512 @opindex EL
12513 Generate little-endian code. This is the default for @samp{mips*el-*-*}
12514 configurations.

12516 @item -march=@var{arch}
12517 @opindex march
12518 Generate code that will run on @var{arch}, which can be the name of a
12519 generic MIPS ISA, or the name of a particular processor.
12520 The ISA names are:
12521 @samp{mips1}, @samp{mips2}, @samp{mips3}, @samp{mips4},
12522 @samp{mips32}, @samp{mips32r2}, @samp{mips64} and @samp{mips64r2}.
12523 The processor names are:
12524 @samp{4kc}, @samp{4km}, @samp{4kp}, @samp{4ksc},
12525 @samp{4kec}, @samp{4kem}, @samp{4kep}, @samp{4ksd},
12526 @samp{5kc}, @samp{5kf},
12527 @samp{20kc},
12528 @samp{24kc}, @samp{24kf2_1}, @samp{24kf1_1},
12529 @samp{24kec}, @samp{24kef2_1}, @samp{24kef1_1},
12530 @samp{34kc}, @samp{34kf2_1}, @samp{34kf1_1},
12531 @samp{74kc}, @samp{74kf2_1}, @samp{74kf1_1}, @samp{74kf3_2},
12532 @samp{loongson2e}, @samp{loongson2f},
12533 @samp{m4k},
12534 @samp{octeon},
12535 @samp{orion},
12536 @samp{r2000}, @samp{r3000}, @samp{r3900}, @samp{r4000}, @samp{r4400},
12537 @samp{r4600}, @samp{r4650}, @samp{r6000}, @samp{r8000},
12538 @samp{rm7000}, @samp{rm9000},
12539 @samp{r10000}, @samp{r12000}, @samp{r14000}, @samp{r16000},
12540 @samp{sbi},
12541 @samp{sr71000},
12542 @samp{vr4100}, @samp{vr4111}, @samp{vr4120}, @samp{vr4130}, @samp{vr4300},
12543 @samp{vr5000}, @samp{vr5400}, @samp{vr5500}
12544 and @samp{xlr}.
12545 The special value @samp{from-abi} selects the
12546 most compatible architecture for the selected ABI (that is,
12547 @samp{mips1} for 32-bit ABIs and @samp{mips3} for 64-bit ABIs).

12549 Native Linux/GNU toolchains also support the value @samp{native},
12550 which selects the best architecture option for the host processor.
12551 @option{-march=native} has no effect if GCC does not recognize
12552 the processor.

12554 In processor names, a final @samp{000} can be abbreviated as @samp{k}
12555 (for example, @samp{-march=r2k}). Prefixes are optional, and
12556 @samp{vr} may be written @samp{r}.

12558 Names of the form @samp{@var{n}f2_1} refer to processors with
12559 FPUs clocked at half the rate of the core, names of the form
12560 @samp{@var{n}f1_1} refer to processors with FPUs clocked at the same
12561 rate as the core, and names of the form @samp{@var{n}f3_2} refer to
12562 processors with FPUs clocked a ratio of 3:2 with respect to the core.
12563 For compatibility reasons, @samp{@var{n}f} is accepted as a synonym
12564 for @samp{@var{n}f2_1} while @samp{@var{n}x} and @samp{@var{b}fx} are

```



12565 accepted as synonyms for `@samp{@var{n}f1_1}`.

12567 GCC defines two macros based on the value of this option. The first  
 12568 is `@samp{MIPS_ARCH}`, which gives the name of target architecture, as  
 12569 a string. The second has the form `@samp{MIPS_ARCH_@var{foo}}`,  
 12570 where `@var{foo}` is the capitalized value of `@samp{MIPS_ARCH}`.  
 12571 For example, `@samp{-march=r2000}` will set `@samp{MIPS_ARCH}`  
 12572 to `@samp{"r2000"}` and define the macro `@samp{MIPS_ARCH_R2000}`.

12574 Note that the `@samp{MIPS_ARCH}` macro uses the processor names given  
 12575 above. In other words, it will have the full prefix and will not  
 12576 abbreviate `@samp{000}` as `@samp{k}`. In the case of `@samp{from-abi}`,  
 12577 the macro names the resolved architecture (either `@samp{"mips1"}` or  
 12578 `@samp{"mips3"}`). It names the default architecture when no  
 12579 `@option{-march}` option is given.

12581 `@item -mtune=@var{arch}`  
 12582 `@opindex mtune`  
 12583 Optimize for `@var{arch}`. Among other things, this option controls  
 12584 the way instructions are scheduled, and the perceived cost of arithmetic  
 12585 operations. The list of `@var{arch}` values is the same as for  
 12586 `@option{-march}`.

12588 When this option is not used, GCC will optimize for the processor  
 12589 specified by `@option{-march}`. By using `@option{-march}` and  
 12590 `@option{-mtune}` together, it is possible to generate code that will  
 12591 run on a family of processors, but optimize the code for one  
 12592 particular member of that family.

12594 `@samp{-mtune}` defines the macros `@samp{MIPS_TUNE}` and  
 12595 `@samp{MIPS_TUNE_@var{foo}}`, which work in the same way as the  
 12596 `@samp{-march}` ones described above.

12598 `@item -mips1`  
 12599 `@opindex mips1`  
 12600 Equivalent to `@samp{-march=mips1}`.

12602 `@item -mips2`  
 12603 `@opindex mips2`  
 12604 Equivalent to `@samp{-march=mips2}`.

12606 `@item -mips3`  
 12607 `@opindex mips3`  
 12608 Equivalent to `@samp{-march=mips3}`.

12610 `@item -mips4`  
 12611 `@opindex mips4`  
 12612 Equivalent to `@samp{-march=mips4}`.

12614 `@item -mips32`  
 12615 `@opindex mips32`  
 12616 Equivalent to `@samp{-march=mips32}`.

12618 `@item -mips32r2`  
 12619 `@opindex mips32r2`  
 12620 Equivalent to `@samp{-march=mips32r2}`.

12622 `@item -mips64`  
 12623 `@opindex mips64`  
 12624 Equivalent to `@samp{-march=mips64}`.

12626 `@item -mips64r2`  
 12627 `@opindex mips64r2`  
 12628 Equivalent to `@samp{-march=mips64r2}`.

12630 `@item -mips16`

12631 `@itemx -mno-mips16`  
 12632 `@opindex mips16`  
 12633 `@opindex mno-mips16`  
 12634 Generate (do not generate) MIPS16 code. If GCC is targeting a  
 12635 MIPS32 or MIPS64 architecture, it will make use of the MIPS16e ASE@.

12637 MIPS16 code generation can also be controlled on a per-function basis  
 12638 by means of `@code{mips16}` and `@code{nomips16}` attributes.  
 12639 `@xref{Function Attributes}`, for more information.

12641 `@item -mflip-mips16`  
 12642 `@opindex mflip-mips16`  
 12643 Generate MIPS16 code on alternating functions. This option is provided  
 12644 for regression testing of mixed MIPS16/non-MIPS16 code generation, and is  
 12645 not intended for ordinary use in compiling user code.

12647 `@item -minterlink-mips16`  
 12648 `@itemx -mno-interlink-mips16`  
 12649 `@opindex minterlink-mips16`  
 12650 `@opindex mno-interlink-mips16`  
 12651 Require (do not require) that non-MIPS16 code be link-compatible with  
 12652 MIPS16 code.

12654 For example, non-MIPS16 code cannot jump directly to MIPS16 code;  
 12655 it must either use a call or an indirect jump. `@option{-minterlink-mips16}`  
 12656 therefore disables direct jumps unless GCC knows that the target of the  
 12657 jump is not MIPS16.

12659 `@item -mabi=32`  
 12660 `@itemx -mabi=o64`  
 12661 `@itemx -mabi=n32`  
 12662 `@itemx -mabi=64`  
 12663 `@itemx -mabi=eabi`  
 12664 `@opindex mabi=32`  
 12665 `@opindex mabi=o64`  
 12666 `@opindex mabi=n32`  
 12667 `@opindex mabi=64`  
 12668 `@opindex mabi=eabi`  
 12669 Generate code for the given ABI@.

12671 Note that the EABI has a 32-bit and a 64-bit variant. GCC normally  
 12672 generates 64-bit code when you select a 64-bit architecture, but you  
 12673 can use `@option{-mfp32}` to get 32-bit code instead.

12675 For information about the O64 ABI, see  
 12676 `@w{@uref{http://gcc.gnu.org/projects/mipso64-abi.html}}`.

12678 GCC supports a variant of the o32 ABI in which floating-point registers  
 12679 are 64 rather than 32 bits wide. You can select this combination with  
 12680 `@option{-mabi=32}` `@option{-mfp64}`. This ABI relies on the `@samp{mthcl}`  
 12681 and `@samp{mfhcl}` instructions and is therefore only supported for  
 12682 MIPS32R2 processors.

12684 The register assignments for arguments and return values remain the  
 12685 same, but each scalar value is passed in a single 64-bit register  
 12686 rather than a pair of 32-bit registers. For example, scalar  
 12687 floating-point values are returned in `@samp{f0}` only, not a  
 12688 `@samp{f0}/@samp{f1}` pair. The set of call-saved registers also  
 12689 remains the same, but all 64 bits are saved.

12691 `@item -mabicalls`  
 12692 `@itemx -mno-abicalls`  
 12693 `@opindex mabicalls`  
 12694 `@opindex mno-abicalls`  
 12695 Generate (do not generate) code that is suitable for SVR4-style  
 12696 dynamic objects. `@option{-mabicalls}` is the default for SVR4-based

12697 systems.

12699 @item -mshared  
 12700 @itemx -mno-shared  
 12701 Generate (do not generate) code that is fully position-independent,  
 12702 and that can therefore be linked into shared libraries. This option  
 12703 only affects @option{-mabicalls}.

12705 All @option{-mabicalls} code has traditionally been position-independent,  
 12706 regardless of options like @option{-fPIC} and @option{-fpic}. However,  
 12707 as an extension, the GNU toolchain allows executables to use absolute  
 12708 accesses for locally-binding symbols. It can also use shorter GP  
 12709 initialization sequences and generate direct calls to locally-defined  
 12710 functions. This mode is selected by @option{-mno-shared}.

12712 @option{-mno-shared} depends on binutils 2.16 or higher and generates  
 12713 objects that can only be linked by the GNU linker. However, the option  
 12714 does not affect the ABI of the final executable; it only affects the ABI  
 12715 of relocatable objects. Using @option{-mno-shared} will generally make  
 12716 executables both smaller and quicker.

12718 @option{-mshared} is the default.

12720 @item -mplt  
 12721 @itemx -mno-plt  
 12722 @opindex plt  
 12723 @opindex mno-plt  
 12724 Assume (do not assume) that the static and dynamic linkers  
 12725 support PLTs and copy relocations. This option only affects  
 12726 @samp{-mno-shared -mabicalls}. For the n64 ABI, this option  
 12727 has no effect without @samp{-msym32}.

12729 You can make @option{-mplt} the default by configuring  
 12730 GCC with @option{--with-mips-plt}. The default is  
 12731 @option{-mno-plt} otherwise.

12733 @item -mxgot  
 12734 @itemx -mno-xgot  
 12735 @opindex mxgot  
 12736 @opindex mno-xgot  
 12737 Lift (do not lift) the usual restrictions on the size of the global  
 12738 offset table.

12740 GCC normally uses a single instruction to load values from the GOT@.  
 12741 While this is relatively efficient, it will only work if the GOT  
 12742 is smaller than about 64k. Anything larger will cause the linker  
 12743 to report an error such as:

12745 @cindex relocation truncated to fit (MIPS)  
 12746 @smallexample  
 12747 relocation truncated to fit: R\_MIPS\_GOT16 foobar  
 12748 @end smallexample

12750 If this happens, you should recompile your code with @option{-mxgot}.  
 12751 It should then work with very large GOTs, although it will also be  
 12752 less efficient, since it will take three instructions to fetch the  
 12753 value of a global symbol.

12755 Note that some linkers can create multiple GOTs. If you have such a  
 12756 linker, you should only need to use @option{-mxgot} when a single object  
 12757 file accesses more than 64k's worth of GOT entries. Very few do.

12759 These options have no effect unless GCC is generating position  
 12760 independent code.

12762 @item -mfp32

12763 @opindex mfp32  
 12764 Assume that general-purpose registers are 32 bits wide.

12766 @item -mfp64  
 12767 @opindex mfp64  
 12768 Assume that general-purpose registers are 64 bits wide.

12770 @item -mfp32  
 12771 @opindex mfp32  
 12772 Assume that floating-point registers are 32 bits wide.

12774 @item -mfp64  
 12775 @opindex mfp64  
 12776 Assume that floating-point registers are 64 bits wide.

12778 @item -mhard-float  
 12779 @opindex mhard-float  
 12780 Use floating-point coprocessor instructions.

12782 @item -msoft-float  
 12783 @opindex msoft-float  
 12784 Do not use floating-point coprocessor instructions. Implement  
 12785 floating-point calculations using library calls instead.

12787 @item -msingle-float  
 12788 @opindex msingle-float  
 12789 Assume that the floating-point coprocessor only supports single-precision  
 12790 operations.

12792 @item -mdouble-float  
 12793 @opindex mdouble-float  
 12794 Assume that the floating-point coprocessor supports double-precision  
 12795 operations. This is the default.

12797 @item -mllsc  
 12798 @itemx -mno-llsc  
 12799 @opindex llsc  
 12800 @opindex mno-llsc  
 12801 Use (do not use) @samp{ll}, @samp{sc}, and @samp{sync} instructions to  
 12802 implement atomic memory built-in functions. When neither option is  
 12803 specified, GCC will use the instructions if the target architecture  
 12804 supports them.

12806 @option{-mllsc} is useful if the runtime environment can emulate the  
 12807 instructions and @option{-mno-llsc} can be useful when compiling for  
 12808 nonstandard ISAs. You can make either option the default by  
 12809 configuring GCC with @option{--with-llsc} and @option{--without-llsc}  
 12810 respectively. @option{--with-llsc} is the default for some  
 12811 configurations; see the installation documentation for details.

12813 @item -mdsp  
 12814 @itemx -mno-dsp  
 12815 @opindex mdsp  
 12816 @opindex mno-dsp  
 12817 Use (do not use) revision 1 of the MIPS DSP ASE@.  
 12818 @xref{MIPS DSP Built-in Functions}. This option defines the  
 12819 preprocessor macro @samp{\_\_mips\_dsp}. It also defines  
 12820 @samp{\_\_mips\_dsp\_rev} to 1.

12822 @item -mdspr2  
 12823 @itemx -mno-dspr2  
 12824 @opindex dspr2  
 12825 @opindex mno-dspr2  
 12826 Use (do not use) revision 2 of the MIPS DSP ASE@.  
 12827 @xref{MIPS DSP Built-in Functions}. This option defines the  
 12828 preprocessor macros @samp{\_\_mips\_dsp} and @samp{\_\_mips\_dspr2}.

12829 It also defines `@samp{__mips_dsp_rev}` to 2.

12831 `@item -msmartmips`  
 12832 `@itemx -mno-smartmips`  
 12833 `@opindex msmartmips`  
 12834 `@opindex mno-smartmips`  
 12835 Use (do not use) the MIPS SmartMIPS ASE.

12837 `@item -mpaired-single`  
 12838 `@itemx -mno-paired-single`  
 12839 `@opindex mpaired-single`  
 12840 `@opindex mno-paired-single`  
 12841 Use (do not use) paired-single floating-point instructions.  
 12842 `@xref{MIPS Paired-Single Support}`. This option requires  
 12843 hardware floating-point support to be enabled.

12845 `@item -mdmx`  
 12846 `@itemx -mno-mdmx`  
 12847 `@opindex mdmx`  
 12848 `@opindex mno-mdmx`  
 12849 Use (do not use) MIPS Digital Media Extension instructions.  
 12850 This option can only be used when generating 64-bit code and requires  
 12851 hardware floating-point support to be enabled.

12853 `@item -mips3d`  
 12854 `@itemx -mno-mips3d`  
 12855 `@opindex mips3d`  
 12856 `@opindex mno-mips3d`  
 12857 Use (do not use) the MIPS-3D ASE@. `@xref{MIPS-3D Built-in Functions}`.  
 12858 The option `@option{-mips3d}` implies `@option{-mpaired-single}`.

12860 `@item -mmt`  
 12861 `@itemx -mno-mt`  
 12862 `@opindex mmt`  
 12863 `@opindex mno-mt`  
 12864 Use (do not use) MT Multithreading instructions.

12866 `@item -mlong64`  
 12867 `@opindex mlong64`  
 12868 Force `@code{long}` types to be 64 bits wide. See `@option{-mlong32}` for  
 12869 an explanation of the default and the way that the pointer size is  
 12870 determined.

12872 `@item -mlong32`  
 12873 `@opindex mlong32`  
 12874 Force `@code{long}`, `@code{int}`, and pointer types to be 32 bits wide.

12876 The default size of `@code{int}`s, `@code{long}`s and pointers depends on  
 12877 the ABI@. All the supported ABIs use 32-bit `@code{int}`s. The n64 ABI  
 12878 uses 64-bit `@code{long}`s, as does the 64-bit EABI; the others use  
 12879 32-bit `@code{long}`s. Pointers are the same size as `@code{long}`s,  
 12880 or the same size as integer registers, whichever is smaller.

12882 `@item -msym32`  
 12883 `@itemx -mno-sym32`  
 12884 `@opindex msym32`  
 12885 `@opindex mno-sym32`  
 12886 Assume (do not assume) that all symbols have 32-bit values, regardless  
 12887 of the selected ABI@. This option is useful in combination with  
 12888 `@option{-mabi=64}` and `@option{-mno-abicalls}` because it allows GCC  
 12889 to generate shorter and faster references to symbolic addresses.

12891 `@item -G @var{num}`  
 12892 `@opindex G`  
 12893 Put definitions of externally-visible data in a small data section  
 12894 if that data is no bigger than `@var{num}` bytes. GCC can then access

12895 the data more efficiently; see `@option{-mgpopt}` for details.

12897 The default `@option{-G}` option depends on the configuration.

12899 `@item -mlocal-sdata`  
 12900 `@itemx -mno-local-sdata`  
 12901 `@opindex mlocal-sdata`  
 12902 `@opindex mno-local-sdata`  
 12903 Extend (do not extend) the `@option{-G}` behavior to local data too,  
 12904 such as to static variables in C@. `@option{-mlocal-sdata}` is the  
 12905 default for all configurations.

12907 If the linker complains that an application is using too much small data,  
 12908 you might want to try rebuilding the less performance-critical parts with  
 12909 `@option{-mno-local-sdata}`. You might also want to build large  
 12910 libraries with `@option{-mno-local-sdata}`, so that the libraries leave  
 12911 more room for the main program.

12913 `@item -mextern-sdata`  
 12914 `@itemx -mno-extern-sdata`  
 12915 `@opindex mextern-sdata`  
 12916 `@opindex mno-extern-sdata`  
 12917 Assume (do not assume) that externally-defined data will be in  
 12918 a small data section if that data is within the `@option{-G}` limit.  
 12919 `@option{-mextern-sdata}` is the default for all configurations.

12921 If you compile a module `@var{Mod}` with `@option{-mextern-sdata}` `@option{-G`  
 12922 `@var{num}` `@option{-mgpopt}`, and `@var{Mod}` references a variable `@var{Var}`  
 12923 that is no bigger than `@var{num}` bytes, you must make sure that `@var{Var}`  
 12924 is placed in a small data section. If `@var{Var}` is defined by another  
 12925 module, you must either compile that module with a high-enough  
 12926 `@option{-G}` setting or attach a `@code{section}` attribute to `@var{Var}`'s  
 12927 definition. If `@var{Var}` is common, you must link the application  
 12928 with a high-enough `@option{-G}` setting.

12930 The easiest way of satisfying these restrictions is to compile  
 12931 and link every module with the same `@option{-G}` option. However,  
 12932 you may wish to build a library that supports several different  
 12933 small data limits. You can do this by compiling the library with  
 12934 the highest supported `@option{-G}` setting and additionally using  
 12935 `@option{-mno-extern-sdata}` to stop the library from making assumptions  
 12936 about externally-defined data.

12938 `@item -mgpopt`  
 12939 `@itemx -mno-gpopt`  
 12940 `@opindex mgpopt`  
 12941 `@opindex mno-gpopt`  
 12942 Use (do not use) GP-relative accesses for symbols that are known to be  
 12943 in a small data section; see `@option{-G}`, `@option{-mlocal-sdata}` and  
 12944 `@option{-mextern-sdata}`. `@option{-mgpopt}` is the default for all  
 12945 configurations.

12947 `@option{-mno-gpopt}` is useful for cases where the `@code{$_gp}` register  
 12948 might not hold the value of `@code{__gp}`. For example, if the code is  
 12949 part of a library that might be used in a boot monitor, programs that  
 12950 call boot monitor routines will pass an unknown value in `@code{$_gp}`.  
 12951 (In such situations, the boot monitor itself would usually be compiled  
 12952 with `@option{-G0}`.)

12954 `@option{-mno-gpopt}` implies `@option{-mno-local-sdata}` and  
 12955 `@option{-mno-extern-sdata}`.

12957 `@item -membedded-data`  
 12958 `@itemx -mno-embedded-data`  
 12959 `@opindex membedded-data`  
 12960 `@opindex mno-embedded-data`

12961 Allocate variables to the read-only data section first if possible, then  
 12962 next in the small data section if possible, otherwise in data. This gives  
 12963 slightly slower code than the default, but reduces the amount of RAM required  
 12964 when executing, and thus may be preferred for some embedded systems.

12966 @item -muninit-const-in-rodada  
 12967 @itemx -mno-uninit-const-in-rodada  
 12968 @opindex muninit-const-in-rodada  
 12969 @opindex mno-uninit-const-in-rodada  
 12970 Put uninitialized @code{const} variables in the read-only data section.  
 12971 This option is only meaningful in conjunction with @option{-membedded-data}.

12973 @item -mcode-readable=@var{setting}  
 12974 @opindex mcode-readable  
 12975 Specify whether GCC may generate code that reads from executable sections.  
 12976 There are three possible settings:

12978 @table @gcctabopt  
 12979 @item -mcode-readable=yes  
 12980 Instructions may freely access executable sections. This is the  
 12981 default setting.

12983 @item -mcode-readable=pcrel  
 12984 MIPS16 PC-relative load instructions can access executable sections,  
 12985 but other instructions must not do so. This option is useful on 4KSc  
 12986 and 4KSd processors when the code TLBs have the Read Inhibit bit set.  
 12987 It is also useful on processors that can be configured to have a dual  
 12988 instruction/data SRAM interface and that, like the M4K, automatically  
 12989 redirect PC-relative loads to the instruction RAM.

12991 @item -mcode-readable=no  
 12992 Instructions must not access executable sections. This option can be  
 12993 useful on targets that are configured to have a dual instruction/data  
 12994 SRAM interface but that (unlike the M4K) do not automatically redirect  
 12995 PC-relative loads to the instruction RAM.  
 12996 @end table

12998 @item -msplit-addresses  
 12999 @itemx -mno-split-addresses  
 13000 @opindex msplit-addresses  
 13001 @opindex mno-split-addresses  
 13002 Enable (disable) use of the @code{%hi()} and @code{%lo()} assembler  
 13003 relocation operators. This option has been superseded by  
 13004 @option{-mexplicit-relocs} but is retained for backwards compatibility.

13006 @item -mexplicit-relocs  
 13007 @itemx -mno-explicit-relocs  
 13008 @opindex mexplicit-relocs  
 13009 @opindex mno-explicit-relocs  
 13010 Use (do not use) assembler relocation operators when dealing with symbolic  
 13011 addresses. The alternative, selected by @option{-mno-explicit-relocs},  
 13012 is to use assembler macros instead.

13014 @option{-mexplicit-relocs} is the default if GCC was configured  
 13015 to use an assembler that supports relocation operators.

13017 @item -mcheck-zero-division  
 13018 @itemx -mno-check-zero-division  
 13019 @opindex mcheck-zero-division  
 13020 @opindex mno-check-zero-division  
 13021 Trap (do not trap) on integer division by zero.

13023 The default is @option{-mcheck-zero-division}.

13025 @item -mdivide-traps  
 13026 @itemx -mdivide-breaks

13027 @opindex mdivide-traps  
 13028 @opindex mdivide-breaks  
 13029 MIPS systems check for division by zero by generating either a  
 13030 conditional trap or a break instruction. Using traps results in  
 13031 smaller code, but is only supported on MIPS II and later. Also, some  
 13032 versions of the Linux kernel have a bug that prevents trap from  
 13033 generating the proper signal (@code{SIGFPE}). Use @option{-mdivide-traps} to  
 13034 allow conditional traps on architectures that support them and  
 13035 @option{-mdivide-breaks} to force the use of breaks.

13037 The default is usually @option{-mdivide-traps}, but this can be  
 13038 overridden at configure time using @option{--with-divide-breaks}.  
 13039 Divide-by-zero checks can be completely disabled using  
 13040 @option{-mno-check-zero-division}.

13042 @item -mmemcpy  
 13043 @itemx -mno-memcpy  
 13044 @opindex memcpy  
 13045 @opindex mno-memcpy  
 13046 Force (do not force) the use of @code{memcpy()} for non-trivial block  
 13047 moves. The default is @option{-mno-memcpy}, which allows GCC to inline  
 13048 most constant-sized copies.

13050 @item -mlong-calls  
 13051 @itemx -mno-long-calls  
 13052 @opindex mlong-calls  
 13053 @opindex mno-long-calls  
 13054 Disable (do not disable) use of the @code{jal} instruction. Calling  
 13055 functions using @code{jal} is more efficient but requires the caller  
 13056 and callee to be in the same 256 megabyte segment.

13058 This option has no effect on abicalls code. The default is  
 13059 @option{-mno-long-calls}.

13061 @item -mmad  
 13062 @itemx -mno-mad  
 13063 @opindex mad  
 13064 @opindex mno-mad  
 13065 Enable (disable) use of the @code{mad}, @code{madu} and @code{mul}  
 13066 instructions, as provided by the R4650 ISA@.

13068 @item -mfused-madd  
 13069 @itemx -mno-fused-madd  
 13070 @opindex mfused-madd  
 13071 @opindex mno-fused-madd  
 13072 Enable (disable) use of the floating point multiply-accumulate  
 13073 instructions, when they are available. The default is  
 13074 @option{-mfused-madd}.

13076 When multiply-accumulate instructions are used, the intermediate  
 13077 product is calculated to infinite precision and is not subject to  
 13078 the FCSR Flush to Zero bit. This may be undesirable in some  
 13079 circumstances.

13081 @item -nocpp  
 13082 @opindex nocpp  
 13083 Tell the MIPS assembler to not run its preprocessor over user  
 13084 assembler files (with a @samp{.s} suffix) when assembling them.

13086 @item -mfix-r4000  
 13087 @itemx -mno-fix-r4000  
 13088 @opindex mfix-r4000  
 13089 @opindex mno-fix-r4000  
 13090 Work around certain R4000 CPU errata:  
 13091 @itemize @minus  
 13092 @item

