

The GNU D Compiler

For GCC version 16.0.0 (pre-release)

(GCC)

David Friedman, Iain Buclaw

Published by the Free Software Foundation
51 Franklin Street, Fifth Floor
Boston, MA 02110-1301, USA

Copyright © 2006-2025 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

Table of Contents

1	Invoking gdc	1
1.1	Input and Output files	1
1.2	Runtime Options	1
1.3	Options for Directory Search	5
1.4	Code Generation	6
1.5	Warnings	7
1.6	Options for Linking	9
1.7	Developer Options	10
2	Language Reference	11
2.1	Attributes	11
2.1.1	Attribute Syntax	11
2.1.2	Common Attributes	12
2.1.3	Other Attributes	16
2.1.4	Target-specific Attributes	17
2.2	Built-in Functions	17
2.2.1	Built-in Types	18
2.2.2	Querying Available Built-ins	18
2.2.3	Other Built-in Functions	19
2.3	Importing C Sources into D	20
2.4	Inline Assembly	21
2.5	Intrinsics	22
2.5.1	Bit Operation Intrinsics	22
2.5.2	Integer Overflow Intrinsics	24
2.5.3	Math Intrinsics	25
2.5.4	Variadic Intrinsics	27
2.5.5	Volatile Intrinsics	27
2.5.6	CTFE Intrinsics	28
2.6	Predefined Pragmas	31
2.7	Predefined Versions	32
2.7.1	Standard Predefined Versions	32
2.7.2	Common Predefined Versions	34
2.7.3	Target-specific Predefined Versions	35
2.8	Special Enums	37
2.9	Traits	37
2.10	Vector Extensions	38
2.11	Vector Intrinsics	39
2.12	Missing Features and Deviations	41
	GNU General Public License	44
	GNU Free Documentation License	55
	ADDENDUM: How to use this License for your documents	62

Option Index	63
Keyword Index	65

1 Invoking `gdc`

The `gdc` command is the GNU compiler for the D language and supports many of the same options as `gcc`. See Section “Option Summary” in *Using the GNU Compiler Collection (GCC)*. This manual only documents the options specific to `gdc`.

1.1 Input and Output files

For any given input file, the file name suffix determines what kind of compilation is done. The following kinds of input file names are supported:

`file.d` D source files.
`file.dd` Ddoc source files.
`file.di` D interface files.

You can specify more than one input file on the `gdc` command line, each being compiled separately in the compilation process. If you specify a `-o file` option, all the input files are compiled together, producing a single output file, named `file`. This is allowed even when using `-S` or `-c`.

A D interface file contains only what an import of the module needs, rather than the whole implementation of that module. They can be created by `gdc` from a D source file by using the `-H` option. When the compiler resolves an import declaration, it searches for matching `.di` files first, then for `.d`.

A Ddoc source file contains code in the D macro processor language. It is primarily designed for use in producing user documentation from embedded comments, with a slight affinity towards HTML generation. If a `.d` source file starts with the string `Ddoc` then it is treated as general purpose documentation, not as a D source file.

1.2 Runtime Options

These options affect the runtime behavior of programs compiled with `gdc`.

`-fall-instantiations`

Generate code for all template instantiations. The default template emission strategy is to not generate code for declarations that were either instantiated speculatively, such as from `__traits(compiles, ...)`, or that come from an imported module not being compiled.

`-fno-assert`

Turn off code generation for `assert` contracts.

`-fno-bounds-check`

Turns off array bounds checking for all functions, which can improve performance for code that uses arrays extensively. Note that this can result in unpredictable behavior if the code in question actually does violate array bounds constraints. It is safe to use this option if you are sure that your code never throws a `RangeError`.

