



GNAT User's Guide for Native Platforms

Release 13.0.0

AdaCore

Nov 10, 2022

This page is intentionally left blank.

GNAT, The GNU Ada Development Environment

GPL

Version 13.0.0

Date: Nov 10, 2022

AdaCore

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, with the Front-Cover Texts being “GNAT Reference Manual”, and with no Back-Cover Texts. A copy of the license is included in the section entitled GNU Free Documentation License.

This page is intentionally left blank.

CONTENTS

1	About This Guide	13
1.1	What This Guide Contains	13
1.2	What You Should Know before Reading This Guide	14
1.3	Related Information	14
1.4	Conventions	14
2	Getting Started with GNAT	17
2.1	System Requirements	17
2.2	Running GNAT	18
2.3	Running a Simple Ada Program	18
2.4	Running a Program with Multiple Units	19
3	The GNAT Compilation Model	21
3.1	Source Representation	21
3.2	Foreign Language Representation	22
3.2.1	Latin-1	22
3.2.2	Other 8-Bit Codes	23
3.2.3	Wide_Character Encodings	24
3.2.4	Wide_Wide_Character Encodings	25
3.3	File Naming Topics and Utilities	26
3.3.1	File Naming Rules	26
3.3.2	Using Other File Names	27
3.3.3	Alternative File Naming Schemes	27
3.3.4	Handling Arbitrary File Naming Conventions with <code>gname</code>	29
	Arbitrary File Naming Conventions	29
	Running <code>gname</code>	30
	Switches for <code>gname</code>	30
	Examples of <code>gname</code> Usage	32
3.3.5	File Name Krunching with <code>gnatkr</code>	32
	About <code>gnatkr</code>	33
	Using <code>gnatkr</code>	33
	Krunching Method	34
	Examples of <code>gnatkr</code> Usage	35
3.3.6	Renaming Files with <code>gnatchop</code>	35
	Handling Files with Multiple Units	35

	Operating gnatcatch in Compilation Mode	36
	Command Line for gnatcatch	37
	Switches for gnatcatch	38
	Examples of gnatcatch Usage	39
3.4	Configuration Pragmas	40
3.4.1	Handling of Configuration Pragmas	41
3.4.2	The Configuration Pragmas Files	42
3.5	Generating Object Files	42
3.6	Source Dependencies	43
3.7	The Ada Library Information Files	44
3.8	Binding an Ada Program	45
3.9	GNAT and Libraries	46
3.9.1	Introduction to Libraries in GNAT	46
3.9.2	General Ada Libraries	47
	Building a library	47
	Installing a library	48
	Using a library	49
3.9.3	Stand-alone Ada Libraries	50
	Introduction to Stand-alone Libraries	50
	Building a Stand-alone Library	51
	Creating a Stand-alone Library to be used in a non-Ada context	52
	Restrictions in Stand-alone Libraries	54
3.9.4	Rebuilding the GNAT Run-Time Library	54
3.10	Conditional Compilation	55
3.10.1	Modeling Conditional Compilation in Ada	55
	Use of Boolean Constants	55
	Debugging - A Special Case	56
	Conditionalizing Declarations	57
	Use of Alternative Implementations	58
	Preprocessing	60
3.10.2	Preprocessing with gnatprep	60
	Preprocessing Symbols	61
	Using gnatprep	61
	Switches for gnatprep	61
	Form of Definitions File	62
	Form of Input Text for gnatprep	63
3.10.3	Integrated Preprocessing	65
3.11	Mixed Language Programming	68
3.11.1	Interfacing to C	68
3.11.2	Calling Conventions	71
3.11.3	Building Mixed Ada and C++ Programs	74
	Interfacing to C++	74
	Linking a Mixed C++ & Ada Program	74
	A Simple Example	76
	Interfacing with C++ constructors	78
	Interfacing with C++ at the Class Level	81
3.11.4	Generating Ada Bindings for C and C++ headers	85
	Running the Binding Generator	85

Generating Bindings for C++ Headers	86
Switches	88
3.11.5 Generating C Headers for Ada Specifications	89
Running the C Header Generator	89
3.12 GNAT and Other Compilation Models	90
3.12.1 Comparison between GNAT and C/C++ Compilation Models	90
3.12.2 Comparison between GNAT and Conventional Ada Library Models	91
3.13 Using GNAT Files with External Tools	91
3.13.1 Using Other Utility Programs with GNAT	92
3.13.2 The External Symbol Naming Scheme of GNAT	92
4 Building Executable Programs with GNAT	93
4.1 Building with <code>gnatmake</code>	93
4.1.1 Running <code>gnatmake</code>	94
4.1.2 Switches for <code>gnatmake</code>	94
4.1.3 Mode Switches for <code>gnatmake</code>	101
4.1.4 Notes on the Command Line	101
4.1.5 How <code>gnatmake</code> Works	102
4.1.6 Examples of <code>gnatmake</code> Usage	103
4.2 Compiling with <code>gcc</code>	103
4.2.1 Compiling Programs	103
4.2.2 Search Paths and the Run-Time Library (RTL)	105
4.2.3 Order of Compilation Issues	106
4.2.4 Examples	106
4.3 Compiler Switches	106
4.3.1 Alphabetical List of All Switches	107
4.3.2 Output and Error Message Control	120
4.3.3 Warning Message Control	123
4.3.4 Debugging and Assertion Control	144
4.3.5 Validity Checking	145
4.3.6 Style Checking	148
4.3.7 Run-Time Checks	155
4.3.8 Using <code>gcc</code> for Syntax Checking	157
4.3.9 Using <code>gcc</code> for Semantic Checking	158
4.3.10 Compiling Different Versions of Ada	159
4.3.11 Character Set Control	160
4.3.12 File Naming Control	161
4.3.13 Subprogram Inlining Control	162
4.3.14 Auxiliary Output Control	162
4.3.15 Debugging Control	163
4.3.16 Exception Handling Control	167
4.3.17 Units to Sources Mapping Files	167
4.3.18 Code Generation Control	168
4.4 Linker Switches	169
4.5 Binding with <code>gnatbind</code>	169
4.5.1 Running <code>gnatbind</code>	169
4.5.2 Switches for <code>gnatbind</code>	170
Consistency-Checking Modes	176

	Binder Error Message Control	176
	Elaboration Control	177
	Output Control	178
	Dynamic Allocation Control	179
	Binding with Non-Ada Main Programs	179
	Binding Programs with No Main Subprogram	180
4.5.3	Command-Line Access	181
4.5.4	Search Paths for <code>gnatbind</code>	181
4.5.5	Examples of <code>gnatbind</code> Usage	182
4.6	Linking with <code>gnatlink</code>	182
4.6.1	Running <code>gnatlink</code>	182
4.6.2	Switches for <code>gnatlink</code>	183
4.7	Using the GNU <code>make</code> Utility	185
4.7.1	Using <code>gnatmake</code> in a Makefile	185
4.7.2	Automatically Creating a List of Directories	187
4.7.3	Generating the Command Line Switches	188
4.7.4	Overcoming Command Line Length Limits	188
5	GNAT Utility Programs	191
5.1	The File Cleanup Utility <code>gnatclean</code>	191
5.1.1	Running <code>gnatclean</code>	191
5.1.2	Switches for <code>gnatclean</code>	192
5.2	The GNAT Library Browser <code>gnatls</code>	193
5.2.1	Running <code>gnatls</code>	193
5.2.2	Switches for <code>gnatls</code>	194
5.2.3	Example of <code>gnatls</code> Usage	195
6	GNAT and Program Execution	197
6.1	Running and Debugging Ada Programs	197
6.1.1	The GNAT Debugger GDB	198
6.1.2	Running GDB	198
6.1.3	Introduction to GDB Commands	199
6.1.4	Using Ada Expressions	201
6.1.5	Calling User-Defined Subprograms	201
6.1.6	Using the <i>next</i> Command in a Function	202
6.1.7	Stopping When Ada Exceptions Are Raised	202
6.1.8	Ada Tasks	203
6.1.9	Debugging Generic Units	203
6.1.10	Remote Debugging with <code>gdbserver</code>	204
6.1.11	GNAT Abnormal Termination or Failure to Terminate	205
6.1.12	Naming Conventions for GNAT Source Files	206
6.1.13	Getting Internal Debugging Information	207
6.1.14	Stack Traceback	207
	Non-Symbolic Traceback	207
	Symbolic Traceback	212
6.1.15	Pretty-Printers for the GNAT runtime	214
6.2	Profiling	215
6.2.1	Profiling an Ada Program with <code>gprof</code>	215

	Compilation for profiling	216
	Program execution	216
	Running gprof	216
	Interpretation of profiling results	217
6.3	Improving Performance	218
6.3.1	Performance Considerations	218
	Controlling Run-Time Checks	218
	Use of Restrictions	219
	Optimization Levels	219
	Debugging Optimized Code	220
	Inlining of Subprograms	222
	Floating Point Operations	223
	Vectorization of loops	224
	Other Optimization Switches	226
	Optimization and Strict Aliasing	226
	Aliased Variables and Optimization	229
	Atomic Variables and Optimization	230
	Passive Task Optimization	231
6.3.2	Text_IO Suggestions	232
6.3.3	Reducing Size of Executables with Unused Subprogram/Data Elimination	232
	About unused subprogram/data elimination	232
	Compilation options	232
	Example of unused subprogram/data elimination	233
6.4	Overflow Check Handling in GNAT	234
6.4.1	Background	234
6.4.2	Management of Overflows in GNAT	236
6.4.3	Specifying the Desired Mode	237
6.4.4	Default Settings	238
6.4.5	Implementation Notes	239
6.5	Performing Dimensionality Analysis in GNAT	239
6.6	Stack Related Facilities	244
6.6.1	Stack Overflow Checking	244
6.6.2	Static Stack Usage Analysis	244
6.6.3	Dynamic Stack Usage Analysis	245
6.7	Memory Management Issues	246
6.7.1	Some Useful Memory Pools	246
6.7.2	The GNAT Debug Pool Facility	247
A	Platform-Specific Information	251
A.1	Run-Time Libraries	251
A.1.1	Summary of Run-Time Configurations	251
A.2	Specifying a Run-Time Library	252
A.2.1	Choosing the Scheduling Policy	253
A.3	GNU/Linux Topics	254
A.3.1	Required Packages on GNU/Linux	254
A.3.2	A GNU/Linux Debug Quirk	254
A.4	Microsoft Windows Topics	254
A.4.1	Using GNAT on Windows	255

A.4.2	Using a network installation of GNAT	256
A.4.3	CONSOLE and WINDOWS subsystems	256
A.4.4	Temporary Files	256
A.4.5	Disabling Command Line Argument Expansion	256
A.4.6	Windows Socket Timeouts	257
A.4.7	Mixed-Language Programming on Windows	258
	Windows Calling Conventions	259
	Introduction to Dynamic Link Libraries (DLLs)	261
	Using DLLs with GNAT	262
	Building DLLs with GNAT Project files	266
	Building DLLs with GNAT	266
	Building DLLs with gnatdll	267
	Ada DLLs and Finalization	270
	Creating a Spec for Ada DLLs	270
	GNAT and Windows Resources	274
	Using GNAT DLLs from Microsoft Visual Studio Applications	276
	Debugging a DLL	277
	Setting Stack Size from gnatlink	280
	Setting Heap Size from gnatlink	280
A.4.8	Windows Specific Add-Ons	281
	Win32Ada	281
	wPOSIX	281
A.5	Mac OS Topics	282
A.5.1	Codesigning the Debugger	282
B	Example of Binder Output File	283
C	Elaboration Order Handling in GNAT	299
C.1	Elaboration Code	299
C.2	Elaboration Order	301
C.3	Checking the Elaboration Order	303
C.4	Controlling the Elaboration Order in Ada	303
C.5	Controlling the Elaboration Order in GNAT	308
C.6	Mixing Elaboration Models	309
C.7	ABE Diagnostics	310
C.8	SPARK Diagnostics	311
C.9	Elaboration Circularities	312
C.10	Resolving Elaboration Circularities	313
C.11	Elaboration-related Compiler Switches	316
C.12	Summary of Procedures for Elaboration Control	318
C.13	Inspecting the Chosen Elaboration Order	319
D	Inline Assembler	323
D.1	Basic Assembler Syntax	323
D.2	A Simple Example of Inline Assembler	325
D.3	Output Variables in Inline Assembler	326
D.4	Input Variables in Inline Assembler	330
D.5	Inlining Inline Assembler Code	331

D.6	Other Asm Functionality	332
D.6.1	The Clobber Parameter	332
D.6.2	The Volatile Parameter	333
E	GNU Free Documentation License	335
E.1	Preamble	335
E.2	1. APPLICABILITY AND DEFINITIONS	335
E.3	2. VERBATIM COPYING	337
E.4	3. COPYING IN QUANTITY	337
E.5	4. MODIFICATIONS	338
E.6	5. COMBINING DOCUMENTS	339
E.7	6. COLLECTIONS OF DOCUMENTS	340
E.8	7. AGGREGATION WITH INDEPENDENT WORKS	340
E.9	8. TRANSLATION	340
E.10	9. TERMINATION	341
E.11	10. FUTURE REVISIONS OF THIS LICENSE	341
E.12	11. RELICENSING	341
E.13	ADDENDUM: How to use this License for your documents	342
	Index	343

This page is intentionally left blank.

ABOUT THIS GUIDE

This guide describes the use of GNAT, a compiler and software development toolset for the full Ada programming language. It documents the features of the compiler and tools, and explains how to use them to build Ada applications.

GNAT implements Ada 95, Ada 2005, Ada 2012, and Ada 202x, and it may also be invoked in Ada 83 compatibility mode. By default, GNAT assumes Ada 2012, but you can override with a compiler switch ([Compiling Different Versions of Ada](#)) to explicitly specify the language version. Throughout this manual, references to ‘Ada’ without a year suffix apply to all Ada versions of the language, starting with Ada 95.

1.1 What This Guide Contains

This guide contains the following chapters:

- [Getting Started with GNAT](#) describes how to get started compiling and running Ada programs with the GNAT Ada programming environment.
- [The GNAT Compilation Model](#) describes the compilation model used by GNAT.
- [Building Executable Programs with GNAT](#) describes how to use the main GNAT tools to build executable programs, and it also gives examples of using the GNU make utility with GNAT.
- [GNAT Utility Programs](#) explains the various utility programs that are included in the GNAT environment
- [GNAT and Program Execution](#) covers a number of topics related to running, debugging, and tuning the performance of programs developed with GNAT

Appendices cover several additional topics:

- [Platform-Specific Information](#) describes the different run-time library implementations and also presents information on how to use GNAT on several specific platforms
- [Example of Binder Output File](#) shows the source code for the binder output file for a sample program.
- [Elaboration Order Handling in GNAT](#) describes how GNAT helps you deal with elaboration order issues.

- [Inline Assembler](#) shows how to use the inline assembly facility in an Ada program.

1.2 What You Should Know before Reading This Guide

This guide assumes a basic familiarity with the Ada 95 language, as described in the International Standard ANSI/ISO/IEC-8652:1995, January 1995. Reference manuals for Ada 95, Ada 2005, and Ada 2012 are included in the GNAT documentation package.

1.3 Related Information

For further information about Ada and related tools, please refer to the following documents:

- *Ada 95 Reference Manual*, *Ada 2005 Reference Manual*, and *Ada 2012 Reference Manual*, which contain reference material for the several revisions of the Ada language standard.
- *GNAT Reference Manual*, which contains all reference material for the GNAT implementation of Ada.
- *Using GNAT Studio*, which describes the GNAT Studio Integrated Development Environment.
- *GNAT Studio Tutorial*, which introduces the main GNAT Studio features through examples.
- *Debugging with GDB*, for all details on the use of the GNU source-level debugger.
- *GNU Emacs Manual*, for full information on the extensible editor and programming environment Emacs.

1.4 Conventions

Following are examples of the typographical and graphic conventions used in this guide:

- Functions, utility program names, standard names, and classes.
- Option flags
- File names
- Variables
- *Emphasis*
- [optional information or parameters]
- Examples are described by text

and then shown this way.

- Commands that are entered by the user are shown as preceded by a prompt string comprising the \$ character followed by a space.

- Full file names are shown with the '/' character as the directory separator; e.g., `parent-dir/subdir/myfile.adb`. If you are using GNAT on a Windows platform, please note that the '\ ' character should be used instead.

This page is intentionally left blank.

GETTING STARTED WITH GNAT

This chapter describes how to use GNAT's command line interface to build executable Ada programs. On most platforms a visually oriented Integrated Development Environment is also available: GNAT Studio. GNAT Studio offers a graphical “look and feel”, support for development in other programming languages, comprehensive browsing features, and many other capabilities. For information on GNAT Studio please refer to the *GNAT Studio documentation*.

2.1 System Requirements

Even though any machine can run the GNAT toolset and GNAT Studio IDE, in order to get the best experience, we recommend using a machine with as many cores as possible since all individual compilations can run in parallel. A comfortable setup for a compiler server is a machine with 24 physical cores or more, with at least 48 GB of memory (2 GB per core).

For a desktop machine, a minimum of 4 cores is recommended (8 preferred), with at least 2GB per core (so 8 to 16GB).

In addition, for running and navigating sources in GNAT Studio smoothly, we recommend at least 1.5 GB plus 3 GB of RAM per 1 million source line of code. In other words, we recommend at least 3 GB for 500K lines of code and 7.5 GB for 2 million lines of code.

Note that using local and fast drives will also make a difference in terms of build and link time. Network drives such as NFS, SMB, or worse, configuration management filesystems (such as ClearCase dynamic views) should be avoided as much as possible and will produce very degraded performance (typically 2 to 3 times slower than on local fast drives). If such slow drives cannot be avoided for accessing the source code, then you should at least configure your project file so that the result of the compilation is stored on a drive local to the machine performing the run. This can be achieved by setting the `Object_Dir` project file attribute.

2.2 Running GNAT

Three steps are needed to create an executable file from an Ada source file:

- The source file(s) must be compiled.
- The file(s) must be bound using the GNAT binder.
- All appropriate object files must be linked to produce an executable.

All three steps are most commonly handled by using the `gnatmake` utility program that, given the name of the main program, automatically performs the necessary compilation, binding and linking steps.