```

13093 A double-word or a variable shift may give an incorrect result if executed
13094 immediately after starting an integer division.
13095 @item
13096 A double-word or a variable shift may give an incorrect result if executed
13097 while an integer multiplication is in progress.
13098 @item
13099 An integer division may give an incorrect result if started in a delay slot
13100 of a taken branch or a jump.
13101 @end itemize

13103 @item -mfix-r4400
13104 @itemx -mno-fix-r4400
13105 @opindex mfix-r4400
13106 @opindex mno-fix-r4400
13107 Work around certain R4400 CPU errata:
13108 @itemize @minus
13109 @item
13110 A double-word or a variable shift may give an incorrect result if executed
13111 immediately after starting an integer division.
13112 @end itemize

13114 @item -mfix-r10000
13115 @itemx -mno-fix-r10000
13116 @opindex mfix-r10000
13117 @opindex mno-fix-r10000
13118 Work around certain R10000 errata:
13119 @itemize @minus
13120 @item
13121 @code{ll}/@code{sc} sequences may not behave atomically on revisions
13122 prior to 3.0. They may deadlock on revisions 2.6 and earlier.
13123 @end itemize

13125 This option can only be used if the target architecture supports
13126 branch-likely instructions. @option{-mfix-r10000} is the default when
13127 @option{-march=r10000} is used; @option{-mno-fix-r10000} is the default
13128 otherwise.

13130 @item -mfix-vr4120
13131 @itemx -mno-fix-vr4120
13132 @opindex mfix-vr4120
13133 Work around certain VR4120 errata:
13134 @itemize @minus
13135 @item
13136 @code{dmultu} does not always produce the correct result.
13137 @item
13138 @code{div} and @code{ddiv} do not always produce the correct result if one
13139 of the operands is negative.
13140 @end itemize
13141 The workarounds for the division errata rely on special functions in
13142 @file{libgcc.a}. At present, these functions are only provided by
13143 the @code{mips64vr*-elf} configurations.

13145 Other VR4120 errata require a nop to be inserted between certain pairs of
13146 instructions. These errata are handled by the assembler, not by GCC itself.

13148 @item -mfix-vr4130
13149 @opindex mfix-vr4130
13150 Work around the VR4130 @code{mflo}/@code{mfhi} errata. The
13151 workarounds are implemented by the assembler rather than by GCC,
13152 although GCC will avoid using @code{mflo} and @code{mfhi} if the
13153 VR4130 @code{macc}, @code{macchi}, @code{dmacc} and @code{dmacchi}
13154 instructions are available instead.

13156 @item -mfix-sb1
13157 @itemx -mno-fix-sb1
13158 @opindex mfix-sb1

```

```

13159 Work around certain SB-1 CPU core errata.
13160 (This flag currently works around the SB-1 revision 2
13161 ``F1'' and ``F2'' floating point errata.)

13163 @item -mr10k-cache-barrier=@var{setting}
13164 @opindex mr10k-cache-barrier
13165 Specify whether GCC should insert cache barriers to avoid the
13166 side-effects of speculation on R10K processors.

13168 In common with many processors, the R10K tries to predict the outcome
13169 of a conditional branch and speculatively executes instructions from
13170 the ``taken'' branch. It later aborts these instructions if the
13171 predicted outcome was wrong. However, on the R10K, even aborted
13172 instructions can have side effects.

13174 This problem only affects kernel stores and, depending on the system,
13175 kernel loads. As an example, a speculatively-executed store may load
13176 the target memory into cache and mark the cache line as dirty, even if
13177 the store itself is later aborted. If a DMA operation writes to the
13178 same area of memory before the ``dirty'' line is flushed, the cached
13179 data will overwrite the DMA-ed data. See the R10K processor manual
13180 for a full description, including other potential problems.

13182 One workaround is to insert cache barrier instructions before every memory
13183 access that might be speculatively executed and that might have side
13184 effects even if aborted. @option{-mr10k-cache-barrier=@var{setting}}
13185 controls GCC's implementation of this workaround. It assumes that
13186 aborted accesses to any byte in the following regions will not have
13187 side effects:

13189 @enumerate
13190 @item
13191 the memory occupied by the current function's stack frame;

13193 @item
13194 the memory occupied by an incoming stack argument;

13196 @item
13197 the memory occupied by an object with a link-time-constant address.
13198 @end enumerate

13200 It is the kernel's responsibility to ensure that speculative
13201 accesses to these regions are indeed safe.

13203 If the input program contains a function declaration such as:

13205 @smallexample
13206 void foo (void);
13207 @end smallexample

13209 then the implementation of @code{foo} must allow @code{j foo} and
13210 @code{jal foo} to be executed speculatively. GCC honors this
13211 restriction for functions it compiles itself. It expects non-GCC
13212 functions (such as hand-written assembly code) to do the same.

13214 The option has three forms:

13216 @table @gcctabopt
13217 @item -mr10k-cache-barrier=load-store
13218 Insert a cache barrier before a load or store that might be
13219 speculatively executed and that might have side effects even
13220 if aborted.

13222 @item -mr10k-cache-barrier=store
13223 Insert a cache barrier before a store that might be speculatively
13224 executed and that might have side effects even if aborted.

```

```

13226 @item -mr10k-cache-barrier=none
13227 Disable the insertion of cache barriers. This is the default setting.
13228 @end table

13230 @item -mflush-func=@var{func}
13231 @itemx -mno-flush-func
13232 @opindex mflush-func
13233 Specifies the function to call to flush the I and D caches, or to not
13234 call any such function. If called, the function must take the same
13235 arguments as the common @code{flush_func()}, that is, the address of the
13236 memory range for which the cache is being flushed, the size of the
13237 memory range, and the number 3 (to flush both caches). The default
13238 depends on the target GCC was configured for, but commonly is either
13239 @samp{flush_func} or @samp{cpu_flush}.

13241 @item mbranch-cost=@var{num}
13242 @opindex mbranch-cost
13243 Set the cost of branches to roughly @var{num} ``simple'' instructions.
13244 This cost is only a heuristic and is not guaranteed to produce
13245 consistent results across releases. A zero cost redundantly selects
13246 the default, which is based on the @option{-mtune} setting.

13248 @item -mbranch-likely
13249 @itemx -mno-branch-likely
13250 @opindex mbranch-likely
13251 @opindex mno-branch-likely
13252 Enable or disable use of Branch Likely instructions, regardless of the
13253 default for the selected architecture. By default, Branch Likely
13254 instructions may be generated if they are supported by the selected
13255 architecture. An exception is for the MIPS32 and MIPS64 architectures
13256 and processors which implement those architectures; for those, Branch
13257 Likely instructions will not be generated by default because the MIPS32
13258 and MIPS64 architectures specifically deprecate their use.

13260 @item -mfp-exceptions
13261 @itemx -mno-fp-exceptions
13262 @opindex mfp-exceptions
13263 Specifies whether FP exceptions are enabled. This affects how we schedule
13264 FP instructions for some processors. The default is that FP exceptions are
13265 enabled.

13267 For instance, on the SB-1, if FP exceptions are disabled, and we are emitting
13268 64-bit code, then we can use both FP pipes. Otherwise, we can only use one
13269 FP pipe.

13271 @item -mvr4130-align
13272 @itemx -mno-vr4130-align
13273 @opindex mvr4130-align
13274 The VR4130 pipeline is two-way superscalar, but can only issue two
13275 instructions together if the first one is 8-byte aligned. When this
13276 option is enabled, GCC will align pairs of instructions that it
13277 thinks should execute in parallel.

13279 This option only has an effect when optimizing for the VR4130.
13280 It normally makes code faster, but at the expense of making it bigger.
13281 It is enabled by default at optimization level @option{-O3}.
13282 @end table

13284 @node MMIX Options
13285 @subsection MMIX Options
13286 @cindex MMIX Options

13288 These options are defined for the MMIX:

13290 @table @gctabopt

```

```

13291 @item -mlibfuncs
13292 @itemx -mno-libfuncs
13293 @opindex mlibfuncs
13294 @opindex mno-libfuncs
13295 Specify that intrinsic library functions are being compiled, passing all
13296 values in registers, no matter the size.

13298 @item -mepsilon
13299 @itemx -mno-epsilon
13300 @opindex mepsilon
13301 @opindex mno-epsilon
13302 Generate floating-point comparison instructions that compare with respect
13303 to the @code{rE} epsilon register.

13305 @item -mabi=mmixware
13306 @itemx -mabi=gnu
13307 @opindex mabi=mmixware
13308 @opindex mabi=gnu
13309 Generate code that passes function parameters and return values that (in
13310 the called function) are seen as registers @code{$0} and up, as opposed to
13311 the GNU ABI which uses global registers @code{$231} and up.

13313 @item -mzero-extend
13314 @itemx -mno-zero-extend
13315 @opindex mzero-extend
13316 @opindex mno-zero-extend
13317 When reading data from memory in sizes shorter than 64 bits, use (do not
13318 use) zero-extending load instructions by default, rather than
13319 sign-extending ones.

13321 @item -mknuthdiv
13322 @itemx -mno-knuthdiv
13323 @opindex mknuthdiv
13324 @opindex mno-knuthdiv
13325 Make the result of a division yielding a remainder have the same sign as
13326 the divisor. With the default, @option{-mno-knuthdiv}, the sign of the
13327 remainder follows the sign of the dividend. Both methods are
13328 arithmetically valid, the latter being almost exclusively used.

13330 @item -mtoplevel-symbols
13331 @itemx -mno-toplevel-symbols
13332 @opindex mtoplevel-symbols
13333 @opindex mno-toplevel-symbols
13334 Prepend (do not prepend) a @samp{:} to all global symbols, so the assembly
13335 code can be used with the @code{PREFIX} assembly directive.

13337 @item -melf
13338 @opindex melf
13339 Generate an executable in the ELF format, rather than the default
13340 @samp{mmo} format used by the @command{mmix} simulator.

13342 @item -mbranch-predict
13343 @itemx -mno-branch-predict
13344 @opindex mbranch-predict
13345 @opindex mno-branch-predict
13346 Use (do not use) the probable-branch instructions, when static branch
13347 prediction indicates a probable branch.

13349 @item -mbase-addresses
13350 @itemx -mno-base-addresses
13351 @opindex mbase-addresses
13352 @opindex mno-base-addresses
13353 Generate (do not generate) code that uses @emph{base addresses}. Using a
13354 base address automatically generates a request (handled by the assembler
13355 and the linker) for a constant to be set up in a global register. The
13356 register is used for one or more base address requests within the range 0

```

13357 to 255 from the value held in the register. The generally leads to short  
 13358 and fast code, but the number of different data items that can be  
 13359 addressed is limited. This means that a program that uses lots of static  
 13360 data may require `@option{-mno-base-addresses}`.

13362 `@item -msingle-exit`  
 13363 `@itemx -mno-single-exit`  
 13364 `@opindex msingle-exit`  
 13365 `@opindex mno-single-exit`  
 13366 Force (do not force) generated code to have a single exit point in each  
 13367 function.  
 13368 `@end table`

13370 `@node MN10300 Options`  
 13371 `@subsection MN10300 Options`  
 13372 `@cindex MN10300 options`

13374 These `@option{-m}` options are defined for Matsushita MN10300 architectures:

13376 `@table @gcctabopt`  
 13377 `@item -mmult-bug`  
 13378 `@opindex mmult-bug`  
 13379 Generate code to avoid bugs in the multiply instructions for the MN10300  
 13380 processors. This is the default.

13382 `@item -mno-mult-bug`  
 13383 `@opindex mno-mult-bug`  
 13384 Do not generate code to avoid bugs in the multiply instructions for the  
 13385 MN10300 processors.

13387 `@item -mam33`  
 13388 `@opindex mam33`  
 13389 Generate code which uses features specific to the AM33 processor.

13391 `@item -mno-am33`  
 13392 `@opindex mno-am33`  
 13393 Do not generate code which uses features specific to the AM33 processor. This  
 13394 is the default.

13396 `@item -mreturn-pointer-on-d0`  
 13397 `@opindex mreturn-pointer-on-d0`  
 13398 When generating a function which returns a pointer, return the pointer  
 13399 in both `@code{a0}` and `@code{d0}`. Otherwise, the pointer is returned  
 13400 only in `a0`, and attempts to call such functions without a prototype  
 13401 would result in errors. Note that this option is on by default; use  
 13402 `@option{-mno-return-pointer-on-d0}` to disable it.

13404 `@item -mno-crt0`  
 13405 `@opindex mno-crt0`  
 13406 Do not link in the C run-time initialization object file.

13408 `@item -mrelax`  
 13409 `@opindex mrelax`  
 13410 Indicate to the linker that it should perform a relaxation optimization pass  
 13411 to shorten branches, calls and absolute memory addresses. This option only  
 13412 has an effect when used on the command line for the final link step.

13414 This option makes symbolic debugging impossible.  
 13415 `@end table`

13417 `@node PDP-11 Options`  
 13418 `@subsection PDP-11 Options`  
 13419 `@cindex PDP-11 Options`

13421 These options are defined for the PDP-11:

13423 `@table @gcctabopt`  
 13424 `@item -mfpu`  
 13425 `@opindex mfpu`  
 13426 Use hardware FPP floating point. This is the default. (FIS floating  
 13427 point on the PDP-11/40 is not supported.)

13429 `@item -msoft-float`  
 13430 `@opindex msoft-float`  
 13431 Do not use hardware floating point.

13433 `@item -mac0`  
 13434 `@opindex mac0`  
 13435 Return floating-point results in `ac0` (`fr0` in Unix assembler syntax).

13437 `@item -mno-ac0`  
 13438 `@opindex mno-ac0`  
 13439 Return floating-point results in memory. This is the default.

13441 `@item -m40`  
 13442 `@opindex m40`  
 13443 Generate code for a PDP-11/40.

13445 `@item -m45`  
 13446 `@opindex m45`  
 13447 Generate code for a PDP-11/45. This is the default.

13449 `@item -m10`  
 13450 `@opindex m10`  
 13451 Generate code for a PDP-11/10.

13453 `@item -mbcopy-builtin`  
 13454 `@opindex bcopy-builtin`  
 13455 Use inline `@code{movmemhi}` patterns for copying memory. This is the  
 13456 default.

13458 `@item -mbcopy`  
 13459 `@opindex mbcopy`  
 13460 Do not use inline `@code{movmemhi}` patterns for copying memory.

13462 `@item -mint16`  
 13463 `@itemx -mno-int32`  
 13464 `@opindex mint16`  
 13465 `@opindex mno-int32`  
 13466 Use 16-bit `@code{int}`. This is the default.

13468 `@item -mint32`  
 13469 `@itemx -mno-int16`  
 13470 `@opindex mint32`  
 13471 `@opindex mno-int16`  
 13472 Use 32-bit `@code{int}`.

13474 `@item -mfloat64`  
 13475 `@itemx -mno-float32`  
 13476 `@opindex mfloat64`  
 13477 `@opindex mno-float32`  
 13478 Use 64-bit `@code{float}`. This is the default.

13480 `@item -mfloat32`  
 13481 `@itemx -mno-float64`  
 13482 `@opindex mfloat32`  
 13483 `@opindex mno-float64`  
 13484 Use 32-bit `@code{float}`.

13486 `@item -mabshi`  
 13487 `@opindex mabshi`  
 13488 Use `@code{abshi2}` pattern. This is the default.

```

13490 @item -mno-abshi
13491 @opindex mno-abshi
13492 Do not use @code{abshi2} pattern.

13494 @item -mbranch-expensive
13495 @opindex mbranch-expensive
13496 Pretend that branches are expensive. This is for experimenting with
13497 code generation only.

13499 @item -mbranch-cheap
13500 @opindex mbranch-cheap
13501 Do not pretend that branches are expensive. This is the default.

13503 @item -msplit
13504 @opindex msplit
13505 Generate code for a system with split I&D@.

13507 @item -mno-split
13508 @opindex mno-split
13509 Generate code for a system without split I&D@. This is the default.

13511 @item -munix-asm
13512 @opindex munix-asm
13513 Use Unix assembler syntax. This is the default when configured for
13514 @samp{pdp11-*-bsd}.

13516 @item -mdec-asm
13517 @opindex mdec-asm
13518 Use DEC assembler syntax. This is the default when configured for any
13519 PDP-11 target other than @samp{pdp11-*-bsd}.
13520 @end table

13522 @node picoChip Options
13523 @subsection picoChip Options
13524 @cindex picoChip options

13526 These @samp{-m} options are defined for picoChip implementations:

13528 @table @gcctabopt

13530 @item -mae=@var{ae_type}
13531 @opindex mcpu
13532 Set the instruction set, register set, and instruction scheduling
13533 parameters for array element type @var{ae_type}. Supported values
13534 for @var{ae_type} are @samp{ANY}, @samp{MUL}, and @samp{MAC}.

13536 @option{-mae=ANY} selects a completely generic AE type. Code
13537 generated with this option will run on any of the other AE types. The
13538 code will not be as efficient as it would be if compiled for a specific
13539 AE type, and some types of operation (e.g., multiplication) will not
13540 work properly on all types of AE.

13542 @option{-mae=MUL} selects a MUL AE type. This is the most useful AE type
13543 for compiled code, and is the default.

13545 @option{-mae=MAC} selects a DSP-style MAC AE. Code compiled with this
13546 option may suffer from poor performance of byte (char) manipulation,
13547 since the DSP AE does not provide hardware support for byte load/stores.

13549 @item -msymbol-as-address
13550 Enable the compiler to directly use a symbol name as an address in a
13551 load/store instruction, without first loading it into a
13552 register. Typically, the use of this option will generate larger
13553 programs, which run faster than when the option isn't used. However, the
13554 results vary from program to program, so it is left as a user option,

```

```

13555 rather than being permanently enabled.

13557 @item -mno-inefficient-warnings
13558 Disables warnings about the generation of inefficient code. These
13559 warnings can be generated, for example, when compiling code which
13560 performs byte-level memory operations on the MAC AE type. The MAC AE has
13561 no hardware support for byte-level memory operations, so all byte
13562 load/stores must be synthesized from word load/store operations. This is
13563 inefficient and a warning will be generated indicating to the programmer
13564 that they should rewrite the code to avoid byte operations, or to target
13565 an AE type which has the necessary hardware support. This option enables
13566 the warning to be turned off.

13568 @end table

13570 @node PowerPC Options
13571 @subsection PowerPC Options
13572 @cindex PowerPC options

13574 These are listed under @xref{RS/6000 and PowerPC Options}.

13576 @node RS/6000 and PowerPC Options
13577 @subsection IBM RS/6000 and PowerPC Options
13578 @cindex RS/6000 and PowerPC Options
13579 @cindex IBM RS/6000 and PowerPC Options