- `'c++23'` Sets `__traits(getTargetInfo, "cppStd")` to 202302.
- `-finclude-imports`
Include imported modules in the compilation, as if they were given on the command line. When this option is enabled, all imported modules are compiled except those that are part of libphobos.
- `-fno-invariants`
Turns off code generation for class `invariant` contracts.
- `-fmain` Generates a default `main()` function when compiling. This is useful when unittesting a library, as it enables running the unittests in a library without having to manually define an entry-point function. This option does nothing when `main` is already defined in user code.
- `-fno-moduleinfo`
Turns off generation of the `ModuleInfo` and related functions that would become unreferenced without it, which may allow linking to programs not written in D. Functions that are not be generated include module constructors and destructors (`static this` and `static ~this`), `unittest` code, and DSO registry functions for dynamically linked code.
- `-fonly=filename`
Tells the compiler to parse and run semantic analysis on all modules on the command line, but only generate code for the module specified by *filename*.
- `-fno-postconditions`
Turns off code generation for postcondition `out` contracts.
- `-fno-preconditions`
Turns off code generation for precondition `in` contracts.
- `-fpreview=id`
Turns on an upcoming D language change identified by *id*. The following values are supported:
- `'all'` Turns on all upcoming D language features.
 - `'bitfields'`
Implements bit-fields in D.
 - `'dip1000'` Implements <https://github.com/dlang/DIPs/blob/master/DIPs/other/DIP1000.md> (Scoped pointers).
 - `'dip1008'` Implements <https://github.com/dlang/DIPs/blob/master/DIPs/other/DIP1008.md> (Allow exceptions in `@nogc` code).
 - `'dip1021'` Implements <https://github.com/dlang/DIPs/blob/master/DIPs/accepted/DIP1021.md> (Mutable function arguments).
 - `'dip25'` Implements <https://github.com/dlang/DIPs/blob/master/DIPs/archive/DIP25.md> (Sealed references).
 - `'dtorfields'`
Turns on generation for destructing fields of partially constructed objects.

<code>'fieldwise'</code>	Turns on generation of struct equality to use field-wise comparisons.
<code>'fixaliasthis'</code>	Implements new lookup rules that check the current scope for <code>alias this</code> before searching in upper scopes.
<code>'fiximmutableconv'</code>	Disallows unsound immutable conversions that were formerly incorrectly permitted.
<code>'in'</code>	Implements <code>in</code> parameters to mean <code>scope const [ref]</code> and accepts rvalues.
<code>'inclusiveincontracts'</code>	Implements <code>in</code> contracts of overridden methods to be a superset of parent contract.
<code>'nosharedaccess'</code>	Turns off and disallows all access to shared memory objects.
<code>'rvaluerefparam'</code>	Implements rvalue arguments to <code>ref</code> parameters.
<code>'systemvariables'</code>	Disables access to variables marked <code>@system</code> from <code>@safe</code> code.

-frelease

Turns on compiling in release mode, which means not emitting runtime checks for contracts and asserts. Array bounds checking is not done for `@system` and `@trusted` functions, and assertion failures are undefined behavior.

This is equivalent to compiling with the following options:

```
gdc -fno-assert -fbounds-check=safe -fno-invariants \
    -fno-postconditions -fno-preconditions -fno-switch-errors
```

-frevert=

Turns off a D language feature identified by *id*. The following values are supported:

<code>'all'</code>	Turns off all revertable D language features.
<code>'dip1000'</code>	Reverts https://github.com/dlang/DIPs/blob/master/DIPs/other/DIP1000.md (Scoped pointers).
<code>'dip25'</code>	Reverts https://github.com/dlang/DIPs/blob/master/DIPs/archive/DIP25.md (Sealed references).
<code>'dtorfields'</code>	Turns off generation for destructing fields of partially constructed objects.
<code>'intpromote'</code>	Turns off C-style integral promotion for unary <code>+</code> , <code>-</code> and <code>~</code> expressions.

- fno-rtti**
Turns off generation of run-time type information for all user defined types. Any code that uses features of the language that require access to this information will result in an error.
- fno-switch-errors**
This option controls what code is generated when no case is matched in a **final switch** statement. The default run time behavior is to throw a **SwitchError**. Turning off **-fswitch-errors** means that instead the execution of the program is immediately halted.
- funittest**
Turns on compilation of **unittest** code, and turns on the **version(unittest)** identifier. This implies **-fassert**.
- fversion=value**
Turns on compilation of conditional **version** code into the program identified by any of the following values:
 - 'ident'** Turns on compilation of **version** code identified by *ident*.
- fno-weak-templates**
Turns off emission of declarations that can be defined in multiple objects as weak symbols. The default is to emit all public symbols as weak, unless the target lacks support for weak symbols. Disabling this option means that common symbols are instead put in COMDAT or become private.

1.3 Options for Directory Search

These options specify directories to search for files, libraries, and other parts of the compiler:

- Idir** Specify a directory to use when searching for imported modules at compile time. Multiple **-I** options can be used, and the paths are searched in the same order.
- Jdir** Specify a directory to use when searching for files in string imports at compile time. This switch is required in order to use **import(file)** expressions. Multiple **-J** options can be used, and the paths are searched in the same order.
- Ldir** When linking, specify a library search directory, as with **gcc**.
- Bdir** This option specifies where to find the executables, libraries, source files, and data files of the compiler itself, as with **gcc**.
- fmodule-file=module=spec**
This option manipulates file paths of imported modules, such that if an imported module matches all or the leftmost part of *module*, the file path in *spec* is used as the location to search for D sources. This is used when the source file path and names are not the same as the package and module hierarchy. Consider the following examples:

```
gdc test.d -fmodule-file=A.B=foo.d -fmodule-file=C=bar
```

This will tell the compiler to search in all import paths for the source file *foo.d* when importing *A.B*, and the directory *bar/* when importing *C*, as annotated in the following D code:

```
module test;
```

```
import A.B;      // Matches A.B, searches for foo.d
import C.D.E;    // Matches C, searches for bar/D/E.d
import A.B.C;    // No match, searches for A/B/C.d
```

-imultilib *dir*

Use *dir* as a subdirectory of the gcc directory containing target-specific D sources and interfaces.