2.3 Running a Simple Ada Program

Any text editor may be used to prepare an Ada program. (If Emacs is used, the optional Ada mode may be helpful in laying out the program.) The program text is a normal text file. We will assume in our initial example that you have used your editor to prepare the following standard format text file:

```
with Ada.Text_IO; use Ada.Text_IO;
procedure Hello is
begin
  Put_Line ("Hello WORLD!");
end Hello;
```

This file should be named `hello.adb`. With the normal default file naming conventions, GNAT requires that each file contain a single compilation unit whose file name is the unit name, with periods replaced by hyphens; the extension is `ads` for a spec and `adb` for a body. You can override this default file naming convention by use of the special pragma `Source_File_Name` (for further information please see [Using Other File Names](#)). Alternatively, if you want to rename your files according to this default convention, which is probably more convenient if you will be using GNAT for all your compilations, then the `gnatchop` utility can be used to generate correctly-named source files (see [Renaming Files with gnatchop](#)).

You can compile the program using the following command (`$` is used as the command prompt in the examples in this document):

```
$ gcc -c hello.adb
```

`gcc` is the command used to run the compiler. This compiler is capable of compiling programs in several languages, including Ada and C. It assumes that you have given it an Ada program if the file extension is either `.ads` or `.adb`, and it will then call the GNAT compiler to compile the specified file.

The `-c` switch is required. It tells `gcc` to only do a compilation. (For C programs, `gcc` can also do linking, but this capability is not used directly for Ada programs, so the `-c` switch must always be present.)

This compile command generates a file `hello.o`, which is the object file corresponding to your Ada program. It also generates an 'Ada Library Information' file `hello.ali`, which contains additional information used to check that an Ada program is consistent.

To build an executable file, use either `gnatmake` or `gprbuild` with the name of the main file: these tools are builders that will take care of all the necessary build steps in the correct order. In particular, these builders automatically recompile any sources that have been modified since they were last compiled, or sources that depend on such modified sources, so that 'version skew' is avoided.

```
$ gnatmake hello.adb
```

The result is an executable program called `hello`, which can be run by entering:

```
$ hello
```

assuming that the current directory is on the search path for executable programs.

and, if all has gone well, you will see:

```
Hello WORLD!
```

appear in response to this command.

2.4 Running a Program with Multiple Units

Consider a slightly more complicated example that has three files: a main program, and the spec and body of a package:

```
package Greetings is
  procedure Hello;
  procedure Goodbye;
end Greetings;

with Ada.Text_IO; use Ada.Text_IO;
package body Greetings is
  procedure Hello is
  begin
    Put_Line ("Hello WORLD!");
  end Hello;

  procedure Goodbye is
  begin
    Put_Line ("Goodbye WORLD!");
  end Goodbye;
end Greetings;

with Greetings;
procedure Gmain is
begin
```

(continues on next page)

(continued from previous page)

```
Greetings.Hello;  
Greetings.Goodbye;  
end Gmain;
```

Following the one-unit-per-file rule, place this program in the following three separate files:

greetings.ads

spec of package Greetings

greetings.adb

body of package Greetings

gmain.adb

body of main program

Note that there is no required order of compilation when using GNAT. In particular it is perfectly fine to compile the main program first. Also, it is not necessary to compile package specs in the case where there is an accompanying body; you only need to compile the body. If you want to submit these files to the compiler for semantic checking and not code generation, then use the `-gnatc` switch:

```
$ gcc -c greetings.ads -gnatc
```

Although the compilation can be done in separate steps, in practice it is almost always more convenient to use the `gnatmake` or `gprbuild` tools:

```
$ gnatmake gmain.adb
```

THE GNAT COMPILATION MODEL

This chapter describes the compilation model used by GNAT. Although similar to that used by other languages such as C and C++, this model is substantially different from the traditional Ada compilation models, which are based on a centralized program library. The chapter covers the following material:

- Topics related to source file makeup and naming
 - Source Representation
 - Foreign Language Representation
 - File Naming Topics and Utilities
- Configuration Pragmas
- Generating Object Files
- Source Dependencies
- The Ada Library Information Files
- Binding an Ada Program
- GNAT and Libraries
- Conditional Compilation
- Mixed Language Programming
- GNAT and Other Compilation Models
- Using GNAT Files with External Tools

3.1 Source Representation

Ada source programs are represented in standard text files, using Latin-1 coding. Latin-1 is an 8-bit code that includes the familiar 7-bit ASCII set, plus additional characters used for representing foreign languages (see [Foreign Language Representation](#) for support of non-USA character sets). The format effector characters are represented using their standard ASCII encodings, as follows:

Character	Effect	Code
VT	Vertical tab	16#0B#
HT	Horizontal tab	16#09#
CR	Carriage return	16#0D#
LF	Line feed	16#0A#
FF	Form feed	16#0C#

Source files are in standard text file format. In addition, GNAT will recognize a wide variety of stream formats, in which the end of physical lines is marked by any of the following sequences: LF, CR, CR-LF, or LF-CR. This is useful in accommodating files that are imported from other operating systems.

The end of a source file is normally represented by the physical end of file. However, the control character 16#1A# (SUB) is also recognized as signalling the end of the source file. Again, this is provided for compatibility with other operating systems where this code is used to represent the end of file.

Each file contains a single Ada compilation unit, including any pragmas associated with the unit. For example, this means you must place a package declaration (a package *spec*) and the corresponding body in separate files. An Ada *compilation* (which is a sequence of compilation units) is represented using a sequence of files. Similarly, you will place each subunit or child unit in a separate file.

3.2 Foreign Language Representation

GNAT supports the standard character sets defined in Ada as well as several other non-standard character sets for use in localized versions of the compiler ([Character Set Control](#)).

3.2.1 Latin-1

The basic character set is Latin-1. This character set is defined by ISO standard 8859, part 1. The lower half (character codes 16#00# ... 16#7F#) is identical to standard ASCII coding, but the upper half is used to represent additional characters. These include extended letters used by European languages, such as French accents, the vowels with umlauts used in German, and the extra letter A-ring used in Swedish.

For a complete list of Latin-1 codes and their encodings, see the source file of library unit `Ada.Characters.Latin_1` in file `a-chlat1.ads`. You may use any of these extended characters freely in character or string literals. In addition, the extended characters that represent letters can be used in identifiers.

3.2.2 Other 8-Bit Codes

GNAT also supports several other 8-bit coding schemes:

ISO 8859-2 (Latin-2)

Latin-2 letters allowed in identifiers, with uppercase and lowercase equivalence.

ISO 8859-3 (Latin-3)

Latin-3 letters allowed in identifiers, with uppercase and lowercase equivalence.

ISO 8859-4 (Latin-4)

Latin-4 letters allowed in identifiers, with uppercase and lowercase equivalence.

ISO 8859-5 (Cyrillic)

ISO 8859-5 letters (Cyrillic) allowed in identifiers, with uppercase and lowercase equivalence.

ISO 8859-15 (Latin-9)

ISO 8859-15 (Latin-9) letters allowed in identifiers, with uppercase and lowercase equivalence

IBM PC (code page 437)

This code page is the normal default for PCs in the U.S. It corresponds to the original IBM PC character set. This set has some, but not all, of the extended Latin-1 letters, but these letters do not have the same encoding as Latin-1. In this mode, these letters are allowed in identifiers with uppercase and lowercase equivalence.

IBM PC (code page 850)

This code page is a modification of 437 extended to include all the Latin-1 letters, but still not with the usual Latin-1 encoding. In this mode, all these letters are allowed in identifiers with uppercase and lowercase equivalence.

Full Upper 8-bit

Any character in the range 80-FF allowed in identifiers, and all are considered distinct. In other words, there are no uppercase and lowercase equivalences in this range. This is useful in conjunction with certain encoding schemes used for some foreign character sets (e.g., the typical method of representing Chinese characters on the PC).

No Upper-Half

No upper-half characters in the range 80-FF are allowed in identifiers. This gives Ada 83 compatibility for identifier names.

For precise data on the encodings permitted, and the uppercase and lowercase equivalences that are recognized, see the file `csets.adb` in the GNAT compiler sources. You will need to obtain a full source release of GNAT to obtain this file.

3.2.3 Wide_Character Encodings

GNAT allows wide character codes to appear in character and string literals, and also optionally in identifiers, by means of the following possible encoding schemes:

Hex Coding

In this encoding, a wide character is represented by the following five character sequence:

```
ESC a b c d
```

where *a*, *b*, *c*, *d* are the four hexadecimal characters (using uppercase letters) of the wide character code. For example, ESC A345 is used to represent the wide character with code 16#A345#. This scheme is compatible with use of the full Wide_Character set.

Upper-Half Coding

The wide character with encoding 16#abcd# where the upper bit is on (in other words, 'a' is in the range 8-F) is represented as two bytes, 16#ab# and 16#cd#. The second byte cannot be a format control character, but is not required to be in the upper half. This method can be also used for shift-JIS or EUC, where the internal coding matches the external coding.

Shift JIS Coding

A wide character is represented by a two-character sequence, 16#ab# and 16#cd#, with the restrictions described for upper-half encoding as described above. The internal character code is the corresponding JIS character according to the standard algorithm for Shift-JIS conversion. Only characters defined in the JIS code set table can be used with this encoding method.

EUC Coding

A wide character is represented by a two-character sequence 16#ab# and 16#cd#, with both characters being in the upper half. The internal character code is the corresponding JIS character according to the EUC encoding algorithm. Only characters defined in the JIS code set table can be used with this encoding method.

UTF-8 Coding

A wide character is represented using UCS Transformation Format 8 (UTF-8) as defined in Annex R of ISO 10646-1/Am.2. Depending on the character value, the representation is a one, two, or three byte sequence:

```
16#0000#-16#007f#: 2#0xxxxxxx#
16#0080#-16#07ff#: 2#110xxxxx# 2#10xxxxxx#
16#0800#-16#ffff#: 2#1110xxxx# 2#10xxxxxx# 2#10xxxxxx#
```

where the xxx bits correspond to the left-padded bits of the 16-bit character value. Note that all lower half ASCII characters are represented as ASCII bytes and all upper half characters and other wide characters are represented as sequences of upper-half (The full UTF-8 scheme allows for encoding 31-bit characters as 6-byte sequences, and in the following section on wide wide characters, the use of these sequences is documented).

Brackets Coding

In this encoding, a wide character is represented by the following eight character sequence:

```
[ " a b c d " ]
```

where `a`, `b`, `c`, `d` are the four hexadecimal characters (using uppercase letters) of the wide character code. For example, `['A345']` is used to represent the wide character with code `16#A345#`. It is also possible (though not required) to use the Brackets coding for upper half characters. For example, the code `16#A3#` can be represented as `['A3']`.

This scheme is compatible with use of the full `Wide_Character` set, and is also the method used for wide character encoding in some standard ACATS (Ada Conformity Assessment Test Suite) test suite distributions.

Note: Some of these coding schemes do not permit the full use of the Ada character set. For example, neither Shift JIS nor EUC allow the use of the upper half of the Latin-1 set.

3.2.4 Wide_Wide_Character Encodings

GNAT allows wide wide character codes to appear in character and string literals, and also optionally in identifiers, by means of the following possible encoding schemes:

UTF-8 Coding

A wide character is represented using UCS Transformation Format 8 (UTF-8) as defined in Annex R of ISO 10646-1/Am.2. Depending on the character value, the representation of character codes with values greater than `16#FFFF#` is a four, five, or six byte sequence:

```
16#01_0000#-16#10_FFFF#:  11110xxx 10xxxxxx 10xxxxxx
                          10xxxxxx
16#0020_0000#-16#03FF_FFFF#: 111110xx 10xxxxxx 10xxxxxx
                          10xxxxxx 10xxxxxx
16#0400_0000#-16#7FFF_FFFF#: 1111110x 10xxxxxx 10xxxxxx
                          10xxxxxx 10xxxxxx 10xxxxxx
```

where the `xxx` bits correspond to the left-padded bits of the 32-bit character value.

Brackets Coding

In this encoding, a wide wide character is represented by the following ten or twelve byte character sequence:

```
[ " a b c d e f " ]
[ " a b c d e f g h " ]
```

where `a-h` are the six or eight hexadecimal characters (using uppercase letters) of the wide wide character code. For example, `["1F4567"]` is used to represent the wide wide character with code `16#001F_4567#`.

This scheme is compatible with use of the full `Wide_Wide_Character` set, and is also the method used for wide wide character encoding in some standard ACATS (Ada Conformity Assessment Test Suite) test suite distributions.

3.3 File Naming Topics and Utilities

GNAT has a default file naming scheme and also provides the user with a high degree of control over how the names and extensions of the source files correspond to the Ada compilation units that they contain.

3.3.1 File Naming Rules

The default file name is determined by the name of the unit that the file contains. The name is formed by taking the full expanded name of the unit and replacing the separating dots with hyphens and using lowercase for all letters.

An exception arises if the file name generated by the above rules starts with one of the characters `a`, `g`, `i`, or `s`, and the second character is a minus. In this case, the character tilde is used in place of the minus. The reason for this special rule is to avoid clashes with the standard names for child units of the packages `System`, `Ada`, `Interfaces`, and `GNAT`, which use the prefixes `s-`, `a-`, `i-`, and `g-`, respectively.

The file extension is `.ads` for a spec and `.adb` for a body. The following table shows some examples of these rules.

Source File	Ada Compilation Unit
<code>main.ads</code>	Main (spec)
<code>main.adb</code>	Main (body)
<code>arith_functions.ads</code>	Arith_Functions (package spec)
<code>arith_functions.adb</code>	Arith_Functions (package body)
<code>func-spec.ads</code>	Func.Spec (child package spec)
<code>func-spec.adb</code>	Func.Spec (child package body)
<code>main-sub.adb</code>	Sub (subunit of Main)
<code>a~bad.adb</code>	A.Bad (child package body)

Following these rules can result in excessively long file names if corresponding unit names are long (for example, if child units or subunits are heavily nested). An option is available to shorten such long file names (called file name ‘krunching’). This may be particularly useful when programs being developed with GNAT are to be used on operating systems with limited file name lengths. [Using `gnatkr`](#).

Of course, no file shortening algorithm can guarantee uniqueness over all possible unit names; if file name krunching is used, it is your responsibility to ensure no name clashes occur. Alternatively you can specify the exact file names that you want used, as described in the next section. Finally, if your Ada programs are migrating from a compiler with a different naming convention, you can use the `gnatchop` utility to produce source files that follow the GNAT naming conventions. (For details see [Renaming Files with `gnatchop`](#).)

Note: in the case of Windows or Mac OS operating systems, case is not significant. So for example on Windows if the canonical name is `main-sub.adb`, you can use the file name `Main-Sub.adb` instead. However, case is significant for other operating systems, so for example, if you want to use other

than canonically cased file names on a Unix system, you need to follow the procedures described in the next section.

3.3.2 Using Other File Names

In the previous section, we have described the default rules used by GNAT to determine the file name in which a given unit resides. It is often convenient to follow these default rules, and if you follow them, the compiler knows without being explicitly told where to find all the files it needs.

However, in some cases, particularly when a program is imported from another Ada compiler environment, it may be more convenient for the programmer to specify which file names contain which units. GNAT allows arbitrary file names to be used by means of the `Source_File_Name` pragma. The form of this pragma is as shown in the following examples:

```
pragma Source_File_Name (My_Uilities.Stacks,  
  Spec_File_Name => "myutilst_a.ada");  
pragma Source_File_name (My_Uilities.Stacks,  
  Body_File_Name => "myutilst.ada");
```

As shown in this example, the first argument for the pragma is the unit name (in this example a child unit). The second argument has the form of a named association. The identifier indicates whether the file name is for a spec or a body; the file name itself is given by a string literal.

The source file name pragma is a configuration pragma, which means that normally it will be placed in the `gnat.adc` file used to hold configuration pragmas that apply to a complete compilation environment. For more details on how the `gnat.adc` file is created and used see [Handling of Configuration Pragmas](#).

GNAT allows completely arbitrary file names to be specified using the source file name pragma. However, if the file name specified has an extension other than `.ads` or `.adb` it is necessary to use a special syntax when compiling the file. The name in this case must be preceded by the special sequence `-x` followed by a space and the name of the language, here `ada`, as in:

```
$ gcc -c -x ada peculiar_file_name.sim
```

`gnatmake` handles non-standard file names in the usual manner (the non-standard file name for the main program is simply used as the argument to `gnatmake`). Note that if the extension is also non-standard, then it must be included in the `gnatmake` command, it may not be omitted.

3.3.3 Alternative File Naming Schemes

The previous section described the use of the `Source_File_Name` pragma to allow arbitrary names to be assigned to individual source files. However, this approach requires one pragma for each file, and especially in large systems can result in very long `gnat.adc` files, and also create a maintenance problem.

GNAT also provides a facility for specifying systematic file naming schemes other than the standard default naming scheme previously described. An alternative scheme for naming is specified by the use of `Source_File_Name` pragmas having the following format:

```

pragma Source_File_Name (
  Spec_File_Name => FILE_NAME_PATTERN
  [ , Casing      => CASING_SPEC ]
  [ , Dot_Replacement => STRING_LITERAL ] );

pragma Source_File_Name (
  Body_File_Name => FILE_NAME_PATTERN
  [ , Casing      => CASING_SPEC ]
  [ , Dot_Replacement => STRING_LITERAL ] );

pragma Source_File_Name (
  Subunit_File_Name => FILE_NAME_PATTERN
  [ , Casing        => CASING_SPEC ]
  [ , Dot_Replacement => STRING_LITERAL ] );

FILE_NAME_PATTERN ::= STRING_LITERAL
CASING_SPEC ::= Lowercase | Uppercase | Mixedcase

```

The `FILE_NAME_PATTERN` string shows how the file name is constructed. It contains a single asterisk character, and the unit name is substituted systematically for this asterisk. The optional parameter `Casing` indicates whether the unit name is to be all upper-case letters, all lower-case letters, or mixed-case. If no `Casing` parameter is used, then the default is all lower-case.

The optional `Dot_Replacement` string is used to replace any periods that occur in subunit or child unit names. If no `Dot_Replacement` argument is used then separating dots appear unchanged in the resulting file name. Although the above syntax indicates that the `Casing` argument must appear before the `Dot_Replacement` argument, but it is also permissible to write these arguments in the opposite order.

As indicated, it is possible to specify different naming schemes for bodies, specs, and subunits. Quite often the rule for subunits is the same as the rule for bodies, in which case, there is no need to give a separate `Subunit_File_Name` rule, and in this case the `Body_File_name` rule is used for subunits as well.

The separate rule for subunits can also be used to implement the rather unusual case of a compilation environment (e.g., a single directory) which contains a subunit and a child unit with the same unit name. Although both units cannot appear in the same partition, the Ada Reference Manual allows (but does not require) the possibility of the two units coexisting in the same environment.