13581 These @samp{-m} options are defined for the IBM RS/6000 and PowerPC:
13582 @table @gcctabopt
13583 @item -mpower
13584 @itemx -mno-power
13585 @itemx -mpower2
13586 @itemx -mno-power2
13587 @itemx -mpowerpc
13588 @itemx -mno-powerpc
13589 @itemx -mpowerpc-gpopt
13590 @itemx -mno-powerpc-gpopt
13591 @itemx -mpowerpc-gfxopt
13592 @itemx -mno-powerpc-gfxopt
13593 @itemx -mpowerpc64
13594 @itemx -mno-powerpc64
13595 @itemx -mmfcrf
13596 @itemx -mno-mfcrf
13597 @itemx -mpopcntb
13598 @itemx -mno-popcntb
13599 @itemx -mfprnd
13600 @itemx -mno-fprnd
13601 @itemx -mcmpb
13602 @itemx -mno-cmpb
13603 @itemx -mmfpgpr
13604 @itemx -mno-mfpgpr
13605 @itemx -mhard-dfp
13606 @itemx -mno-hard-dfp
13607 @opindex mpower
13608 @opindex mno-power
13609 @opindex mpower2
13610 @opindex mno-power2
13611 @opindex mpowerpc
13612 @opindex mno-powerpc
13613 @opindex mpowerpc-gpopt
13614 @opindex mno-powerpc-gpopt
13615 @opindex mpowerpc-gfxopt
13616 @opindex mno-powerpc-gfxopt
13617 @opindex mpowerpc64
13618 @opindex mno-powerpc64
13619 @opindex mmfcrf
13620 @opindex mno-mfcrf

```



13621 @opindex mpopcntb  
 13622 @opindex mno-popcntb  
 13623 @opindex mfprnd  
 13624 @opindex mno-fprnd  
 13625 @opindex mcmpb  
 13626 @opindex mno-cmpb  
 13627 @opindex mmfpgpr  
 13628 @opindex mno-mfpgpr  
 13629 @opindex mhard-dfp  
 13630 @opindex mno-hard-dfp  
 13631 GCC supports two related instruction set architectures for the  
 13632 RS/6000 and PowerPC@. The @dfn{POWER} instruction set are those  
 13633 instructions supported by the @samp{rios} chip set used in the original  
 13634 RS/6000 systems and the @dfn{PowerPC} instruction set is the  
 13635 architecture of the Freescale MPC5xx, MPC6xx, MPC8xx microprocessors, and  
 13636 the IBM 4xx, 6xx, and follow-on microprocessors.

13638 Neither architecture is a subset of the other. However there is a  
 13639 large common subset of instructions supported by both. An MQ  
 13640 register is included in processors supporting the POWER architecture.

13642 You use these options to specify which instructions are available on the  
 13643 processor you are using. The default value of these options is  
 13644 determined when configuring GCC@. Specifying the  
 13645 @option{-mcpu=@var{cpu\_type}} overrides the specification of these  
 13646 options. We recommend you use the @option{-mcpu=@var{cpu\_type}} option  
 13647 rather than the options listed above.

13649 The @option{-mpower} option allows GCC to generate instructions that  
 13650 are found only in the POWER architecture and to use the MQ register.  
 13651 Specifying @option{-mpower2} implies @option{-power} and also allows GCC  
 13652 to generate instructions that are present in the POWER2 architecture but  
 13653 not the original POWER architecture.

13655 The @option{-mpowerpc} option allows GCC to generate instructions that  
 13656 are found only in the 32-bit subset of the PowerPC architecture.  
 13657 Specifying @option{-mpowerpc-gpopt} implies @option{-mpowerpc} and also allows  
 13658 GCC to use the optional PowerPC architecture instructions in the  
 13659 General Purpose group, including floating-point square root. Specifying  
 13660 @option{-mpowerpc-gfxopt} implies @option{-mpowerpc} and also allows GCC to  
 13661 use the optional PowerPC architecture instructions in the Graphics  
 13662 group, including floating-point select.

13664 The @option{-mmfcrf} option allows GCC to generate the move from  
 13665 condition register field instruction implemented on the POWER4  
 13666 processor and other processors that support the PowerPC V2.01  
 13667 architecture.

13668 The @option{-mpopcntb} option allows GCC to generate the popcount and  
 13669 double precision FP reciprocal estimate instruction implemented on the  
 13670 POWER5 processor and other processors that support the PowerPC V2.02  
 13671 architecture.

13672 The @option{-mfprnd} option allows GCC to generate the FP round to  
 13673 integer instructions implemented on the POWER5+ processor and other  
 13674 processors that support the PowerPC V2.03 architecture.

13675 The @option{-mcmpb} option allows GCC to generate the compare bytes  
 13676 instruction implemented on the POWER6 processor and other processors  
 13677 that support the PowerPC V2.05 architecture.

13678 The @option{-mmfpgpr} option allows GCC to generate the FP move to/from  
 13679 general purpose register instructions implemented on the POWER6X  
 13680 processor and other processors that support the extended PowerPC V2.05  
 13681 architecture.

13682 The @option{-mhard-dfp} option allows GCC to generate the decimal floating  
 13683 point instructions implemented on some POWER processors.

13685 The @option{-mpowerpc64} option allows GCC to generate the additional  
 13686 64-bit instructions that are found in the full PowerPC64 architecture

13687 and to treat GPRs as 64-bit, doubleword quantities. GCC defaults to  
 13688 @option{-mno-powerpc64}.

13690 If you specify both @option{-mno-power} and @option{-mno-powerpc}, GCC  
 13691 will use only the instructions in the common subset of both  
 13692 architectures plus some special AIX common-mode calls, and will not use  
 13693 the MQ register. Specifying both @option{-mpower} and @option{-mpowerpc}  
 13694 permits GCC to use any instruction from either architecture and to  
 13695 allow use of the MQ register; specify this for the Motorola MPC601.

13697 @item -mnew-mnemonics  
 13698 @itemx -mold-mnemonics  
 13699 @opindex mnew-mnemonics  
 13700 @opindex mold-mnemonics  
 13701 Select which mnemonics to use in the generated assembler code. With  
 13702 @option{-mnew-mnemonics}, GCC uses the assembler mnemonics defined for  
 13703 the PowerPC architecture. With @option{-mold-mnemonics} it uses the  
 13704 assembler mnemonics defined for the POWER architecture. Instructions  
 13705 defined in only one architecture have only one mnemonic; GCC uses that  
 13706 mnemonic irrespective of which of these options is specified.

13708 GCC defaults to the mnemonics appropriate for the architecture in  
 13709 use. Specifying @option{-mcpu=@var{cpu\_type}} sometimes overrides the  
 13710 value of these option. Unless you are building a cross-compiler, you  
 13711 should normally not specify either @option{-mnew-mnemonics} or  
 13712 @option{-mold-mnemonics}, but should instead accept the default.

13714 @item -mcpu=@var{cpu\_type}  
 13715 @opindex mcpu  
 13716 Set architecture type, register usage, choice of mnemonics, and  
 13717 instruction scheduling parameters for machine type @var{cpu\_type}.  
 13718 Supported values for @var{cpu\_type} are @samp{401}, @samp{403},  
 13719 @samp{405}, @samp{405fp}, @samp{440}, @samp{440fp}, @samp{464}, @samp{464fp},  
 13720 @samp{505}, @samp{601}, @samp{602}, @samp{603}, @samp{603e}, @samp{604},  
 13721 @samp{604e}, @samp{620}, @samp{630}, @samp{740}, @samp{7400},  
 13722 @samp{7450}, @samp{750}, @samp{801}, @samp{821}, @samp{823},  
 13723 @samp{860}, @samp{970}, @samp{8540}, @samp{e300c2}, @samp{e300c3},  
 13724 @samp{e500mc}, @samp{ec603e}, @samp{G3}, @samp{G4}, @samp{G5},  
 13725 @samp{power}, @samp{power2}, @samp{power3}, @samp{power4},  
 13726 @samp{power5}, @samp{power5+}, @samp{power6}, @samp{power6x}, @samp{power7},  
 13727 @samp{common}, @samp{powerpc}, @samp{powerpc64}, @samp{rios},  
 13728 @samp{rios1}, @samp{rios2}, @samp{rsc}, and @samp{rs64}.

13730 @option{-mcpu=common} selects a completely generic processor. Code  
 13731 generated under this option will run on any POWER or PowerPC processor.  
 13732 GCC will use only the instructions in the common subset of both  
 13733 architectures, and will not use the MQ register. GCC assumes a generic  
 13734 processor model for scheduling purposes.

13736 @option{-mcpu=power}, @option{-mcpu=power2}, @option{-mcpu=powerpc}, and  
 13737 @option{-mcpu=powerpc64} specify generic POWER, POWER2, pure 32-bit  
 13738 PowerPC (i.e., not MPC601), and 64-bit PowerPC architecture machine  
 13739 types, with an appropriate, generic processor model assumed for  
 13740 scheduling purposes.

13742 The other options specify a specific processor. Code generated under  
 13743 those options will run best on that processor, and may not run at all on  
 13744 others.

13746 The @option{-mcpu} options automatically enable or disable the  
 13747 following options:

13749 @gccoptlist{-maltivec -mfprnd -mhard-float -mmfcrf -mmultiple @gol  
 13750 -mnew-mnemonics -mpopcntb -mpower -mpower2 -mpowerpc64 @gol  
 13751 -mpowerpc-gpopt -mpowerpc-gfxopt -msingle-float -mdouble-float @gol  
 13752 -msimple-fpu -mstring -mmulhw -mdlmzb -mmfpgpr}

13754 The particular options set for any particular CPU will vary between  
 13755 compiler versions, depending on what setting seems to produce optimal  
 13756 code for that CPU; it doesn't necessarily reflect the actual hardware's  
 13757 capabilities. If you wish to set an individual option to a particular  
 13758 value, you may specify it after the @option{-mcpu} option, like  
 13759 @samp{-mcpu=970 -mno-altivec}.

13761 On AIX, the @option{-maltivec} and @option{-mpowerpc64} options are  
 13762 not enabled or disabled by the @option{-mcpu} option at present because  
 13763 AIX does not have full support for these options. You may still  
 13764 enable or disable them individually if you're sure it'll work in your  
 13765 environment.

13767 @item -mtune=@var{cpu\_type}  
 13768 @opindex mtune  
 13769 Set the instruction scheduling parameters for machine type  
 13770 @var{cpu\_type}, but do not set the architecture type, register usage, or  
 13771 choice of mnemonics, as @option{-mcpu=@var{cpu\_type}} would. The same  
 13772 values for @var{cpu\_type} are used for @option{-mtune} as for  
 13773 @option{-mcpu}. If both are specified, the code generated will use the  
 13774 architecture, registers, and mnemonics set by @option{-mcpu}, but the  
 13775 scheduling parameters set by @option{-mtune}.

13777 @item -mswdiv  
 13778 @itemx -mno-swdiv  
 13779 @opindex mswdiv  
 13780 @opindex mno-swdiv  
 13781 Generate code to compute division as reciprocal estimate and iterative  
 13782 refinement, creating opportunities for increased throughput. This  
 13783 feature requires: optional PowerPC Graphics instruction set for single  
 13784 precision and FRE instruction for double precision, assuming divides  
 13785 cannot generate user-visible traps, and the domain values not include  
 13786 Infinities, denormals or zero denominator.

13788 @item -maltivec  
 13789 @itemx -mno-altivec  
 13790 @opindex altivec  
 13791 @opindex mno-altivec  
 13792 Generate code that uses (does not use) AltiVec instructions, and also  
 13793 enable the use of built-in functions that allow more direct access to  
 13794 the AltiVec instruction set. You may also need to set  
 13795 @option{-mabi=altivec} to adjust the current ABI with AltiVec ABI  
 13796 enhancements.

13798 @item -mvrsave  
 13799 @itemx -mno-vrsave  
 13800 @opindex mvrsave  
 13801 @opindex mno-vrsave  
 13802 Generate VRSAVE instructions when generating AltiVec code.

13804 @item -mgen-cell-microcode  
 13805 @opindex mgen-cell-microcode  
 13806 Generate Cell microcode instructions

13808 @item -mwarn-cell-microcode  
 13809 @opindex mwarn-cell-microcode  
 13810 Warning when a Cell microcode instruction is going to emitted. An example  
 13811 of a Cell microcode instruction is a variable shift.

13813 @item -msecure-plt  
 13814 @opindex msecure-plt  
 13815 Generate code that allows ld and ld.so to build executables and shared  
 13816 libraries with non-exec .plt and .got sections. This is a PowerPC  
 13817 32-bit SYSV ABI option.

13819 @item -mbss-plt  
 13820 @opindex mbss-plt  
 13821 Generate code that uses a BSS .plt section that ld.so fills in, and  
 13822 requires .plt and .got sections that are both writable and executable.  
 13823 This is a PowerPC 32-bit SYSV ABI option.

13825 @item -misel  
 13826 @itemx -mno-isel  
 13827 @opindex misel  
 13828 @opindex mno-isel  
 13829 This switch enables or disables the generation of ISEL instructions.

13831 @item -misel=@var{yes/no}  
 13832 This switch has been deprecated. Use @option{-misel} and  
 13833 @option{-mno-isel} instead.

13835 @item -mspe  
 13836 @itemx -mno-spe  
 13837 @opindex mspe  
 13838 @opindex mno-spe  
 13839 This switch enables or disables the generation of SPE simd  
 13840 instructions.

13842 @item -mpaired  
 13843 @itemx -mno-paired  
 13844 @opindex mpaired  
 13845 @opindex mno-paired  
 13846 This switch enables or disables the generation of PAIRED simd  
 13847 instructions.

13849 @item -mspe=@var{yes/no}  
 13850 This option has been deprecated. Use @option{-mspe} and  
 13851 @option{-mno-spe} instead.

13853 @item -mfloat-gprs=@var{yes/single/double/no}  
 13854 @itemx -mfloat-gprs  
 13855 @opindex mfloat-gprs  
 13856 This switch enables or disables the generation of floating point  
 13857 operations on the general purpose registers for architectures that  
 13858 support it.

13860 The argument @var{yes} or @var{single} enables the use of  
 13861 single-precision floating point operations.

13863 The argument @var{double} enables the use of single and  
 13864 double-precision floating point operations.

13866 The argument @var{no} disables floating point operations on the  
 13867 general purpose registers.

13869 This option is currently only available on the MPC854x.

13871 @item -m32  
 13872 @itemx -m64  
 13873 @opindex m32  
 13874 @opindex m64  
 13875 Generate code for 32-bit or 64-bit environments of Darwin and SVR4  
 13876 targets (including GNU/Linux). The 32-bit environment sets int, long  
 13877 and pointer to 32 bits and generates code that runs on any PowerPC  
 13878 variant. The 64-bit environment sets int to 32 bits and long and  
 13879 pointer to 64 bits, and generates code for PowerPC64, as for  
 13880 @option{-mpowerpc64}.

13882 @item -mfull-toc  
 13883 @itemx -mno-fp-in-toc  
 13884 @itemx -mno-sum-in-toc

13885 @itemx -mminimal-toc  
 13886 @opindex mfull-toc  
 13887 @opindex mno-fp-in-toc  
 13888 @opindex mno-sum-in-toc  
 13889 @opindex mminimal-toc  
 13890 Modify generation of the TOC (Table Of Contents), which is created for  
 13891 every executable file. The @option{-mfull-toc} option is selected by  
 13892 default. In that case, GCC will allocate at least one TOC entry for  
 13893 each unique non-automatic variable reference in your program. GCC  
 13894 will also place floating-point constants in the TOC@. However, only  
 13895 16,384 entries are available in the TOC@.

13897 If you receive a linker error message that saying you have overflowed  
 13898 the available TOC space, you can reduce the amount of TOC space used  
 13899 with the @option{-mno-fp-in-toc} and @option{-mno-sum-in-toc} options.  
 13900 @option{-mno-fp-in-toc} prevents GCC from putting floating-point  
 13901 constants in the TOC and @option{-mno-sum-in-toc} forces GCC to  
 13902 generate code to calculate the sum of an address and a constant at  
 13903 run-time instead of putting that sum into the TOC@. You may specify one  
 13904 or both of these options. Each causes GCC to produce very slightly  
 13905 slower and larger code at the expense of conserving TOC space.

13907 If you still run out of space in the TOC even when you specify both of  
 13908 these options, specify @option{-mminimal-toc} instead. This option causes  
 13909 GCC to make only one TOC entry for every file. When you specify this  
 13910 option, GCC will produce code that is slower and larger but which  
 13911 uses extremely little TOC space. You may wish to use this option  
 13912 only on files that contain less frequently executed code.

13914 @item -maix64  
 13915 @itemx -maix32  
 13916 @opindex maix64  
 13917 @opindex maix32  
 13918 Enable 64-bit AIX ABI and calling convention: 64-bit pointers, 64-bit  
 13919 @code{long} type, and the infrastructure needed to support them.  
 13920 Specifying @option{-maix64} implies @option{-mpowerpc64} and  
 13921 @option{-mpowerpc}, while @option{-maix32} disables the 64-bit ABI and  
 13922 implies @option{-mno-powerpc64}. GCC defaults to @option{-maix32}.

13924 @item -mxl-compat  
 13925 @itemx -mno-xl-compat  
 13926 @opindex mxl-compat  
 13927 @opindex mno-xl-compat  
 13928 Produce code that conforms more closely to IBM XL compiler semantics  
 13929 when using AIX-compatible ABI@. Pass floating-point arguments to  
 13930 prototyped functions beyond the register save area (RSA) on the stack  
 13931 in addition to argument FPRs. Do not assume that most significant  
 13932 double in 128-bit long double value is properly rounded when comparing  
 13933 values and converting to double. Use XL symbol names for long double  
 13934 support routines.

13936 The AIX calling convention was extended but not initially documented to  
 13937 handle an obscure K&R C case of calling a function that takes the  
 13938 address of its arguments with fewer arguments than declared. IBM XL  
 13939 compilers access floating point arguments which do not fit in the  
 13940 RSA from the stack when a subroutine is compiled without  
 13941 optimization. Because always storing floating-point arguments on the  
 13942 stack is inefficient and rarely needed, this option is not enabled by  
 13943 default and only is necessary when calling subroutines compiled by IBM  
 13944 XL compilers without optimization.

13946 @item -mpe  
 13947 @opindex mpe  
 13948 Support @dfn{IBM RS/6000 SP} @dfn{Parallel Environment} (PE)@. Link an  
 13949 application written to use message passing with special startup code to  
 13950 enable the application to run. The system must have PE installed in the

13951 standard location (@file{/usr/lpp/ppe.poe/}), or the @file{specs} file  
 13952 must be overridden with the @option{-specs=} option to specify the  
 13953 appropriate directory location. The Parallel Environment does not  
 13954 support threads, so the @option{-mpe} option and the @option{-pthread}  
 13955 option are incompatible.

13957 @item -malign-natural  
 13958 @itemx -malign-power  
 13959 @opindex malign-natural  
 13960 @opindex malign-power  
 13961 On AIX, 32-bit Darwin, and 64-bit PowerPC GNU/Linux, the option  
 13962 @option{-malign-natural} overrides the ABI-defined alignment of larger  
 13963 types, such as floating-point doubles, on their natural size-based boundary.  
 13964 The option @option{-malign-power} instructs GCC to follow the ABI-specified  
 13965 alignment rules. GCC defaults to the standard alignment defined in the ABI@.

13967 On 64-bit Darwin, natural alignment is the default, and @option{-malign-power}  
 13968 is not supported.

13970 @item -msoft-float  
 13971 @itemx -mhard-float  
 13972 @opindex msoft-float  
 13973 @opindex mhard-float  
 13974 Generate code that does not use (uses) the floating-point register set.  
 13975 Software floating point emulation is provided if you use the  
 13976 @option{-msoft-float} option, and pass the option to GCC when linking.

13978 @item -msingle-float  
 13979 @itemx -mdouble-float  
 13980 @opindex msingle-float  
 13981 @opindex mdouble-float  
 13982 Generate code for single or double-precision floating point operations.  
 13983 @option{-mdouble-float} implies @option{-msingle-float}.

13985 @item -msimple-fpu  
 13986 @opindex msimple-fpu  
 13987 Do not generate sqrt and div instructions for hardware floating point unit.

13989 @item -mfpu  
 13990 @opindex mfp  
 13991 Specify type of floating point unit. Valid values are @var{sp\_lite}  
 13992 (equivalent to -msingle-float -msimple-fpu), @var{dp\_lite} (equivalent  
 13993 to -mdouble-float -msimple-fpu), @var{sp\_full} (equivalent to -msingle-float),  
 13994 and @var{dp\_full} (equivalent to -mdouble-float).

13996 @item -mxilinx-fpu  
 13997 @opindex mxilinx-fpu  
 13998 Perform optimizations for floating point unit on Xilinx PPC 405/440.

14000 @item -mmultiple  
 14001 @itemx -mno-multiple  
 14002 @opindex mmultiple  
 14003 @opindex mno-multiple  
 14004 Generate code that uses (does not use) the load multiple word  
 14005 instructions and the store multiple word instructions. These  
 14006 instructions are generated by default on POWER systems, and not  
 14007 generated on PowerPC systems. Do not use @option{-mmultiple} on little  
 14008 endian PowerPC systems, since those instructions do not work when the  
 14009 processor is in little endian mode. The exceptions are PPC740 and  
 14010 PPC750 which permit the instructions usage in little endian mode.

14012 @item -mstring  
 14013 @itemx -mno-string  
 14014 @opindex mstring  
 14015 @opindex mno-string  
 14016 Generate code that uses (does not use) the load string instructions

14017 and the store string word instructions to save multiple registers and  
 14018 do small block moves. These instructions are generated by default on  
 14019 POWER systems, and not generated on PowerPC systems. Do not use  
 14020 `@option{-mstring}` on little endian PowerPC systems, since those  
 14021 instructions do not work when the processor is in little endian mode.  
 14022 The exceptions are PPC740 and PPC750 which permit the instructions  
 14023 usage in little endian mode.

14025 `@item -mupdate`  
 14026 `@itemx -mno-update`  
 14027 `@opindex mupdate`  
 14028 `@opindex mno-update`  
 14029 Generate code that uses (does not use) the load or store instructions  
 14030 that update the base register to the address of the calculated memory  
 14031 location. These instructions are generated by default. If you use  
 14032 `@option{-mno-update}`, there is a small window between the time that the  
 14033 stack pointer is updated and the address of the previous frame is  
 14034 stored, which means code that walks the stack frame across interrupts or  
 14035 signals may get corrupted data.

14037 `@item -mavoid-indexed-addresses`  
 14038 `@itemx -mno-avoid-indexed-addresses`  
 14039 `@opindex mavoid-indexed-addresses`  
 14040 `@opindex mno-avoid-indexed-addresses`  
 14041 Generate code that tries to avoid (not avoid) the use of indexed load  
 14042 or store instructions. These instructions can incur a performance  
 14043 penalty on Power6 processors in certain situations, such as when  
 14044 stepping through large arrays that cross a 16M boundary. This option  
 14045 is enabled by default when targetting Power6 and disabled otherwise.

14047 `@item -mfused-madd`  
 14048 `@itemx -mno-fused-madd`  
 14049 `@opindex mfused-madd`  
 14050 `@opindex mno-fused-madd`  
 14051 Generate code that uses (does not use) the floating point multiply and  
 14052 accumulate instructions. These instructions are generated by default if  
 14053 hardware floating is used.

14055 `@item -mmulhw`  
 14056 `@itemx -mno-mulhw`  
 14057 `@opindex mmulhw`  
 14058 `@opindex mno-mulhw`  
 14059 Generate code that uses (does not use) the half-word multiply and  
 14060 multiply-accumulate instructions on the IBM 405, 440 and 464 processors.  
 14061 These instructions are generated by default when targetting those  
 14062 processors.

14064 `@item -mdlmzb`  
 14065 `@itemx -mno-dlmzb`  
 14066 `@opindex mdlmzb`  
 14067 `@opindex mno-dlmzb`  
 14068 Generate code that uses (does not use) the string-search `@samp{dlmzb}`  
 14069 instruction on the IBM 405, 440 and 464 processors. This instruction is  
 14070 generated by default when targetting those processors.

14072 `@item -mno-bit-align`  
 14073 `@itemx -mbit-align`  
 14074 `@opindex mno-bit-align`  
 14075 `@opindex mbit-align`  
 14076 On System V.4 and embedded PowerPC systems do not (do) force structures  
 14077 and unions that contain bit-fields to be aligned to the base type of the  
 14078 bit-field.

14080 For example, by default a structure containing nothing but 8  
 14081 `@code{unsigned}` bit-fields of length 1 would be aligned to a 4 byte  
 14082 boundary and have a size of 4 bytes. By using `@option{-mno-bit-align}`,

14083 the structure would be aligned to a 1 byte boundary and be one byte in  
 14084 size.

14086 `@item -mno-strict-align`  
 14087 `@itemx -mstrict-align`  
 14088 `@opindex mno-strict-align`  
 14089 `@opindex mstrict-align`  
 14090 On System V.4 and embedded PowerPC systems do not (do) assume that  
 14091 unaligned memory references will be handled by the system.

14093 `@item -mrelocatable`  
 14094 `@itemx -mno-relocatable`  
 14095 `@opindex mrelocatable`  
 14096 `@opindex mno-relocatable`  
 14097 On embedded PowerPC systems generate code that allows (does not allow)  
 14098 the program to be relocated to a different address at runtime. If you  
 14099 use `@option{-mrelocatable}` on any module, all objects linked together must  
 14100 be compiled with `@option{-mrelocatable}` or `@option{-mrelocatable-lib}`.

14102 `@item -mrelocatable-lib`  
 14103 `@itemx -mno-relocatable-lib`  
 14104 `@opindex mrelocatable-lib`  
 14105 `@opindex mno-relocatable-lib`  
 14106 On embedded PowerPC systems generate code that allows (does not allow)  
 14107 the program to be relocated to a different address at runtime. Modules  
 14108 compiled with `@option{-mrelocatable-lib}` can be linked with either modules  
 14109 compiled without `@option{-mrelocatable}` and `@option{-mrelocatable-lib}` or  
 14110 with modules compiled with the `@option{-mrelocatable}` options.

14112 `@item -mno-toc`  
 14113 `@itemx -mtoc`  
 14114 `@opindex mno-toc`  
 14115 `@opindex mtoc`  
 14116 On System V.4 and embedded PowerPC systems do not (do) assume that  
 14117 register 2 contains a pointer to a global area pointing to the addresses  
 14118 used in the program.

14120 `@item -mlittle`  
 14121 `@itemx -mlittle-endian`  
 14122 `@opindex mlittle`  
 14123 `@opindex mlittle-endian`  
 14124 On System V.4 and embedded PowerPC systems compile code for the  
 14125 processor in little endian mode. The `@option{-mlittle-endian}` option is  
 14126 the same as `@option{-mlittle}`.

14128 `@item -mbig`  
 14129 `@itemx -mbig-endian`  
 14130 `@opindex mbig`  
 14131 `@opindex mbig-endian`  
 14132 On System V.4 and embedded PowerPC systems compile code for the  
 14133 processor in big endian mode. The `@option{-mbig-endian}` option is  
 14134 the same as `@option{-mbig}`.

14136 `@item -mdynamic-no-pic`  
 14137 `@opindex mdynamic-no-pic`  
 14138 On Darwin and Mac OS X systems, compile code so that it is not  
 14139 relocatable, but that its external references are relocatable. The  
 14140 resulting code is suitable for applications, but not shared  
 14141 libraries.

14143 `@item -mprioritize-restricted-insns=@var{priority}`  
 14144 `@opindex mprioritize-restricted-insns`  
 14145 This option controls the priority that is assigned to  
 14146 dispatch-slot restricted instructions during the second scheduling  
 14147 pass. The argument `@var{priority}` takes the value `@var{0/1/2}` to assign  
 14148 `@var{no/highest/second-highest}` priority to dispatch slot restricted

14149 instructions.

14151 @item -msched-costly-dep=@var{dependence\_type}  
 14152 @opindex msched-costly-dep  
 14153 This option controls which dependences are considered costly  
 14154 by the target during instruction scheduling. The argument  
 14155 @var{dependence\_type} takes one of the following values:  
 14156 @var{no}: no dependence is costly,  
 14157 @var{all}: all dependences are costly,  
 14158 @var{true\_store\_to\_load}: a true dependence from store to load is costly,  
 14159 @var{store\_to\_load}: any dependence from store to load is costly,  