-iprefix *prefix*

Specify *prefix* as the prefix for the gcc directory containing target-specific D sources and interfaces. If the *prefix* represents a directory, you should include the final '/'.

-nostdinc

Do not search the standard system directories for D source and interface files. Only the directories that have been specified with **-I** options (and the directory of the current file, if appropriate) are searched.

1.4 Code Generation

In addition to the many gcc options controlling code generation, gdc has several options specific to itself.

-H Generates D interface files for all modules being compiled. The compiler determines the output file based on the name of the input file, removes any directory components and suffix, and applies the **.di** suffix.

-Hd *dir* Same as **-H**, but writes interface files to directory *dir*. This option can be used with **-Hf *file*** to independently set the output file and directory path.

-Hf *file* Same as **-H** but writes interface files to *file*. This option can be used with **-Hd *dir*** to independently set the output file and directory path.

-M Output the module dependencies of all source files being compiled in a format suitable for **make**. The compiler outputs one **make** rule containing the object file name for that source file, a colon, and the names of all imported files.

-MM Like **-M** but does not mention imported modules from the D standard library package directories.

-MF *file* When used with **-M** or **-MM**, specifies a *file* to write the dependencies to. When used with the driver options **-MD** or **-MMD**, **-MF** overrides the default dependency output file.

-MG This option is for compatibility with gcc, and is ignored by the compiler.

-MP Outputs a phony target for each dependency other than the modules being compiled, causing each to depend on nothing.

-MT *target*

Change the *target* of the rule emitted by dependency generation to be exactly the string you specify. If you want multiple targets, you can specify them as a single argument to **-MT**, or use multiple **-MT** options.

- MQ *target***
Same as **-MT**, but it quotes any characters which are special to **make**.
- MD**
This option is equivalent to **-M -MF *file***. The driver determines *file* by removing any directory components and suffix from the input file, and then adding a **.deps** suffix.
- MMD**
Like **-MD** but does not mention imported modules from the D standard library package directories.
- X**
Output information describing the contents of all source files being compiled in JSON format to a file. The driver determines *file* by removing any directory components and suffix from the input file, and then adding a **.json** suffix.
- Xf *file***
Same as **-X**, but writes all JSON contents to the specified *file*.
- fdoc**
Generates Ddoc documentation and writes it to a file. The compiler determines *file* by removing any directory components and suffix from the input file, and then adding a **.html** suffix.
- fdoc-dir=*dir***
Same as **-fdoc**, but writes documentation to directory *dir*. This option can be used with **-fdoc-file=*file*** to independently set the output file and directory path.
- fdoc-file=*file***
Same as **-fdoc**, but writes documentation to *file*. This option can be used with **-fdoc-dir=*dir*** to independently set the output file and directory path.
- fdoc-inc=*file***
Specify *file* as a Ddoc macro file to be read. Multiple **-fdoc-inc** options can be used, and files are read and processed in the same order.
- fdump-c++-spec=*file***
For D source files, generate corresponding C++ declarations in *file*.
- fdump-c++-spec-verbose**
In conjunction with **-fdump-c++-spec=** above, add comments for ignored declarations in the generated C++ header.
- fsave-mixins=*file***
Generates code expanded from D **mixin** statements and writes the processed sources to *file*. This is useful to debug errors in compilation and provides source for debuggers to show when requested.

1.5 Warnings

Warnings are diagnostic messages that report constructions that are not inherently erroneous but that are risky or suggest there is likely to be a bug in the program. Unless **-Werror** is specified, they do not prevent compilation of the program.

- Wall**
Turns on all warnings messages. Warnings are not a defined part of the D language, and all constructs for which this may generate a warning message are valid code.