The file name translation works in the following steps:

- If there is a specific `Source_File_Name` pragma for the given unit, then this is always used, and any general pattern rules are ignored.
- If there is a pattern type `Source_File_Name` pragma that applies to the unit, then the resulting file name will be used if the file exists. If more than one pattern matches, the latest one will be tried first, and the first attempt resulting in a reference to a file that exists will be used.
- If no pattern type `Source_File_Name` pragma that applies to the unit for which the corresponding file exists, then the standard GNAT default naming rules are used.

As an example of the use of this mechanism, consider a commonly used scheme in which file names

are all lower case, with separating periods copied unchanged to the resulting file name, and specs end with `.1.ad`, and bodies end with `.2.ad`. GNAT will follow this scheme if the following two pragmas appear:

```
pragma Source_File_Name
  (Spec_File_Name => ".1.ad");
pragma Source_File_Name
  (Body_File_Name => ".2.ad");
```

The default GNAT scheme is actually implemented by providing the following default pragmas internally:

```
pragma Source_File_Name
  (Spec_File_Name => ".ads", Dot_Replacement => "-");
pragma Source_File_Name
  (Body_File_Name => ".adb", Dot_Replacement => "-");
```

Our final example implements a scheme typically used with one of the Ada 83 compilers, where the separator character for subunits was ‘`__`’ (two underscores), specs were identified by adding `_.ADA`, bodies by adding `.ADA`, and subunits by adding `.SEP`. All file names were upper case. Child units were not present of course since this was an Ada 83 compiler, but it seems reasonable to extend this scheme to use the same double underscore separator for child units.

```
pragma Source_File_Name
  (Spec_File_Name => "_.ADA",
   Dot_Replacement => "__",
   Casing = Uppercase);
pragma Source_File_Name
  (Body_File_Name => ".ADA",
   Dot_Replacement => "__",
   Casing = Uppercase);
pragma Source_File_Name
  (Subunit_File_Name => ".SEP",
   Dot_Replacement => "__",
   Casing = Uppercase);
```

3.3.4 Handling Arbitrary File Naming Conventions with `gnatname`

Arbitrary File Naming Conventions

The GNAT compiler must be able to know the source file name of a compilation unit. When using the standard GNAT default file naming conventions (`.ads` for specs, `.adb` for bodies), the GNAT compiler does not need additional information.

When the source file names do not follow the standard GNAT default file naming conventions, the GNAT compiler must be given additional information through a configuration pragmas file ([Configuration Pragmas](#)) or a project file. When the non-standard file naming conventions are well-defined, a small number of pragmas `Source_File_Name` specifying a naming pattern ([Alternative File Naming Schemes](#)) may be sufficient. However, if the file naming conventions are irregular or arbitrary, a number of pragma `Source_File_Name` for individual compilation units must be defined.

To help maintain the correspondence between compilation unit names and source file names within the compiler, GNAT provides a tool `gnatname` to generate the required pragmas for a set of files.

Running `gnatname`

The usual form of the `gnatname` command is:

```
$ gnatname [ switches ] naming_pattern [ naming_patterns ]
  [--and [ switches ] naming_pattern [ naming_patterns ]]
```

All of the arguments are optional. If invoked without any argument, `gnatname` will display its usage.

When used with at least one naming pattern, `gnatname` will attempt to find all the compilation units in files that follow at least one of the naming patterns. To find these compilation units, `gnatname` will use the GNAT compiler in syntax-check-only mode on all regular files.

One or several Naming Patterns may be given as arguments to `gnatname`. Each Naming Pattern is enclosed between double quotes (or single quotes on Windows). A Naming Pattern is a regular expression similar to the wildcard patterns used in file names by the Unix shells or the DOS prompt.

`gnatname` may be called with several sections of directories/patterns. Sections are separated by the switch `--and`. In each section, there must be at least one pattern. If no directory is specified in a section, the current directory (or the project directory if `-P` is used) is implied. The options other than the directory switches and the patterns apply globally even if they are in different sections.

Examples of Naming Patterns are:

```
"*. [12].ada"
"*.ad[sb]*"
"body_*"    "spec_*
```

For a more complete description of the syntax of Naming Patterns, see the second kind of regular expressions described in `g-regexp.ads` (the ‘Glob’ regular expressions).

When invoked without the switch `-P`, `gnatname` will create a configuration pragmas file `gnat.adc` in the current working directory, with pragmas `Source_File_Name` for each file that contains a valid Ada unit.

Switches for `gnatname`

Switches for `gnatname` must precede any specified Naming Pattern.

You may specify any of the following switches to `gnatname`:

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

--subdirs=*dir*

Real object, library or exec directories are subdirectories `<dir>` of the specified ones.

--no-backup

Do not create a backup copy of an existing project file.

--and

Start another section of directories/patterns.

-cfilename

Create a configuration pragmas file `filename` (instead of the default `gnat.adc`). There may be zero, one or more space between `-c` and `filename`. `filename` may include directory information. `filename` must be writable. There may be only one switch `-c`. When a switch `-c` is specified, no switch `-P` may be specified (see below).

-ddir

Look for source files in directory `dir`. There may be zero, one or more spaces between `-d` and `dir`. `dir` may end with `/**`, that is it may be of the form `root_dir/**`. In this case, the directory `root_dir` and all of its subdirectories, recursively, have to be searched for sources. When a switch `-d` is specified, the current working directory will not be searched for source files, unless it is explicitly specified with a `-d` or `-D` switch. Several switches `-d` may be specified. If `dir` is a relative path, it is relative to the directory of the configuration pragmas file specified with switch `-c`, or to the directory of the project file specified with switch `-P` or, if neither switch `-c` nor switch `-P` are specified, it is relative to the current working directory. The directory specified with switch `-d` must exist and be readable.

-Dfilename

Look for source files in all directories listed in text file `filename`. There may be zero, one or more spaces between `-D` and `filename`. `filename` must be an existing, readable text file. Each nonempty line in `filename` must be a directory. Specifying switch `-D` is equivalent to specifying as many switches `-d` as there are nonempty lines in `file`.

-eL

Follow symbolic links when processing project files.

-fpattern

Foreign patterns. Using this switch, it is possible to add sources of languages other than Ada to the list of sources of a project file. It is only useful if a `-P` switch is used. For example,

```
gnatname -Pprj -f"*.c" "*.ada"
```

will look for Ada units in all files with the `.ada` extension, and will add to the list of file for project `prj.gpr` the C files with extension `.c`.

-h

Output usage (help) information. The output is written to `stdout`.

-Pproj

Create or update project file `proj`. There may be zero, one or more space between `-P` and `proj`. `proj` may include directory information. `proj` must be writable. There may be only one switch `-P`. When a switch `-P` is specified, no switch `-c` may be specified. On all platforms, except on VMS, when `gnatname` is invoked for an existing project file `<proj>.gpr`, a backup copy of the project file is created in the project directory with file name `<proj>.gpr.saved_x`. 'x' is the first non negative number that makes this backup copy a new file.

-v

Verbose mode. Output detailed explanation of behavior to `stdout`. This includes name of the file written, the name of the directories to search and, for each file in those directories whose name matches at least one of the Naming Patterns, an indication of whether the file contains a unit, and if so the name of the unit.

-v -v

Very Verbose mode. In addition to the output produced in verbose mode, for each file in the searched directories whose name matches none of the Naming Patterns, an indication is given that there is no match.

-xpattern

Excluded patterns. Using this switch, it is possible to exclude some files that would match the name patterns. For example,

```
gnatname -x "*_nt.ada" "*.ada"
```

will look for Ada units in all files with the `.ada` extension, except those whose names end with `_nt.ada`.

Examples of gnatname Usage

```
$ gnatname -c /home/me/names.adc -d sources "[a-z]*.ada"
```

In this example, the directory `/home/me` must already exist and be writable. In addition, the directory `/home/me/sources` (specified by `-d sources`) must exist and be readable.

Note the optional spaces after `-c` and `-d`.

```
$ gnatname -P/home/me/proj -x "*_nt_body.ada"  
-dsources -dsources/plus -Dcommon_dirs.txt "body_*" "spec_*"
```

Note that several switches `-d` may be used, even in conjunction with one or several switches `-D`. Several Naming Patterns and one excluded pattern are used in this example.

3.3.5 File Name Krunching with gnatkr

This section discusses the method used by the compiler to shorten the default file names chosen for Ada units so that they do not exceed the maximum length permitted. It also describes the `gnatkr` utility that can be used to determine the result of applying this shortening.

About `gnatkr`

The default file naming rule in GNAT is that the file name must be derived from the unit name. The exact default rule is as follows:

- Take the unit name and replace all dots by hyphens.
- If such a replacement occurs in the second character position of a name, and the first character is `a`, `g`, `s`, or `i`, then replace the dot by the character `~` (tilde) instead of a minus.

The reason for this exception is to avoid clashes with the standard names for children of System, Ada, Interfaces, and GNAT, which use the prefixes `s-`, `a-`, `i-`, and `g-`, respectively.

The `-gnatkrn` switch of the compiler activates a 'krunching' circuit that limits file names to `nn` characters (where `nn` is a decimal integer).

The `gnatkr` utility can be used to determine the krunched name for a given file, when krunched to a specified maximum length.

Using `gnatkr`

The `gnatkr` command has the form:

```
$ gnatkr name [ length ]
```

`name` is the uncrunched file name, derived from the name of the unit in the standard manner described in the previous section (i.e., in particular all dots are replaced by hyphens). The file name may or may not have an extension (defined as a suffix of the form period followed by arbitrary characters other than period). If an extension is present then it will be preserved in the output. For example, when krunching `hellofile.ads` to eight characters, the result will be `hellofil.ads`.

Note: for compatibility with previous versions of `gnatkr` dots may appear in the name instead of hyphens, but the last dot will always be taken as the start of an extension. So if `gnatkr` is given an argument such as `Hello.World.adb` it will be treated exactly as if the first period had been a hyphen, and for example krunching to eight characters gives the result `hellworl.adb`.

Note that the result is always all lower case. Characters of the other case are folded as required.

`length` represents the length of the krunched name. The default when no argument is given is 8 characters. A length of zero stands for unlimited, in other words do not chop except for system files where the implied krunching length is always eight characters.

The output is the krunched name. The output has an extension only if the original argument was a file name with an extension.

Krunching Method

The initial file name is determined by the name of the unit that the file contains. The name is formed by taking the full expanded name of the unit and replacing the separating dots with hyphens and using lowercase for all letters, except that a hyphen in the second character position is replaced by a tilde if the first character is `a`, `i`, `g`, or `s`. The extension is `.ads` for a spec and `.adb` for a body. Krunching does not affect the extension, but the file name is shortened to the specified length by following these rules:

- The name is divided into segments separated by hyphens, tildes or underscores and all hyphens, tildes, and underscores are eliminated. If this leaves the name short enough, we are done.
- If the name is too long, the longest segment is located (left-most if there are two of equal length), and shortened by dropping its last character. This is repeated until the name is short enough.

As an example, consider the krunching of `our-strings-wide_fixed.adb` to fit the name into 8 characters as required by some operating systems:

```
our-strings-wide_fixed 22
our strings wide fixed 19
our string  wide fixed 18
our strin  wide fixed 17
our stri   wide fixed 16
our stri   wide fixe 15
our str    wide fixe 14
our str    wid  fixe 13
our str    wid  fix  12
ou str     wid  fix  11
ou st      wid  fix  10
ou st      wi   fix  9
ou st      wi   fi   8
Final file name: oustwifi.adb
```

- The file names for all predefined units are always krunched to eight characters. The krunching of these predefined units uses the following special prefix replacements:

Prefix	Replacement
ada-	a-
gnat-	g-
interfac es-	i-
system-	s-

These system files have a hyphen in the second character position. That is why normal user files replace such a character with a tilde, to avoid confusion with system file names.

As an example of this special rule, consider `ada-strings-wide_fixed.adb`, which gets krunched as follows:

```

ada-strings-wide_fixed 22
a- strings wide fixed 18
a- string  wide fixed 17
a- strin  wide fixed 16
a- stri   wide fixed 15
a- stri   wide fixe 14
a- str    wide fixe 13
a- str    wid  fixe 12
a- str    wid  fix  11
a- st     wid  fix  10
a- st     wi   fix  9
a- st     wi   fi   8
Final file name: a-stwifi.adb

```

Of course no file shortening algorithm can guarantee uniqueness over all possible unit names, and if file name krunching is used then it is your responsibility to ensure that no name clashes occur. The utility program `gnatkr` is supplied for conveniently determining the krunched name of a file.

Examples of `gnatkr` Usage

```

$ gnatkr very_long_unit_name.ads --> velounna.ads
$ gnatkr grandparent-parent-child.ads --> grparchi.ads
$ gnatkr Grandparent.Parent.Child.ads --> grparchi.ads
$ gnatkr grandparent-parent-child --> grparchi
$ gnatkr very_long_unit_name.ads/count=6 --> vlunna.ads
$ gnatkr very_long_unit_name.ads/count=0 --> very_long_unit_name.ads

```

3.3.6 Renaming Files with `gnatchop`

This section discusses how to handle files with multiple units by using the `gnatchop` utility. This utility is also useful in renaming files to meet the standard GNAT default file naming conventions.

Handling Files with Multiple Units

The basic compilation model of GNAT requires that a file submitted to the compiler have only one unit and there be a strict correspondence between the file name and the unit name.

If you want to keep your files with multiple units, perhaps to maintain compatibility with some other Ada compilation system, you can use `gnatname` to generate or update your project files. Generated or modified project files can be processed by GNAT.

See [Handling Arbitrary File Naming Conventions with `gnatname`](#) for more details on how to use `gnatname`.

Alternatively, if you want to permanently restructure a set of ‘foreign’ files so that they match the GNAT rules, and do the remaining development using the GNAT structure, you can simply use `gnatchop` once, generate the new set of files and work with them from that point on.

Note that if your file containing multiple units starts with a byte order mark (BOM) specifying UTF-8 encoding, then the files generated by `gnatchop` will each start with a copy of this BOM, meaning that they can be compiled automatically in UTF-8 mode without needing to specify an explicit encoding.

Operating `gnatchop` in Compilation Mode

The basic function of `gnatchop` is to take a file with multiple units and split it into separate files. The boundary between files is reasonably clear, except for the issue of comments and pragmas. In default mode, the rule is that any pragmas between units belong to the previous unit, except that configuration pragmas always belong to the following unit. Any comments belong to the following unit. These rules almost always result in the right choice of the split point without needing to mark it explicitly and most users will find this default to be what they want. In this default mode it is incorrect to submit a file containing only configuration pragmas, or one that ends in configuration pragmas, to `gnatchop`.

However, using a special option to activate ‘compilation mode’, `gnatchop` can perform another function, which is to provide exactly the semantics required by the RM for handling of configuration pragmas in a compilation. In the absence of configuration pragmas (at the main file level), this option has no effect, but it causes such configuration pragmas to be handled in a quite different manner.

First, in compilation mode, if `gnatchop` is given a file that consists of only configuration pragmas, then this file is appended to the `gnat.adc` file in the current directory. This behavior provides the required behavior described in the RM for the actions to be taken on submitting such a file to the compiler, namely that these pragmas should apply to all subsequent compilations in the same compilation environment. Using GNAT, the current directory, possibly containing a `gnat.adc` file is the representation of a compilation environment. For more information on the `gnat.adc` file, see [Handling of Configuration Pragmas](#).

Second, in compilation mode, if `gnatchop` is given a file that starts with configuration pragmas, and contains one or more units, then these configuration pragmas are prepended to each of the chopped files. This behavior provides the required behavior described in the RM for the actions to be taken on compiling such a file, namely that the pragmas apply to all units in the compilation, but not to subsequently compiled units.

Finally, if configuration pragmas appear between units, they are appended to the previous unit. This results in the previous unit being illegal, since the compiler does not accept configuration pragmas that follow a unit. This provides the required RM behavior that forbids configuration pragmas other than those preceding the first compilation unit of a compilation.

For most purposes, `gnatchop` will be used in default mode. The compilation mode described above is used only if you need exactly accurate behavior with respect to compilations, and you have files that contain multiple units and configuration pragmas. In this circumstance the use of `gnatchop` with the compilation mode switch provides the required behavior, and is for example the mode in which GNAT processes the ACVC tests.

Command Line for gnat Chop

The `gnat Chop` command has the form:

```
$ gnat Chop switches file_name [file_name ...]
    [directory]
```

The only required argument is the file name of the file to be chopped. There are no restrictions on the form of this file name. The file itself contains one or more Ada units, in normal GNAT format, concatenated together. As shown, more than one file may be presented to be chopped.

When run in default mode, `gnat Chop` generates one output file in the current directory for each unit in each of the files.

`directory`, if specified, gives the name of the directory to which the output files will be written. If it is not specified, all files are written to the current directory.

For example, given a file called `hellofiles` containing

```
procedure Hello;

with Ada.Text_IO; use Ada.Text_IO;
procedure Hello is
begin
    Put_Line ("Hello");
end Hello;
```

the command

```
$ gnat Chop hellofiles
```

generates two files in the current directory, one called `hello.ads` containing the single line that is the procedure spec, and the other called `hello.adb` containing the remaining text. The original file is not affected. The generated files can be compiled in the normal manner.

When `gnat Chop` is invoked on a file that is empty or that contains only empty lines and/or comments, `gnat Chop` will not fail, but will not produce any new sources.

For example, given a file called `toto.txt` containing

```
-- Just a comment
```

the command

```
$ gnat Chop toto.txt
```

will not produce any new file and will result in the following warnings:

```
toto.txt:1:01: warning: empty file, contains no compilation units
no compilation units found
no source files written
```

Switches for `gnatchop`

`gnatchop` recognizes the following switches:

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

-c

Causes `gnatchop` to operate in compilation mode, in which configuration pragmas are handled according to strict RM rules. See previous section for a full description of this mode.

-gnatxxx

This passes the given `-gnatxxx` switch to `gnat` which is used to parse the given file. Not all `xxx` options make sense, but for example, the use of `-gnati2` allows `gnatchop` to process a source file that uses Latin-2 coding for identifiers.

-h

Causes `gnatchop` to generate a brief help summary to the standard output file showing usage information.

-kmm

Limit generated file names to the specified number `mm` of characters. This is useful if the resulting set of files is required to be interoperable with systems which limit the length of file names. No space is allowed between the `-k` and the numeric value. The numeric value may be omitted in which case a default of `-k8`, suitable for use with DOS-like file systems, is used. If no `-k` switch is present then there is no limit on the length of file names.