 14160 @var{number}: any dependence which latency >= @var{number} is costly.

14162 @item -minsert-sched-nops=@var{scheme}  
 14163 @opindex minsert-sched-nops  
 14164 This option controls which nop insertion scheme will be used during  
 14165 the second scheduling pass. The argument @var{scheme} takes one of the  
 14166 following values:  
 14167 @var{no}: Don't insert nops.  
 14168 @var{pad}: Pad with nops any dispatch group which has vacant issue slots,  
 14169 according to the scheduler's grouping.  
 14170 @var{regroup\_exact}: Insert nops to force costly dependent insns into  
 14171 separate groups. Insert exactly as many nops as needed to force an insn  
 14172 to a new group, according to the estimated processor grouping.  
 14173 @var{number}: Insert nops to force costly dependent insns into  
 14174 separate groups. Insert @var{number} nops to force an insn to a new group.

14176 @item -mcall-sysv  
 14177 @opindex mcall-sysv  
 14178 On System V.4 and embedded PowerPC systems compile code using calling  
 14179 conventions that adheres to the March 1995 draft of the System V  
 14180 Application Binary Interface, PowerPC processor supplement. This is the  
 14181 default unless you configured GCC using @samp{powerpc-\*-eablaix}.

14183 @item -mcall-sysv-eabi  
 14184 @opindex mcall-sysv-eabi  
 14185 Specify both @option{-mcall-sysv} and @option{-meabi} options.

14187 @item -mcall-sysv-noeabi  
 14188 @opindex mcall-sysv-noeabi  
 14189 Specify both @option{-mcall-sysv} and @option{-mno-eabi} options.

14191 @item -mcall-solaris  
 14192 @opindex mcall-solaris  
 14193 On System V.4 and embedded PowerPC systems compile code for the Solaris  
 14194 operating system.

14196 @item -mcall-linux  
 14197 @opindex mcall-linux  
 14198 On System V.4 and embedded PowerPC systems compile code for the  
 14199 Linux-based GNU system.

14201 @item -mcall-gnu  
 14202 @opindex mcall-gnu  
 14203 On System V.4 and embedded PowerPC systems compile code for the  
 14204 Hurd-based GNU system.

14206 @item -mcall-netbsd  
 14207 @opindex mcall-netbsd  
 14208 On System V.4 and embedded PowerPC systems compile code for the  
 14209 NetBSD operating system.

14211 @item -maix-struct-return  
 14212 @opindex maix-struct-return  
 14213 Return all structures in memory (as specified by the AIX ABI)@.

14215 @item -msvr4-struct-return  
 14216 @opindex msvr4-struct-return  
 14217 Return structures smaller than 8 bytes in registers (as specified by the  
 14218 SVR4 ABI)@.

14220 @item -mabi=@var{abi-type}  
 14221 @opindex mabi  
 14222 Extend the current ABI with a particular extension, or remove such extension.  
 14223 Valid values are @var{altivec}, @var{no-altivec}, @var{spe},  
 14224 @var{no-spe}, @var{ibmldouble}, @var{ieeeldouble}@.

14226 @item -mabi=spe  
 14227 @opindex mabi=spe  
 14228 Extend the current ABI with SPE ABI extensions. This does not change  
 14229 the default ABI, instead it adds the SPE ABI extensions to the current  
 14230 ABI@.

14232 @item -mabi=no-spe  
 14233 @opindex mabi=no-spe  
 14234 Disable BookE SPE ABI extensions for the current ABI@.

14236 @item -mabi=ibmldouble  
 14237 @opindex mabi=ibmldouble  
 14238 Change the current ABI to use IBM extended precision long double.  
 14239 This is a PowerPC 32-bit SYSV ABI option.

14241 @item -mabi=ieeeldouble  
 14242 @opindex mabi=ieeeldouble  
 14243 Change the current ABI to use IEEE extended precision long double.  
 14244 This is a PowerPC 32-bit Linux ABI option.

14246 @item -mprototype  
 14247 @itemx -mno-prototype  
 14248 @opindex mprototype  
 14249 @opindex mno-prototype  
 14250 On System V.4 and embedded PowerPC systems assume that all calls to  
 14251 variable argument functions are properly prototyped. Otherwise, the  
 14252 compiler must insert an instruction before every non prototyped call to  
 14253 set or clear bit 6 of the condition code register (@var{CR}) to  
 14254 indicate whether floating point values were passed in the floating point  
 14255 registers in case the function takes a variable arguments. With  
 14256 @option{-mprototype}, only calls to prototyped variable argument functions  
 14257 will set or clear the bit.

14259 @item -msim  
 14260 @opindex msim  
 14261 On embedded PowerPC systems, assume that the startup module is called  
 14262 @file{sim-crt0.o} and that the standard C libraries are @file{libsim.a} and  
 14263 @file{libc.a}. This is the default for @samp{powerpc-\*-eabisim}  
 14264 configurations.

14266 @item -mmvme  
 14267 @opindex mmvme  
 14268 On embedded PowerPC systems, assume that the startup module is called  
 14269 @file{crt0.o} and the standard C libraries are @file{libmvme.a} and  
 14270 @file{libc.a}.

14272 @item -mads  
 14273 @opindex mads  
 14274 On embedded PowerPC systems, assume that the startup module is called  
 14275 @file{crt0.o} and the standard C libraries are @file{libads.a} and  
 14276 @file{libc.a}.

14278 @item -myellowknife  
 14279 @opindex myellowknife  
 14280 On embedded PowerPC systems, assume that the startup module is called

14281 @file{crt0.o} and the standard C libraries are @file{libyk.a} and  
 14282 @file{libc.a}.

14284 @item -mvxworks  
 14285 @opindex mvxworks  
 14286 On System V.4 and embedded PowerPC systems, specify that you are  
 14287 compiling for a VxWorks system.

14289 @item -memb  
 14290 @opindex memb  
 14291 On embedded PowerPC systems, set the @var{PPC\_EMB} bit in the ELF flags  
 14292 header to indicate that @samp{eabi} extended relocations are used.

14294 @item -meabi  
 14295 @itemx -mno-eabi  
 14296 @opindex meabi  
 14297 @opindex mno-eabi  
 14298 On System V.4 and embedded PowerPC systems do (do not) adhere to the  
 14299 Embedded Applications Binary Interface (eabi) which is a set of  
 14300 modifications to the System V.4 specifications. Selecting @option{-meabi}  
 14301 means that the stack is aligned to an 8 byte boundary, a function  
 14302 @code{\_\_eabi} is called to from @code{main} to set up the eabi  
 14303 environment, and the @option{-msdata} option can use both @code{r2} and  
 14304 @code{r13} to point to two separate small data areas. Selecting  
 14305 @option{-mno-eabi} means that the stack is aligned to a 16 byte boundary,  
 14306 do not call an initialization function from @code{main}, and the  
 14307 @option{-msdata} option will only use @code{r13} to point to a single  
 14308 small data area. The @option{-meabi} option is on by default if you  
 14309 configured GCC using one of the @samp{powerpc\*-\*-eabi\*} options.

14311 @item -msdata=eabi  
 14312 @opindex msdata=eabi  
 14313 On System V.4 and embedded PowerPC systems, put small initialized  
 14314 @code{const} global and static data in the @samp{.sdata2} section, which  
 14315 is pointed to by register @code{r2}. Put small initialized  
 14316 non-@code{const} global and static data in the @samp{.sdata} section,  
 14317 which is pointed to by register @code{r13}. Put small uninitialized  
 14318 global and static data in the @samp{.sbss} section, which is adjacent to  
 14319 the @samp{.sdata} section. The @option{-msdata=eabi} option is  
 14320 incompatible with the @option{-mrelocatable} option. The  
 14321 @option{-msdata=eabi} option also sets the @option{-memb} option.

14323 @item -msdata=sysv  
 14324 @opindex msdata=sysv  
 14325 On System V.4 and embedded PowerPC systems, put small global and static  
 14326 data in the @samp{.sdata} section, which is pointed to by register  
 14327 @code{r13}. Put small uninitialized global and static data in the  
 14328 @samp{.sbss} section, which is adjacent to the @samp{.sdata} section.  
 14329 The @option{-msdata=sysv} option is incompatible with the  
 14330 @option{-mrelocatable} option.

14332 @item -msdata=default  
 14333 @itemx -msdata  
 14334 @opindex msdata=default  
 14335 @opindex msdata  
 14336 On System V.4 and embedded PowerPC systems, if @option{-meabi} is used,  
 14337 compile code the same as @option{-msdata=eabi}, otherwise compile code the  
 14338 same as @option{-msdata=sysv}.

14340 @item -msdata=data  
 14341 @opindex msdata=data  
 14342 On System V.4 and embedded PowerPC systems, put small global  
 14343 data in the @samp{.sdata} section. Put small uninitialized global  
 14344 data in the @samp{.sbss} section. Do not use register @code{r13}  
 14345 to address small data however. This is the default behavior unless  
 14346 other @option{-msdata} options are used.

14348 @item -msdata=none  
 14349 @itemx -mno-sdata  
 14350 @opindex msdata=none  
 14351 @opindex mno-sdata  
 14352 On embedded PowerPC systems, put all initialized global and static data  
 14353 in the @samp{.data} section, and all uninitialized data in the  
 14354 @samp{.bss} section.

14356 @item -G @var{num}  
 14357 @opindex G  
 14358 @cindex smaller data references (PowerPC)  
 14359 @cindex .sdata/.sdata2 references (PowerPC)  
 14360 On embedded PowerPC systems, put global and static items less than or  
 14361 equal to @var{num} bytes into the small data or bss sections instead of  
 14362 the normal data or bss section. By default, @var{num} is 8. The  
 14363 @option{-G @var{num}} switch is also passed to the linker.  
 14364 All modules should be compiled with the same @option{-G @var{num}} value.

14366 @item -mregnames  
 14367 @itemx -mno-regnames  
 14368 @opindex mregnames  
 14369 @opindex mno-regnames  
 14370 On System V.4 and embedded PowerPC systems do (do not) emit register  
 14371 names in the assembly language output using symbolic forms.

14373 @item -mlongcall  
 14374 @itemx -mno-longcall  
 14375 @opindex mlongcall  
 14376 @opindex mno-longcall  
 14377 By default assume that all calls are far away so that a longer more  
 14378 expensive calling sequence is required. This is required for calls  
 14379 further than 32 megabytes (33,554,432 bytes) from the current location.  
 14380 A short call will be generated if the compiler knows  
 14381 the call cannot be that far away. This setting can be overridden by  
 14382 the @code{shortcall} function attribute, or by @code{#pragma  
 14383 longcall(0)}.

14385 Some linkers are capable of detecting out-of-range calls and generating  
 14386 glue code on the fly. On these systems, long calls are unnecessary and  
 14387 generate slower code. As of this writing, the AIX linker can do this,  
 14388 as can the GNU linker for PowerPC/64. It is planned to add this feature  
 14389 to the GNU linker for 32-bit PowerPC systems as well.

14391 On Darwin/PPC systems, @code{#pragma longcall} will generate ``jbsr  
 14392 callee, L42'', plus a ``branch island'' (glue code). The two target  
 14393 addresses represent the callee and the ``branch island''. The  
 14394 Darwin/PPC linker will prefer the first address and generate a ``bl  
 14395 callee'' if the PPC ``bl'' instruction will reach the callee directly;  
 14396 otherwise, the linker will generate ``bl L42'' to call the ``branch  
 14397 island''. The ``branch island'' is appended to the body of the  
 14398 calling function; it computes the full 32-bit address of the callee  
 14399 and jumps to it.

14401 On Mach-O (Darwin) systems, this option directs the compiler emit to  
 14402 the glue for every direct call, and the Darwin linker decides whether  
 14403 to use or discard it.

14405 In the future, we may cause GCC to ignore all longcall specifications  
 14406 when the linker is known to generate glue.

14408 @item -pthread  
 14409 @opindex pthread  
 14410 Adds support for multithreading with the @dfn{pthreads} library.  
 14411 This option sets flags for both the preprocessor and linker.

14413 @end table

14415 @node S/390 and zSeries Options  
 14416 @subsection S/390 and zSeries Options  
 14417 @cindex S/390 and zSeries Options

14419 These are the @samp{-m} options defined for the S/390 and zSeries architecture.

14421 @table @gcctabopt  
 14422 @item -mhard-float  
 14423 @itemx -msoft-float  
 14424 @opindex mhard-float  
 14425 @opindex msoft-float  
 14426 Use (do not use) the hardware floating-point instructions and registers  
 14427 for floating-point operations. When @option{-msoft-float} is specified,  
 14428 functions in @file{libgcc.a} will be used to perform floating-point  
 14429 operations. When @option{-mhard-float} is specified, the compiler  
 14430 generates IEEE floating-point instructions. This is the default.

14432 @item -mhard-dfp  
 14433 @itemx -mno-hard-dfp  
 14434 @opindex mhard-dfp  
 14435 @opindex mno-hard-dfp  
 14436 Use (do not use) the hardware decimal-floating-point instructions for  
 14437 decimal-floating-point operations. When @option{-mno-hard-dfp} is  
 14438 specified, functions in @file{libgcc.a} will be used to perform  
 14439 decimal-floating-point operations. When @option{-mhard-dfp} is  
 14440 specified, the compiler generates decimal-floating-point hardware  
 14441 instructions. This is the default for @option{-march=z9-ec} or higher.

14443 @item -mlong-double-64  
 14444 @itemx -mlong-double-128  
 14445 @opindex mlong-double-64  
 14446 @opindex mlong-double-128  
 14447 These switches control the size of @code{long double} type. A size  
 14448 of 64bit makes the @code{long double} type equivalent to the @code{double}  
 14449 type. This is the default.

14451 @item -mbackchain  
 14452 @itemx -mno-backchain  
 14453 @opindex mbackchain  
 14454 @opindex mno-backchain  
 14455 Store (do not store) the address of the caller's frame as backchain pointer  
 14456 into the callee's stack frame.  
 14457 A backchain may be needed to allow debugging using tools that do not understand  
 14458 DWARF-2 call frame information.  
 14459 When @option{-mno-packed-stack} is in effect, the backchain pointer is stored  
 14460 at the bottom of the stack frame; when @option{-mpacked-stack} is in effect,  
 14461 the backchain is placed into the topmost word of the 96/160 byte register  
 14462 save area.

14464 In general, code compiled with @option{-mbackchain} is call-compatible with  
 14465 code compiled with @option{-mno-backchain}; however, use of the backchain  
 14466 for debugging purposes usually requires that the whole binary is built with  
 14467 @option{-mbackchain}. Note that the combination of @option{-mbackchain},  
 14468 @option{-mpacked-stack} and @option{-mhard-float} is not supported. In order  
 14469 to build a linux kernel use @option{-msoft-float}.

14471 The default is to not maintain the backchain.

14473 @item -mpacked-stack  
 14474 @itemx -mno-packed-stack  
 14475 @opindex mpacked-stack  
 14476 @opindex mno-packed-stack  
 14477 Use (do not use) the packed stack layout. When @option{-mno-packed-stack} is  
 14478 specified, the compiler uses the all fields of the 96/160 byte register save

14479 area only for their default purpose; unused fields still take up stack space.  
 14480 When @option{-mpacked-stack} is specified, register save slots are densely  
 14481 packed at the top of the register save area; unused space is reused for other  
 14482 purposes, allowing for more efficient use of the available stack space.  
 14483 However, when @option{-mbackchain} is also in effect, the topmost word of  
 14484 the save area is always used to store the backchain, and the return address  
 14485 register is always saved two words below the backchain.

14487 As long as the stack frame backchain is not used, code generated with  
 14488 @option{-mpacked-stack} is call-compatible with code generated with  
 14489 @option{-mno-packed-stack}. Note that some non-FSF releases of GCC 2.95 for  
 14490 S/390 or zSeries generated code that uses the stack frame backchain at run  
 14491 time, not just for debugging purposes. Such code is not call-compatible  
 14492 with code compiled with @option{-mpacked-stack}. Also, note that the  
 14493 combination of @option{-mbackchain},  
 14494 @option{-mpacked-stack} and @option{-mhard-float} is not supported. In order  
 14495 to build a linux kernel use @option{-msoft-float}.

14497 The default is to not use the packed stack layout.

14499 @item -msmall-exec  
 14500 @itemx -mno-small-exec  
 14501 @opindex msmall-exec  
 14502 @opindex mno-small-exec  
 14503 Generate (or do not generate) code using the @code{bras} instruction  
 14504 to do subroutine calls.  
 14505 This only works reliably if the total executable size does not  
 14506 exceed 64k. The default is to use the @code{basr} instruction instead,  
 14507 which does not have this limitation.

14509 @item -m64  
 14510 @itemx -m31  
 14511 @opindex m64  
 14512 @opindex m31  
 14513 When @option{-m31} is specified, generate code compliant to the  
 14514 GNU/Linux for S/390 ABI@. When @option{-m64} is specified, generate  
 14515 code compliant to the GNU/Linux for zSeries ABI@. This allows GCC in  
 14516 particular to generate 64-bit instructions. For the @samp{s390}  
 14517 targets, the default is @option{-m31}, while the @samp{s390x}  
 14518 targets default to @option{-m64}.

14520 @item -mzarch  
 14521 @itemx -mesa  
 14522 @opindex mzarch  
 14523 @opindex mesa  
 14524 When @option{-mzarch} is specified, generate code using the  
 14525 instructions available on z/Architecture.  
 14526 When @option{-mesa} is specified, generate code using the  
 14527 instructions available on ESA/390. Note that @option{-mesa} is  
 14528 not possible with @option{-m64}.  
 14529 When generating code compliant to the GNU/Linux for S/390 ABI,  
 14530 the default is @option{-mesa}. When generating code compliant  
 14531 to the GNU/Linux for zSeries ABI, the default is @option{-mzarch}.

14533 @item -mmvcl  
 14534 @itemx -mno-mvcl  
 14535 @opindex mmvcl  
 14536 @opindex mno-mvcl  
 14537 Generate (or do not generate) code using the @code{mvcl} instruction  
 14538 to perform block moves. When @option{-mno-mvcl} is specified,  
 14539 use a @code{mvc} loop instead. This is the default unless optimizing for  
 14540 size.

14542 @item -mdebug  
 14543 @itemx -mno-debug  
 14544 @opindex mdebug

14545 @opindex mno-debug  
 14546 Print (or do not print) additional debug information when compiling.  
 14547 The default is to not print debug information.

14549 @item -march=@var{cpu-type}  
 14550 @opindex march  
 14551 Generate code that will run on @var{cpu-type}, which is the name of a system  
 14552 representing a certain processor type. Possible values for  
 14553 @var{cpu-type} are @samp{g5}, @samp{g6}, @samp{z900}, @samp{z990},  
 14554 @samp{z9-109}, @samp{z9-ec} and @samp{z10}.  
 14555 When generating code using the instructions available on z/Architecture,  
 14556 the default is @option{-march=z900}. Otherwise, the default is  
 14557 @option{-march=g5}.

14559 @item -mtune=@var{cpu-type}  
 14560 @opindex mtune  
 14561 Tune to @var{cpu-type} everything applicable about the generated code,  
 14562 except for the ABI and the set of available instructions.  
 14563 The list of @var{cpu-type} values is the same as for @option{-march}.  
 14564 The default is the value used for @option{-march}.

14566 @item -mtpf-trace  
 14567 @itemx -mno-tpf-trace  
 14568 @opindex mtpf-trace  
 14569 @opindex mno-tpf-trace  
 14570 Generate code that adds (does not add) in TPF OS specific branches to trace  
 14571 routines in the operating system. This option is off by default, even  
 14572 when compiling for the TPF OS@.

14574 @item -mfused-madd  
 14575 @itemx -mno-fused-madd  
 14576 @opindex mfused-madd  
 14577 @opindex mno-fused-madd  
 14578 Generate code that uses (does not use) the floating point multiply and  
 14579 accumulate instructions. These instructions are generated by default if  
 14580 hardware floating point is used.

14582 @item -mwarn-framesize=@var{framesize}  
 14583 @opindex mwarn-framesize  
 14584 Emit a warning if the current function exceeds the given frame size. Because  
 14585 this is a compile time check it doesn't need to be a real problem when the progr  
 14586 runs. It is intended to identify functions which most probably cause  
 14587 a stack overflow. It is useful to be used in an environment with limited stack  
 14588 size e.g.@: the linux kernel.

14590 @item -mwarn-dynamicstack  
 14591 @opindex mwarn-dynamicstack  
 14592 Emit a warning if the function calls alloca or uses dynamically  
 14593 sized arrays. This is generally a bad idea with a limited stack size.

14595 @item -mstack-guard=@var{stack-guard}  
 14596 @itemx -mstack-size=@var{stack-size}  
 14597 @opindex mstack-guard  
 14598 @opindex mstack-size  
 14599 If these options are provided the s390 back end emits additional instructions in  
 14600 the function prologue which trigger a trap if the stack size is @var{stack-guard}  
 14601 bytes above the @var{stack-size} (remember that the stack on s390 grows downward  
 14602 If the @var{stack-guard} option is omitted the smallest power of 2 larger than  
 14603 the frame size of the compiled function is chosen.  
 14604 These options are intended to be used to help debugging stack overflow problems.  
 14605 The additionally emitted code causes only little overhead and hence can also be  
 14606 used in production like systems without greater performance degradation. The gi  
 14607 values have to be exact powers of 2 and @var{stack-size} has to be greater than  
 14608 @var{stack-guard} without exceeding 64k.  
 14609 In order to be efficient the extra code makes the assumption that the stack star  
 14610 at an address aligned to the value given by @var{stack-size}.

14611 The @var{stack-guard} option can only be used in conjunction with @var{stack-siz  
 14612 @end table

14614 @node Score Options  
 14615 @subsection Score Options  
 14616 @cindex Score Options

14618 These options are defined for Score implementations:

14620 @table @gcctabopt  
 14621 @item -meb  
 14622 @opindex meb  
 14623 Compile code for big endian mode. This is the default.

14625 @item -mel  
 14626 @opindex mel  
 14627 Compile code for little endian mode.

14629 @item -mnhwloop  
 14630 @opindex mnhwloop  
 14631 Disable generate bcnz instruction.

14633 @item -muls  
 14634 @opindex muls  
 14635 Enable generate unaligned load and store instruction.

14637 @item -mmac  
 14638 @opindex mmac  
 14639 Enable the use of multiply-accumulate instructions. Disabled by default.

14641 @item -mscore5  
 14642 @opindex mscore5  
 14643 Specify the SCORE5 as the target architecture.

14645 @item -mscore5u  
 14646 @opindex mscore5u  
 14647 Specify the SCORE5U of the target architecture.

14649 @item -mscore7  
 14650 @opindex mscore7  
 14651 Specify the SCORE7 as the target architecture. This is the default.

14653 @item -mscore7d  
 14654 @opindex mscore7d  
 14655 Specify the SCORE7D as the target architecture.  
 14656 @end table

14658 @node SH Options  
 14659 @subsection SH Options

14661 These @samp{-m} options are defined for the SH implementations:

14663 @table @gcctabopt  
 14664 @item -m1  
 14665 @opindex m1  
 14666 Generate code for the SH1.

14668 @item -m2  
 14669 @opindex m2  
 14670 Generate code for the SH2.

14672 @item -m2e  
 14673 Generate code for the SH2e.

14675 @item -m3  
 14676 @opindex m3



14677 Generate code for the SH3.

14679 @item -m3e  
14680 @opindex m3e  
14681 Generate code for the SH3e.

14683 @item -m4-nofpu  
14684 @opindex m4-nofpu  
14685 Generate code for the SH4 without a floating-point unit.

14687 @item -m4-single-only  
14688 @opindex m4-single-only  
14689 Generate code for the SH4 with a floating-point unit that only supports single-precision arithmetic.

14692 @item -m4-single  
14693 @opindex m4-single  
14694 Generate code for the SH4 assuming the floating-point unit is in single-precision mode by default.

14697 @item -m4  
14698 @opindex m4  
14699 Generate code for the SH4.

14701 @item -m4a-nofpu  
14702 @opindex m4a-nofpu  
14703 Generate code for the SH4al-dsp, or for a SH4a in such a way that the floating-point unit is not used.

14706 @item -m4a-single-only  
14707 @opindex m4a-single-only  
14708 Generate code for the SH4a, in such a way that no double-precision floating point operations are used.

14711 @item -m4a-single  
14712 @opindex m4a-single  
14713 Generate code for the SH4a assuming the floating-point unit is in single-precision mode by default.

14716 @item -m4a  
14717 @opindex m4a  
14718 Generate code for the SH4a.

14720 @item -m4al  
14721 @opindex m4al  
14722 Same as @option{-m4a-nofpu}, except that it implicitly passes @option{-dsp} to the assembler. GCC doesn't generate any DSP instructions at the moment.

14726 @item -mb  
14727 @opindex mb  
14728 Compile code for the processor in big endian mode.

14730 @item -ml  
14731 @opindex ml  
14732 Compile code for the processor in little endian mode.

14734 @item -mdalign  
14735 @opindex mdalign  
14736 Align doubles at 64-bit boundaries. Note that this changes the calling conventions, and thus some functions from the standard C library will not work unless you recompile it first with @option{-mdalign}.

14740 @item -mrelax  
14741 @opindex mrelax  
14742 Shorten some address references at link time, when possible; uses the

14743 linker option @option{-relax}.

14745 @item -mbigtable  
14746 @opindex mbigtable  
14747 Use 32-bit offsets in @code{switch} tables. The default is to use 14748 16-bit offsets.

14750 @item -mbitops  
14751 @opindex mbitops  
14752 Enable the use of bit manipulation instructions on SH2A.

14754 @item -mfmovd  
14755 @opindex fmovd  
14756 Enable the use of the instruction @code{fmovd}.

14758 @item -mhitachi  
14759 @opindex mhitachi  
14760 Comply with the calling conventions defined by Renesas.

14762 @item -mrenesas  
14763 @opindex mhitachi  
14764 Comply with the calling conventions defined by Renesas.

14766 @item -mno-renesas  
14767 @opindex mhitachi  
14768 Comply with the calling conventions defined for GCC before the Renesas 14769 conventions were available. This option is the default for all 14770 targets of the SH toolchain except for @samp{sh-symbianelf}.

14772 @item -mnomacsave  
14773 @opindex mnomacsave  
14774 Mark the @code{MAC} register as call-clobbered, even if 14775 @option{-mhitachi} is given.

14777 @item -mieee  
14778 @opindex mieee  
14779 Increase IEEE-compliance of floating-point code.  
14780 At the moment, this is equivalent to @option{-fno-finite-math-only}.  
14781 When generating 16 bit SH opcodes, getting IEEE-conforming results for 14782 comparisons of NaNs / infinities incurs extra overhead in every 14783 floating point comparison, therefore the default is set to 14784 @option{-ffinite-math-only}.

14786 @item -minline-ic\_invalidate  
14787 @opindex minline-ic\_invalidate  
14788 Inline code to invalidate instruction cache entries after setting up 14789 nested function trampolines.  
14790 This option has no effect if -musermode is in effect and the selected 14791 code generation option (e.g. -m4) does not allow the use of the icbi 14792 instruction.  
14793 If the selected code generation option does not allow the use of the icbi 14794 instruction, and -musermode is not in effect, the inlined code will 14795 manipulate the instruction cache address array directly with an associative 14796 write. This not only requires privileged mode, but it will also 14797 fail if the cache line had been mapped via the TLB and has become unmapped.

14799 @item -misize  
14800 @opindex misize  
14801 Dump instruction size and location in the assembly code.

14803 @item -mpadstruct  
14804 @opindex mpadstruct  
14805 This option is deprecated. It pads structures to multiple of 4 bytes, 14806 which is incompatible with the SH ABI@.

14808 @item -mspace