- Walloca** This option warns on all uses of "alloca" in the source.
- Walloca-larger-than=*n***
Warn on unbounded uses of `alloca`, and on bounded uses of `alloca` whose bound can be larger than *n* bytes. **-Wno-alloc-larger-than** disables **-Walloca-larger-than** warning and is equivalent to **-Walloca-larger-than=SIZE_MAX** or larger.
- Wno-builtin-declaration-mismatch**
Warn if a built-in function is declared with an incompatible signature.
- Wcast-result**
Warn about casts that will produce a null or zero result. Currently this is only done for casting between an imaginary and non-imaginary data type, or casting between a D and C++ class.
- Wno-deprecated**
Do not warn about usage of deprecated features and symbols with **deprecated** attributes.
- Werror** Turns all warnings into errors.
- Wextra** This enables some extra warning flags that are not enabled by **-Wall**.
 - Waddress**
 - Wcast-result**
 - Wmismatched-special-enum**
 - Wunknown-pragmas**
- Wmismatched-special-enum**
Warn when an enum the compiler recognizes as special is declared with a different size to the built-in type it is representing.
- Wspeculative**
List all error messages from speculative compiles, such as `__traits(compiles, ...)`. This option does not report messages as warnings, and these messages therefore never become errors when the **-Werror** option is also used.
- Wunknown-pragmas**
Warn when a `pragma()` is encountered that is not understood by `gdc`. This differs from **-fignore-unknown-pragmas** where a `pragma` that is part of the D language, but not implemented by the compiler, won't get reported.
- Wno-varargs**
Do not warn upon questionable usage of the macros used to handle variable arguments like `va_start`.
- fno-ignore-unknown-pragmas**
Do not recognize unsupported pragmas. Any `pragma()` encountered that is not part of the D language will result in an error. This option is now deprecated and will be removed in a future release.
- fmax-errors=*n***
Limits the maximum number of error messages to *n*, at which point `gdc` bails out rather than attempting to continue processing the source code. If *n* is 0 (the default), there is no limit on the number of error messages produced.

-fsyntax-only

Check the code for syntax errors, but do not actually compile it. This can be used in conjunction with **-fdoc** or **-H** to generate files for each module present on the command-line, but no other output file.

-ftransition=id

Report additional information about D language changes identified by *id*. The following values are supported:

- 'all'** List information on all D language transitions.
- 'complex'** List all usages of complex or imaginary types.
- 'field'** List all non-mutable fields which occupy an object instance.
- 'in'** List all usages of **in** on parameter.
- 'nogc'** List all hidden GC allocations.
- 'templates'** List statistics on template instantiations.
- 'tls'** List all variables going into thread local storage.

1.6 Options for Linking

These options come into play when the compiler links object files into an executable output file. They are meaningless if the compiler is not doing a link step.

-defaultlib=libname

Specify the library to use instead of **libphobos** when linking. Options specifying the linkage of **libphobos**, such as **-static-libphobos** or **-shared-libphobos**, are ignored.

-debuglib=libname

Specify the debug library to use instead of **libphobos** when linking. This option has no effect unless the **-g** option was also given on the command line. Options specifying the linkage of **libphobos**, such as **-static-libphobos** or **-shared-libphobos**, are ignored.

-nophoboslib

Do not use the Phobos or D runtime library when linking. Options specifying the linkage of **libphobos**, such as **-static-libphobos** or **-shared-libphobos**, are ignored. The standard system libraries are used normally, unless **-nostdlib** or **-nodefaultlibs** is used.

-shared-libphobos

On systems that provide **libgphobos** and **libgdruntime** as a shared and a static library, this option forces the use of the shared version. If no shared version was built when the compiler was configured, this option has no effect.

-static-libphobos

On systems that provide **libgphobos** and **libgdruntime** as a shared and a static library, this option forces the use of the static version. If no static version was built when the compiler was configured, this option has no effect.

1.7 Developer Options

This section describes command-line options that are primarily of interest to developers or language tooling.

-fdump-d-original

Output the internal front-end AST after the **semantic3** stage. This option is only useful for debugging the GNU D compiler itself.

-v

Dump information about the compiler language processing stages as the source program is being compiled. This includes listing all modules that are processed through the **parse**, **semantic**, **semantic2**, and **semantic3** stages; all **import** modules and their file paths; and all **function** bodies that are being compiled.

2 Language Reference

The implementation of the D programming language used by the GNU D compiler is shared with parts of the front-end for the Digital Mars D compiler, hosted at <https://github.com/dlang/dmd/>. This common front-end covers lexical analysis, parsing, and semantic analysis of the D programming language defined in the documents at <https://dlang.org/>.