-p

Causes the file modification time stamp of the input file to be preserved and used for the time stamp of the output file(s). This may be useful for preserving coherency of time stamps in an environment where `gnatchop` is used as part of a standard build process.

-q

Causes output of informational messages indicating the set of generated files to be suppressed. Warnings and error messages are unaffected.

-r

Generate `Source_Reference` pragmas. Use this switch if the output files are regarded as temporary and development is to be done in terms of the original unchopped file. This switch causes `Source_Reference` pragmas to be inserted into each of the generated files to refer back to the original file name and line number. The result is that all error messages refer back to the original unchopped file. In addition, the debugging information placed into the object file (when the `-g` switch of `gcc` or `gnatmake` is specified) also refers back to this original file so that tools like profilers and debuggers will give information in terms of the original unchopped file.

If the original file to be chopped itself contains a `Source_Reference` pragma referencing a third file, then `gnatchop` respects this pragma, and the generated `Source_Reference` pragmas in the chopped file refer to the original file, with appropriate line numbers. This is particularly useful

when `gnatchop` is used in conjunction with `gnatprep` to compile files that contain preprocessing statements and multiple units.

-v

Causes `gnatchop` to operate in verbose mode. The version number and copyright notice are output, as well as exact copies of the `gnat1` commands spawned to obtain the chop control information.

-w

Overwrite existing file names. Normally `gnatchop` regards it as a fatal error if there is already a file with the same name as a file it would otherwise output, in other words if the files to be chopped contain duplicated units. This switch bypasses this check, and causes all but the last instance of such duplicated units to be skipped.

--GCC=xxxx

Specify the path of the GNAT parser to be used. When this switch is used, no attempt is made to add the prefix to the GNAT parser executable.

Examples of `gnatchop` Usage

```
$ gnatchop -w hello_s.ada prerelease/files
```

Chops the source file `hello_s.ada`. The output files will be placed in the directory `prerelease/files`, overwriting any files with matching names in that directory (no files in the current directory are modified).

```
$ gnatchop archive
```

Chops the source file `archive` into the current directory. One useful application of `gnatchop` is in sending sets of sources around, for example in email messages. The required sources are simply concatenated (for example, using a Unix `cat` command), and then `gnatchop` is used at the other end to reconstitute the original file names.

```
$ gnatchop file1 file2 file3 direc
```

Chops all units in files `file1`, `file2`, `file3`, placing the resulting files in the directory `direc`. Note that if any units occur more than once anywhere within this set of files, an error message is generated, and no files are written. To override this check, use the `-w` switch, in which case the last occurrence in the last file will be the one that is output, and earlier duplicate occurrences for a given unit will be skipped.

3.4 Configuration Pragmas

Configuration pragmas include those pragmas described as such in the Ada Reference Manual, as well as implementation-dependent pragmas that are configuration pragmas. See the `Implementation_Defined_Pragmas` chapter in the *GNAT Reference Manual* for details on these additional GNAT-specific configuration pragmas. Most notably, the pragma `Source_File_Name`, which allows specifying non-default names for source files, is a configuration pragma. The following is a complete list of configuration pragmas recognized by GNAT:

```
Ada_83
Ada_95
Ada_05
Ada_2005
Ada_12
Ada_2012
Ada_2022
Aggregate_Individually_Assign
Allow_Integer_Address
Annotate
Assertion_Policy
Assume_No_Invalid_Values
C_Pass_By_Copy
Check_Float_Overflow
Check_Name
Check_Policy
Component_Alignment
Convention_Identifier
Debug_Policy
Default_Scalar_Storage_Order
Default_Storage_Pool
Detect_Blocking
Disable_Atomic_Synchronization
Discard_Names
Elaboration_Checks
Eliminate
Enable_Atomic_Synchronization
Extend_System
Extensions_Allowed
External_Name_Casing
Fast_Math
Favor_Top_Level
Ignore_Pragma
Implicit_Packing
InitializeScalars
Interrupt_State
License
Locking_Policy
No_Component_Reordering
No_Heap_Finalization
No_Strict_Aliasing
NormalizeScalars
Optimize_Alignment
```

(continues on next page)

(continued from previous page)

```
Overflow_Mode
Overriding_Renamings
Partition_Elaboration_Policy
Persistent_BSS
Prefix_Exception_Messages
Priority_Specific_Dispatching
Profile
Profile_Warnings
Queuing_Policy
Rename_Pragma
Restrictions
Restriction_Warnings
Reviewable
Short_Circuit_And_Or
Source_File_Name
Source_File_Name_Project
SPARK_Mode
Style_Checks
Suppress
Suppress_Exception_Locations
Task_Dispatching_Policy
Unevaluated_Use_Of_Old
Unsuppress
Use_VADS_Size
Validity_Checks
Warning_As_Error
Warnings
Wide_Character-Encoding
```

3.4.1 Handling of Configuration Pragmas

Configuration pragmas may either appear at the start of a compilation unit, or they can appear in a configuration pragma file to apply to all compilations performed in a given compilation environment.

GNAT also provides the `gnatchop` utility to provide an automatic way to handle configuration pragmas following the semantics for compilations (that is, files with multiple units), described in the RM. See [Operating `gnatchop` in Compilation Mode](#) for details. However, for most purposes, it will be more convenient to edit the `gnat.adc` file that contains configuration pragmas directly, as described in the following section.

In the case of `Restrictions` pragmas appearing as configuration pragmas in individual compilation units, the exact handling depends on the type of restriction.

Restrictions that require partition-wide consistency (like `No_Tasking`) are recognized wherever they appear and can be freely inherited, e.g. from a *withed* unit to the *withing* unit. This makes sense since the binder will in any case insist on seeing consistent use, so any unit not conforming to any restrictions that are anywhere in the partition will be rejected, and you might as well find that out at compile time rather than at bind time.

For restrictions that do not require partition-wide consistency, e.g. `SPARK` or

`No_Implementation_Attributes`, in general the restriction applies only to the unit in which the pragma appears, and not to any other units.

The exception is `No_Elaboration_Code` which always applies to the entire object file from a compilation, i.e. to the body, spec, and all subunits. This restriction can be specified in a configuration pragma file, or it can be on the body and/or the spec (in either case it applies to all the relevant units). It can appear on a subunit only if it has previously appeared in the body of spec.

3.4.2 The Configuration Pragmas Files

In GNAT a compilation environment is defined by the current directory at the time that a compile command is given. This current directory is searched for a file whose name is `gnat.adc`. If this file is present, it is expected to contain one or more configuration pragmas that will be applied to the current compilation. However, if the switch `-gnatA` is used, `gnat.adc` is not considered. When taken into account, `gnat.adc` is added to the dependencies, so that if `gnat.adc` is modified later, an invocation of `gnatmake` will recompile the source.

Configuration pragmas may be entered into the `gnat.adc` file either by running `gnatchop` on a source file that consists only of configuration pragmas, or more conveniently by direct editing of the `gnat.adc` file, which is a standard format source file.

Besides `gnat.adc`, additional files containing configuration pragmas may be applied to the current compilation using the switch `-gnatec=path` where `path` must designate an existing file that contains only configuration pragmas. These configuration pragmas are in addition to those found in `gnat.adc` (provided `gnat.adc` is present and switch `-gnatA` is not used).

It is allowable to specify several switches `-gnatec=`, all of which will be taken into account.

Files containing configuration pragmas specified with switches `-gnatec=` are added to the dependencies, unless they are temporary files. A file is considered temporary if its name ends in `.tmp` or `.TMP`. Certain tools follow this naming convention because they pass information to `gcc` via temporary files that are immediately deleted; it doesn't make sense to depend on a file that no longer exists. Such tools include `gprbuild`, `gnatmake`, and `gnatcheck`.

By default, configuration pragma files are stored by their absolute paths in ALI files. You can use the `-gnateb` switch in order to store them by their basename instead.

If you are using project file, a separate mechanism is provided using project attributes.

3.5 Generating Object Files

An Ada program consists of a set of source files, and the first step in compiling the program is to generate the corresponding object files. These are generated by compiling a subset of these source files. The files you need to compile are the following:

- If a package spec has no body, compile the package spec to produce the object file for the package.

- If a package has both a spec and a body, compile the body to produce the object file for the package. The source file for the package spec need not be compiled in this case because there is only one object file, which contains the code for both the spec and body of the package.
- For a subprogram, compile the subprogram body to produce the object file for the subprogram. The spec, if one is present, is as usual in a separate file, and need not be compiled.
- In the case of subunits, only compile the parent unit. A single object file is generated for the entire subunit tree, which includes all the subunits.
- Compile child units independently of their parent units (though, of course, the spec of all the ancestor unit must be present in order to compile a child unit).
- Compile generic units in the same manner as any other units. The object files in this case are small dummy files that contain at most the flag used for elaboration checking. This is because GNAT always handles generic instantiation by means of macro expansion. However, it is still necessary to compile generic units, for dependency checking and elaboration purposes.

The preceding rules describe the set of files that must be compiled to generate the object files for a program. Each object file has the same name as the corresponding source file, except that the extension is `.o` as usual.

You may wish to compile other files for the purpose of checking their syntactic and semantic correctness. For example, in the case where a package has a separate spec and body, you would not normally compile the spec. However, it is convenient in practice to compile the spec to make sure it is error-free before compiling clients of this spec, because such compilations will fail if there is an error in the spec.

GNAT provides an option for compiling such files purely for the purposes of checking correctness; such compilations are not required as part of the process of building a program. To compile a file in this checking mode, use the `-gnatc` switch.

3.6 Source Dependencies

A given object file clearly depends on the source file which is compiled to produce it. Here we are using “depends” in the sense of a typical `make` utility; in other words, an object file depends on a source file if changes to the source file require the object file to be recompiled. In addition to this basic dependency, a given object may depend on additional source files as follows:

- If a file being compiled *withs* a unit `X`, the object file depends on the file containing the spec of unit `X`. This includes files that are *withed* implicitly either because they are parents of *withed* child units or they are run-time units required by the language constructs used in a particular unit.
- If a file being compiled instantiates a library level generic unit, the object file depends on both the spec and body files for this generic unit.
- If a file being compiled instantiates a generic unit defined within a package, the object file depends on the body file for the package as well as the spec file.
- If a file being compiled contains a call to a subprogram for which `pragma Inline` applies and inlining is activated with the `-gnatn` switch, the object file depends on the file containing the

body of this subprogram as well as on the file containing the spec. Note that for inlining to actually occur as a result of the use of this switch, it is necessary to compile in optimizing mode.

The use of `-gnatN` activates inlining optimization that is performed by the front end of the compiler. This inlining does not require that the code generation be optimized. Like `-gnatn`, the use of this switch generates additional dependencies.

When using a gcc-based back end, then the use of `-gnatN` is deprecated, and the use of `-gnatn` is preferred. Historically front end inlining was more extensive than the gcc back end inlining, but that is no longer the case.

- If an object file `O` depends on the proper body of a subunit through inlining or instantiation, it depends on the parent unit of the subunit. This means that any modification of the parent unit or one of its subunits affects the compilation of `O`.
- The object file for a parent unit depends on all its subunit body files.
- The previous two rules meant that for purposes of computing dependencies and recompilation, a body and all its subunits are treated as an indivisible whole.

These rules are applied transitively: if unit `A` *with*s unit `B`, whose elaboration calls an inlined procedure in package `C`, the object file for unit `A` will depend on the body of `C`, in file `c.adb`.

The set of dependent files described by these rules includes all the files on which the unit is semantically dependent, as dictated by the Ada language standard. However, it is a superset of what the standard describes, because it includes generic, inline, and subunit dependencies.

An object file must be recreated by recompiling the corresponding source file if any of the source files on which it depends are modified. For example, if the `make` utility is used to control compilation, the rule for an Ada object file must mention all the source files on which the object file depends, according to the above definition. The determination of the necessary recompilations is done automatically when one uses `gnatmake`.

3.7 The Ada Library Information Files

Each compilation actually generates two output files. The first of these is the normal object file that has a `.o` extension. The second is a text file containing full dependency information. It has the same name as the source file, but an `.ali` extension. This file is known as the Ada Library Information (ALI) file. The following information is contained in the ALI file.

- Version information (indicates which version of GNAT was used to compile the unit(s) in question)
- Main program information (including priority and time slice settings, as well as the wide character encoding used during compilation).
- List of arguments used in the `gcc` command for the compilation
- Attributes of the unit, including configuration pragmas used, an indication of whether the compilation was successful, exception model used etc.
- A list of relevant restrictions applying to the unit (used for consistency) checking.

- Categorization information (e.g., use of pragma `Pure`).
- Information on all *withed* units, including presence of `Elaborate` or `Elaborate_All` pragmas.
- Information from any `Linker_Options` pragmas used in the unit
- Information on the use of `Body_Version` or `Version` attributes in the unit.
- Dependency information. This is a list of files, together with time stamp and checksum information. These are files on which the unit depends in the sense that recompilation is required if any of these units are modified.
- Cross-reference data. Contains information on all entities referenced in the unit. Used by some tools to provide cross-reference information.

For a full detailed description of the format of the ALI file, see the source of the body of unit `Lib.Writ`, contained in file `lib-writ.adb` in the GNAT compiler sources.

3.8 Binding an Ada Program

When using languages such as C and C++, once the source files have been compiled the only remaining step in building an executable program is linking the object modules together. This means that it is possible to link an inconsistent version of a program, in which two units have included different versions of the same header.

The rules of Ada do not permit such an inconsistent program to be built. For example, if two clients have different versions of the same package, it is illegal to build a program containing these two clients. These rules are enforced by the GNAT binder, which also determines an elaboration order consistent with the Ada rules.

The GNAT binder is run after all the object files for a program have been created. It is given the name of the main program unit, and from this it determines the set of units required by the program, by reading the corresponding ALI files. It generates error messages if the program is inconsistent or if no valid order of elaboration exists.

If no errors are detected, the binder produces a main program, in Ada by default, that contains calls to the elaboration procedures of those compilation unit that require them, followed by a call to the main program. This Ada program is compiled to generate the object file for the main program. The name of the Ada file is `b-xxx.adb` (with the corresponding spec `b-xxx.ads`) where `xxx` is the name of the main program unit.

Finally, the linker is used to build the resulting executable program, using the object from the main program from the bind step as well as the object files for the Ada units of the program.

3.9 GNAT and Libraries

This section describes how to build and use libraries with GNAT, and also shows how to recompile the GNAT run-time library. You should be familiar with the Project Manager facility (see the *GNAT_Project_Manager* chapter of the *GPRbuild User's Guide*) before reading this chapter.

3.9.1 Introduction to Libraries in GNAT

A library is, conceptually, a collection of objects which does not have its own main thread of execution, but rather provides certain services to the applications that use it. A library can be either statically linked with the application, in which case its code is directly included in the application, or, on platforms that support it, be dynamically linked, in which case its code is shared by all applications making use of this library.

GNAT supports both types of libraries. In the static case, the compiled code can be provided in different ways. The simplest approach is to provide directly the set of objects resulting from compilation of the library source files. Alternatively, you can group the objects into an archive using whatever commands are provided by the operating system. For the latter case, the objects are grouped into a shared library.

In the GNAT environment, a library has three types of components:

- Source files,
- ALI files (see [The Ada Library Information Files](#)), and
- Object files, an archive or a shared library.

A GNAT library may expose all its source files, which is useful for documentation purposes. Alternatively, it may expose only the units needed by an external user to make use of the library. That is to say, the specs reflecting the library services along with all the units needed to compile those specs, which can include generic bodies or any body implementing an inlined routine. In the case of *stand-alone libraries* those exposed units are called *interface units* ([Stand-alone Ada Libraries](#)).

All compilation units comprising an application, including those in a library, need to be elaborated in an order partially defined by Ada's semantics. GNAT computes the elaboration order from the ALI files and this is why they constitute a mandatory part of GNAT libraries. *Stand-alone libraries* are the exception to this rule because a specific library elaboration routine is produced independently of the application(s) using the library.

3.9.2 General Ada Libraries

Building a library

The easiest way to build a library is to use the Project Manager, which supports a special type of project called a *Library Project* (see the *Library Projects* section in the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*).

A project is considered a library project, when two project-level attributes are defined in it: `Library_Name` and `Library_Dir`. In order to control different aspects of library configuration, additional optional project-level attributes can be specified:

- **Library_Kind**
This attribute controls whether the library is to be static or dynamic
- **Library_Version**
This attribute specifies the library version; this value is used during dynamic linking of shared libraries to determine if the currently installed versions of the binaries are compatible.
- `Library_Options`
- **Library_GCC**
These attributes specify additional low-level options to be used during library generation, and redefine the actual application used to generate library.

The GNAT Project Manager takes full care of the library maintenance task, including recompilation of the source files for which objects do not exist or are not up to date, assembly of the library archive, and installation of the library (i.e., copying associated source, object and ALI files to the specified location).

Here is a simple library project file:

```
project My_Lib is
  for Source_Dirs use ("src1", "src2");
  for Object_Dir use "obj";
  for Library_Name use "mylib";
  for Library_Dir use "lib";
  for Library_Kind use "dynamic";
end My_Lib;
```

and the compilation command to build and install the library:

```
$ gnatmake -Pmy_lib
```

It is not entirely trivial to perform manually all the steps required to produce a library. We recommend that you use the GNAT Project Manager for this task. In special cases where this is not desired, the necessary steps are discussed below.

There are various possibilities for compiling the units that make up the library: for example with a Makefile ([Using the GNU make Utility](#)) or with a conventional script. For simple libraries, it is also possible to create a dummy main program which depends upon all the packages that comprise

the interface of the library. This dummy main program can then be given to `gnatmake`, which will ensure that all necessary objects are built.

After this task is accomplished, you should follow the standard procedure of the underlying operating system to produce the static or shared library.

Here is an example of such a dummy program:

```
with My_Lib.Service1;
with My_Lib.Service2;
with My_Lib.Service3;
procedure My_Lib_Dummy is
begin
  null;
end;
```

Here are the generic commands that will build an archive or a shared library.

```
# compiling the library
$ gnatmake -c my_lib_dummy.adb

# we don't need the dummy object itself
$ rm my_lib_dummy.o my_lib_dummy.ali

# create an archive with the remaining objects
$ ar rc libmy_lib.a *.o
# some systems may require "ranlib" to be run as well

# or create a shared library
$ gcc -shared -o libmy_lib.so *.o
# some systems may require the code to have been compiled with -fPIC

# remove the object files that are now in the library
$ rm *.o

# Make the ALI files read-only so that gnatmake will not try to
# regenerate the objects that are in the library
$ chmod -w *.ali
```

Please note that the library must have a name of the form `libxxx.a` or `libxxx.so` (or `libxxx.dll` on Windows) in order to be accessed by the directive `-lxxx` at link time.