```

14809 @opindex mspace
14810 Optimize for space instead of speed.  Implied by @option{-Os}.

14812 @item -mprefergot
14813 @opindex mprefergot
14814 When generating position-independent code, emit function calls using
14815 the Global Offset Table instead of the Procedure Linkage Table.

14817 @item -musermode
14818 @opindex musermode
14819 Don't generate privileged mode only code; implies -mno-inline-ic_invalidate
14820 if the inlined code would not work in user mode.
14821 This is the default when the target is @code{sh-*-linux*}.

14823 @item -multcost=@var{number}
14824 @opindex multcost=@var{number}
14825 Set the cost to assume for a multiply insn.

14827 @item -mdiv=@var{strategy}
14828 @opindex mdiv=@var{strategy}
14829 Set the division strategy to use for SHmedia code.  @var{strategy} must be
14830 one of: call, call2, fp, inv, inv:minlat, inv20u, inv20l, inv:call,
14831 inv:call2, inv:fp .
14832 "fp" performs the operation in floating point.  This has a very high latency,
14833 but needs only a few instructions, so it might be a good choice if
14834 your code has enough easily exploitable ILP to allow the compiler to
14835 schedule the floating point instructions together with other instructions.
14836 Division by zero causes a floating point exception.
14837 "inv" uses integer operations to calculate the inverse of the divisor,
14838 and then multiplies the dividend with the inverse.  This strategy allows
14839 cse and hoisting of the inverse calculation.  Division by zero calculates
14840 an unspecified result, but does not trap.
14841 "inv:minlat" is a variant of "inv" where if no cse / hoisting opportunities
14842 have been found, or if the entire operation has been hoisted to the same
14843 place, the last stages of the inverse calculation are intertwined with the
14844 final multiply to reduce the overall latency, at the expense of using a few
14845 more instructions, and thus offering fewer scheduling opportunities with
14846 other code.
14847 "call" calls a library function that usually implements the inv:minlat
14848 strategy.
14849 This gives high code density for m5-*media-nofpu compilations.
14850 "call2" uses a different entry point of the same library function, where it
14851 assumes that a pointer to a lookup table has already been set up, which
14852 exposes the pointer load to cse / code hoisting optimizations.
14853 "inv:call", "inv:call2" and "inv:fp" all use the "inv" algorithm for initial
14854 code generation, but if the code stays unoptimized, revert to the "call",
14855 "call2", or "fp" strategies, respectively.  Note that the
14856 potentially-trapping side effect of division by zero is carried by a
14857 separate instruction, so it is possible that all the integer instructions
14858 are hoisted out, but the marker for the side effect stays where it is.
14859 A recombination to fp operations or a call is not possible in that case.
14860 "inv20u" and "inv20l" are variants of the "inv:minlat" strategy.  In the case
14861 that the inverse calculation was not separated from the multiply, they speed
14862 up division where the dividend fits into 20 bits (plus sign where applicable),
14863 by inserting a test to skip a number of operations in this case; this test
14864 slows down the case of larger dividends.  inv20u assumes the case of a such
14865 a small dividend to be unlikely, and inv20l assumes it to be likely.

14867 @item -mdivsi3_libfunc=@var{name}
14868 @opindex mdivsi3_libfunc=@var{name}
14869 Set the name of the library function used for 32 bit signed division to
14870 @var{name}.  This only affect the name used in the call and inv:call
14871 division strategies, and the compiler will still expect the same
14872 sets of input/output/clobbered registers as if this option was not present.

14874 @item -mfixed-range=@var{register-range}

```

```

14875 @opindex mfixed-range
14876 Generate code treating the given register range as fixed registers.
14877 A fixed register is one that the register allocator can not use.  This is
14878 useful when compiling kernel code.  A register range is specified as
14879 two registers separated by a dash.  Multiple register ranges can be
14880 specified separated by a comma.

14882 @item -madjust-unroll
14883 @opindex madjust-unroll
14884 Throttle unrolling to avoid thrashing target registers.
14885 This option only has an effect if the gcc code base supports the
14886 TARGET_ADJUST_UNROLL_MAX target hook.

14888 @item -mindexed-addressing
14889 @opindex mindexed-addressing
14890 Enable the use of the indexed addressing mode for SHmedia32/SHcompact.
14891 This is only safe if the hardware and/or OS implement 32 bit wrap-around
14892 semantics for the indexed addressing mode.  The architecture allows the
14893 implementation of processors with 64 bit MMU, which the OS could use to
14894 get 32 bit addressing, but since no current hardware implementation supports
14895 this or any other way to make the indexed addressing mode safe to use in
14896 the 32 bit ABI, the default is -mno-indexed-addressing.

14898 @item -mgettrcost=@var{number}
14899 @opindex mgettrcost=@var{number}
14900 Set the cost assumed for the gettr instruction to @var{number}.
14901 The default is 2 if @option{-mpt-fixed} is in effect, 100 otherwise.

14903 @item -mpt-fixed
14904 @opindex mpt-fixed
14905 Assume pt* instructions won't trap.  This will generally generate better
14906 scheduled code, but is unsafe on current hardware.  The current architecture
14907 definition says that ptabs and ptreload trap when the target anded with 3 is 3.
14908 This has the unintentional effect of making it unsafe to schedule ptabs /
14909 ptreload before a branch, or hoist it out of a loop.  For example,
14910 __do_global_ctors, a part of libgcc that runs constructors at program
14911 startup, calls functions in a list which is delimited by @minus{1}.  With the
14912 -mpt-fixed option, the ptabs will be done before testing against @minus{1}.
14913 That means that all the constructors will be run a bit quicker, but when
14914 the loop comes to the end of the list, the program crashes because ptabs
14915 loads @minus{1} into a target register.  Since this option is unsafe for any
14916 hardware implementing the current architecture specification, the default
14917 is -mno-pt-fixed.  Unless the user specifies a specific cost with
14918 @option{-mgettrcost}, -mno-pt-fixed also implies @option{-mgettrcost=100};
14919 this deters register allocation using target registers for storing
14920 ordinary integers.

14922 @item -minvalid-symbols
14923 @opindex minvalid-symbols
14924 Assume symbols might be invalid.  Ordinary function symbols generated by
14925 the compiler will always be valid to load with movi/shori/ptabs or
14926 movi/shori/ptrel, but with assembler and/or linker tricks it is possible
14927 to generate symbols that will cause ptabs / ptreload to trap.
14928 This option is only meaningful when @option{-mno-pt-fixed} is in effect.
14929 It will then prevent cross-basic-block cse, hoisting and most scheduling
14930 of symbol loads.  The default is @option{-mno-invalid-symbols}.
14931 @end table

14933 @node SPARC Options
14934 @subsection SPARC Options
14935 @cindex SPARC options

14937 These @samp{-m} options are supported on the SPARC:

14939 @table @gcctabopt
14940 @item -mno-app-regs

```

14941 @itemx -mapp-regs  
 14942 @opindex mno-app-regs  
 14943 @opindex mapp-regs  
 14944 Specify @option{-mapp-regs} to generate output using the global registers  
 14945 2 through 4, which the SPARC SVR4 ABI reserves for applications. This  
 14946 is the default.

14948 To be fully SVR4 ABI compliant at the cost of some performance loss,  
 14949 specify @option{-mno-app-regs}. You should compile libraries and system  
 14950 software with this option.

14952 @item -mfp  
 14953 @itemx -mhard-float  
 14954 @opindex mfp  
 14955 @opindex mhard-float  
 14956 Generate output containing floating point instructions. This is the  
 14957 default.

14959 @item -mno-fpu  
 14960 @itemx -msoft-float  
 14961 @opindex mno-fpu  
 14962 @opindex msoft-float  
 14963 Generate output containing library calls for floating point.  
 14964 @strong{Warning;} the requisite libraries are not available for all SPARC  
 14965 targets. Normally the facilities of the machine's usual C compiler are  
 14966 used, but this cannot be done directly in cross-compilation. You must make  
 14967 your own arrangements to provide suitable library functions for  
 14968 cross-compilation. The embedded targets @samp{sparc-\*-aout} and  
 14969 @samp{sparclite-\*-aout} do provide software floating point support.

14971 @option{-msoft-float} changes the calling convention in the output file;  
 14972 therefore, it is only useful if you compile @emph{all} of a program with  
 14973 this option. In particular, you need to compile @file{libgcc.a}, the  
 14974 library that comes with GCC, with @option{-msoft-float} in order for  
 14975 this to work.

14977 @item -mhard-quad-float  
 14978 @opindex mhard-quad-float  
 14979 Generate output containing quad-word (long double) floating point  
 14980 instructions.

14982 @item -msoft-quad-float  
 14983 @opindex msoft-quad-float  
 14984 Generate output containing library calls for quad-word (long double)  
 14985 floating point instructions. The functions called are those specified  
 14986 in the SPARC ABI@. This is the default.

14988 As of this writing, there are no SPARC implementations that have hardware  
 14989 support for the quad-word floating point instructions. They all invoke  
 14990 a trap handler for one of these instructions, and then the trap handler  
 14991 emulates the effect of the instruction. Because of the trap handler overhead,  
 14992 this is much slower than calling the ABI library routines. Thus the  
 14993 @option{-msoft-quad-float} option is the default.

14995 @item -mno-unaligned-doubles  
 14996 @itemx -munaligned-doubles  
 14997 @opindex mno-unaligned-doubles  
 14998 @opindex munaligned-doubles  
 14999 Assume that doubles have 8 byte alignment. This is the default.

15001 With @option{-munaligned-doubles}, GCC assumes that doubles have 8 byte  
 15002 alignment only if they are contained in another type, or if they have an  
 15003 absolute address. Otherwise, it assumes they have 4 byte alignment.  
 15004 Specifying this option avoids some rare compatibility problems with code  
 15005 generated by other compilers. It is not the default because it results  
 15006 in a performance loss, especially for floating point code.

15008 @item -mno-faster-structs  
 15009 @itemx -mfaster-structs  
 15010 @opindex mno-faster-structs  
 15011 @opindex mfaster-structs  
 15012 With @option{-mfaster-structs}, the compiler assumes that structures  
 15013 should have 8 byte alignment. This enables the use of pairs of  
 15014 @code{ldd} and @code{std} instructions for copies in structure  
 15015 assignment, in place of twice as many @code{ld} and @code{st} pairs.  
 15016 However, the use of this changed alignment directly violates the SPARC  
 15017 ABI@. Thus, it's intended only for use on targets where the developer  
 15018 acknowledges that their resulting code will not be directly in line with  
 15019 the rules of the ABI@.

15021 @item -mimpure-text  
 15022 @opindex mimpure-text  
 15023 @option{-mimpure-text}, used in addition to @option{-shared}, tells  
 15024 the compiler to not pass @option{-z text} to the linker when linking a  
 15025 shared object. Using this option, you can link position-dependent  
 15026 code into a shared object.

15028 @option{-mimpure-text} suppresses the ``relocations remain against  
 15029 allocatable but non-writable sections'' linker error message.  
 15030 However, the necessary relocations will trigger copy-on-write, and the  
 15031 shared object is not actually shared across processes. Instead of  
 15032 using @option{-mimpure-text}, you should compile all source code with  
 15033 @option{-fpic} or @option{-fPIC}.

15035 This option is only available on SunOS and Solaris.

15037 @item -mcpu=@var{cpu\_type}  
 15038 @opindex mcpu  
 15039 Set the instruction set, register set, and instruction scheduling parameters  
 15040 for machine type @var{cpu\_type}. Supported values for @var{cpu\_type} are  
 15041 @samp{v7}, @samp{cypress}, @samp{v8}, @samp{supersparc}, @samp{sparclite},  
 15042 @samp{f930}, @samp{f934}, @samp{hypersparc}, @samp{sparclite86x},  
 15043 @samp{sparclet}, @samp{tsc701}, @samp{v9}, @samp{ultrasparc},  
 15044 @samp{ultrasparc3}, @samp{niagara} and @samp{niagara2}.

15046 Default instruction scheduling parameters are used for values that select  
 15047 an architecture and not an implementation. These are @samp{v7}, @samp{v8},  
 15048 @samp{sparclite}, @samp{sparclet}, @samp{v9}.

15050 Here is a list of each supported architecture and their supported  
 15051 implementations.

15053 @smallexample  
 15054 v7: cypress  
 15055 v8: supersparc, hypersparc  
 15056 sparclite: f930, f934, sparclite86x  
 15057 sparclet: tsc701  
 15058 v9: ultrasparc, ultrasparc3, niagara, niagara2  
 15059 @end smallexample

15061 By default (unless configured otherwise), GCC generates code for the V7  
 15062 variant of the SPARC architecture. With @option{-mcpu=cypress}, the compiler  
 15063 additionally optimizes it for the Cypress CY7C602 chip, as used in the  
 15064 SPARCStation/SPARCServer 3xx series. This is also appropriate for the older  
 15065 SPARCStation 1, 2, IPX etc.

15067 With @option{-mcpu=v8}, GCC generates code for the V8 variant of the SPARC  
 15068 architecture. The only difference from V7 code is that the compiler emits  
 15069 the integer multiply and integer divide instructions which exist in SPARC-V8  
 15070 but not in SPARC-V7. With @option{-mcpu=supersparc}, the compiler additionally  
 15071 optimizes it for the SuperSPARC chip, as used in the SPARCStation 10, 1000 and  
 15072 2000 series.

15074 With `@option{-mcpu=sparclite}`, GCC generates code for the SPARClite variant of  
 15075 the SPARC architecture. This adds the integer multiply, integer divide step  
 15076 and scan (`@code{ffs}`) instructions which exist in SPARClite but not in SPARC-V7.  
 15077 With `@option{-mcpu=f930}`, the compiler additionally optimizes it for the  
 15078 Fujitsu MB86930 chip, which is the original SPARClite, with no FPU@. With  
 15079 `@option{-mcpu=f934}`, the compiler additionally optimizes it for the Fujitsu  
 15080 MB86934 chip, which is the more recent SPARClite with FPU@.

15082 With `@option{-mcpu=sparclet}`, GCC generates code for the SPARClet variant of  
 15083 the SPARC architecture. This adds the integer multiply, multiply/accumulate,  
 15084 integer divide step and scan (`@code{ffs}`) instructions which exist in SPARClet  
 15085 but not in SPARC-V7. With `@option{-mcpu=tsc701}`, the compiler additionally  
 15086 optimizes it for the TEMIC SPARClet chip.

15088 With `@option{-mcpu=v9}`, GCC generates code for the V9 variant of the SPARC  
 15089 architecture. This adds 64-bit integer and floating-point move instructions,  
 15090 3 additional floating-point condition code registers and conditional move  
 15091 instructions. With `@option{-mcpu=ultrasparc}`, the compiler additionally  
 15092 optimizes it for the Sun UltraSPARC I/II/III chips. With  
 15093 `@option{-mcpu=ultrasparc3}`, the compiler additionally optimizes it for the  
 15094 Sun UltraSPARC III/III+/IIIi/IIIi+/IV/IV+ chips. With  
 15095 `@option{-mcpu=niagara}`, the compiler additionally optimizes it for  
 15096 Sun UltraSPARC T1 chips. With `@option{-mcpu=niagara2}`, the compiler  
 15097 additionally optimizes it for Sun UltraSPARC T2 chips.

15099 `@item -mtune=@var{cpu_type}`  
 15100 `@opindex mtune`  
 15101 Set the instruction scheduling parameters for machine type  
 15102 `@var{cpu_type}`, but do not set the instruction set or register set that the  
 15103 option `@option{-mcpu=@var{cpu_type}}` would.

15105 The same values for `@option{-mcpu=@var{cpu_type}}` can be used for  
 15106 `@option{-mtune=@var{cpu_type}}`, but the only useful values are those  
 15107 that select a particular cpu implementation. Those are `@samp{cypress}`,  
 15108 `@samp{supersparc}`, `@samp{hypersparc}`, `@samp{f930}`, `@samp{f934}`,  
 15109 `@samp{sparclite86x}`, `@samp{tsc701}`, `@samp{ultrasparc}`,  
 15110 `@samp{ultrasparc3}`, `@samp{niagara}`, and `@samp{niagara2}`.

15112 `@item -mv8plus`  
 15113 `@itemx -mno-v8plus`  
 15114 `@opindex mv8plus`  
 15115 `@opindex mno-v8plus`  
 15116 With `@option{-mv8plus}`, GCC generates code for the SPARC-V8+ ABI@. The  
 15117 difference from the V8 ABI is that the global and out registers are  
 15118 considered 64-bit wide. This is enabled by default on Solaris in 32-bit  
 15119 mode for all SPARC-V9 processors.

15121 `@item -mvis`  
 15122 `@itemx -mno-vis`  
 15123 `@opindex mvis`  
 15124 `@opindex mno-vis`  
 15125 With `@option{-mvis}`, GCC generates code that takes advantage of the UltraSPARC  
 15126 Visual Instruction Set extensions. The default is `@option{-mno-vis}`.

15128 `@item -mno-integer-ldd-std`  
 15129 `@opindex mno-integer-ldd-std`  
 15130 With `@option{-mno-integer-ldd-std}`, GCC does not use the `@code{ldd}`  
 15131 and `@code{std}` instructions for integer operands in 32-bit mode. This  
 15132 is for use with legacy code using 64-bit quantities which are not  
 15133 64-bit aligned.

15135 `@item -massume-32bit-callers`  
 15136 `@opindex massume-32bit-callers`  
 15137 With `@option{-massume-32bit-callers}`, The type promotion of function  
 15138 arguments is altered such that integer arguments smaller than the word

15139 size are extended in the callee rather than the caller. This is  
 15140 necessary for system calls from 32bit processes to 64bit kernels in  
 15141 certain systems. This option should not be used in any situation  
 15142 other than compiling the kernels of such systems, and has not been  
 15143 tested outside of that scenario.  
 15144 `@end table`

15146 These `@samp{-m}` options are supported in addition to the above  
 15147 on SPARC-V9 processors in 64-bit environments:

15149 `@table @gcctabopt`  
 15150 `@item -mlittle-endian`  
 15151 `@opindex mlittle-endian`  
 15152 Generate code for a processor running in little-endian mode. It is only  
 15153 available for a few configurations and most notably not on Solaris and Linux.

15155 `@item -m32`  
 15156 `@itemx -m64`  
 15157 `@opindex m32`  
 15158 `@opindex m64`  
 15159 Generate code for a 32-bit or 64-bit environment.  
 15160 The 32-bit environment sets `int`, `long` and `pointer` to 32 bits.  
 15161 The 64-bit environment sets `int` to 32 bits and `long` and `pointer`  
 15162 to 64 bits.

15164 `@item -mcmmodel=medlow`  
 15165 `@opindex mcmmodel=medlow`  
 15166 Generate code for the Medium/Low code model: 64-bit addresses, programs  
 15167 must be linked in the low 32 bits of memory. Programs can be statically  
 15168 or dynamically linked.

15170 `@item -mcmmodel=medmid`  
 15171 `@opindex mcmmodel=medmid`  
 15172 Generate code for the Medium/Middle code model: 64-bit addresses, programs  
 15173 must be linked in the low 44 bits of memory, the text and data segments must  
 15174 be less than 2GB in size and the data segment must be located within 2GB of  
 15175 the text segment.

15177 `@item -mcmmodel=medany`  
 15178 `@opindex mcmmodel=medany`  
 15179 Generate code for the Medium/Anywhere code model: 64-bit addresses, programs  
 15180 may be linked anywhere in memory, the text and data segments must be less  
 15181 than 2GB in size and the data segment must be located within 2GB of the  
 15182 text segment.

15184 `@item -mcmmodel=embmedany`  
 15185 `@opindex mcmmodel=embmedany`  
 15186 Generate code for the Medium/Anywhere code model for embedded systems:  
 15187 64-bit addresses, the text and data segments must be less than 2GB in  
 15188 size, both starting anywhere in memory (determined at link time). The  
 15189 global register `%g4` points to the base of the data segment. Programs  
 15190 are statically linked and PIC is not supported.

15192 `@item -mstack-bias`  
 15193 `@itemx -mno-stack-bias`  
 15194 `@opindex mstack-bias`  
 15195 `@opindex mno-stack-bias`  
 15196 With `@option{-mstack-bias}`, GCC assumes that the stack pointer, and  
 15197 frame pointer if present, are offset by `@minus{}2047` which must be added back  
 15198 when making stack frame references. This is the default in 64-bit mode.  
 15199 Otherwise, assume no such offset is present.  
 15200 `@end table`

15202 These switches are supported in addition to the above on Solaris:

15204 `@table @gcctabopt`

15205 @item -threads  
 15206 @opindex threads  
 15207 Add support for multithreading using the Solaris threads library. This  
 15208 option sets flags for both the preprocessor and linker. This option does  
 15209 not affect the thread safety of object code produced by the compiler or  
 15210 that of libraries supplied with it.

15212 @item -pthreads  
 15213 @opindex pthreads  
 15214 Add support for multithreading using the POSIX threads library. This  
 15215 option sets flags for both the preprocessor and linker. This option does  
 15216 not affect the thread safety of object code produced by the compiler or  
 15217 that of libraries supplied with it.

15219 @item -pthread  
 15220 @opindex pthread  
 15221 This is a synonym for @option{-pthreads}.  
 15222 @end table

15224 @node SPU Options  
 15225 @subsection SPU Options  
 15226 @cindex SPU options

15228 These @samp{-m} options are supported on the SPU:

15230 @table @gcctabopt  
 15231 @item -mwarn-reloc  
 15232 @itemx -merror-reloc  
 15233 @opindex mwarn-reloc  
 15234 @opindex merror-reloc

15236 The loader for SPU does not handle dynamic relocations. By default, GCC  
 15237 will give an error when it generates code that requires a dynamic  
 15238 relocation. @option{-mno-error-reloc} disables the error,  
 15239 @option{-mwarn-reloc} will generate a warning instead.

15241 @item -msafe-dma  
 15242 @itemx -munsafe-dma  
 15243 @opindex msafe-dma  
 15244 @opindex unsafe-dma

15246 Instructions which initiate or test completion of DMA must not be  
 15247 reordered with respect to loads and stores of the memory which is being  
 15248 accessed. Users typically address this problem using the volatile  
 15249 keyword, but that can lead to inefficient code in places where the  
 15250 memory is known to not change. Rather than mark the memory as volatile  
 15251 we treat the DMA instructions as potentially effecting all memory. With  
 15252 @option{-munsafe-dma} users must use the volatile keyword to protect  
 15253 memory accesses.

15255 @item -mbranch-hints  
 15256 @opindex mbranch-hints

15258 By default, GCC will generate a branch hint instruction to avoid  
 15259 pipeline stalls for always taken or probably taken branches. A hint  
 15260 will not be generated closer than 8 instructions away from its branch.  
 15261 There is little reason to disable them, except for debugging purposes,  
 15262 or to make an object a little bit smaller.

15264 @item -msmall-mem  
 15265 @itemx -mlarge-mem  
 15266 @opindex msmall-mem  
 15267 @opindex mlarge-mem

15269 By default, GCC generates code assuming that addresses are never larger  
 15270 than 18 bits. With @option{-mlarge-mem} code is generated that assumes

15271 a full 32 bit address.

15273 @item -mstdmain  
 15274 @opindex mstdmain

15276 By default, GCC links against startup code that assumes the SPU-style  
 15277 main function interface (which has an unconventional parameter list).  
 15278 With @option{-mstdmain}, GCC will link your program against startup  
 15279 code that assumes a C99-style interface to @code{main}, including a  
 15280 local copy of @code{argv} strings.

15282 @item -mfixed-range=@var{register-range}  
 15283 @opindex mfixed-range  
 15284 Generate code treating the given register range as fixed registers.  
 15285 A fixed register is one that the register allocator can not use. This is  
 15286 useful when compiling kernel code. A register range is specified as  
 15287 two registers separated by a dash. Multiple register ranges can be  
 15288 specified separated by a comma.

15290 @item -mdual-nops  
 15291 @itemx -mdual-nops=@var{n}  
 15292 @opindex mdual-nops  
 15293 By default, GCC will insert nops to increase dual issue when it expects  
 15294 it to increase performance. @var{n} can be a value from 0 to 10. A  
 15295 smaller @var{n} will insert fewer nops. 10 is the default, 0 is the  
 15296 same as @option{-mno-dual-nops}. Disabled with @option{-Os}.

15298 @item -mhint-max-nops=@var{n}  
 15299 @opindex mhint-max-nops  
 15300 Maximum number of nops to insert for a branch hint. A branch hint must  
 15301 be at least 8 instructions away from the branch it is effecting. GCC  
 15302 will insert up to @var{n} nops to enforce this, otherwise it will not  
 15303 generate the branch hint.

15305 @item -mhint-max-distance=@var{n}  
 15306 @opindex mhint-max-distance  
 15307 The encoding of the branch hint instruction limits the hint to be within  
 15308 256 instructions of the branch it is effecting. By default, GCC makes  
 15309 sure it is within 125.

15311 @item -msafe-hints  
 15312 @opindex msafe-hints  
 15313 Work around a hardware bug which causes the SPU to stall indefinitely.  
 15314 By default, GCC will insert the @code{hrp} instruction to make sure  
 15315 this stall won't happen.

15317 @end table

15319 @node System V Options  
 15320 @subsection Options for System V

15322 These additional options are available on System V Release 4 for  
 15323 compatibility with other compilers on those systems:

15325 @table @gcctabopt  
 15326 @item -G  
 15327 @opindex G  
 15328 Create a shared object.  
 15329 It is recommended that @option{-symbolic} or @option{-shared} be used instead.

15331 @item -Qy  
 15332 @opindex Qy  
 15333 Identify the versions of each tool used by the compiler, in a  
 15334 @code{.ident} assembler directive in the output.

15336 @item -Qn

15337 @opindex Qn  
 15338 Refrain from adding @code{.ident} directives to the output file (this is  
 15339 the default).

15341 @item -YP,@var{dirs}  
 15342 @opindex YP  
 15343 Search the directories @var{dirs}, and no others, for libraries  
 15344 specified with @option{-l}.

15346 @item -Ym,@var{dir}  
 15347 @opindex Ym  
 15348 Look in the directory @var{dir} to find the M4 preprocessor.  
 15349 The assembler uses this option.  
 15350 @c This is supposed to go with a -Yd for predefined M4 macro files, but  
 15351 @c the generic assembler that comes with Solaris takes just -Ym.  
 15352 @end table

15354 @node V850 Options  
 15355 @subsection V850 Options  
 15356 @cindex V850 Options

15358 These @samp{-m} options are defined for V850 implementations:

15360 @table @gcctabopt  
 15361 @item -mlong-calls  
 15362 @itemx -mno-long-calls  
 15363 @opindex mlong-calls  
 15364 @opindex mno-long-calls  
 15365 Treat all calls as being far away (near). If calls are assumed to be  
 15366 far away, the compiler will always load the functions address up into a  
 15367 register, and call indirect through the pointer.

15369 @item -mno-ep  
 15370 @itemx -mep  
 15371 @opindex mno-ep  
 15372 @opindex mep  
 15373 Do not optimize (do optimize) basic blocks that use the same index  
 15374 pointer 4 or more times to copy pointer into the @code{ep} register, and  
 15375 use the shorter @code{sld} and @code{sst} instructions. The @option{-mep}  
 15376 option is on by default if you optimize.

15378 @item -mno-prolog-function  
 15379 @itemx -mprolog-function  
 15380 @opindex mno-prolog-function  
 15381 @opindex mprolog-function  
 15382 Do not use (do use) external functions to save and restore registers  
 15383 at the prologue and epilogue of a function. The external functions  
 15384 are slower, but use less code space if more than one function saves  
 15385 the same number of registers. The @option{-mprolog-function} option  
 15386 is on by default if you optimize.

15388 @item -mspace  
 15389 @opindex mspace  
 15390 Try to make the code as small as possible. At present, this just turns  
 15391 on the @option{-mep} and @option{-mprolog-function} options.

15393 @item -mtda=@var{n}  
 15394 @opindex mtda  
 15395 Put static or global variables whose size is @var{n} bytes or less into  
 15396 the tiny data area that register @code{ep} points to. The tiny data  
 15397 area can hold up to 256 bytes in total (128 bytes for byte references).

15399 @item -msda=@var{n}  
 15400 @opindex msda  
 15401 Put static or global variables whose size is @var{n} bytes or less into  
 15402 the small data area that register @code{gp} points to. The small data

15403 area can hold up to 64 kilobytes.

15405 @item -mzda=@var{n}  
 15406 @opindex mzda  
 15407 Put static or global variables whose size is @var{n} bytes or less into  
 15408 the first 32 kilobytes of memory.

15410 @item -mv850  
 15411 @opindex mv850  
 15412 Specify that the target processor is the V850.

15414 @item -mbig-switch  
 15415 @opindex mbig-switch  
 15416 Generate code suitable for big switch tables. Use this option only if  
 15417 the assembler/linker complain about out of range branches within a switch  
 15418 table.

15420 @item -mapp-regs  
 15421 @opindex mapp-regs  
 15422 This option will cause r2 and r5 to be used in the code generated by  
 15423 the compiler. This setting is the default.

15425 @item -mno-app-regs  
 15426 @opindex mno-app-regs  
 15427 This option will cause r2 and r5 to be treated as fixed registers.

15429 @item -mv850e1  
 15430 @opindex mv850e1  
 15431 Specify that the target processor is the V850E1. The preprocessor  
 15432 constants @samp{\_\_v850e1\_\_} and @samp{\_\_v850e\_\_} will be defined if  
 15433 this option is used.

15435 @item -mv850e  
 15436 @opindex mv850e  
 15437 Specify that the target processor is the V850E@. The preprocessor  
 15438 constant @samp{\_\_v850e\_\_} will be defined if this option is used.

15440 If neither @option{-mv850} nor @option{-mv850e} nor @option{-mv850e1}  
 15441 are defined then a default target processor will be chosen and the  
 15442 relevant @samp{\_\_v850\*\_\_} preprocessor constant will be defined.

15444 The preprocessor constants @samp{\_\_v850} and @samp{\_\_v851\_\_} are always  
 15445 defined, regardless of which processor variant is the target.

15447 @item -mdisable-callt  
 15448 @opindex mdisable-callt  
 15449 This option will suppress generation of the CALLT instruction for the  
 15450 v850e and v850e1 flavors of the v850 architecture. The default is  
 15451 @option{-mno-disable-callt} which allows the CALLT instruction to be used.

15453 @end table

15455 @node VAX Options  
 15456 @subsection VAX Options  
 15457 @cindex VAX options

15459 These @samp{-m} options are defined for the VAX:

15461 @table @gcctabopt  
 15462 @item -munix  
 15463 @opindex munix  
 15464 Do not output certain jump instructions (@code{aobleq} and so on)  
 15465 that the Unix assembler for the VAX cannot handle across long  
 15466 ranges.

15468 @item -mgnu