The implementation details described in this manual are GNU D extensions to the D programming language. If you want to write code that checks whether these features are available, you can test for the predefined version `GNU`, or you can check whether a specific feature is compilable using `__traits(compiles)`.

```
version (GNU)
{
    import gcc.builtins;
    return __builtin_atan2(x, y);
}

static if (__traits(compiles, { asm {"";} }))
{
    asm { "magic instruction"; }
}
```

2.1 Attributes

User-Defined Attributes (UDA) are compile-time expressions introduced by the `@` token that can be attached to a declaration. These attributes can then be queried, extracted, and manipulated at compile time.

GNU D provides a number of extra special attributes to control specific compiler behavior that may help the compiler optimize or check code more carefully for correctness. The attributes are defined in the `gcc.attributes` module.

There is some overlap between the purposes of attributes and pragmas. It has been found more convenient to use `@attribute` to achieve a natural attachment of attributes to their corresponding declarations, whereas `pragma` is of use for compatibility with other compilers or constructs that do not naturally form part of the grammar.

2.1.1 Attribute Syntax

`@(gcc.attributes.attribute)` is the generic entrypoint for applying GCC attributes to a function, variable, or type. There is no type checking done, as well as no deprecation path for attributes removed from the compiler. So the recommendation is to use any of the other UDAs available as described in Section 2.1.2 [Common Attributes], page 12, unless it is a target-specific attribute (See Section 2.1.4 [Target Attributes], page 17).

Function attributes introduced by the `@attribute` UDA are used in the declaration of a function, followed by an attribute name string and any arguments separated by commas enclosed in parentheses.

```
import gcc.attributes;
@attribute("regparm", 1) int func(int size);
```

Multiple attributes can be applied to a single declaration either with multiple `@attribute` attributes, or passing all attributes as a comma-separated list enclosed by parentheses.

```
// Both func1 and func2 have the same attributes applied.
```


can be provided, separated by commas to specify multiple options. Each numeric argument specifies an optimization level. Each string argument that begins with the letter `O` refers to an optimization option such as `-O0` or `-Os`. Other options are taken as suffixes to the `-f` prefix jointly forming the name of an optimization option.

Not every optimization option that starts with the `-f` prefix specified by the attribute necessarily has an effect on the function. The `@optimize` attribute should be used for debugging purposes only. It is not suitable in production code.

```
@optimize(2) double fn0(double x);
@optimize("2") double fn1(double x);
@optimize("s") double fn2(double x);
@optimize("Ofast") double fn3(double x);
@optimize("-O2") double fn4(double x);
@optimize("tree-vectorize") double fn5(double x);
@optimize("-ftree-vectorize") double fn6(double x);
@optimize("no-finite-math-only", 3) double fn7(double x);
```

`@(gcc.attributes.register ("registerName"))`

The `@register` attribute specifies that a local or `__gshared` variable is to be given a register storage-class in the C99 sense of the term, and will be placed into a register named *registerName*.

The variable needs to be boiled down to a data type that fits the target register. It also cannot have either thread-local or `extern` storage. It is an error to take the address of a register variable.

```
@register("ebx") __gshared int ebx = void;
void func() { @register("r10") long r10 = 0x2a; }
```

`@(gcc.attributes.restrict)`

The `@restrict` attribute specifies that a function parameter is to be restrict-qualified in the C99 sense of the term. The parameter needs to boil down to either a pointer or reference type, such as a D pointer, class reference, or a `ref` parameter.

```
void func(@restrict ref const float[16] array);
```

`@(gcc.attributes.section ("sectionName"))`

The `@section` attribute specifies that a function or variable lives in a particular section. For when you need certain particular functions to appear in special sections.

Some file formats do not support arbitrary sections so the section attribute is not available on all platforms. If you need to map the entire contents of a module to a particular section, consider using the facilities of the linker instead.

```
@section("bar") extern void func();
@section("stack") ubyte[10000] stack;
```

`@(gcc.attributes.simd)`

The `@simd` attribute enables creation of one or more function versions that can process multiple arguments using SIMD instructions from a single invocation. Specifying this attribute allows compiler to assume that such versions are available at link time (provided in the same or another module). Generated versions are target-dependent and described in the corresponding Vector ABI document.


```
@(gcc.attributes.noSanitize ("sanitize_option"))
```

This attribute is a synonym for `@no_sanitize("sanitize_option")`.

```
@(gcc.attributes.optStrategy ("strategy"))
```

This attribute is a synonym for `@optimize("O0")` and `@optimize("Os")`. Sets the optimization strategy for a function. Valid strategies are "none", "optsize", "minsize". The strategies are mutually exclusive.

```
@(gcc.attributes.polly)
```

This attribute is a synonym for `@optimize("loop-parallelize-all", "loop-nest-optimize")`. Only effective when GDC was built with ISL included.