Installing a library

If you use project files, library installation is part of the library build process (see the *Installing a Library with Project Files* section of the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*).

When project files are not an option, it is also possible, but not recommended, to install the library so that the sources needed to use the library are on the Ada source path and the ALI files & libraries be on the Ada Object path (see [Search Paths](#) and the [Run-Time Library \(RTL\)](#)). Alternatively, the system administrator can place general-purpose libraries in the default compiler paths, by

specifying the libraries' location in the configuration files `ada_source_path` and `ada_object_path`. These configuration files must be located in the GNAT installation tree at the same place as the gcc spec file. The location of the gcc spec file can be determined as follows:

```
$ gcc -v
```

The configuration files mentioned above have a simple format: each line must contain one unique directory name. Those names are added to the corresponding path in their order of appearance in the file. The names can be either absolute or relative; in the latter case, they are relative to where these files are located.

The files `ada_source_path` and `ada_object_path` might not be present in a GNAT installation, in which case, GNAT will look for its run-time library in the directories `adainclude` (for the sources) and `adalib` (for the objects and ALI files). When the files exist, the compiler does not look in `adainclude` and `adalib`, and thus the `ada_source_path` file must contain the location for the GNAT run-time sources (which can simply be `adainclude`). In the same way, the `ada_object_path` file must contain the location for the GNAT run-time objects (which can simply be `adalib`).

You can also specify a new default path to the run-time library at compilation time with the switch `--RTS=rts-path`. You can thus choose / change the run-time library you want your program to be compiled with. This switch is recognized by `gcc`, `gnatmake`, `gnatbind`, `gnatls`, and all project aware tools.

It is possible to install a library before or after the standard GNAT library, by reordering the lines in the configuration files. In general, a library must be installed before the GNAT library if it redefines any part of it.

Using a library

Once again, the project facility greatly simplifies the use of libraries. In this context, using a library is just a matter of adding a *with* clause in the user project. For instance, to make use of the library `My_Lib` shown in examples in earlier sections, you can write:

```
with "my_lib";
project My_Proj is
  ...
end My_Proj;
```

Even if you have a third-party, non-Ada library, you can still use GNAT's Project Manager facility to provide a wrapper for it. For example, the following project, when *withed* by your main project, will link with the third-party library `liba.a`:

```
project Liba is
  for Externally_Built use "true";
  for Source_Files use ();
  for Library_Dir use "lib";
  for Library_Name use "a";
  for Library_Kind use "static";
end Liba;
```

This is an alternative to the use of `pragma Linker_Options`. It is especially interesting in the context of systems with several interdependent static libraries where finding a proper linker order is not easy and best be left to the tools having visibility over project dependence information.

In order to use an Ada library manually, you need to make sure that this library is on both your source and object path (see [Search Paths and the Run-Time Library \(RTL\)](#) and [Search Paths for `gnatbind`](#)). Furthermore, when the objects are grouped in an archive or a shared library, you need to specify the desired library at link time.

For example, you can use the library `mylib` installed in `/dir/my_lib_src` and `/dir/my_lib_obj` with the following commands:

```
$ gnatmake -aI/dir/my_lib_src -aO/dir/my_lib_obj my_appl \\  
-largS -lmy_lib
```

This can be expressed more simply:

```
$ gnatmake my_appl
```

when the following conditions are met:

- `/dir/my_lib_src` has been added by the user to the environment variable `ADA_INCLUDE_PATH`, or by the administrator to the file `ada_source_path`
- `/dir/my_lib_obj` has been added by the user to the environment variable `ADA_OBJECTS_PATH`, or by the administrator to the file `ada_object_path`
- a `pragma Linker_Options` has been added to one of the sources. For example:

```
pragma Linker_Options ("-lmy_lib");
```

Note that you may also load a library dynamically at run time given its filename, as illustrated in the GNAT `plugins` example in the directory `share/examples/gnat/plugins` within the GNAT install area.

3.9.3 Stand-alone Ada Libraries

Introduction to Stand-alone Libraries

A Stand-alone Library (abbreviated ‘SAL’) is a library that contains the necessary code to elaborate the Ada units that are included in the library. In contrast with an ordinary library, which consists of all sources, objects and ALI files of the library, a SAL may specify a restricted subset of compilation units to serve as a library interface. In this case, the fully self-sufficient set of files will normally consist of an objects archive, the sources of interface units’ specs, and the ALI files of interface units. If an interface spec contains a generic unit or an inlined subprogram, the body’s source must also be provided; if the units that must be provided in the source form depend on other units, the source and ALI files of those must also be provided.

The main purpose of a SAL is to minimize the recompilation overhead of client applications when a new version of the library is installed. Specifically, if the interface sources have not changed, client applications do not need to be recompiled. If, furthermore, a SAL is provided in the shared form

and its version, controlled by `Library_Version` attribute, is not changed, then the clients do not need to be relinked.

SALs also allow the library providers to minimize the amount of library source text exposed to the clients. Such 'information hiding' might be useful or necessary for various reasons.

Stand-alone libraries are also well suited to be used in an executable whose main routine is not written in Ada.

Building a Stand-alone Library

GNAT's Project facility provides a simple way of building and installing stand-alone libraries; see the *Stand-alone Library Projects* section in the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*. To be a Stand-alone Library Project, in addition to the two attributes that make a project a Library Project (`Library_Name` and `Library_Dir`; see the *Library Projects* section in the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*), the attribute `Library_Interface` must be defined. For example:

```
for Library_Dir use "lib_dir";
for Library_Name use "dummy";
for Library_Interface use ("intl", "intl.child");
```

Attribute `Library_Interface` has a non-empty string list value, each string in the list designating a unit contained in an immediate source of the project file.

When a Stand-alone Library is built, first the binder is invoked to build a package whose name depends on the library name (`b~dummy.ads/b` in the example above). This binder-generated package includes initialization and finalization procedures whose names depend on the library name (`dummyinit` and `dummyfinal` in the example above). The object corresponding to this package is included in the library.

You must ensure timely (e.g., prior to any use of interfaces in the SAL) calling of these procedures if a static SAL is built, or if a shared SAL is built with the project-level attribute `Library_Auto_Init` set to `"false"`.

For a Stand-Alone Library, only the ALI files of the Interface Units (those that are listed in attribute `Library_Interface`) are copied to the Library Directory. As a consequence, only the Interface Units may be imported from Ada units outside of the library. If other units are imported, the binding phase will fail.

It is also possible to build an encapsulated library where not only the code to elaborate and finalize the library is embedded but also ensuring that the library is linked only against static libraries. So an encapsulated library only depends on system libraries, all other code, including the GNAT runtime, is embedded. To build an encapsulated library the attribute `Library_StandAlone` must be set to `encapsulated`:

```
for Library_Dir use "lib_dir";
for Library_Name use "dummy";
for Library_Kind use "dynamic";
for Library_Interface use ("intl", "intl.child");
for Library_StandAlone use "encapsulated";
```

The default value for this attribute is `standard` in which case a stand-alone library is built.

The attribute `Library_Src_Dir` may be specified for a Stand-Alone Library. `Library_Src_Dir` is a simple attribute that has a single string value. Its value must be the path (absolute or relative to the project directory) of an existing directory. This directory cannot be the object directory or one of the source directories, but it can be the same as the library directory. The sources of the Interface Units of the library that are needed by an Ada client of the library will be copied to the designated directory, called the Interface Copy directory. These sources include the specs of the Interface Units, but they may also include bodies and subunits, when pragmas `Inline` or `Inline_Always` are used, or when there is a generic unit in the spec. Before the sources are copied to the Interface Copy directory, an attempt is made to delete all files in the Interface Copy directory.

Building stand-alone libraries by hand is somewhat tedious, but for those occasions when it is necessary here are the steps that you need to perform:

- Compile all library sources.
- Invoke the binder with the switch `-n` (No Ada main program), with all the ALI files of the interfaces, and with the switch `-L` to give specific names to the `init` and `final` procedures. For example:

```
$ gnatbind -n int1.ali int2.ali -Lsal1
```

- Compile the binder generated file:

```
$ gcc -c b-int2.adb
```

- Link the dynamic library with all the necessary object files, indicating to the linker the names of the `init` (and possibly `final`) procedures for automatic initialization (and finalization). The built library should be placed in a directory different from the object directory.
- Copy the ALI files of the interface to the library directory, add in this copy an indication that it is an interface to a SAL (i.e., add a word `SL` on the line in the ALI file that starts with letter 'P') and make the modified copy of the ALI file read-only.

Using SALs is not different from using other libraries (see [Using a library](#)).

Creating a Stand-alone Library to be used in a non-Ada context

It is easy to adapt the SAL build procedure discussed above for use of a SAL in a non-Ada context.

The only extra step required is to ensure that library interface subprograms are compatible with the main program, by means of `pragma Export` or `pragma Convention`.

Here is an example of simple library interface for use with C main program:

```
package My_Package is
    procedure Do_Something;
    pragma Export (C, Do_Something, "do_something");
    procedure Do_Something_Else;
```

(continues on next page)

(continued from previous page)

```

pragma Export (C, Do_Something_Else, "do_something_else");

end My_Package;

```

On the foreign language side, you must provide a ‘foreign’ view of the library interface; remember that it should contain elaboration routines in addition to interface subprograms.

The example below shows the content of `mylib_interface.h` (note that there is no rule for the naming of this file, any name can be used)

```

/* the library elaboration procedure */
extern void mylibinit (void);

/* the library finalization procedure */
extern void mylibfinal (void);

/* the interface exported by the library */
extern void do_something (void);
extern void do_something_else (void);

```

Libraries built as explained above can be used from any program, provided that the elaboration procedures (named `mylibinit` in the previous example) are called before the library services are used. Any number of libraries can be used simultaneously, as long as the elaboration procedure of each library is called.

Below is an example of a C program that uses the `mylib` library.

```

#include "mylib_interface.h"

int
main (void)
{
    /* First, elaborate the library before using it */
    mylibinit ();

    /* Main program, using the library exported entities */
    do_something ();
    do_something_else ();

    /* Library finalization at the end of the program */
    mylibfinal ();
    return 0;
}

```

Note that invoking any library finalization procedure generated by `gnatbind` shuts down the Ada run-time environment. Consequently, the finalization of all Ada libraries must be performed at the end of the program. No call to these libraries or to the Ada run-time library should be made after the finalization phase.

Note also that special care must be taken with multi-tasks applications. The initialization and finalization routines are not protected against concurrent access. If such requirement is needed it

must be ensured at the application level using a specific operating system services like a mutex or a critical-section.

Restrictions in Stand-alone Libraries

The pragmas listed below should be used with caution inside libraries, as they can create incompatibilities with other Ada libraries:

- pragma Locking_Policy
- pragma Partition_Elaboration_Policy
- pragma Queuing_Policy
- pragma Task_Dispatching_Policy
- pragma Unreserve_All_Interrupts

When using a library that contains such pragmas, the user must make sure that all libraries use the same pragmas with the same values. Otherwise, `Program_Error` will be raised during the elaboration of the conflicting libraries. The usage of these pragmas and its consequences for the user should therefore be well documented.

Similarly, the traceback in the exception occurrence mechanism should be enabled or disabled in a consistent manner across all libraries. Otherwise, `Program_Error` will be raised during the elaboration of the conflicting libraries.

If the `Version` or `Body_Version` attributes are used inside a library, then you need to perform a `gnatbind` step that specifies all ALI files in all libraries, so that version identifiers can be properly computed. In practice these attributes are rarely used, so this is unlikely to be a consideration.

3.9.4 Rebuilding the GNAT Run-Time Library

It may be useful to recompile the GNAT library in various debugging or experimentation contexts. A project file called `libada.gpr` is provided to that effect and can be found in the directory containing the GNAT library. The location of this directory depends on the way the GNAT environment has been installed and can be determined by means of the command:

```
$ gnatls -v
```

The last entry in the source search path usually contains the gnat library (the `adainclude` directory). This project file contains its own documentation and in particular the set of instructions needed to rebuild a new library and to use it.

Note that rebuilding the GNAT Run-Time is only recommended for temporary experiments or debugging, and is not supported.

3.10 Conditional Compilation

This section presents some guidelines for modeling conditional compilation in Ada and describes the gnatprep preprocessor utility.

3.10.1 Modeling Conditional Compilation in Ada

It is often necessary to arrange for a single source program to serve multiple purposes, where it is compiled in different ways to achieve these different goals. Some examples of the need for this feature are

- Adapting a program to a different hardware environment
- Adapting a program to a different target architecture
- Turning debugging features on and off
- Arranging for a program to compile with different compilers

In C, or C++, the typical approach would be to use the preprocessor that is defined as part of the language. The Ada language does not contain such a feature. This is not an oversight, but rather a very deliberate design decision, based on the experience that overuse of the preprocessing features in C and C++ can result in programs that are extremely difficult to maintain. For example, if we have ten switches that can be on or off, this means that there are a thousand separate programs, any one of which might not even be syntactically correct, and even if syntactically correct, the resulting program might not work correctly. Testing all combinations can quickly become impossible.

Nevertheless, the need to tailor programs certainly exists, and in this section we will discuss how this can be achieved using Ada in general, and GNAT in particular.

Use of Boolean Constants

In the case where the difference is simply which code sequence is executed, the cleanest solution is to use Boolean constants to control which code is executed.

```
FP_Initialize_Required : constant Boolean := True;
...
if FP_Initialize_Required then
...
end if;
```

Not only will the code inside the `if` statement not be executed if the constant Boolean is `False`, but it will also be completely deleted from the program. However, the code is only deleted after the `if` statement has been checked for syntactic and semantic correctness. (In contrast, with preprocessors the code is deleted before the compiler ever gets to see it, so it is not checked until the switch is turned on.)

Typically the Boolean constants will be in a separate package, something like:

```
package Config is
  FP_Initialize_Required : constant Boolean := True;
  Reset_Available       : constant Boolean := False;
  ...
end Config;
```

The Config package exists in multiple forms for the various targets, with an appropriate script selecting the version of Config needed. Then any other unit requiring conditional compilation can do a *with* of Config to make the constants visible.

Debugging - A Special Case

A common use of conditional code is to execute statements (for example dynamic checks, or output of intermediate results) under control of a debug switch, so that the debugging behavior can be turned on and off. This can be done using a Boolean constant to control whether the code is active:

```
if Debugging then
  Put_Line ("got to the first stage!");
end if;
```

or

```
if Debugging and then Temperature > 999.0 then
  raise Temperature_Crazy;
end if;
```

Since this is a common case, there are special features to deal with this in a convenient manner. For the case of tests, Ada 2005 has added a pragma `Assert` that can be used for such tests. This pragma is modeled on the `Assert` pragma that has always been available in GNAT, so this feature may be used with GNAT even if you are not using Ada 2005 features. The use of pragma `Assert` is described in the *GNAT_Reference_Manual*, but as an example, the last test could be written:

```
pragma Assert (Temperature <= 999.0, "Temperature Crazy");
```

or simply

```
pragma Assert (Temperature <= 999.0);
```

In both cases, if assertions are active and the temperature is excessive, the exception `Assert_Failure` will be raised, with the given string in the first case or a string indicating the location of the pragma in the second case used as the exception message.

You can turn assertions on and off by using the `Assertion_Policy` pragma.

This is an Ada 2005 pragma which is implemented in all modes by GNAT. Alternatively, you can use the `-gnata` switch to enable assertions from the command line, which applies to all versions of Ada.

For the example above with the `Put_Line`, the GNAT-specific pragma `Debug` can be used:

```
pragma Debug (Put_Line ("got to the first stage!"));
```

If debug pragmas are enabled, the argument, which must be of the form of a procedure call, is executed (in this case, `Put_Line` will be called). Only one call can be present, but of course a special debugging procedure containing any code you like can be included in the program and then called in a `pragma Debug` argument as needed.

One advantage of `pragma Debug` over the `if Debugging then` construct is that `pragma Debug` can appear in declarative contexts, such as at the very beginning of a procedure, before local declarations have been elaborated.

Debug pragmas are enabled using either the `-gnata` switch that also controls assertions, or with a separate `Debug_Policy` pragma.

The latter pragma is new in the Ada 2005 versions of GNAT (but it can be used in Ada 95 and Ada 83 programs as well), and is analogous to `pragma Assertion_Policy` to control assertions.

`Assertion_Policy` and `Debug_Policy` are configuration pragmas, and thus they can appear in `gnat.adc` if you are not using a project file, or in the file designated to contain configuration pragmas in a project file. They then apply to all subsequent compilations. In practice the use of the `-gnata` switch is often the most convenient method of controlling the status of these pragmas.

Note that a pragma is not a statement, so in contexts where a statement sequence is required, you can't just write a pragma on its own. You have to add a `null` statement.

```
if ... then
  ... -- some statements
else
  pragma Assert (Num_Cases < 10);
  null;
end if;
```

Conditionalizing Declarations

In some cases it may be necessary to conditionalize declarations to meet different requirements. For example we might want a bit string whose length is set to meet some hardware message requirement.

This may be possible using declare blocks controlled by conditional constants:

```
if Small_Machine then
  declare
    X : Bit_String (1 .. 10);
  begin
    ...
  end;
else
  declare
    X : Large_Bit_String (1 .. 1000);
  begin
    ...
  end;
end if;
```

(continues on next page)

(continued from previous page)

```
end;  
end if;
```

Note that in this approach, both declarations are analyzed by the compiler so this can only be used where both declarations are legal, even though one of them will not be used.

Another approach is to define integer constants, e.g., `Bits_Per_Word`, or Boolean constants, e.g., `Little_Endian`, and then write declarations that are parameterized by these constants. For example

```
for Rec use  
  Field1 at 0 range Boolean'Pos (Little_Endian) * 10 .. Bits_Per_Word;  
end record;
```

If `Bits_Per_Word` is set to 32, this generates either

```
for Rec use  
  Field1 at 0 range 0 .. 32;  
end record;
```

for the big endian case, or

```
for Rec use record  
  Field1 at 0 range 10 .. 32;  
end record;
```

for the little endian case. Since a powerful subset of Ada expression notation is usable for creating static constants, clever use of this feature can often solve quite difficult problems in conditionalizing compilation (note incidentally that in Ada 95, the little endian constant was introduced as `System.Default_Bit_Order`, so you do not need to define this one yourself).

Use of Alternative Implementations

In some cases, none of the approaches described above are adequate. This can occur for example if the set of declarations required is radically different for two different configurations.