```

15469 @opindex mgnu
15470 Do output those jump instructions, on the assumption that you
15471 will assemble with the GNU assembler.

15473 @item -mg
15474 @opindex mg
15475 Output code for g-format floating point numbers instead of d-format.
15476 @end table

15478 @node VxWorks Options
15479 @subsection VxWorks Options
15480 @cindex VxWorks Options

15482 The options in this section are defined for all VxWorks targets.
15483 Options specific to the target hardware are listed with the other
15484 options for that target.

15486 @table @gcctabopt
15487 @item -mrtp
15488 @opindex mrtp
15489 GCC can generate code for both VxWorks kernels and real time processes
15490 (RTPs). This option switches from the former to the latter. It also
15491 defines the preprocessor macro @code{__RTP__}.

15493 @item -non-static
15494 @opindex non-static
15495 Link an RTP executable against shared libraries rather than static
15496 libraries. The options @option{-static} and @option{-shared} can
15497 also be used for RTPs (@pxref{Link Options}); @option{-static}
15498 is the default.

15500 @item -Bstatic
15501 @itemx -Bdynamic
15502 @opindex Bstatic
15503 @opindex Bdynamic
15504 These options are passed down to the linker. They are defined for
15505 compatibility with Diab.

15507 @item -Xbind-lazy
15508 @opindex Xbind-lazy
15509 Enable lazy binding of function calls. This option is equivalent to
15510 @option{-Wl,-z,now} and is defined for compatibility with Diab.

15512 @item -Xbind-now
15513 @opindex Xbind-now
15514 Disable lazy binding of function calls. This option is the default and
15515 is defined for compatibility with Diab.
15516 @end table

15518 @node x86-64 Options
15519 @subsection x86-64 Options
15520 @cindex x86-64 options

15522 These are listed under @xref{i386 and x86-64 Options}.

15524 @node i386 and x86-64 Windows Options
15525 @subsection i386 and x86-64 Windows Options
15526 @cindex i386 and x86-64 Windows Options

15528 These additional options are available for Windows targets:

15530 @table @gcctabopt
15531 @item -mconsole
15532 @opindex mconsole
15533 This option is available for Cygwin and MinGW targets. It
15534 specifies that a console application is to be generated, by

```

```

15535 instructing the linker to set the PE header subsystem type
15536 required for console applications.
15537 This is the default behaviour for Cygwin and MinGW targets.

15539 @item -mcygwin
15540 @opindex mcygwin
15541 This option is available for Cygwin targets. It specifies that
15542 the Cygwin internal interface is to be used for predefined
15543 preprocessor macros, C runtime libraries and related linker
15544 paths and options. For Cygwin targets this is the default behaviour.
15545 This option is deprecated and will be removed in a future release.

15547 @item -mno-cygwin
15548 @opindex mno-cygwin
15549 This option is available for Cygwin targets. It specifies that
15550 the MinGW internal interface is to be used instead of Cygwin's, by
15551 setting MinGW-related predefined macros and linker paths and default
15552 library options.
15553 This option is deprecated and will be removed in a future release.

15555 @item -mdll
15556 @opindex mdll
15557 This option is available for Cygwin and MinGW targets. It
15558 specifies that a DLL - a dynamic link library - is to be
15559 generated, enabling the selection of the required runtime
15560 startup object and entry point.

15562 @item -mnop-fun-dllimport
15563 @opindex mnop-fun-dllimport
15564 This option is available for Cygwin and MinGW targets. It
15565 specifies that the dllimport attribute should be ignored.

15567 @item -mthread
15568 @opindex mthread
15569 This option is available for MinGW targets. It specifies
15570 that MinGW-specific thread support is to be used.

15572 @item -mwin32
15573 @opindex mwin32
15574 This option is available for Cygwin and MinGW targets. It
15575 specifies that the typical Windows pre-defined macros are to
15576 be set in the pre-processor, but does not influence the choice
15577 of runtime library/startup code.

15579 @item -mwindows
15580 @opindex mwindows
15581 This option is available for Cygwin and MinGW targets. It
15582 specifies that a GUI application is to be generated by
15583 instructing the linker to set the PE header subsystem type
15584 appropriately.
15585 @end table

15587 See also under @ref{i386 and x86-64 Options} for standard options.

15589 @node Xstormy16 Options
15590 @subsection Xstormy16 Options
15591 @cindex Xstormy16 Options