2.1.4 Target-specific Attributes

Many targets have their own target-specific attributes. These are also exposed via the `gcc.attributes` module with use of the generic `@(gcc.attributes.attribute)` UDA function.

See Section 2.1.1 [Attribute Syntax], page 11, for details of the exact syntax for using attributes.

See the function and variable attribute documentation in the GCC manual for more information about what attributes are available on each target.

Examples of using x86-specific target attributes are shown as follows:

```
import gcc.attributes;

@attribute("cdecl")
@attribute("fastcall")
@attribute("ms_abi")
@attribute("sysv_abi")
@attribute("callee_pop_aggregate_return", 1)
@attribute("ms_hook_prologue")
@attribute("naked")
@attribute("regparm", 2)
@attribute("sseregparm")
@attribute("force_align_arg_pointer")
@attribute("stdcall")
@attribute("no_caller_saved_registers")
@attribute("interrupt")
@attribute("indirect_branch", "thunk")
@attribute("function_return", "keep"))
@attribute("nof_check")
@attribute("cf_check")
@attribute("indirect_return")
@attribute("fentry_name", "nop")
@attribute("fentry_section", "__entry_loc")
@attribute("nodirect_extern_access")
```

2.2 Built-in Functions

GCC provides a large number of built-in functions that are made available in GNU D by importing the `gcc.builtins` module. Declarations in this module are automatically created by the compiler. All declarations start with `__builtin_`. Refer to the built-in function documentation in the GCC manual for a full list of functions that are available.

2.2.1 Built-in Types

In addition to built-in functions, the following types are defined in the `gcc.builtins` module.

```

__builtin_clong
    The D equivalent of the target's C long type.

__builtin_clonglong
    The D equivalent of the target's C long long type.

__builtin_culong
    The D equivalent of the target's C unsigned long type.

__builtin_culonglong
    The D equivalent of the target's C unsigned long long type.

__builtin_machine_byte
    Signed unit-sized integer type.

__builtin_machine_int
    Signed word-sized integer type.

__builtin_machine_ubyte
    Unsigned unit-sized integer type.

__builtin_machine_uint
    Unsigned word-sized integer type.

__builtin_pointer_int
    Signed pointer-sized integer type.

__builtin_pointer_uint
    Unsigned pointer-sized integer type.

__builtin_unwind_int
    The D equivalent of the target's C _Unwind_Sword type.

__builtin_unwind_uint
    The D equivalent of the target's C _Unwind_Word type.

__builtin_va_list
    The target's va_list type.

```

2.2.2 Querying Available Built-ins

Not all of the functions are supported, and some target-specific functions may only be available when compiling for a particular ISA. One way of finding out what is exposed by the built-ins module is by generating a D interface file. Assuming you have no file `builtins.d`, the command

```
echo "module gcc.builtins;" > builtins.d; gdc -H -fsyntax-only builtins.d
```

will save all built-in declarations to the file `builtins.di`.

Another way to determine whether a specific built-in is available is by using compile-time reflection.

```
enum X86_HAVE_SSE3 = __traits(compiles, __builtin_ia32_haddps);
```


ImportC does not have a preprocessor. It is designed to compile C files after they have been first run through the C preprocessor. If the C file has a `.i` extension, the file is presumed to be already preprocessed. Preprocessing can be run manually:

```
gcc -E file.c > file.i
```

ImportC collects all the `#define` macros from the preprocessor run when it is run automatically. The macros that look like manifest constants, such as:

```
#define COLOR 0x123456
```

are interpreted as D manifest constant declarations of the form:

```
enum COLOR = 0x123456;
```

The variety of macros that can be interpreted as D declarations may be expanded, but will never encompass all the metaprogramming uses of C macros.

GNU D does not directly compile C files into modules that can be linked in with D code to form an executable. When given a source file with the suffix `.c`, the compiler driver program `gdc` instead runs the subprogram `cc1`.

```
gdc file1.d file2.c // d2i file1.d -o file1.s
                  // cc1 file2.c -o file2.s
                  // as file1.s -o file1.o
                  // as file2.s -o file2.o
                  // ld file1.o file2.o
```

2.4 Inline Assembly

The `asm` keyword allows you to embed assembler instructions within D code. GNU D provides two forms of inline `asm` statements. A *basic* `asm` statement is one with no operands, while an *extended* `asm` statement includes one or more operands.

```
asm FunctionAttributes {
    AssemblerInstruction ;
}

asm FunctionAttributes {
    AssemblerTemplate
    : OutputOperands
    [ : InputOperands
    [ : Clobbers
    [ : GotoLabels ] ] ] ;
}
```