In this situation, the official Ada way of dealing with conditionalizing such code is to write separate units for the different cases. As long as this does not result in excessive duplication of code, this can be done without creating maintenance problems. The approach is to share common code as far as possible, and then isolate the code and declarations that are different. Subunits are often a convenient method for breaking out a piece of a unit that is to be conditionalized, with separate files for different versions of the subunit for different targets, where the build script selects the right one to give to the compiler.

As an example, consider a situation where a new feature in Ada 2005 allows something to be done in a really nice way. But your code must be able to compile with an Ada 95 compiler. Conceptually you want to say:

```
if Ada_2005 then  
  ... neat Ada 2005 code
```

(continues on next page)

(continued from previous page)

```
else
  ... not quite as neat Ada 95 code
end if;
```

where `Ada_2005` is a Boolean constant.

But this won't work when `Ada_2005` is set to `False`, since the `then` clause will be illegal for an Ada 95 compiler. (Recall that although such unreachable code would eventually be deleted by the compiler, it still needs to be legal. If it uses features introduced in Ada 2005, it will be illegal in Ada 95.)

So instead we write

```
procedure Insert is separate;
```

Then we have two files for the subunit `Insert`, with the two sets of code. If the package containing this is called `File_Queries`, then we might have two files

- `file_queries-insert-2005.adb`
- `file_queries-insert-95.adb`

and the build script renames the appropriate file to `file_queries-insert.adb` and then carries out the compilation.

This can also be done with project files' naming schemes. For example:

```
for body ("File_Queries.Insert") use "file_queries-insert-2005.adb";
```

Note also that with project files it is desirable to use a different extension than `ads` / `adb` for alternative versions. Otherwise a naming conflict may arise through another commonly used feature: to declare as part of the project a set of directories containing all the sources obeying the default naming scheme.

The use of alternative units is certainly feasible in all situations, and for example the Ada part of the GNAT run-time is conditionalized based on the target architecture using this approach. As a specific example, consider the implementation of the AST feature in VMS. There is one spec: `s-asthan.ads` which is the same for all architectures, and three bodies:

- **`s-asthan.adb`**
used for all non-VMS operating systems
- **`s-asthan-vms-alpha.adb`**
used for VMS on the Alpha
- **`s-asthan-vms-ia64.adb`**
used for VMS on the ia64

The dummy version `s-asthan.adb` simply raises exceptions noting that this operating system feature is not available, and the two remaining versions interface with the corresponding versions of VMS to provide VMS-compatible AST handling. The GNAT build script knows the architecture and operating system, and automatically selects the right version, renaming it if necessary to `s-asthan.adb` before the run-time build.

Another style for arranging alternative implementations is through Ada's access-to-subprogram facility. In case some functionality is to be conditionally included, you can declare an access-to-procedure variable `Ref` that is initialized to designate a 'do nothing' procedure, and then invoke `Ref.all` when appropriate. In some library package, set `Ref` to `Proc'Access` for some procedure `Proc` that performs the relevant processing. The initialization only occurs if the library package is included in the program. The same idea can also be implemented using tagged types and dispatching calls.

Preprocessing

Although it is quite possible to conditionalize code without the use of C-style preprocessing, as described earlier in this section, it is nevertheless convenient in some cases to use the C approach. Moreover, older Ada compilers have often provided some preprocessing capability, so legacy code may depend on this approach, even though it is not standard.

To accommodate such use, GNAT provides a preprocessor (modeled to a large extent on the various preprocessors that have been used with legacy code on other compilers, to enable easier transition).

The preprocessor may be used in two separate modes. It can be used quite separately from the compiler, to generate a separate output source file that is then fed to the compiler as a separate step. This is the `gnatprep` utility, whose use is fully described in [Preprocessing with gnatprep](#).

The preprocessing language allows such constructs as

```
#if DEBUG or else (PRIORITY > 4) then
  sequence of declarations
#else
  completely different sequence of declarations
#end if;
```

The values of the symbols `DEBUG` and `PRIORITY` can be defined either on the command line or in a separate file.

The other way of running the preprocessor is even closer to the C style and often more convenient. In this approach the preprocessing is integrated into the compilation process. The compiler is given the preprocessor input which includes `#if` lines etc, and then the compiler carries out the preprocessing internally and processes the resulting output. For more details on this approach, see [Integrated Preprocessing](#).

3.10.2 Preprocessing with gnatprep

This section discusses how to use GNAT's `gnatprep` utility for simple preprocessing. Although designed for use with GNAT, `gnatprep` does not depend on any special GNAT features. For further discussion of conditional compilation in general, see [Conditional Compilation](#).

Preprocessing Symbols

Preprocessing symbols are defined in *definition files* and referenced in the sources to be preprocessed. A preprocessing symbol is an identifier, following normal Ada (case-insensitive) rules for its syntax, with the restriction that all characters need to be in the ASCII set (no accented letters).

Using gnatprep

To call gnatprep use:

```
$ gnatprep [ switches ] infile outfile [ deffile ]
```

where

- *switches*
is an optional sequence of switches as described in the next section.
- *infile*
is the full name of the input file, which is an Ada source file containing preprocessor directives.
- *outfile*
is the full name of the output file, which is an Ada source in standard Ada form. When used with GNAT, this file name will normally have an `ads` or `adb` suffix.
- *deffile*
is the full name of a text file containing definitions of preprocessing symbols to be referenced by the preprocessor. This argument is optional, and can be replaced by the use of the `-D` switch.

Switches for gnatprep

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage and then exit disregarding all other options.

-b

Causes both preprocessor lines and the lines deleted by preprocessing to be replaced by blank lines in the output source file, preserving line numbers in the output file.

-c

Causes both preprocessor lines and the lines deleted by preprocessing to be retained in the output source as comments marked with the special string `--!` . This option will result in line numbers being preserved in the output file.

-C

Causes comments to be scanned. Normally comments are ignored by gnatprep. If this option is specified, then comments are scanned and any `$symbol` substitutions performed as in program text. This is particularly useful when structured comments are used (e.g., for

programs written in a pre-2014 version of the SPARK Ada subset). Note that this switch is not available when doing integrated preprocessing (it would be useless in this context since comments are ignored by the compiler in any case).

-Dsymbol[=value]

Defines a new preprocessing symbol with the specified value. If no value is given on the command line, then symbol is considered to be True. This switch can be used in place of a definition file.

-r

Causes a `Source_Reference` pragma to be generated that references the original input file, so that error messages will use the file name of this original file. The use of this switch implies that preprocessor lines are not to be removed from the file, so its use will force `-b` mode if `-c` has not been specified explicitly.

Note that if the file to be preprocessed contains multiple units, then it will be necessary to `gnatchop` the output file from `gnatprep`. If a `Source_Reference` pragma is present in the preprocessed file, it will be respected by `gnatchop -r` so that the final chopped files will correctly refer to the original input source file for `gnatprep`.

-s

Causes a sorted list of symbol names and values to be listed on the standard output file.

-T

Use LF as line terminators when writing files. By default the line terminator of the host (LF under unix, CR/LF under Windows) is used.

-u

Causes undefined symbols to be treated as having the value FALSE in the context of a preprocessor test. In the absence of this option, an undefined symbol in a `#if` or `#elsif` test will be treated as an error.

-v

Verbose mode: generates more output about work done.

Note: if neither `-b` nor `-c` is present, then preprocessor lines and deleted lines are completely removed from the output, unless `-r` is specified, in which case `-b` is assumed.

Form of Definitions File

The definitions file contains lines of the form:

```
symbol := value
```

where `symbol` is a preprocessing symbol, and `value` is one of the following:

- Empty, corresponding to a null substitution,
- A string literal using normal Ada syntax, or
- Any sequence of characters from the set {letters, digits, period, underline}.

Comment lines may also appear in the definitions file, starting with the usual `--`, and comments may be added to the definitions lines.

Form of Input Text for `gnatprep`

The input text may contain preprocessor conditional inclusion lines, as well as general symbol substitution sequences.

The preprocessor conditional inclusion commands have the form:

```
#if <expression> [then]
  lines
#elif <expression> [then]
  lines
#elif <expression> [then]
  lines
...
#else
  lines
#endif;
```

In this example, `<expression>` is defined by the following grammar:

```
<expression> ::= <symbol>
<expression> ::= <symbol> = "<value>"
<expression> ::= <symbol> = <symbol>
<expression> ::= <symbol> = <integer>
<expression> ::= <symbol> > <integer>
<expression> ::= <symbol> >= <integer>
<expression> ::= <symbol> < <integer>
<expression> ::= <symbol> <= <integer>
<expression> ::= <symbol> 'Defined
<expression> ::= not <expression>
<expression> ::= <expression> and <expression>
<expression> ::= <expression> or <expression>
<expression> ::= <expression> and then <expression>
<expression> ::= <expression> or else <expression>
<expression> ::= ( <expression> )
```

Note the following restriction: it is not allowed to have “and” or “or” following “not” in the same expression without parentheses. For example, this is not allowed:

```
not X or Y
```

This can be expressed instead as one of the following forms:

```
(not X) or Y
not (X or Y)
```

For the first test (`<expression> ::= <symbol>`) the symbol must have either the value true or false, that is to say the right-hand of the symbol definition must be one of the (case-insensitive) literals

True or **False**. If the value is true, then the corresponding lines are included, and if the value is false, they are excluded.

When comparing a symbol to an integer, the integer is any non negative literal integer as defined in the Ada Reference Manual, such as 3, 16#FF# or 2#11#. The symbol value must also be a non negative integer. Integer values in the range 0 .. 2**31-1 are supported.

The test (<expression> ::= <symbol>'Defined) is true only if the symbol has been defined in the definition file or by a **-D** switch on the command line. Otherwise, the test is false.

The equality tests are case insensitive, as are all the preprocessor lines.

If the symbol referenced is not defined in the symbol definitions file, then the effect depends on whether or not switch **-u** is specified. If so, then the symbol is treated as if it had the value **false** and the test fails. If this switch is not specified, then it is an error to reference an undefined symbol. It is also an error to reference a symbol that is defined with a value other than **True** or **False**.

The use of the **not** operator inverts the sense of this logical test. The **not** operator cannot be combined with the **or** or **and** operators, without parentheses. For example, "if not X or Y then" is not allowed, but "if (not X) or Y then" and "if not (X or Y) then" are.

The **then** keyword is optional as shown

The **#** must be the first non-blank character on a line, but otherwise the format is free form. Spaces or tabs may appear between the **#** and the keyword. The keywords and the symbols are case insensitive as in normal Ada code. Comments may be used on a preprocessor line, but other than that, no other tokens may appear on a preprocessor line. Any number of **elsif** clauses can be present, including none at all. The **else** is optional, as in Ada.

The **#** marking the start of a preprocessor line must be the first non-blank character on the line, i.e., it must be preceded only by spaces or horizontal tabs.

Symbol substitution outside of preprocessor lines is obtained by using the sequence:

```
$symbol
```

anywhere within a source line, except in a comment or within a string literal. The identifier following the **\$** must match one of the symbols defined in the symbol definition file, and the result is to substitute the value of the symbol in place of **\$symbol** in the output file.

Note that although the substitution of strings within a string literal is not possible, it is possible to have a symbol whose defined value is a string literal. So instead of setting XYZ to **hello** and writing:

```
Header : String := "$XYZ";
```

you should set XYZ to **"hello"** and write:

```
Header : String := $XYZ;
```

and then the substitution will occur as desired.

3.10.3 Integrated Preprocessing

As noted above, a file to be preprocessed consists of Ada source code in which preprocessing lines have been inserted. However, instead of using `gnatprep` to explicitly preprocess a file as a separate step before compilation, you can carry out the preprocessing implicitly as part of compilation. Such *integrated preprocessing*, which is the common style with C, is performed when either or both of the following switches are passed to the compiler:

- `-gnatp`, which specifies the *preprocessor data file*. This file dictates how the source files will be preprocessed (e.g., which symbol definition files apply to which sources).
- `-gnateD`, which defines values for preprocessing symbols.

Integrated preprocessing applies only to Ada source files, it is not available for configuration pragma files.

With integrated preprocessing, the output from the preprocessor is not, by default, written to any external file. Instead it is passed internally to the compiler. To preserve the result of preprocessing in a file, either run `gnatprep` in standalone mode or else supply the `-gnateG` switch (described below) to the compiler.

When using project files:

- the builder switch `-x` should be used if any Ada source is compiled with `gnatp=`, so that the compiler finds the *preprocessor data file*.
- the preprocessing data file and the symbol definition files should be located in the source directories of the project.

Note that the `gnatmake` switch `-m` will almost always trigger recompilation for sources that are preprocessed, because `gnatmake` cannot compute the checksum of the source after preprocessing.

The actual preprocessing function is described in detail in [Preprocessing with gnatprep](#). This section explains the switches that relate to integrated preprocessing.

`-gnatp=preprocessor_data_file`

This switch specifies the file name (without directory information) of the preprocessor data file. Either place this file in one of the source directories, or, when using project files, reference the project file's directory via the `project_name'Project_Dir` project attribute; e.g:

```
project Prj is
  package Compiler is
    for Switches ("Ada") use
      ("-gnatp=" & Prj'Project_Dir & "prep.def");
    end Compiler;
end Prj;
```

A preprocessor data file is a text file that contains *preprocessor control lines*. A preprocessor control line directs the preprocessing of either a particular source file, or, analogous to `others` in Ada, all sources not specified elsewhere in the preprocessor data file. A preprocessor control line can optionally identify a *definition file* that assigns values to preprocessor symbols, as well as a list of switches that relate to preprocessing. Empty lines and comments (using Ada syntax) are also permitted, with no semantic effect.

Here's an example of a preprocessor data file:

```
"toto.adb" "prep.def" -u
-- Preprocess toto.adb, using definition file prep.def
-- Undefined symbols are treated as False

* -c -DVERSION=V101
-- Preprocess all other sources without using a definition file
-- Suppressed lined are commented
-- Symbol VERSION has the value V101

"tata.adb" "prep2.def" -s
-- Preprocess tata.adb, using definition file prep2.def
-- List all symbols with their values
```

A preprocessor control line has the following syntax:

```
<preprocessor_control_line> ::=
  <preprocessor_input> [ <definition_file_name> ] { <switch> }

<preprocessor_input> ::= <source_file_name> | '*'

<definition_file_name> ::= <string_literal>

<source_file_name> ::= <string_literal>

<switch> := (See below for list)
```

Thus each preprocessor control line starts with either a literal string or the character '*':

- A literal string is the file name (without directory information) of the source file that will be input to the preprocessor.
- The character '*' is a wild-card indicator; the additional parameters on the line indicate the preprocessing for all the sources that are not specified explicitly on other lines (the order of the lines is not significant).

It is an error to have two lines with the same file name or two lines starting with the character '*'.

After the file name or '*', an optional literal string specifies the name of the definition file to be used for preprocessing ([Form of Definitions File](#)). The definition files are found by the compiler in one of the source directories. In some cases, when compiling a source in a directory other than the current directory, if the definition file is in the current directory, it may be necessary to add the current directory as a source directory through the `-I` switch; otherwise the compiler would not find the definition file.

Finally, switches similar to those of `gnatprep` may optionally appear:

-b

Causes both preprocessor lines and the lines deleted by preprocessing to be replaced by blank lines, preserving the line number. This switch is always implied; however, if specified after `-c` it cancels the effect of `-c`.

-c

Causes both preprocessor lines and the lines deleted by preprocessing to be retained as comments marked with the special string `'-!'`.

-Dsymbol=new_value

Define or redefine `symbol` to have `new_value` as its value. The permitted form for `symbol` is either an Ada identifier, or any Ada reserved word aside from `if`, `else`, `elsif`, `end`, `and`, or `and then`. The permitted form for `new_value` is a literal string, an Ada identifier or any Ada reserved word. A symbol declared with this switch replaces a symbol with the same name defined in a definition file.

-s

Causes a sorted list of symbol names and values to be listed on the standard output file.

-u

Causes undefined symbols to be treated as having the value `FALSE` in the context of a preprocessor test. In the absence of this option, an undefined symbol in a `#if` or `#elsif` test will be treated as an error.

-gnateDsymbol[=new_value]

Define or redefine `symbol` to have `new_value` as its value. If no value is supplied, then the value of `symbol` is `True`. The form of `symbol` is an identifier, following normal Ada (case-insensitive) rules for its syntax, and `new_value` is either an arbitrary string between double quotes or any sequence (including an empty sequence) of characters from the set (letters, digits, period, underline). Ada reserved words may be used as symbols, with the exceptions of `if`, `else`, `elsif`, `end`, `and`, or `and then`.

Examples:

```
-gnateDToto=Tata
-gnateDFoo
-gnateDFoo=\"Foo-Bar\"
```

A symbol declared with this switch on the command line replaces a symbol with the same name either in a definition file or specified with a switch `-D` in the preprocessor data file.

This switch is similar to switch `-D` of `gnatprep`.

-gnateG

When integrated preprocessing is performed on source file `filename.extension`, create or overwrite `filename.extension.prep` to contain the result of the preprocessing. For example if the source file is `foo.adb` then the output file will be `foo.adb.prep`.

3.11 Mixed Language Programming

This section describes how to develop a mixed-language program, with a focus on combining Ada with C or C++.

3.11.1 Interfacing to C

Interfacing Ada with a foreign language such as C involves using compiler directives to import and/or export entity definitions in each language – using `extern` statements in C, for instance, and the `Import`, `Export`, and `Convention` pragmas in Ada. A full treatment of these topics is provided in Appendix B, section 1 of the Ada Reference Manual.