15593 These options are defined for Xstormy16:

15595 @table @gcctabopt
15596 @item -msim
15597 @opindex msim
15598 Choose startup files and linker script suitable for the simulator.
15599 @end table

```

15601 @node Xtensa Options  
 15602 @subsection Xtensa Options  
 15603 @cindex Xtensa Options

15605 These options are supported for Xtensa targets:

15607 @table @gcctabopt  
 15608 @item -mconst16  
 15609 @itemx -mno-const16  
 15610 @opindex mconst16  
 15611 @opindex mno-const16  
 15612 Enable or disable use of @code{CONST16} instructions for loading  
 15613 constant values. The @code{CONST16} instruction is currently not a  
 15614 standard option from Tensilica. When enabled, @code{CONST16}  
 15615 instructions are always used in place of the standard @code{L32R}  
 15616 instructions. The use of @code{CONST16} is enabled by default only if  
 15617 the @code{L32R} instruction is not available.

15619 @item -mfused-madd  
 15620 @itemx -mno-fused-madd  
 15621 @opindex mfused-madd  
 15622 @opindex mno-fused-madd  
 15623 Enable or disable use of fused multiply/add and multiply/subtract  
 15624 instructions in the floating-point option. This has no effect if the  
 15625 floating-point option is not also enabled. Disabling fused multiply/add  
 15626 and multiply/subtract instructions forces the compiler to use separate  
 15627 instructions for the multiply and add/subtract operations. This may be  
 15628 desirable in some cases where strict IEEE 754-compliant results are  
 15629 required: the fused multiply add/subtract instructions do not round the  
 15630 intermediate result, thereby producing results with @emph{more} bits of  
 15631 precision than specified by the IEEE standard. Disabling fused multiply  
 15632 add/subtract instructions also ensures that the program output is not  
 15633 sensitive to the compiler's ability to combine multiply and add/subtract  
 15634 operations.

15636 @item -mserialize-volatile  
 15637 @itemx -mno-serialize-volatile  
 15638 @opindex mserialize-volatile  
 15639 @opindex mno-serialize-volatile  
 15640 When this option is enabled, GCC inserts @code{MEMW} instructions before  
 15641 @code{volatile} memory references to guarantee sequential consistency.  
 15642 The default is @option{-mserialize-volatile}. Use  
 15643 @option{-mno-serialize-volatile} to omit the @code{MEMW} instructions.

15645 @item -mtext-section-literals  
 15646 @itemx -mno-text-section-literals  
 15647 @opindex mtext-section-literals  
 15648 @opindex mno-text-section-literals  
 15649 Control the treatment of literal pools. The default is  
 15650 @option{-mno-text-section-literals}, which places literals in a separate  
 15651 section in the output file. This allows the literal pool to be placed  
 15652 in a data RAM/ROM, and it also allows the linker to combine literal  
 15653 pools from separate object files to remove redundant literals and  
 15654 improve code size. With @option{-mtext-section-literals}, the literals  
 15655 are interspersed in the text section in order to keep them as close as  
 15656 possible to their references. This may be necessary for large assembly  
 15657 files.

15659 @item -mtarget-align  
 15660 @itemx -mno-target-align  
 15661 @opindex mtarget-align  
 15662 @opindex mno-target-align  
 15663 When this option is enabled, GCC instructs the assembler to  
 15664 automatically align instructions to reduce branch penalties at the  
 15665 expense of some code density. The assembler attempts to widen density  
 15666 instructions to align branch targets and the instructions following call

15667 instructions. If there are not enough preceding safe density  
 15668 instructions to align a target, no widening will be performed. The  
 15669 default is @option{-mtarget-align}. These options do not affect the  
 15670 treatment of auto-aligned instructions like @code{LOOP}, which the  
 15671 assembler will always align, either by widening density instructions or  
 15672 by inserting no-op instructions.

15674 @item -mlongcalls  
 15675 @itemx -mno-longcalls  
 15676 @opindex mlongcalls  
 15677 @opindex mno-longcalls  
 15678 When this option is enabled, GCC instructs the assembler to translate  
 15679 direct calls to indirect calls unless it can determine that the target  
 15680 of a direct call is in the range allowed by the call instruction. This  
 15681 translation typically occurs for calls to functions in other source  
 15682 files. Specifically, the assembler translates a direct @code{CALL}  
 15683 instruction into an @code{L32R} followed by a @code{CALLX} instruction.  
 15684 The default is @option{-mno-longcalls}. This option should be used in  
 15685 programs where the call target can potentially be out of range. This  
 15686 option is implemented in the assembler, not the compiler, so the  
 15687 assembly code generated by GCC will still show direct call  
 15688 instructions---look at the disassembled object code to see the actual  
 15689 instructions. Note that the assembler will use an indirect call for  
 15690 every cross-file call, not just those that really will be out of range.  
 15691 @end table

15693 @node zSeries Options  
 15694 @subsection zSeries Options  
 15695 @cindex zSeries options

15697 These are listed under @xref{S/390 and zSeries Options}.

15699 @node Code Gen Options  
 15700 @section Options for Code Generation Conventions  
 15701 @cindex code generation conventions  
 15702 @cindex options, code generation  
 15703 @cindex run-time options

15705 These machine-independent options control the interface conventions  
 15706 used in code generation.

15708 Most of them have both positive and negative forms; the negative form  
 15709 of @option{-ffoo} would be @option{-fno-foo}. In the table below, only  
 15710 one of the forms is listed---the one which is not the default. You  
 15711 can figure out the other form by either removing @samp{no-} or adding  
 15712 it.

15714 @table @gcctabopt  
 15715 @item -fbounds-check  
 15716 @opindex fbounds-check  
 15717 For front-ends that support it, generate additional code to check that  
 15718 indices used to access arrays are within the declared range. This is  
 15719 currently only supported by the Java and Fortran front-ends, where  
 15720 this option defaults to true and false respectively.

15722 @item -ftrapv  
 15723 @opindex ftrapv  
 15724 This option generates traps for signed overflow on addition, subtraction,  
 15725 multiplication operations.

15727 @item -fwrapv  
 15728 @opindex fwrapv  
 15729 This option instructs the compiler to assume that signed arithmetic  
 15730 overflow of addition, subtraction and multiplication wraps around  
 15731 using twos-complement representation. This flag enables some optimizations  
 15732 and disables others. This option is enabled by default for the Java



15733 front-end, as required by the Java language specification.

15735 @item -fexceptions  
 15736 @opindex fexceptions  
 15737 Enable exception handling. Generates extra code needed to propagate  
 15738 exceptions. For some targets, this implies GCC will generate frame  
 15739 unwind information for all functions, which can produce significant data  
 15740 size overhead, although it does not affect execution. If you do not  
 15741 specify this option, GCC will enable it by default for languages like  
 15742 C++ which normally require exception handling, and disable it for  
 15743 languages like C that do not normally require it. However, you may need  
 15744 to enable this option when compiling C code that needs to interoperate  
 15745 properly with exception handlers written in C++. You may also wish to  
 15746 disable this option if you are compiling older C++ programs that don't  
 15747 use exception handling.

15749 @item -fnon-call-exceptions  
 15750 @opindex fnon-call-exceptions  
 15751 Generate code that allows trapping instructions to throw exceptions.  
 15752 Note that this requires platform-specific runtime support that does  
 15753 not exist everywhere. Moreover, it only allows @emph{trapping}  
 15754 instructions to throw exceptions, i.e. @code{memory references or floating  
 15755 point instructions}. It does not allow exceptions to be thrown from  
 15756 arbitrary signal handlers such as @code{SIGALRM}.

15758 @item -funwind-tables  
 15759 @opindex funwind-tables  
 15760 Similar to @option{-fexceptions}, except that it will just generate any needed  
 15761 static data, but will not affect the generated code in any other way.  
 15762 You will normally not enable this option; instead, a language processor  
 15763 that needs this handling would enable it on your behalf.

15765 @item -fasynchronous-unwind-tables  
 15766 @opindex fasynchronous-unwind-tables  
 15767 Generate unwind table in dwarf2 format, if supported by target machine. The  
 15768 table is exact at each instruction boundary, so it can be used for stack  
 15769 unwinding from asynchronous events (such as debugger or garbage collector).

15771 @item -fpcc-struct-return  
 15772 @opindex fpcc-struct-return  
 15773 Return ``short'' @code{struct} and @code{union} values in memory like  
 15774 longer ones, rather than in registers. This convention is less  
 15775 efficient, but it has the advantage of allowing intercallability between  
 15776 GCC-compiled files and files compiled with other compilers, particularly  
 15777 the Portable C Compiler (pcc).

15779 The precise convention for returning structures in memory depends  
 15780 on the target configuration macros.

15782 Short structures and unions are those whose size and alignment match  
 15783 that of some integer type.

15785 @strong{Warning;} code compiled with the @option{-fpcc-struct-return}  
 15786 switch is not binary compatible with code compiled with the  
 15787 @option{-freg-struct-return} switch.  
 15788 Use it to conform to a non-default application binary interface.

15790 @item -freg-struct-return  
 15791 @opindex freg-struct-return  
 15792 Return @code{struct} and @code{union} values in registers when possible.  
 15793 This is more efficient for small structures than  
 15794 @option{-fpcc-struct-return}.

15796 If you specify neither @option{-fpcc-struct-return} nor  
 15797 @option{-freg-struct-return}, GCC defaults to whichever convention is  
 15798 standard for the target. If there is no standard convention, GCC

15799 defaults to @option{-fpcc-struct-return}, except on targets where GCC is  
 15800 the principal compiler. In those cases, we can choose the standard, and  
 15801 we chose the more efficient register return alternative.

15803 @strong{Warning;} code compiled with the @option{-freg-struct-return}  
 15804 switch is not binary compatible with code compiled with the  
 15805 @option{-fpcc-struct-return} switch.  
 15806 Use it to conform to a non-default application binary interface.

15808 @item -fshort-enums  
 15809 @opindex fshort-enums  
 15810 Allocate to an @code{enum} type only as many bytes as it needs for the  
 15811 declared range of possible values. Specifically, the @code{enum} type  
 15812 will be equivalent to the smallest integer type which has enough room.

15814 @strong{Warning;} the @option{-fshort-enums} switch causes GCC to generate  
 15815 code that is not binary compatible with code generated without that switch.  
 15816 Use it to conform to a non-default application binary interface.

15818 @item -fshort-double  
 15819 @opindex fshort-double  
 15820 Use the same size for @code{double} as for @code{float}.

15822 @strong{Warning;} the @option{-fshort-double} switch causes GCC to generate  
 15823 code that is not binary compatible with code generated without that switch.  
 15824 Use it to conform to a non-default application binary interface.

15826 @item -fshort-wchar  
 15827 @opindex fshort-wchar  
 15828 Override the underlying type for @code{wchar\_t} to be @code{short  
 15829 unsigned int} instead of the default for the target. This option is  
 15830 useful for building programs to run under WINE@.

15832 @strong{Warning;} the @option{-fshort-wchar} switch causes GCC to generate  
 15833 code that is not binary compatible with code generated without that switch.  
 15834 Use it to conform to a non-default application binary interface.

15836 @item -fno-common  
 15837 @opindex fno-common  
 15838 In C code, controls the placement of uninitialized global variables.  
 15839 Unix C compilers have traditionally permitted multiple definitions of  
 15840 such variables in different compilation units by placing the variables  
 15841 in a common block.  
 15842 This is the behavior specified by @option{-fcommon}, and is the default  
 15843 for GCC on most targets.  
 15844 On the other hand, this behavior is not required by ISO C, and on some  
 15845 targets may carry a speed or code size penalty on variable references.  
 15846 The @option{-fno-common} option specifies that the compiler should place  
 15847 uninitialized global variables in the data section of the object file,  
 15848 rather than generating them as common blocks.  
 15849 This has the effect that if the same variable is declared  
 15850 (without @code{extern}) in two different compilations,  
 15851 you will get a multiple-definition error when you link them.  
 15852 In this case, you must compile with @option{-fcommon} instead.  
 15853 Compiling with @option{-fno-common} is useful on targets for which  
 15854 it provides better performance, or if you wish to verify that the  
 15855 program will work on other systems which always treat uninitialized  
 15856 variable declarations this way.

15858 @item -fno-ident  
 15859 @opindex fno-ident  
 15860 Ignore the @code{#ident} directive.

15862 @item -finhibit-size-directive  
 15863 @opindex finhibit-size-directive  
 15864 Don't output a @code{.size} assembler directive, or anything else that

15865 would cause trouble if the function is split in the middle, and the  
 15866 two halves are placed at locations far apart in memory. This option is  
 15867 used when compiling @file{crtstuff.c}; you should not need to use it  
 15868 for anything else.

15870 @item -fverbose-asm  
 15871 @opindex fverbose-asm  
 15872 Put extra commentary information in the generated assembly code to  
 15873 make it more readable. This option is generally only of use to those  
 15874 who actually need to read the generated assembly code (perhaps while  
 15875 debugging the compiler itself).

15877 @option{-fno-verbose-asm}, the default, causes the  
 15878 extra information to be omitted and is useful when comparing two assembler  
 15879 files.

15881 @item -frecord-gcc-switches  
 15882 @opindex frecord-gcc-switches  
 15883 This switch causes the command line that was used to invoke the  
 15884 compiler to be recorded into the object file that is being created.  
 15885 This switch is only implemented on some targets and the exact format  
 15886 of the recording is target and binary file format dependent, but it  
 15887 usually takes the form of a section containing ASCII text. This  
 15888 switch is related to the @option{-fverbose-asm} switch, but that  
 15889 switch only records information in the assembler output file as  
 15890 comments, so it never reaches the object file.

15892 @item -fpic  
 15893 @opindex fpic  
 15894 @cindex global offset table  
 15895 @cindex PIC  
 15896 Generate position-independent code (PIC) suitable for use in a shared  
 15897 library, if supported for the target machine. Such code accesses all  
 15898 constant addresses through a global offset table (GOT)@. The dynamic  
 15899 loader resolves the GOT entries when the program starts (the dynamic  
 15900 loader is not part of GCC; it is part of the operating system). If  
 15901 the GOT size for the linked executable exceeds a machine-specific  
 15902 maximum size, you get an error message from the linker indicating that  
 15903 @option{-fpic} does not work; in that case, recompile with @option{-fPIC}  
 15904 instead. (These maximums are 8k on the SPARC and 32k  
 15905 on the m68k and RS/6000. The 386 has no such limit.)

15907 Position-independent code requires special support, and therefore works  
 15908 only on certain machines. For the 386, GCC supports PIC for System V  
 15909 but not for the Sun 386i. Code generated for the IBM RS/6000 is always  
 15910 position-independent.

15912 When this flag is set, the macros @code{\_\_pic\_\_} and @code{\_\_PIC\_\_}  
 15913 are defined to 1.

15915 @item -fPIC  
 15916 @opindex fPIC  
 15917 If supported for the target machine, emit position-independent code,  
 15918 suitable for dynamic linking and avoiding any limit on the size of the  
 15919 global offset table. This option makes a difference on the m68k,  
 15920 PowerPC and SPARC@.

15922 Position-independent code requires special support, and therefore works  
 15923 only on certain machines.

15925 When this flag is set, the macros @code{\_\_pic\_\_} and @code{\_\_PIC\_\_}  
 15926 are defined to 2.

15928 @item -fpie  
 15929 @itemx -fPIE  
 15930 @opindex fpie

15931 @opindex fPIE  
 15932 These options are similar to @option{-fpic} and @option{-fPIC}, but  
 15933 generated position independent code can be only linked into executables.  
 15934 Usually these options are used when @option{-pie} GCC option will be  
 15935 used during linking.

15937 @option{-fpie} and @option{-fPIE} both define the macros  
 15938 @code{\_\_pie\_\_} and @code{\_\_PIE\_\_}. The macros have the value 1  
 15939 for @option{-fpie} and 2 for @option{-fPIE}.

15941 @item -fno-jump-tables  
 15942 @opindex fno-jump-tables  
 15943 Do not use jump tables for switch statements even where it would be  
 15944 more efficient than other code generation strategies. This option is  
 15945 of use in conjunction with @option{-fpic} or @option{-fPIC} for  
 15946 building code which forms part of a dynamic linker and cannot  
 15947 reference the address of a jump table. On some targets, jump tables  
 15948 do not require a GOT and this option is not needed.

15950 @item -ffixed-@var{reg}  
 15951 @opindex ffixed  
 15952 Treat the register named @var{reg} as a fixed register; generated code  
 15953 should never refer to it (except perhaps as a stack pointer, frame  
 15954 pointer or in some other fixed role).

15956 @var{reg} must be the name of a register. The register names accepted  
 15957 are machine-specific and are defined in the @code{REGISTER\_NAMES}  
 15958 macro in the machine description macro file.

15960 This flag does not have a negative form, because it specifies a  
 15961 three-way choice.

15963 @item -fcall-used-@var{reg}  
 15964 @opindex fcall-used  
 15965 Treat the register named @var{reg} as an allocable register that is  
 15966 clobbered by function calls. It may be allocated for temporaries or  
 15967 variables that do not live across a call. Functions compiled this way  
 15968 will not save and restore the register @var{reg}.

15970 It is an error to used this flag with the frame pointer or stack pointer.  
 15971 Use of this flag for other registers that have fixed pervasive roles in  
 15972 the machine's execution model will produce disastrous results.

15974 This flag does not have a negative form, because it specifies a  
 15975 three-way choice.

15977 @item -fcall-saved-@var{reg}  
 15978 @opindex fcall-saved  
 15979 Treat the register named @var{reg} as an allocable register saved by  
 15980 functions. It may be allocated even for temporaries or variables that  
 15981 live across a call. Functions compiled this way will save and restore  
 15982 the register @var{reg} if they use it.

15984 It is an error to used this flag with the frame pointer or stack pointer.  
 15985 Use of this flag for other registers that have fixed pervasive roles in  
 15986 the machine's execution model will produce disastrous results.

15988 A different sort of disaster will result from the use of this flag for  
 15989 a register in which function values may be returned.

15991 This flag does not have a negative form, because it specifies a  
 15992 three-way choice.

15994 @item -fpack-struct[=@var{n}]  
 15995 @opindex fpack-struct  
 15996 Without a value specified, pack all structure members together without

15997 holes. When a value is specified (which must be a small power of two), pack  
 15998 structure members according to this value, representing the maximum  
 15999 alignment (that is, objects with default alignment requirements larger than  
 16000 this will be output potentially unaligned at the next fitting location.

16002 @strong{Warning:} the @option{-fpack-struct} switch causes GCC to generate  
 16003 code that is not binary compatible with code generated without that switch.  
 16004 Additionally, it makes the code suboptimal.  
 16005 Use it to conform to a non-default application binary interface.

16007 @item -finstrument-functions  
 16008 @opindex finstrument-functions  
 16009 Generate instrumentation calls for entry and exit to functions. Just  
 16010 after function entry and just before function exit, the following  
 16011 profiling functions will be called with the address of the current  
 16012 function and its call site. (On some platforms,  
 16013 @code{\_\_builtin\_return\_address} does not work beyond the current  
 16014 function, so the call site information may not be available to the  
 16015 profiling functions otherwise.)

16017 @smallexample  
 16018 void \_\_cyg\_profile\_func\_enter (void \*this\_fn,  
 16019 void \*call\_site);  
 16020 void \_\_cyg\_profile\_func\_exit (void \*this\_fn,  
 16021 void \*call\_site);  
 16022 @end smallexample

16024 The first argument is the address of the start of the current function,  
 16025 which may be looked up exactly in the symbol table.

16027 This instrumentation is also done for functions expanded inline in other  
 16028 functions. The profiling calls will indicate where, conceptually, the  
 16029 inline function is entered and exited. This means that addressable  
 16030 versions of such functions must be available. If all your uses of a  
 16031 function are expanded inline, this may mean an additional expansion of  
 16032 code size. If you use @samp{extern inline} in your C code, an  
 16033 addressable version of such functions must be provided. (This is  
 16034 normally the case anyways, but if you get lucky and the optimizer always  
 16035 expands the functions inline, you might have gotten away without  
 16036 providing static copies.)

16038 A function may be given the attribute @code{no\_instrument\_function}, in  
 16039 which case this instrumentation will not be done. This can be used, for  
 16040 example, for the profiling functions listed above, high-priority  
 16041 interrupt routines, and any functions from which the profiling functions  
 16042 cannot safely be called (perhaps signal handlers, if the profiling  
 16043 routines generate output or allocate memory).

16045 @item -finstrument-functions-exclude-file-list=@var{file},@var{file},@dots{}  
 16046 @opindex finstrument-functions-exclude-file-list

16048 Set the list of functions that are excluded from instrumentation (see  
 16049 the description of @code{-finstrument-functions}). If the file that  
 16050 contains a function definition matches with one of @var{file}, then  
 16051 that function is not instrumented. The match is done on substrings:  
 16052 if the @var{file} parameter is a substring of the file name, it is  
 16053 considered to be a match.

16055 For example,  
 16056 @code{-finstrument-functions-exclude-file-list=/bits/stl/include/sys}  
 16057 will exclude any inline function defined in files whose pathnames  
 16058 contain @code{/bits/stl} or @code{include/sys}.

16060 If, for some reason, you want to include letter @code{'\,'} in one of  
 16061 @var{sym}, write @code{'\,'}. For example,  
 16062 @code{-finstrument-functions-exclude-file-list='\,\,tmp'}

16063 (note the single quote surrounding the option).

16065 @item -finstrument-functions-exclude-function-list=@var{sym},@var{sym},@dots{}  
 16066 @opindex finstrument-functions-exclude-function-list

16068 This is similar to @code{-finstrument-functions-exclude-file-list},  
 16069 but this option sets the list of function names to be excluded from  
 16070 instrumentation. The function name to be matched is its user-visible  
 16071 name, such as @code{vector<int> blah(const vector<int> &)}, not the  
 16072 internal mangled name (e.g., @code{\_Z4blahRSt6vectorIiSaIEE}). The  
 16073 match is done on substrings: if the @var{sym} parameter is a substring  
 16074 of the function name, it is considered to be a match.

16076 @item -fstack-check  
 16077 @opindex fstack-check  
 16078 Generate code to verify that you do not go beyond the boundary of the  
 16079 stack. You should specify this flag if you are running in an  
 16080 environment with multiple threads, but only rarely need to specify it in  
 16081 a single-threaded environment since stack overflow is automatically  
 16082 detected on nearly all systems if there is only one stack.

16084 Note that this switch does not actually cause checking to be done; the  
 16085 operating system or the language runtime must do that. The switch causes  
 16086 generation of code to ensure that they see the stack being extended.

16088 You can additionally specify a string parameter: @code{no} means no  
 16089 checking, @code{generic} means force the use of old-style checking,  
 16090 @code{specific} means use the best checking method and is equivalent  
 16091 to bare @option{-fstack-check}.

16093 Old-style checking is a generic mechanism that requires no specific  
 16094 target support in the compiler but comes with the following drawbacks:

16096 @enumerate  
 16097 @item  
 16098 Modified allocation strategy for large objects: they will always be  
 16099 allocated dynamically if their size exceeds a fixed threshold.

16101 @item  
 16102 Fixed limit on the size of the static frame of functions: when it is  
 16103 topped by a particular function, stack checking is not reliable and  
 16104 a warning is issued by the compiler.

16106 @item  
 16107 Inefficiency: because of both the modified allocation strategy and the  
 16108 generic implementation, the performances of the code are hampered.  
 16109 @end enumerate

16111 Note that old-style stack checking is also the fallback method for  
 16112 @code{specific} if no target support has been added in the compiler.

16114 @item -fstack-limit-register=@var{reg}  
 16115 @itemx -fstack-limit-symbol=@var{sym}  
 16116 @itemx -fno-stack-limit  
 16117 @opindex fstack-limit-register  
 16118 @opindex fstack-limit-symbol  
 16119 @opindex fno-stack-limit  
 16120 Generate code to ensure that the stack does not grow beyond a certain value,  
 16121 either the value of a register or the address of a symbol. If the stack  
 16122 would grow beyond the value, a signal is raised. For most targets,  
 16123 the signal is raised before the stack overruns the boundary, so  
 16124 it is possible to catch the signal without taking special precautions.

16126 For instance, if the stack starts at absolute address @samp{0x80000000}  
 16127 and grows downwards, you can use the flags  
 16128 @option{-fstack-limit-symbol=\_\_stack\_limit} and

16129 `@option{-Wl,--defsym,__stack_limit=0x7ffe0000}` to enforce a stack limit  
 16130 of 128KB@. Note that this may only work with the GNU linker.

16132 `@cindex` aliasing of parameters  
 16133 `@cindex` parameters, aliased  
 16134 `@item -fargument-alias`  
 16135 `@itemx -fargument-noalias`  
 16136 `@itemx -fargument-noalias-global`  
 16137 `@itemx -fargument-noalias-anything`  
 16138 `@opindex` fargument-alias  
 16139 `@opindex` fargument-noalias  
 16140 `@opindex` fargument-noalias-global  
 16141 `@opindex` fargument-noalias-anything  
 16142 Specify the possible relationships among parameters and between  
 16143 parameters and global data.

16145 `@option{-fargument-alias}` specifies that arguments (parameters) may  
 16146 alias each other and may alias global storage.\*  
 16147 `@option{-fargument-noalias}` specifies that arguments do not alias  
 16148 each other, but may alias global storage.\*  
 16149 `@option{-fargument-noalias-global}` specifies that arguments do not  
 16150 alias each other and do not alias global storage.  
 16151 `@option{-fargument-noalias-anything}` specifies that arguments do not  
 16152 alias any other storage.

16154 Each language will automatically use whatever option is required by  
 16155 the language standard. You should not need to use these options yourself.

16157 `@item -fleading-underscore`  
 16158 `@opindex` fleading-underscore  
 16159 This option and its counterpart, `@option{-fno-leading-underscore}`, forcibly  
 16160 change the way C symbols are represented in the object file. One use  
 16161 is to help link with legacy assembly code.

16163 `@strong{Warning:}` the `@option{-fleading-underscore}` switch causes GCC to  
 16164 generate code that is not binary compatible with code generated without that  
 16165 switch. Use it to conform to a non-default application binary interface.  
 16166 Not all targets provide complete support for this switch.

16168 `@item -ftls-model=@var{model}`  
 16169 `@opindex` ftls-model  
 16170 Alter the thread-local storage model to be used (`@pxref{Thread-Local}`).  
 16171 The `@var{model}` argument should be one of `@code{global-dynamic}`,  
 16172 `@code{local-dynamic}`, `@code{initial-exec}` or `@code{local-exec}`.

16174 The default without `@option{-fpic}` is `@code{initial-exec}`; with  
 16175 `@option{-fpic}` the default is `@code{global-dynamic}`.

16177 `@item -fvisibility=@var{default|internal|hidden|protected}`  
 16178 `@opindex` fvisibility  
 16179 Set the default ELF image symbol visibility to the specified option---all  
 16180 symbols will be marked with this unless overridden within the code.  
 16181 Using this feature can very substantially improve linking and  
 16182 load times of shared object libraries, produce more optimized  
 16183 code, provide near-perfect API export and prevent symbol clashes.  
 16184 It is `@strong{strongly}` recommended that you use this in any shared objects  
 16185 you distribute.

16187 Despite the nomenclature, `@code{default}` always means public ie;  
 16188 available to be linked against from outside the shared object.  
 16189 `@code{protected}` and `@code{internal}` are pretty useless in real-world  
 16190 usage so the only other commonly used option will be `@code{hidden}`.  
 16191 The default if `@option{-fvisibility}` isn't specified is  
 16192 `@code{default}`, i.e., make every  
 16193 symbol public---this causes the same behavior as previous versions of  
 16194 GCC@.

16196 A good explanation of the benefits offered by ensuring ELF  
 16197 symbols have the correct visibility is given by "How To Write  
 16198 Shared Libraries" by Ulrich Drepper (which can be found at  
 16199 `@w{@uref{http://people.redhat.com/~drepper/}}`)---however a superior  
 16200 solution made possible by this option to marking things hidden when  
 16201 the default is public is to make the default hidden and mark things  
 16202 public. This is the norm with DLL's on Windows and with `@option{-fvisibility=hi}`  
 16203 and `@code{__attribute__((visibility("default")))}` instead of  
 16204 `@code{__declspec(dllexport)}` you get almost identical semantics with  
 16205 identical syntax. This is a great boon to those working with  
 16206 cross-platform projects.

16208 For those adding visibility support to existing code, you may find  
 16209 `@samp{#pragma GCC visibility}` of use. This works by you enclosing  
 16210 the declarations you wish to set visibility for with (for example)  
 16211 `@samp{#pragma GCC visibility push(hidden)}` and  
 16212 `@samp{#pragma GCC visibility pop}`.  
 16213 Bear in mind that symbol visibility should be viewed `@strong{as}`  
 16214 part of the API interface contract and thus all new code should  
 16215 always specify visibility when it is not the default ie; declarations  
 16216 only for use within the local DSO should `@strong{always}` be marked explicitly  
 16217 as hidden as so to avoid PLT indirection overheads---making this  
 16218 abundantly clear also aids readability and self-documentation of the code.  
 16219 Note that due to ISO C++ specification requirements, operator new and  
 16220 operator delete must always be of default visibility.

16222 Be aware that headers from outside your project, in particular system  
 16223 headers and headers from any other library you use, may not be  
 16224 expecting to be compiled with visibility other than the default. You  
 16225 may need to explicitly say `@samp{#pragma GCC visibility push(default)}`  
 16226 before including any such headers.

16228 `@samp{extern}` declarations are not affected by `@samp{-fvisibility}`, so  
 16229 a lot of code can be recompiled with `@samp{-fvisibility=hidden}` with  
 16230 no modifications. However, this means that calls to `@samp{extern}`  
 16231 functions with no explicit visibility will use the PLT, so it is more  
 16232 effective to use `@samp{__attribute__((visibility))}` and/or  
 16233 `@samp{#pragma GCC visibility}` to tell the compiler which `@samp{extern}`  
 16234 declarations should be treated as hidden.

16236 Note that `@samp{-fvisibility}` does affect C++ vague linkage  
 16237 entities. This means that, for instance, an exception class that will  
 16238 be thrown between DSOs must be explicitly marked with default  
 16239 visibility so that the `@samp{type_info}` nodes will be unified between  
 16240 the DSOs.

16242 An overview of these techniques, their benefits and how to use them  
 16243 is at `@w{@uref{http://gcc.gnu.org/wiki/Visibility}}`.

16245 `@end table`

16247 `@c man end`

16249 `@node` Environment Variables  
 16250 `@section` Environment Variables Affecting GCC  
 16251 `@cindex` environment variables

16253 `@c man` begin ENVIRONMENT  
 16254 This section describes several environment variables that affect how GCC  
 16255 operates. Some of them work by specifying directories or prefixes to use  
 16256 when searching for various kinds of files. Some are used to specify other  
 16257 aspects of the compilation environment.

16259 Note that you can also specify places to search using options such as  
 16260 `@option{-B}`, `@option{-I}` and `@option{-L}` (`@pxref{Directory Options}`). These

16261 take precedence over places specified using environment variables, which  
 16262 in turn take precedence over those specified by the configuration of GCC@.  
 16263 @xref{Driver,, Controlling the Compilation Driver @file{gcc}}, gccint,  
 16264 GNU Compiler Collection (GCC) Internals).

16266 @table @env  
 16267 @item LANG  
 16268 @itemx LC\_CTYPE  
 16269 @c @itemx LC\_COLLATE  
 16270 @itemx LC\_MESSAGES  
 16271 @c @itemx LC\_MONETARY  
 16272 @c @itemx LC\_NUMERIC  
 16273 @c @itemx LC\_TIME  
 16274 @itemx LC\_ALL  
 16275 @findex LANG  
 16276 @findex LC\_CTYPE  
 16277 @c @findex LC\_COLLATE  
 16278 @findex LC\_MESSAGES  
 16279 @c @findex LC\_MONETARY  
 16280 @c @findex LC\_NUMERIC  
 16281 @c @findex LC\_TIME  
 16282 @findex LC\_ALL  
 16283 @cindex locale  
 16284 These environment variables control the way that GCC uses  
 16285 localization information that allow GCC to work with different  
 16286 national conventions. GCC inspects the locale categories  
 16287 @env{LC\_CTYPE} and @env{LC\_MESSAGES} if it has been configured to do  
 16288 so. These locale categories can be set to any value supported by your  
 16289 installation. A typical value is @samp{en\_GB.UTF-8} for English in the United  
 16290 Kingdom encoded in UTF-8.

16292 The @env{LC\_CTYPE} environment variable specifies character  
 16293 classification. GCC uses it to determine the character boundaries in  
 16294 a string; this is needed for some multibyte encodings that contain quote  
 16295 and escape characters that would otherwise be interpreted as a string  
 16296 end or escape.

16298 The @env{LC\_MESSAGES} environment variable specifies the language to  
 16299 use in diagnostic messages.

16301 If the @env{LC\_ALL} environment variable is set, it overrides the value  
 16302 of @env{LC\_CTYPE} and @env{LC\_MESSAGES}; otherwise, @env{LC\_CTYPE}  
 16303 and @env{LC\_MESSAGES} default to the value of the @env{LANG}  
 16304 environment variable. If none of these variables are set, GCC  
 16305 defaults to traditional C English behavior.

16307 @item TMPDIR  
 16308 @findex TMPDIR  
 16309 If @env{TMPDIR} is set, it specifies the directory to use for temporary  
 16310 files. GCC uses temporary files to hold the output of one stage of  
 16311 compilation which is to be used as input to the next stage: for example,  
 16312 the output of the preprocessor, which is the input to the compiler  
 16313 proper.

16315 @item GCC\_EXEC\_PREFIX  
 16316 @findex GCC\_EXEC\_PREFIX  
 16317 If @env{GCC\_EXEC\_PREFIX} is set, it specifies a prefix to use in the  
 16318 names of the subprograms executed by the compiler. No slash is added  
 16319 when this prefix is combined with the name of a subprogram, but you can  
 16320 specify a prefix that ends with a slash if you wish.

16322 If @env{GCC\_EXEC\_PREFIX} is not set, GCC will attempt to figure out  
 16323 an appropriate prefix to use based on the pathname it was invoked with.

16325 If GCC cannot find the subprogram using the specified prefix, it  
 16326 tries looking in the usual places for the subprogram.

16328 The default value of @env{GCC\_EXEC\_PREFIX} is  
 16329 @file{@var{prefix}/lib/gcc/} where @var{prefix} is the prefix to  
 16330 the installed compiler. In many cases @var{prefix} is the value  
 16331 of @code{prefix} when you ran the @file{configure} script.

16333 Other prefixes specified with @option{-B} take precedence over this prefix.

16335 This prefix is also used for finding files such as @file{crt0.o} that are  
 16336 used for linking.

16338 In addition, the prefix is used in an unusual way in finding the  
 16339 directories to search for header files. For each of the standard  
 16340 directories whose name normally begins with @samp{/usr/local/lib/gcc}  
 16341 (more precisely, with the value of @env{GCC\_INCLUDE\_DIR}), GCC tries  
 16342 replacing that beginning with the specified prefix to produce an  
 16343 alternate directory name. Thus, with @option{-Bfoo/}, GCC will search  
 16344 @file{foo/bar} where it would normally search @file{/usr/local/lib/bar}.  
 16345 These alternate directories are searched first; the standard directories  
 16346 come next. If a standard directory begins with the configured  
 16347 @var{prefix} then the value of @var{prefix} is replaced by  
 16348 @env{GCC\_EXEC\_PREFIX} when looking for header files.

16350 @item COMPILER\_PATH  
 16351 @findex COMPILER\_PATH  
 16352 The value of @env{COMPILER\_PATH} is a colon-separated list of  
 16353 directories, much like @env{PATH}. GCC tries the directories thus  
 16354 specified when searching for subprograms, if it can't find the  
 16355 subprograms using @env{GCC\_EXEC\_PREFIX}.

16357 @item LIBRARY\_PATH  
 16358 @findex LIBRARY\_PATH  
 16359 The value of @env{LIBRARY\_PATH} is a colon-separated list of  
 16360 directories, much like @env{PATH}. When configured as a native compiler,  
 16361 GCC tries the directories thus specified when searching for special  
 16362 linker files, if it can't find them using @env{GCC\_EXEC\_PREFIX}. Linking  
 16363 using GCC also uses these directories when searching for ordinary  
 16364 libraries for the @option{-l} option (but directories specified with  
 16365 @option{-L} come first).

16367 @item LANG  
 16368 @findex LANG  
 16369 @cindex locale definition  
 16370 This variable is used to pass locale information to the compiler. One way in  
 16371 which this information is used is to determine the character set to be used  
 16372 when character literals, string literals and comments are parsed in C and C++.  
 16373 When the compiler is configured to allow multibyte characters,  
 16374 the following values for @env{LANG} are recognized:

16376 @table @samp  
 16377 @item C-JIS  
 16378 Recognize JIS characters.  
 16379 @item C-SJIS  
 16380 Recognize SJIS characters.  
 16381 @item C-EUCJP  
 16382 Recognize EUCJP characters.  
 16383 @end table

16385 If @env{LANG} is not defined, or if it has some other value, then the  
 16386 compiler will use mblen and mbtowc as defined by the default locale to  
 16387 recognize and translate multibyte characters.  
 16388 @end table

16390 @noindent  
 16391 Some additional environments variables affect the behavior of the  
 16392 preprocessor.

16394 @include cppenv.texi

16396 @c man end

16398 @node Precompiled Headers  
 16399 @section Using Precompiled Headers  
 16400 @cindex precompiled headers  
 16401 @cindex speed of compilation

16403 Often large projects have many header files that are included in every  
 16404 source file. The time the compiler takes to process these header files  
 16405 over and over again can account for nearly all of the time required to  
 16406 build the project. To make builds faster, GCC allows users to  
 16407 'precompile' a header file; then, if builds can use the precompiled  
 16408 header file they will be much faster.

16410 To create a precompiled header file, simply compile it as you would any  
 16411 other file, if necessary using the @option{-x} option to make the driver  
 16412 treat it as a C or C++ header file. You will probably want to use a  
 16413 tool like @command{make} to keep the precompiled header up-to-date when  
 16414 the headers it contains change.

16416 A precompiled header file will be searched for when @code{#include} is  
 16417 seen in the compilation. As it searches for the included file  
 16418 (@pxref{Search Path,,Search Path,cpp,The C Preprocessor}) the  
 16419 compiler looks for a precompiled header in each directory just before it  
 16420 looks for the include file in that directory. The name searched for is  
 16421 the name specified in the @code{#include} with @samp{.gch} appended. If  
 16422 the precompiled header file can't be used, it is ignored.

16424 For instance, if you have @code{#include "all.h"}, and you have  
 16425 @file{all.h.gch} in the same directory as @file{all.h}, then the  
 16426 precompiled header file will be used if possible, and the original  
 16427 header will be used otherwise.

16429 Alternatively, you might decide to put the precompiled header file in a  
 16430 directory and use @option{-I} to ensure that directory is searched  
 16431 before (or instead of) the directory containing the original header.  
 16432 Then, if you want to check that the precompiled header file is always  
 16433 used, you can put a file of the same name as the original header in this  
 16434 directory containing an @code{#error} command.

16436 This also works with @option{-include}. So yet another way to use  
 16437 precompiled headers, good for projects not designed with precompiled  
 16438 header files in mind, is to simply take most of the header files used by  
 16439 a project, include them from another header file, precompile that header  
 16440 file, and @option{-include} the precompiled header. If the header files  
 16441 have guards against multiple inclusion, they will be skipped because  
 16442 they've already been included (in the precompiled header).

16444 If you need to precompile the same header file for different  
 16445 languages, targets, or compiler options, you can instead make a  
 16446 @emph{directory} named like @file{all.h.gch}, and put each precompiled  
 16447 header in the directory, perhaps using @option{-o}. It doesn't matter  
 16448 what you call the files in the directory, every precompiled header in  
 16449 the directory will be considered. The first precompiled header  
 16450 encountered in the directory that is valid for this compilation will  
 16451 be used; they're searched in no particular order.

16453 There are many other possibilities, limited only by your imagination,  
 16454 good sense, and the constraints of your build system.

16456 A precompiled header file can be used only when these conditions apply:

16458 @itemize

16459 @item  
 16460 Only one precompiled header can be used in a particular compilation.

16462 @item  
 16463 A precompiled header can't be used once the first C token is seen. You  
 16464 can have preprocessor directives before a precompiled header; you can  
 16465 even include a precompiled header from inside another header, so long as  
 16466 there are no C tokens before the @code{#include}.

16468 @item  
 16469 The precompiled header file must be produced for the same language as  
 16470 the current compilation. You can't use a C precompiled header for a C++  
 16471 compilation.

16473 @item  
 16474 The precompiled header file must have been produced by the same compiler  
 16475 binary as the current compilation is using.

16477 @item  
 16478 Any macros defined before the precompiled header is included must  
 16479 either be defined in the same way as when the precompiled header was  
 16480 generated, or must not affect the precompiled header, which usually  
 16481 means that they don't appear in the precompiled header at all.

16483 The @option{-D} option is one way to define a macro before a  
 16484 precompiled header is included; using a @code{#define} can also do it.  
 16485 There are also some options that define macros implicitly, like  
 16486 @option{-O} and @option{-Wdeprecated}; the same rule applies to macros  
 16487 defined this way.

16489 @item If debugging information is output when using the precompiled  
 16490 header, using @option{-g} or similar, the same kind of debugging information  
 16491 must have been output when building the precompiled header. However,  
 16492 a precompiled header built using @option{-g} can be used in a compilation  
 16493 when no debugging information is being output.

16495 @item The same @option{-m} options must generally be used when building  
 16496 and using the precompiled header. @xref{Submodel Options},  
 16497 for any cases where this rule is relaxed.

16499 @item Each of the following options must be the same when building and using  
 16500 the precompiled header:

16502 @gcoptlist{-fexceptions}

16504 @item  
 16505 Some other command-line options starting with @option{-f},  
 16506 @option{-p}, or @option{-O} must be defined in the same way as when  
 16507 the precompiled header was generated. At present, it's not clear  
 16508 which options are safe to change and which are not; the safest choice  
 16509 is to use exactly the same options when generating and using the  
 16510 precompiled header. The following are known to be safe:

16512 @gcoptlist{-fmessage-length= -fpreprocessed -fsched-interblock @gol  
 16513 -fsched-spec -fsched-spec-load -fsched-spec-load-dangerous @gol  
 16514 -fsched-verbose=<number> -fschedule-insns -fvisibility=@gol  
 16515 -pedantic-errors}

16517 @end itemize

16519 For all of these except the last, the compiler will automatically  
 16520 ignore the precompiled header if the conditions aren't met. If you  
 16521 find an option combination that doesn't work and doesn't cause the  
 16522 precompiled header to be ignored, please consider filing a bug report,  
 16523 see @ref{Bugs}.

16525 If you do use differing options when generating and using the  
 16526 precompiled header, the actual behavior will be a mixture of the  
 16527 behavior for the options. For instance, if you use `@option{-g}` to  
 16528 generate the precompiled header but not when using it, you may or may  
 16529 not get debugging information for routines in the precompiled header.

16531 @node Running Protoize  
 16532 @section Running Protoize

16534 The program `@code{protoize}` is an optional part of GCC@. You can use  
 16535 it to add prototypes to a program, thus converting the program to ISO  
 16536 C in one respect. The companion program `@code{unprotoize}` does the  
 16537 reverse: it removes argument types from any prototypes that are found.

16539 When you run these programs, you must specify a set of source files as  
 16540 command line arguments. The conversion programs start out by compiling  
 16541 these files to see what functions they define. The information gathered  
 16542 about a file `@var{foo}` is saved in a file named `@file{@var{foo}.X}`.

16544 After scanning comes actual conversion. The specified files are all  
 16545 eligible to be converted; any files they include (whether sources or  
 16546 just headers) are eligible as well.

16548 But not all the eligible files are converted. By default,  
 16549 `@code{protoize}` and `@code{unprotoize}` convert only source and header  
 16550 files in the current directory. You can specify additional directories  
 16551 whose files should be converted with the `@option{-d @var{directory}}`  
 16552 option. You can also specify particular files to exclude with the  
 16553 `@option{-x @var{file}}` option. A file is converted if it is eligible, its  
 16554 directory name matches one of the specified directory names, and its  
 16555 name within the directory has not been excluded.

16557 Basic conversion with `@code{protoize}` consists of rewriting most  
 16558 function definitions and function declarations to specify the types of  
 16559 the arguments. The only ones not rewritten are those for varargs  
 16560 functions.

16562 `@code{protoize}` optionally inserts prototype declarations at the  
 16563 beginning of the source file, to make them available for any calls that  
 16564 precede the function's definition. Or it can insert prototype  
 16565 declarations with block scope in the blocks where undeclared functions  
 16566 are called.

16568 Basic conversion with `@code{unprotoize}` consists of rewriting most  
 16569 function declarations to remove any argument types, and rewriting  
 16570 function definitions to the old-style pre-ISO form.

16572 Both conversion programs print a warning for any function declaration or  
 16573 definition that they can't convert. You can suppress these warnings  
 16574 with `@option{-q}`.

16576 The output from `@code{protoize}` or `@code{unprotoize}` replaces the  
 16577 original source file. The original file is renamed to a name ending  
 16578 with `@samp{.save}` (for DOS, the saved filename ends in `@samp{.sav}`  
 16579 without the original `@samp{.c}` suffix). If the `@samp{.save}` (`@samp{.sav}`  
 16580 for DOS) file already exists, then the source file is simply discarded.

16582 `@code{protoize}` and `@code{unprotoize}` both depend on GCC itself to  
 16583 scan the program and collect information about the functions it uses.  
 16584 So neither of these programs will work until GCC is installed.

16586 Here is a table of the options you can use with `@code{protoize}` and  
 16587 `@code{unprotoize}`. Each option works with both programs unless  
 16588 otherwise stated.

16590 @table @code

16591 @item -B `@var{directory}`  
 16592 Look for the file `@file{SYSCALLS.c.X}` in `@var{directory}`, instead of the  
 16593 usual directory (normally `@file{/usr/local/lib}`). This file contains  
 16594 prototype information about standard system functions. This option  
 16595 applies only to `@code{protoize}`.

16597 @item -c `@var{compilation-options}`  
 16598 Use `@var{compilation-options}` as the options when running `@command{gcc}` to  
 16599 produce the `@samp{.X}` files. The special option `@option{-aux-info}` is  
 16600 always passed in addition, to tell `@command{gcc}` to write a `@samp{.X}` file.

16602 Note that the compilation options must be given as a single argument to  
 16603 `@code{protoize}` or `@code{unprotoize}`. If you want to specify several  
 16604 `@command{gcc}` options, you must quote the entire set of compilation options  
 16605 to make them a single word in the shell.

16607 There are certain `@command{gcc}` arguments that you cannot use, because they  
 16608 would produce the wrong kind of output. These include `@option{-g}`,  
 16609 `@option{-O}`, `@option{-c}`, `@option{-S}`, and `@option{-o}` If you include these in  
 16610 the `@var{compilation-options}`, they are ignored.

16612 @item -C  
 16613 Rename files to end in `@samp{.C}` (`@samp{.cc}` for DOS-based file  
 16614 systems) instead of `@samp{.c}`. This is convenient if you are converting  
 16615 a C program to C++. This option applies only to `@code{protoize}`.

16617 @item -g  
 16618 Add explicit global declarations. This means inserting explicit  
 16619 declarations at the beginning of each source file for each function  
 16620 that is called in the file and was not declared. These declarations  
 16621 precede the first function definition that contains a call to an  
 16622 undeclared function. This option applies only to `@code{protoize}`.

16624 @item -i `@var{string}`  
 16625 Indent old-style parameter declarations with the string `@var{string}`.  
 16626 This option applies only to `@code{protoize}`.

16628 `@code{unprotoize}` converts prototyped function definitions to old-style  
 16629 function definitions, where the arguments are declared between the  
 16630 argument list and the initial `@samp{@{}`. By default, `@code{unprotoize}`  
 16631 uses five spaces as the indentation. If you want to indent with just  
 16632 one space instead, use `@option{-i " "}`.

16634 @item -k  
 16635 Keep the `@samp{.X}` files. Normally, they are deleted after conversion  
 16636 is finished.

16638 @item -l  
 16639 Add explicit local declarations. `@code{protoize}` with `@option{-l}` inserts  
 16640 a prototype declaration for each function in each block which calls the  
 16641 function without any declaration. This option applies only to  
 16642 `@code{protoize}`.

16644 @item -n  
 16645 Make no real changes. This mode just prints information about the conversions  
 16646 that would have been done without `@option{-n}`.

16648 @item -N  
 16649 Make no `@samp{.save}` files. The original files are simply deleted.  
 16650 Use this option with caution.

16652 @item -p `@var{program}`  
 16653 Use the program `@var{program}` as the compiler. Normally, the name  
 16654 `@file{gcc}` is used.

16656 @item -q

16657 Work quietly. Most warnings are suppressed.

16659 @item -v  
16660 Print the version number, just like @option{-v} for @command{gcc}.  
16661 @end table

16663 If you need special compiler options to compile one of your program's  
16664 source files, then you should generate that file's @samp{.X} file  
16665 specially, by running @command{gcc} on that source file with the  
16666 appropriate options and the option @option{-aux-info}. Then run  
16667 @code{protoize} on the entire set of files. @code{protoize} will use  
16668 the existing @samp{.X} file because it is newer than the source file.  
16669 For example:

16671 @smallexample  
16672 gcc -Dfoo=bar file1.c -aux-info file1.X  
16673 protoize \*.c  
16674 @end smallexample

16676 @noindent  
16677 You need to include the special files along with the rest in the  
16678 @code{protoize} command, even though their @samp{.X} files already  
16679 exist, because otherwise they won't get converted.

16681 @xref{Protoize Caveats}, for more information on how to use  
16682 @code{protoize} successfully.



new/gcc/testsuite/gcc.target/i386/local.c

1

\*\*\*\*\*

483 Sun Oct 28 20:56:11 2012

new/gcc/testsuite/gcc.target/i386/local.c

Implement `-fstrict-calling-conventions`

Stock GCC is overly willing to violate the ABI when calling local functions, such that it passes arguments in registers on i386. This hampers debugging with anything other than a fully-aware DWARF debugger, and is generally not something we desire.

Implement a flag which disables this behaviour, enabled by default. The flag is global, though only effective on i386, to more easily allow its globalization later which, given the odds, is likely to be necessary.

\*\*\*\*\*