The extended form is preferred for mixing D and assembly language within a function, but to include assembly language in a function declared with the `naked` attribute you must use basic `asm`.

```
uint incr (uint value)
{
    uint result;
    asm { "incl %0"
        : "=a" (result)
        : "a" (value);
    }
    return result;
}
```

Multiple assembler instructions can appear within an `asm` block, or the instruction template can be a multi-line or concatenated string. In both cases, GCC's optimizers won't discard or move any instruction within the statement block.

```
bool hasCUID()
{
    uint flags = 0;
    asm nothrow @nogc {
        "pushfl";
        "pushfl";
        "xorl %0, (%esp)" :: "i" (0x00200000);
        "popfl";
        "pushfl";
        "popl %0" : "=a" (flags);
        "xorl (%esp), %0" : "=a" (flags);
        "popfl";
    }
    return (flags & 0x0020_0000) != 0;
}
```

The instruction templates for both basic and extended `asm` can be any expression that can be evaluated at compile-time to a string, not just string literals.

```
uint invert(uint v)
{
    uint result;
    asm @safe @nogc nothrow pure {
        genAsmInsn(`invert`)
        : [res] `=r` (result)
        : [arg1] `r` (v);
    }
    return result;
}
```

The total number of input + output + goto operands is limited to 30.

2.5 Intrinsics

The D language specification itself does not define any intrinsics that a compatible compiler must implement. Rather, within the D core library there are a number of modules that define primitives with generic implementations. While the generic versions of these functions are computationally expensive relative to the cost of the operation itself, compiler implementations are free to recognize them and generate equivalent and faster code.

The following are the kinds of intrinsics recognized by GNU D.

2.5.1 Bit Operation Intrinsics

The following functions are a collection of intrinsics that do bit-level operations, available by importing the `core.bitop` module.

Although most are named after x86 hardware instructions, it is not guaranteed that they will result in generating equivalent assembly on x86. If the compiler determines there is a better way to get the same result in hardware, then that will be used instead.

`float std.math.rounding.floor (float x)` [Function]
`double std.math.rounding.floor (double x)` [Function]
`real std.math.rounding.floor (real x)` [Function]
 Returns the value of *x* rounded downward to the next integer (toward negative infinity).
 This function is evaluated during CTFE as the GCC built-in function `__builtin_floor`.

`real std.math.rounding.round (real x)` [Function]
 Return the value of *x* rounded to the nearest integer. If the fractional part of *x* is exactly 0.5, the return value is rounded away from zero.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_round`.

`real std.math.rounding.trunc (real x)` [Function]
 Returns the integer portion of *x*, dropping the fractional portion.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_trunc`.

`R std.math.traits.copysign (R, X)(R to, X from)` [Template]
 Returns a value composed of *to* with *from*'s sign bit.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_copysign`.

`bool std.math.traits.isFinite (X)(X x)` [Template]
 Returns true if *x* is finite.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_isfinite`.

`bool std.math.traits.isInfinity (X)(X x)` [Template]
 Returns true if *x* is infinite.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_isinf`.

`bool std.math.traits.isNaN (X)(X x)` [Template]
 Returns true if *x* is NaN.
 This function is evaluated during CTFE as the GCC built-in function `__builtin_isnan`.

`float std.math.trigoometry.tan (float x)` [Function]
`double std.math.trigoometry.tan (double x)` [Function]
`real std.math.trigonometry.tan (real x)` [Function]
 Returns tangent of *x*, where *x* is in radians.
 This intrinsic is the same as the GCC built-in function `__builtin_tan`.


```
extern(C) void body_func();

#pragma(mangle, "function")
extern(C++) struct _function {}
```

`pragma(msg)`

`pragma(msg, "message")` causes the compiler to print an informational message with the text ‘`message`’. The pragma accepts multiple arguments, each to which is evaluated at compile time and then all are combined into one concatenated message.

```
pragma(msg, "compiling...", 6, 1.0); // prints "compiling...61.0"
```

`pragma(sprintf)`

`pragma(scanf)`

`pragma(sprintf)` and `pragma(scanf)` specifies that a function declaration with `printf` or `scanf` style arguments that should be type-checked against a format string.

A `printf`-like or `scanf`-like function can either be an `extern(C)` or `extern(C++)` function with a *format* parameter accepting a pointer to a 0-terminated `char` string, immediately followed by either a ... variadic argument list or a parameter of type `va_list` as the last parameter.

```
extern(C):
#pragma(sprintf)
int printf(scope const char* format, scope const ...);

#pragma(scanf)
int vsprintf(scope const char* format, va_list arg);
```

`pragma(startaddress)`

This pragma is accepted, but has no effect.

```
void foo() { }
#pragma(startaddress, foo);
```

2.7 Predefined Versions

Several conditional version identifiers are predefined; you use them without supplying their definitions. They fall into three classes: standard, common, and target-specific.