There are two ways to build a program using GNAT that contains some Ada sources and some foreign language sources, depending on whether or not the main subprogram is written in Ada. Here is a source example with the main subprogram in Ada:

```
/* file1.c */
#include <stdio.h>

void print_num (int num)
{
    printf ("num is %d.\n", num);
    return;
}
```

```
/* file2.c */

/* num_from_Ada is declared in my_main.adb */
extern int num_from_Ada;

int get_num (void)
{
    return num_from_Ada;
}
```

```
-- my_main.adb
procedure My_Main is

    -- Declare then export an Integer entity called num_from_Ada
    My_Num : Integer := 10;
    pragma Export (C, My_Num, "num_from_Ada");

    -- Declare an Ada function spec for Get_Num, then use
    -- C function get_num for the implementation.
    function Get_Num return Integer;
    pragma Import (C, Get_Num, "get_num");

    -- Declare an Ada procedure spec for Print_Num, then use
    -- C function print_num for the implementation.
    procedure Print_Num (Num : Integer);
```

(continues on next page)

(continued from previous page)

```

pragma Import (C, Print_Num, "print_num");

begin
  Print_Num (Get_Num);
end My_Main;

```

To build this example:

- First compile the foreign language files to generate object files:

```

$ gcc -c file1.c
$ gcc -c file2.c

```

- Then, compile the Ada units to produce a set of object files and ALI files:

```

$ gnatmake -c my_main.adb

```

- Run the Ada binder on the Ada main program:

```

$ gnatbind my_main.ali

```

- Link the Ada main program, the Ada objects and the other language objects:

```

$ gnatlink my_main.ali file1.o file2.o

```

The last three steps can be grouped in a single command:

```

$ gnatmake my_main.adb -larges file1.o file2.o

```

If the main program is in a language other than Ada, then you may have more than one entry point into the Ada subsystem. You must use a special binder option to generate callable routines that initialize and finalize the Ada units ([Binding with Non-Ada Main Programs](#)). Calls to the initialization and finalization routines must be inserted in the main program, or some other appropriate point in the code. The call to initialize the Ada units must occur before the first Ada subprogram is called, and the call to finalize the Ada units must occur after the last Ada subprogram returns. The binder will place the initialization and finalization subprograms into the `b~xxx.adb` file where they can be accessed by your C sources. To illustrate, we have the following example:

```

/* main.c */
extern void adainit (void);
extern void adafinal (void);
extern int add (int, int);
extern int sub (int, int);

int main (int argc, char *argv[])
{
  int a = 21, b = 7;

  adainit();

```

(continues on next page)

(continued from previous page)

```
/* Should print "21 + 7 = 28" */
printf ("%d + %d = %d\\n", a, b, add (a, b));

/* Should print "21 - 7 = 14" */
printf ("%d - %d = %d\\n", a, b, sub (a, b));

adafinal();
}
```

```
-- unit1.ads
package Unit1 is
  function Add (A, B : Integer) return Integer;
  pragma Export (C, Add, "add");
end Unit1;
```

```
-- unit1.adb
package body Unit1 is
  function Add (A, B : Integer) return Integer is
  begin
    return A + B;
  end Add;
end Unit1;
```

```
-- unit2.ads
package Unit2 is
  function Sub (A, B : Integer) return Integer;
  pragma Export (C, Sub, "sub");
end Unit2;
```

```
-- unit2.adb
package body Unit2 is
  function Sub (A, B : Integer) return Integer is
  begin
    return A - B;
  end Sub;
end Unit2;
```

The build procedure for this application is similar to the last example's:

- First, compile the foreign language files to generate object files:

```
$ gcc -c main.c
```

- Next, compile the Ada units to produce a set of object files and ALI files:

```
$ gnatmake -c unit1.adb
$ gnatmake -c unit2.adb
```

- Run the Ada binder on every generated ALI file. Make sure to use the `-n` option to specify a foreign main program:

```
$ gnatbind -n unit1.ali unit2.ali
```

- Link the Ada main program, the Ada objects and the foreign language objects. You need only list the last ALI file here:

```
$ gnatlink unit2.ali main.o -o exec_file
```

This procedure yields a binary executable called `exec_file`.

Depending on the circumstances (for example when your non-Ada main object does not provide symbol `main`), you may also need to instruct the GNAT linker not to include the standard startup objects by passing the `-nostartfiles` switch to `gnatlink`.

3.11.2 Calling Conventions

GNAT follows standard calling sequence conventions and will thus interface to any other language that also follows these conventions. The following Convention identifiers are recognized by GNAT:

Ada

This indicates that the standard Ada calling sequence will be used and all Ada data items may be passed without any limitations in the case where GNAT is used to generate both the caller and callee. It is also possible to mix GNAT generated code and code generated by another Ada compiler. In this case, the data types should be restricted to simple cases, including primitive types. Whether complex data types can be passed depends on the situation. Probably it is safe to pass simple arrays, such as arrays of integers or floats. Records may or may not work, depending on whether both compilers lay them out identically. Complex structures involving variant records, access parameters, tasks, or protected types, are unlikely to be able to be passed.

Note that in the case of GNAT running on a platform that supports HP Ada 83, a higher degree of compatibility can be guaranteed, and in particular records are laid out in an identical manner in the two compilers. Note also that if output from two different compilers is mixed, the program is responsible for dealing with elaboration issues. Probably the safest approach is to write the main program in the version of Ada other than GNAT, so that it takes care of its own elaboration requirements, and then call the GNAT-generated `adainit` procedure to ensure elaboration of the GNAT components. Consult the documentation of the other Ada compiler for further details on elaboration.

However, it is not possible to mix the tasking run time of GNAT and HP Ada 83, All the tasking operations must either be entirely within GNAT compiled sections of the program, or entirely within HP Ada 83 compiled sections of the program.

Assembler

Specifies assembler as the convention. In practice this has the same effect as convention `Ada` (but is not equivalent in the sense of being considered the same convention).

Asm

Equivalent to `Assembler`.

COBOL

Data will be passed according to the conventions described in section B.4 of the Ada Reference Manual.

C

Data will be passed according to the conventions described in section B.3 of the Ada Reference Manual.

A note on interfacing to a C 'varargs' function:

In C, `varargs` allows a function to take a variable number of arguments. There is no direct equivalent in this to Ada. One approach that can be used is to create a C wrapper for each different profile and then interface to this C wrapper. For example, to print an `int` value using `printf`, create a C function `printfi` that takes two arguments, a pointer to a string and an `int`, and calls `printf`. Then in the Ada program, use `pragma Import` to interface to `printfi`.

It may work on some platforms to directly interface to a `varargs` function by providing a specific Ada profile for a particular call. However, this does not work on all platforms, since there is no guarantee that the calling sequence for a two argument normal C function is the same as for calling a `varargs` C function with the same two arguments.

Default

Equivalent to C.

External

Equivalent to C.

C_Plus_Plus (or CPP)

This stands for C++. For most purposes this is identical to C. See the separate description of the specialized GNAT pragmas relating to C++ interfacing for further details.

Fortran

Data will be passed according to the conventions described in section B.5 of the Ada Reference Manual.

Intrinsic

This applies to an intrinsic operation, as defined in the Ada Reference Manual. If a `pragma Import (Intrinsic)` applies to a subprogram, this means that the body of the subprogram is provided by the compiler itself, usually by means of an efficient code sequence, and that the user does not supply an explicit body for it. In an application program, the pragma may be applied to the following sets of names:

- `Rotate_Left`, `Rotate_Right`, `Shift_Left`, `Shift_Right`, `Shift_Right_Arithmetic`. The corresponding subprogram declaration must have two formal parameters. The first one must be a signed integer type or a modular type with a binary modulus, and the second parameter must be of type `Natural`. The return type must be the same as the type of the first argument. The size of this type can only be 8, 16, 32, or 64.
- Binary arithmetic operators: `'+'`, `'-'`, `'*'`, `'/'`. The corresponding operator declaration must have parameters and result type that have the same root numeric type (for example,

all three are `long_float` types). This simplifies the definition of operations that use type checking to perform dimensional checks:

```
type Distance is new Long_Float;
type Time     is new Long_Float;
type Velocity is new Long_Float;
function "/" (D : Distance; T : Time)
  return Velocity;
pragma Import (Intrinsic, "/");
```

This common idiom is often programmed with a generic definition and an explicit body. The `pragma` makes it simpler to introduce such declarations. It incurs no overhead in compilation time or code size, because it is implemented as a single machine instruction.

- General subprogram entities. This is used to bind an Ada subprogram declaration to a compiler builtin by name with back-ends where such interfaces are available. A typical example is the set of `__builtin` functions exposed by the GCC back-end, as in the following example:

```
function builtin_sqrt (F : Float) return Float;
pragma Import (Intrinsic, builtin_sqrt, "__builtin_sqrtf");
```

Most of the GCC builtins are accessible this way, and as for other import conventions (e.g. C), it is the user's responsibility to ensure that the Ada subprogram profile matches the underlying builtin expectations.

Stdcall

This is relevant only to Windows implementations of GNAT, and specifies that the `Stdcall` calling sequence will be used, as defined by the NT API. Nevertheless, to ease building cross-platform bindings this convention will be handled as a C calling convention on non-Windows platforms.

DLL

This is equivalent to `Stdcall`.

Win32

This is equivalent to `Stdcall`.

Stubbed

This is a special convention that indicates that the compiler should provide a stub body that raises `Program_Error`.

GNAT additionally provides a useful `pragma Convention_Identifier` that can be used to parameterize conventions and allow additional synonyms to be specified. For example if you have legacy code in which the convention identifier `Fortran77` was used for Fortran, you can use the configuration `pragma`:

```
pragma Convention_Identifier (Fortran77, Fortran);
```

And from now on the identifier `Fortran77` may be used as a convention identifier (for example in an `Import` `pragma`) with the same meaning as `Fortran`.

3.11.3 Building Mixed Ada and C++ Programs

A programmer inexperienced with mixed-language development may find that building an application containing both Ada and C++ code can be a challenge. This section gives a few hints that should make this task easier.

Interfacing to C++

GNAT supports interfacing with the G++ compiler (or any C++ compiler generating code that is compatible with the G++ Application Binary Interface —see <http://itanium-cxx-abi.github.io/cxx-abi/abi.html>).

Interfacing can be done at 3 levels: simple data, subprograms, and classes. In the first two cases, GNAT offers a specific `Convention C_Plus_Plus` (or `CPP`) that behaves exactly like `Convention C`. Usually, C++ mangles the names of subprograms. To generate proper mangled names automatically, see [Generating Ada Bindings for C and C++ headers](#)). This problem can also be addressed manually in two ways:

- by modifying the C++ code in order to force a C convention using the `extern "C"` syntax.
- by figuring out the mangled name (using e.g. `nm`) and using it as the `Link_Name` argument of the `pragma import`.

Interfacing at the class level can be achieved by using the GNAT specific pragmas such as `CPP_Constructor`. See the *GNAT_Reference_Manual* for additional information.

Linking a Mixed C++ & Ada Program

Usually the linker of the C++ development system must be used to link mixed applications because most C++ systems will resolve elaboration issues (such as calling constructors on global class instances) transparently during the link phase. GNAT has been adapted to ease the use of a foreign linker for the last phase. Three cases can be considered:

- Using GNAT and G++ (GNU C++ compiler) from the same GCC installation: The C++ linker can simply be called by using the C++ specific driver called `g++`.

Note that if the C++ code uses inline functions, you will need to compile your C++ code with the `-fkeep-inline-functions` switch in order to provide an existing function implementation that the Ada code can link with.

```
$ g++ -c -fkeep-inline-functions file1.C
$ g++ -c -fkeep-inline-functions file2.C
$ gnatmake ada_unit -largS file1.o file2.o --LINK=g++
```

- Using GNAT and G++ from two different GCC installations: If both compilers are on the `:envvar'PATH'`, the previous method may be used. It is important to note that environment variables such as `C_INCLUDE_PATH`, `GCC_EXEC_PREFIX`, `BINUTILS_ROOT`, and `GCC_ROOT` will affect both compilers at the same time and may make one of the two compilers operate improperly if set during invocation of the wrong compiler. It is also very important that the linker uses the proper `libgcc.a` GCC library – that is, the one from the C++ compiler installation. The

implicit link command as suggested in the `gnatmake` command from the former example can be replaced by an explicit link command with the full-verbosity option in order to verify which library is used:

```
$ gnatbind ada_unit
$ gnatlink -v -v ada_unit file1.o file2.o --LINK=c++
```

If there is a problem due to interfering environment variables, it can be worked around by using an intermediate script. The following example shows the proper script to use when GNAT has not been installed at its default location and `g++` has been installed at its default location:

```
$ cat ./my_script
#!/bin/sh
unset BINUTILS_ROOT
unset GCC_ROOT
c++ $*
$ gnatlink -v -v ada_unit file1.o file2.o --LINK=./my_script
```

- Using a non-GNU C++ compiler: The commands previously described can be used to insure that the C++ linker is used. Nonetheless, you need to add a few more parameters to the link command line, depending on the exception mechanism used.

If the `setjmp` / `longjmp` exception mechanism is used, only the paths to the `libgcc` libraries are required:

```
$ cat ./my_script
#!/bin/sh
CC $* gcc -print-file-name=libgcc.a gcc -print-file-name=libgcc_eh.a
$ gnatlink ada_unit file1.o file2.o --LINK=./my_script
```

where `CC` is the name of the non-GNU C++ compiler.

If the “zero cost” exception mechanism is used, and the platform supports automatic registration of exception tables (e.g., Solaris), paths to more objects are required:

```
$ cat ./my_script
#!/bin/sh
CC gcc -print-file-name=crtbegin.o $* \
gcc -print-file-name=libgcc.a gcc -print-file-name=libgcc_eh.a \
gcc -print-file-name=crtend.o
$ gnatlink ada_unit file1.o file2.o --LINK=./my_script
```

If the “zero cost exception” mechanism is used, and the platform doesn't support automatic registration of exception tables (e.g., HP-UX or AIX), the simple approach described above will not work and a pre-linking phase using GNAT will be necessary.

Another alternative is to use the `gprbuild` multi-language builder which has a large knowledge base and knows how to link Ada and C++ code together automatically in most cases.

A Simple Example

The following example, provided as part of the GNAT examples, shows how to achieve procedural interfacing between Ada and C++ in both directions. The C++ class A has two methods. The first method is exported to Ada by the means of an extern C wrapper function. The second method calls an Ada subprogram. On the Ada side, the C++ calls are modelled by a limited record with a layout comparable to the C++ class. The Ada subprogram, in turn, calls the C++ method. So, starting from the C++ main program, the process passes back and forth between the two languages.

Here are the compilation commands:

```
$ gnatmake -c simple_cpp_interface
$ g++ -c cpp_main.C
$ g++ -c ex7.C
$ gnatbind -n simple_cpp_interface
$ gnatlink simple_cpp_interface -o cpp_main --LINK=g++ -lstdc++ ex7.o cpp_main.o
```

Here are the corresponding sources:

```
//cpp_main.C

#include "ex7.h"

extern "C" {
    void adainit (void);
    void adafinal (void);
    void method1 (A *t);
}

void method1 (A *t)
{
    t->method1 ();
}

int main ()
{
    A obj;
    adainit ();
    obj.method2 (3030);
    adafinal ();
}
```

```
//ex7.h

class Origin {
public:
    int o_value;
};

class A : public Origin {
public:
    void method1 (void);
    void method2 (int v);
}
```

(continues on next page)

(continued from previous page)

```
A();
  int  a_value;
};
```

```
//ex7.C

#include "ex7.h"
#include <stdio.h>

extern "C" { void ada_method2 (A *t, int v);}

void A::method1 (void)
{
  a_value = 2020;
  printf ("in A::method1, a_value = %d \\n",a_value);
}

void A::method2 (int v)
{
  ada_method2 (this, v);
  printf ("in A::method2, a_value = %d \\n",a_value);
}

A::A(void)
{
  a_value = 1010;
  printf ("in A::A, a_value = %d \\n",a_value);
}
```

```
-- simple_cpp_interface.ads
with System;
package Simple_Cpp_Interface is
  type A is limited
    record
      Vptr      : System.Address;
      O_Value  : Integer;
      A_Value  : Integer;
    end record;
  pragma Convention (C, A);

  procedure Method1 (This : in out A);
  pragma Import (C, Method1);

  procedure Ada_Method2 (This : in out A; V : Integer);
  pragma Export (C, Ada_Method2);
end Simple_Cpp_Interface;
```

```
-- simple_cpp_interface.adb
package body Simple_Cpp_Interface is
```

(continues on next page)

(continued from previous page)

```

procedure Ada_Method2 (This : in out A; V : Integer) is
begin
  Method1 (This);
  This.A_Value := V;
end Ada_Method2;

end Simple_Cpp_Interface;

```

Interfacing with C++ constructors

In order to interface with C++ constructors GNAT provides the pragma `CPP_Constructor` (see the *GNAT_Reference_Manual* for additional information). In this section we present some common uses of C++ constructors in mixed-languages programs in GNAT.

Let us assume that we need to interface with the following C++ class:

```

class Root {
public:
  int a_value;
  int b_value;
  virtual int Get_Value ();
  Root();           // Default constructor
  Root(int v);     // 1st non-default constructor
  Root(int v, int w); // 2nd non-default constructor
};

```

For this purpose we can write the following package spec (further information on how to build this spec is available in [Interfacing with C++ at the Class Level](#) and [Generating Ada Bindings for C and C++ headers](#)).

```

with Interfaces.C; use Interfaces.C;
package Pkg_Root is
  type Root is tagged limited record
    A_Value : int;
    B_Value : int;
  end record;
  pragma Import (CPP, Root);

  function Get_Value (Obj : Root) return int;
  pragma Import (CPP, Get_Value);

  function Constructor return Root;
  pragma Cpp_Constructor (Constructor, "_ZN4RootC1Ev");

  function Constructor (v : Integer) return Root;
  pragma Cpp_Constructor (Constructor, "_ZN4RootC1Ei");

  function Constructor (v, w : Integer) return Root;

```

(continues on next page)

(continued from previous page)

```
pragma Cpp_Constructor (Constructor, "_ZN4RootC1Eii");
end Pkg_Root;
```

On the Ada side the constructor is represented by a function (whose name is arbitrary) that returns the classwide type corresponding to the imported C++ class. Although the constructor is described as a function, it is typically a procedure with an extra implicit argument (the object being initialized) at the implementation level. GNAT issues the appropriate call, whatever it is, to get the object properly initialized.

Constructors can only appear in the following contexts:

- On the right side of an initialization of an object of type T.
- On the right side of an initialization of a record component of type T.
- In an Ada 2005 limited aggregate.
- In an Ada 2005 nested limited aggregate.
- In an Ada 2005 limited aggregate that initializes an object built in place by an extended return statement.

In a declaration of an object whose type is a class imported from C++, either the default C++ constructor is implicitly called by GNAT, or else the required C++ constructor must be explicitly called in the expression that initializes the object. For example:

```
Obj1 : Root;
Obj2 : Root := Constructor;
Obj3 : Root := Constructor (v => 10);
Obj4 : Root := Constructor (30, 40);
```

The first two declarations are equivalent: in both cases the default C++ constructor is invoked (in the former case the call to the constructor is implicit, and in the latter case the call is explicit in the object declaration). `Obj3` is initialized by the C++ non-default constructor that takes an integer argument, and `Obj4` is initialized by the non-default C++ constructor that takes two integers.