```
1 /* { dg-do compile } */
2 /* { dg-options "-O2 -funit-at-a-time -fno-strict-calling-conventions" { target
3 /* { dg-options "-O2 -funit-at-a-time" { target lp64 } } */
2 /* { dg-options "-O2 -funit-at-a-time" } */
4 /* { dg-final { scan-assembler "magic\[\^\n\]*eax" { target ilp32 } } } */
5 /* { dg-final { scan-assembler "magic\[\^\n\]*edi" { target lp64 } } } */
```

```
7 /* Verify that local calling convention is used. */
```

```
8 static t(int) __attribute__((noinline));
```

```
9 m()
```

```
10 {
```

```
11     t(1);
```

```
12 }
```

unchanged\_portion\_omitted

new/gcc/testsuite/gcc.target/i386/save-args-1.c

1

\*\*\*\*\*

376 Sun Oct 28 20:56:11 2012

new/gcc/testsuite/gcc.target/i386/save-args-1.c

testsuite/save-args-1: Actually test the full thing, rather than just the bit I

\*\*\*\*\*

```
1 /* Test -msave-args */
2 /* { dg-do compile { target { { i?86-*-solaris2.* } && lp64 } } } */
3 /* { dg-options "-msave-args" } */
4 /* { dg-final { scan-assembler "movq\t%rsi, -32\\(%rbp\\)" } } */
5 /* { dg-final { scan-assembler "movq\t%rsi, -16\\(%rbp\\)" } } */
6 /* { dg-final { scan-assembler "movq\t%rdi, -8\\(%rbp\\)" } } */
```

```
8 int
9 foo(int argc, char **argv)
10 {
11     return (1);
12 }
13 #endif /* ! codereview */
```

new/gcc/testsuite/gcc.target/i386/strict-cc.c

1

\*\*\*\*\*

373 Sun Oct 28 20:56:11 2012

new/gcc/testsuite/gcc.target/i386/strict-cc.c

Implement `-fstrict-calling-conventions`

Stock GCC is overly willing to violate the ABI when calling local functions, such that it passes arguments in registers on i386. This hampers debugging with anything other than a fully-aware DWARF debugger, and is generally not something we desire.

Implement a flag which disables this behaviour, enabled by default. The flag is global, though only effective on i386, to more easily allow its globalization later which, given the odds, is likely to be necessary.

\*\*\*\*\*

```
1 /* { dg-do compile { target { ilp32 } } } */
2 /* { dg-options "-O2 -funit-at-a-time -fstrict-calling-conventions" } */
3 /* { dg-final { scan-assembler "pushl.*\\$1" } } */

5 #include <stdio.h>

7 /* Verify that local calling convention is not used if strict conventions. */
8 static t(int) __attribute__((noinline));
9 m()
10 {
11     t(1);
12 }

14 static t(int a)
15 {
16     printf("%d\n", a);
17 }
18 #endif /* ! codereview */
```

```

*****
7746 Sun Oct 28 20:56:11 2012
new/libiberty/testsuite/test-demangle.c
libiberty/testsuite: Avoid conflicting getline()
*****
_unchanged_portion_omitted_

42 static unsigned int lineno;

44 /* Safely read a single line of arbitrary length from standard input. */

46 #define LINELEN 80

48 static void
49 _getline(buf)
49 getline(buf)
50 struct line *buf;
51 {
52   char *data = buf->data;
53   size_t alloc = buf->allocated;
54   size_t count = 0;
55   int c;

57   if (data == 0)
58     {
59       data = xmalloc (LINELEN);
60       alloc = LINELEN;
61     }

63   /* Skip comment lines. */
64   while ((c = getchar()) == '#')
65     {
66       while ((c = getchar()) != EOF && c != '\n');
67       lineno++;
68     }

70   /* c is the first character on the line, and it's not a comment
71      line: copy this line into the buffer and return. */
72   while (c != EOF && c != '\n')
73     {
74       if (count + 1 >= alloc)
75         {
76           alloc *= 2;
77           data = xrealloc (data, alloc);
78         }
79       data[count++] = c;
80       c = getchar();
81     }
82   lineno++;
83   data[count] = '\0';

85   buf->data = data;
86   buf->allocated = alloc;
87 }
_unchanged_portion_omitted_

149 /* The tester operates on a data file consisting of groups of lines:
150 options
151 input to be demangled
152 expected output

154 Supported options:
155 --format=<name> Sets the demangling style.
156 --no-params There are two lines of expected output; the first
157 is with DMGL_PARAMS, the second is without it.
158 --is-v3-ctor Calls is_gnu_v3_mangled_ctor on input; expected

```

```

159 output is an integer representing ctor_kind.
160 --is-v3-dtor Likewise, but for dtors.
161 --ret-postfix Passes the DMGL_RET_POSTFIX option

163 For compatibility, just in case it matters, the options line may be
164 empty, to mean --format=auto. If it doesn't start with --, then it
165 may contain only a format name.
166 */

168 int
169 main(argc, argv)
170 int argc;
171 char **argv;
172 {
173   enum demangling_styles style = auto_demangling;
174   int no_params;
175   int is_v3_ctor;
176   int is_v3_dtor;
177   int ret_postfix;
178   struct line format;
179   struct line input;
180   struct line expect;
181   char *result;
182   int failures = 0;
183   int tests = 0;

185   if (argc > 1)
186     {
187       fprintf (stderr, "usage: %s < test-set\n", argv[0]);
188       return 2;
189     }

191   format.data = 0;
192   input.data = 0;
193   expect.data = 0;

195   for (;;)
196     {
197       const char *inp;
198
199       _getline (&format);
200       getline (&format);
201       if (feof (stdin))
202         break;

203       _getline (&input);
204       _getline (&expect);
205       getline (&input);
206       getline (&expect);

207       inp = protect_end (input.data);

208       tests++;

210       no_params = 0;
211       ret_postfix = 0;
212       is_v3_ctor = 0;
213       is_v3_dtor = 0;
214       if (format.data[0] == '\0')
215         style = auto_demangling;
216       else if (format.data[0] != '-')
217         {
218           style = cplus_demangle_name_to_style (format.data);
219           if (style == unknown_demangling)
220             {
221               printf ("FAIL at line %d: unknown demangling style %s\n",

```

```

222         lineno, format.data);
223     failures++;
224     continue;
225 }
226 }
227 else
228 {
229     char *p;
230     char *opt;

232     p = format.data;
233     while (*p != '\0')
234     {
235         char c;

237         opt = p;
238         p += strcspn (p, " \t=");
239         c = *p;
240         *p = '\0';
241         if (strcmp (opt, "--format") == 0 && c == '=')
242         {
243             char *fstyle;

245             *p = c;
246             ++p;
247             fstyle = p;
248             p += strcspn (p, " \t");
249             c = *p;
250             *p = '\0';
251             style = cplus_demangle_name_to_style (fstyle);
252             if (style == unknown_demangling)
253             {
254                 printf ("FAIL at line %d: unknown demangling style %s\n",
255                     lineno, fstyle);
256                 failures++;
257                 continue;
258             }
259         }
260         else if (strcmp (opt, "--no-params") == 0)
261             no_params = 1;
262         else if (strcmp (opt, "--is-v3-ctor") == 0)
263             is_v3_ctor = 1;
264         else if (strcmp (opt, "--is-v3-dtor") == 0)
265             is_v3_dtor = 1;
266         else if (strcmp (opt, "--ret-postfix") == 0)
267             ret_postfix = 1;
268         else
269         {
270             printf ("FAIL at line %d: unrecognized option %s\n",
271                 lineno, opt);
272             failures++;
273             continue;
274         }
275         *p = c;
276         p += strspn (p, " \t");
277     }
278 }

280 if (is_v3_ctor || is_v3_dtor)
281 {
282     char buf[20];

284     if (is_v3_ctor)
285     {
286         enum gnu_v3_ctor_kinds kc;

```

```

288         kc = is_gnu_v3_mangled_ctor (inp);
289         sprintf (buf, "%d", (int) kc);
290     }
291     else
292     {
293         enum gnu_v3_dtor_kinds kd;

295         kd = is_gnu_v3_mangled_dtor (inp);
296         sprintf (buf, "%d", (int) kd);
297     }

299     if (strcmp (buf, expect.data) != 0)
300     {
301         fail (lineno, format.data, input.data, buf, expect.data);
302         failures++;
303     }

305     continue;
306 }

308 cplus_demangle_set_style (style);

310 result = cplus_demangle (inp,
311     DMGL_PARAMS|DMGL_ANSI|DMGL_TYPES
312     |(ret_postfix ? DMGL_RET_POSTFIX : 0));

314 if (result
315     ? strcmp (result, expect.data)
316     : strcmp (input.data, expect.data))
317 {
318     fail (lineno, format.data, input.data, result, expect.data);
319     failures++;
320 }
321 free (result);

323 if (no_params)
324 {
325     getline (&expect);
325     getline (&expect);
326     result = cplus_demangle (inp, DMGL_ANSI|DMGL_TYPES);

328     if (result
329         ? strcmp (result, expect.data)
330         : strcmp (input.data, expect.data))
331     {
332         fail (lineno, format.data, input.data, result, expect.data);
333         failures++;
334     }
335     free (result);
336 }
337 }

339 free (format.data);
340 free (input.data);
341 free (expect.data);

343 printf ("%s: %d tests, %d failures\n", argv[0], tests, failures);
344 return failures ? 1 : 0;
345 }
unchanged_portion_omitted

```