Predefined version identifiers from this list cannot be set from the command line or from version statements. This prevents things like both `Windows` and `linux` being simultaneously set.

2.7.1 Standard Predefined Versions

The standard predefined versions are documented by the D language specification hosted at <https://dlang.org/spec/version.html#predefined-versions>.

`all`

`none` Version `none` is never defined; used to just disable a section of code. Version `all` is always defined; used as the opposite of `none`.

MIPS32
MIPS64
MIPS_EABI
MIPS_HardFloat
MIPS_N32
MIPS_N64
MIPS_O32
MIPS_O64
MIPS_SoftFloat
 Versions relating to the MIPS family of processors.

NetBSD Version relating to NetBSD systems.

OpenBSD Version relating to OpenBSD systems.

OSX Version relating to OSX systems.

Posix Version relating to POSIX systems (includes Linux, FreeBSD, OSX, Solaris, etc).

PPC
PPC64
PPC_HardFloat
PPC_SoftFloat
 Versions relating to the PowerPC family of processors.

RISCV32
RISCV64 Versions relating to the RISC-V family of processors.

S390
SystemZ Versions relating to the S/390 and System Z family of processors.

S390X Deprecated; use `SystemZ` instead.

Solaris Versions relating to Solaris systems.

SPARC
SPARC64
SPARC_HardFloat
SPARC_SoftFloat
SPARC_V8Plus
 Versions relating to the SPARC family of processors.

Thumb Deprecated; use `ARM_Thumb` instead.

D_X32
X86
X86_64 Versions relating to the x86-32 and x86-64 family of processors.

Windows
Win32
Win64 Versions relating to Microsoft Windows systems.

2.8 Special Enums

Special `enum` names are used to represent types that do not have an equivalent basic D type. For example, C++ types used by the C++ name mangler.

Special enums are declared opaque, with a base type explicitly set. Unlike regular opaque enums, special enums can be used as any other value type. They have a default `.init` value, as well as other enum properties available (`.min`, `.max`). Special enums can be declared in any module, and will be recognized by the compiler.

```
import gcc.builtins;
enum __c_long : __builtin_clong;
__c_long var = 0x800A;
```

The following identifiers are recognized by GNU D.

```
__c_complex_double
    C _Complex double type.

__c_complex_float
    C _Complex float type.

__c_complex_real
    C _Complex long double type.

__c_long    C++ long type.

__c_longlong
    C++ long long type.

__c_long_double
    C long double type.

__c_ulong
    C++ unsigned long type.

__c_ulonglong
    C++ unsigned long long type.

__c_wchar_t
    C++ wchar_t type.
```

The `core.stdc.config` module declares the following shorthand alias types for convenience: `c_complex_double`, `c_complex_float`, `c_complex_real`, `cpp_long`, `cpp_longlong`, `c_long_double`, `cpp_ulong`, `cpp_ulonglong`.

It may cause undefined behavior at runtime if a special enum is declared with a base type that has a different size to the target C/C++ type it is representing. The GNU D compiler will catch such declarations and emit a warning when the `-Wmismatched-special-enum` option is seen on the command-line.

2.9 Traits

Traits are extensions to the D programming language to enable programs, at compile time, to get at information internal to the compiler. This is also known as compile time reflection.

GNU D implements a `__traits(getTargetInfo)` trait that receives a string key as its argument. The result is an expression describing the requested target information.

```
version (OSX)
```


On x86 targets, all intrinsics are available as functions in the `gcc.builtins` module, and have predictable equivalents.

```
version (DigitalMars)
{
    __simd(XMM.PSLLW, op1, op2);
    __simd_ib(XMM.PSLLW, op1, imm8);
}
version (GNU)
{
    __builtin_ia32_psllw(op1, op2);
    __builtin_ia32_psllwi(op1, imm8);
}
```

TypeInfo-based `va_arg`

The Digital Mars D compiler implements a version of `core.vararg.va_arg` that accepts a run-time `TypeInfo` argument for use when the static type is not known. This function is not implemented by GNU D. It is more portable to use variadic template functions instead.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a. The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source.

The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient’s use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others’ Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

- be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
 - C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
 - D. Preserve all the copyright notices of the Document.
 - E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
 - F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
 - G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
 - H. Include an unaltered copy of this License.
 - I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
 - J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
 - K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
 - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
 - M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
 - N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
 - O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts. A copy of the license is included in the section entitled ``GNU
Free Documentation License''.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