Let us derive the imported C++ class in the Ada side. For example:

```
type DT is new Root with record
  C_Value : Natural := 2009;
end record;
```

In this case the components DT inherited from the C++ side must be initialized by a C++ constructor, and the additional Ada components of type DT are initialized by GNAT. The initialization of such an object is done either by default, or by means of a function returning an aggregate of type DT, or by means of an extension aggregate.

```
Obj5 : DT;
Obj6 : DT := Function_Returning_DT (50);
Obj7 : DT := (Constructor (30,40) with C_Value => 50);
```

The declaration of `Obj5` invokes the default constructors: the C++ default constructor of the parent type takes care of the initialization of the components inherited from `Root`, and GNAT takes care

of the default initialization of the additional Ada components of type DT (that is, `C_Value` is initialized to value 2009). The order of invocation of the constructors is consistent with the order of elaboration required by Ada and C++. That is, the constructor of the parent type is always called before the constructor of the derived type.

Let us now consider a record that has components whose type is imported from C++. For example:

```
type Rec1 is limited record
  Data1 : Root := Constructor (10);
  Value : Natural := 1000;
end record;

type Rec2 (D : Integer := 20) is limited record
  Rec   : Rec1;
  Data2 : Root := Constructor (D, 30);
end record;
```

The initialization of an object of type `Rec2` will call the non-default C++ constructors specified for the imported components. For example:

```
Obj8 : Rec2 (40);
```

Using Ada 2005 we can use limited aggregates to initialize an object invoking C++ constructors that differ from those specified in the type declarations. For example:

```
Obj9 : Rec2 := (Rec => (Data1 => Constructor (15, 16),
                    others => <>),
              others => <>);
```

The above declaration uses an Ada 2005 limited aggregate to initialize `Obj9`, and the C++ constructor that has two integer arguments is invoked to initialize the `Data1` component instead of the constructor specified in the declaration of type `Rec1`. In Ada 2005 the box in the aggregate indicates that unspecified components are initialized using the expression (if any) available in the component declaration. That is, in this case discriminant `D` is initialized to value 20, `Value` is initialized to value 1000, and the non-default C++ constructor that handles two integers takes care of initializing component `Data2` with values 20, 30.

In Ada 2005 we can use the extended return statement to build the Ada equivalent to C++ non-default constructors. For example:

```
function Constructor (V : Integer) return Rec2 is
begin
  return Obj : Rec2 := (Rec => (Data1 => Constructor (V, 20),
                            others => <>),
                    others => <>) do
    -- Further actions required for construction of
    -- objects of type Rec2
    ...
  end record;
end Constructor;
```

In this example the extended return statement construct is used to build in place the returned object

whose components are initialized by means of a limited aggregate. Any further action associated with the constructor can be placed inside the construct.

Interfacing with C++ at the Class Level

In this section we demonstrate the GNAT features for interfacing with C++ by means of an example making use of Ada 2005 abstract interface types. This example consists of a classification of animals; classes have been used to model our main classification of animals, and interfaces provide support for the management of secondary classifications. We first demonstrate a case in which the types and constructors are defined on the C++ side and imported from the Ada side, and latter the reverse case.

The root of our derivation will be the `Animal` class, with a single private attribute (the `Age` of the animal), a constructor, and two public primitives to set and get the value of this attribute.

```
class Animal {
public:
    virtual void Set_Age (int New_Age);
    virtual int Age ();
    Animal() {Age_Count = 0;};
private:
    int Age_Count;
};
```

Abstract interface types are defined in C++ by means of classes with pure virtual functions and no data members. In our example we will use two interfaces that provide support for the common management of `Carnivore` and `Domestic` animals:

```
class Carnivore {
public:
    virtual int Number_Of_Teeth () = 0;
};

class Domestic {
public:
    virtual void Set_Owner (char* Name) = 0;
};
```

Using these declarations, we can now say that a `Dog` is an animal that is both `Carnivore` and `Domestic`, that is:

```
class Dog : Animal, Carnivore, Domestic {
public:
    virtual int Number_Of_Teeth ();
    virtual void Set_Owner (char* Name);

    Dog(); // Constructor
private:
    int Tooth_Count;
    char *Owner;
};
```

In the following examples we will assume that the previous declarations are located in a file named `animals.h`. The following package demonstrates how to import these C++ declarations from the Ada side:

```
with Interfaces.C.Strings; use Interfaces.C.Strings;
package Animals is
  type Carnivore is limited interface;
  pragma Convention (C_Plus_Plus, Carnivore);
  function Number_Of_Teeth (X : Carnivore)
    return Natural is abstract;

  type Domestic is limited interface;
  pragma Convention (C_Plus_Plus, Domestic);
  procedure Set_Owner
    (X : in out Domestic;
     Name : Chars_Ptr) is abstract;

  type Animal is tagged limited record
    Age : Natural;
  end record;
  pragma Import (C_Plus_Plus, Animal);

  procedure Set_Age (X : in out Animal; Age : Integer);
  pragma Import (C_Plus_Plus, Set_Age);

  function Age (X : Animal) return Integer;
  pragma Import (C_Plus_Plus, Age);

  function New_Animal return Animal;
  pragma CPP_Constructor (New_Animal);
  pragma Import (CPP, New_Animal, "_ZN6AnimalC1Ev");

  type Dog is new Animal and Carnivore and Domestic with record
    Tooth_Count : Natural;
    Owner : Chars_Ptr;
  end record;
  pragma Import (C_Plus_Plus, Dog);

  function Number_Of_Teeth (A : Dog) return Natural;
  pragma Import (C_Plus_Plus, Number_Of_Teeth);

  procedure Set_Owner (A : in out Dog; Name : Chars_Ptr);
  pragma Import (C_Plus_Plus, Set_Owner);

  function New_Dog return Dog;
  pragma CPP_Constructor (New_Dog);
  pragma Import (CPP, New_Dog, "_ZN3DogC2Ev");
end Animals;
```

Thanks to the compatibility between GNAT run-time structures and the C++ ABI, interfacing with these C++ classes is easy. The only requirement is that all the primitives and components must be declared exactly in the same order in the two languages.

Regarding the abstract interfaces, we must indicate to the GNAT compiler by means of a `pragma Convention (C_Plus_Plus)`, the convention used to pass the arguments to the called primitives will be the same as for C++. For the imported classes we use `pragma Import` with convention `C_Plus_Plus` to indicate that they have been defined on the C++ side; this is required because the dispatch table associated with these tagged types will be built in the C++ side and therefore will not contain the predefined Ada primitives which Ada would otherwise expect.

As the reader can see there is no need to indicate the C++ mangled names associated with each subprogram because it is assumed that all the calls to these primitives will be dispatching calls. The only exception is the constructor, which must be registered with the compiler by means of `pragma CPP_Constructor` and needs to provide its associated C++ mangled name because the Ada compiler generates direct calls to it.

With the above packages we can now declare objects of type `Dog` on the Ada side and dispatch calls to the corresponding subprograms on the C++ side. We can also extend the tagged type `Dog` with further fields and primitives, and override some of its C++ primitives on the Ada side. For example, here we have a type derivation defined on the Ada side that inherits all the dispatching primitives of the ancestor from the C++ side.

```
with Animals; use Animals;
package Vaccinated_Animals is
  type Vaccinated_Dog is new Dog with null record;
  function Vaccination_Expired (A : Vaccinated_Dog) return Boolean;
end Vaccinated_Animals;
```

It is important to note that, because of the ABI compatibility, the programmer does not need to add any further information to indicate either the object layout or the dispatch table entry associated with each dispatching operation.

Now let us define all the types and constructors on the Ada side and export them to C++, using the same hierarchy of our previous example:

```
with Interfaces.C.Strings;
use Interfaces.C.Strings;
package Animals is
  type Carnivore is limited interface;
  pragma Convention (C_Plus_Plus, Carnivore);
  function Number_Of_Teeth (X : Carnivore)
    return Natural is abstract;

  type Domestic is limited interface;
  pragma Convention (C_Plus_Plus, Domestic);
  procedure Set_Owner
    (X : in out Domestic;
     Name : Chars_Ptr) is abstract;

  type Animal is tagged record
    Age : Natural;
  end record;
  pragma Convention (C_Plus_Plus, Animal);

  procedure Set_Age (X : in out Animal; Age : Integer);
```

(continues on next page)

(continued from previous page)

```

pragma Export (C_Plus_Plus, Set_Age);

function Age (X : Animal) return Integer;
pragma Export (C_Plus_Plus, Age);

function New_Animal return Animal'Class;
pragma Export (C_Plus_Plus, New_Animal);

type Dog is new Animal and Carnivore and Domestic with record
  Tooth_Count : Natural;
  Owner       : String (1 .. 30);
end record;
pragma Convention (C_Plus_Plus, Dog);

function Number_Of_Teeth (A : Dog) return Natural;
pragma Export (C_Plus_Plus, Number_Of_Teeth);

procedure Set_Owner (A : in out Dog; Name : Chars_Ptr);
pragma Export (C_Plus_Plus, Set_Owner);

function New_Dog return Dog'Class;
pragma Export (C_Plus_Plus, New_Dog);
end Animals;

```

Compared with our previous example the only differences are the use of `pragma Convention` (instead of `pragma Import`), and the use of `pragma Export` to indicate to the GNAT compiler that the primitives will be available to C++. Thanks to the ABI compatibility, on the C++ side there is nothing else to be done; as explained above, the only requirement is that all the primitives and components are declared in exactly the same order.

For completeness, let us see a brief C++ main program that uses the declarations available in `animals.h` (presented in our first example) to import and use the declarations from the Ada side, properly initializing and finalizing the Ada run-time system along the way:

```

#include "animals.h"
#include <iostream>
using namespace std;

void Check_Carnivore (Carnivore *obj) {...}
void Check_Domestic (Domestic *obj) {...}
void Check_Animal (Animal *obj) {...}
void Check_Dog (Dog *obj) {...}

extern "C" {
  void adainit (void);
  void adafinal (void);
  Dog* new_dog ();
}

void test ()
{

```

(continues on next page)

(continued from previous page)

```

Dog *obj = new_dog(); // Ada constructor
Check_Carnivore (obj); // Check secondary DT
Check_Domestic (obj); // Check secondary DT
Check_Animal (obj); // Check primary DT
Check_Dog (obj); // Check primary DT
}

int main ()
{
  adainit (); test(); adafinal ();
  return 0;
}

```

3.11.4 Generating Ada Bindings for C and C++ headers

GNAT includes a binding generator for C and C++ headers which is intended to do 95% of the tedious work of generating Ada specs from C or C++ header files.

Note that this capability is not intended to generate 100% correct Ada specs, and will in some cases require manual adjustments, although it can often be used out of the box in practice.

Some of the known limitations include:

- only very simple character constant macros are translated into Ada constants. Function macros (macros with arguments) are partially translated as comments, to be completed manually if needed.
- some extensions (e.g. vector types) are not supported
- pointers to pointers are mapped to System.Address
- identifiers with identical name (except casing) may generate compilation errors (e.g. `shm_get` vs `SHM_GET`).

The code is generated using Ada 2012 syntax, which makes it easier to interface with other languages. In most cases you can still use the generated binding even if your code is compiled using earlier versions of Ada (e.g. `-gnat95`).

Running the Binding Generator

The binding generator is part of the gcc compiler and can be invoked via the `-fdump-ada-spec` switch, which will generate Ada spec files for the header files specified on the command line, and all header files needed by these files transitively. For example:

```

$ gcc -c -fdump-ada-spec -C /usr/include/time.h
$ gcc -c *.ads

```

will generate, under GNU/Linux, the following files: `time_h.ads`, `bits_time_h.ads`, `stddef_h.ads`, `bits_types_h.ads` which correspond to the files `/usr/include/time.h`, `/usr/include/bits/time.h`, etc..., and then compile these Ada specs. That is to say, the name of the Ada specs is in keeping

with the relative path under `/usr/include/` of the header files. This behavior is specific to paths ending with `/include/`; in all the other cases, the name of the Ada specs is derived from the simple name of the header files instead.

The `-C` switch tells `gcc` to extract comments from headers, and will attempt to generate corresponding Ada comments.

If you want to generate a single Ada file and not the transitive closure, you can use instead the `-fdump-ada-spec-slim` switch.

You can optionally specify a parent unit, of which all generated units will be children, using `-fada-spec-parent=unit`.

The simple `gcc`-based command works only for C headers. For C++ headers you need to use either the `g++` command or the combination `gcc -x c++`.

In some cases, the generated bindings will be more complete or more meaningful when defining some macros, which you can do via the `-D` switch. This is for example the case with `Xlib.h` under GNU/Linux:

```
$ gcc -c -fdump-ada-spec -DXLIB_ILLEGAL_ACCESS -C /usr/include/X11/Xlib.h
```

The above will generate more complete bindings than a straight call without the `-DXLIB_ILLEGAL_ACCESS` switch.

In other cases, it is not possible to parse a header file in a stand-alone manner, because other include files need to be included first. In this case, the solution is to create a small header file including the needed `#include` and possible `#define` directives. For example, to generate Ada bindings for `readline/readline.h`, you need to first include `stdio.h`, so you can create a file with the following two lines in e.g. `readline1.h`:

```
#include <stdio.h>
#include <readline/readline.h>
```

and then generate Ada bindings from this file:

```
$ gcc -c -fdump-ada-spec readline1.h
```

Generating Bindings for C++ Headers

Generating bindings for C++ headers is done using the same options, always with the `g++` compiler. Note that generating Ada spec from C++ headers is a much more complex job and support for C++ headers is much more limited than support for C headers. As a result, you will need to modify the resulting bindings by hand more extensively when using C++ headers.

In this mode, C++ classes will be mapped to Ada tagged types, constructors will be mapped using the `CPP_Constructor` pragma, and when possible, multiple inheritance of abstract classes will be mapped to Ada interfaces (see the *Interfacing to C++* section in the *GNAT Reference Manual* for additional information on interfacing to C++).

For example, given the following C++ header file:

```

class Carnivore {
public:
    virtual int Number_Of_Teeth () = 0;
};

class Domestic {
public:
    virtual void Set_Owner (char* Name) = 0;
};

class Animal {
public:
    int Age_Count;
    virtual void Set_Age (int New_Age);
};

class Dog : Animal, Carnivore, Domestic {
public:
    int Tooth_Count;
    char *Owner;

    virtual int Number_Of_Teeth ();
    virtual void Set_Owner (char* Name);

    Dog();
};

```

The corresponding Ada code is generated:

```

package Class_Carnivore is
    type Carnivore is limited interface;
    pragma Import (CPP, Carnivore);

    function Number_Of_Teeth (this : access Carnivore) return int is abstract;
end;
use Class_Carnivore;

package Class_Domestic is
    type Domestic is limited interface;
    pragma Import (CPP, Domestic);

    procedure Set_Owner
        (this : access Domestic;
         Name : Interfaces.C.Strings.chars_ptr) is abstract;
end;
use Class_Domestic;

package Class_Animal is
    type Animal is tagged limited record
        Age_Count : aliased int;
    end record;
    pragma Import (CPP, Animal);

```

(continues on next page)

(continued from previous page)

```
procedure Set_Age (this : access Animal; New_Age : int);
pragma Import (CPP, Set_Age, "_ZN6Animal7Set_AgeEi");
end;
use Class_Animal;

package Class_Dog is
  type Dog is new Animal and Carnivore and Domestic with record
    Tooth_Count : aliased int;
    Owner : Interfaces.C.Strings.chars_ptr;
  end record;
  pragma Import (CPP, Dog);

  function Number_Of_Teeth (this : access Dog) return int;
  pragma Import (CPP, Number_Of_Teeth, "_ZN3Dog15Number_Of_TeethEv");

  procedure Set_Owner
    (this : access Dog; Name : Interfaces.C.Strings.chars_ptr);
  pragma Import (CPP, Set_Owner, "_ZN3Dog9Set_OwnerEPc");

  function New_Dog return Dog;
  pragma CPP_Constructor (New_Dog);
  pragma Import (CPP, New_Dog, "_ZN3DogC1Ev");
end;
use Class_Dog;
```

Switches

-fdump-ada-spec

Generate Ada spec files for the given header files transitively (including all header files that these headers depend upon).

-fdump-ada-spec-slim

Generate Ada spec files for the header files specified on the command line only.

-fada-spec-parent=unit

Specifies that all files generated by `-fdump-ada-spec` are to be child units of the specified parent unit.

-c

Extract comments from headers and generate Ada comments in the Ada spec files.

3.11.5 Generating C Headers for Ada Specifications

GNAT includes a C header generator for Ada specifications which supports Ada types that have a direct mapping to C types. This includes in particular support for:

- Scalar types
- Constrained arrays
- Records (untagged)
- Composition of the above types
- Constant declarations
- Object declarations
- Subprogram declarations

Running the C Header Generator

The C header generator is part of the GNAT compiler and can be invoked via the `-gnatceg` combination of switches, which will generate a `.h` file corresponding to the given input file (Ada spec or body). Note that only spec files are processed in any case, so giving a spec or a body file as input is equivalent. For example:

```
$ gcc -c -gnatceg pack1.ads
```

will generate a self-contained file called `pack1.h` including common definitions from the Ada Standard package, followed by the definitions included in `pack1.ads`, as well as all the other units withed by this file.

For instance, given the following Ada files:

```
package Pack2 is
  type Int is range 1 .. 10;
end Pack2;
```

```
with Pack2;

package Pack1 is
  type Rec is record
    Field1, Field2 : Pack2.Int;
  end record;

  Global : Rec := (1, 2);

  procedure Proc1 (R : Rec);
  procedure Proc2 (R : in out Rec);
end Pack1;
```

The above gcc command will generate the following `pack1.h` file:

```
/* Standard definitions skipped */
#ifndef PACK2_ADS
#define PACK2_ADS
typedef short_short_integer pack2__TintB;
typedef pack2__TintB pack2__int;
#endif /* PACK2_ADS */

#ifndef PACK1_ADS
#define PACK1_ADS
typedef struct _pack1__rec {
    pack2__int field1;
    pack2__int field2;
} pack1__rec;
extern pack1__rec pack1__global;
extern void pack1__proc1(const pack1__rec r);
extern void pack1__proc2(pack1__rec *r);
#endif /* PACK1_ADS */
```

You can then include `pack1.h` from a C source file and use the types, call subprograms, reference objects, and constants.

3.12 GNAT and Other Compilation Models

This section compares the GNAT model with the approaches taken in other environments, first the C/C++ model and then the mechanism that has been used in other Ada systems, in particular those traditionally used for Ada 83.

3.12.1 Comparison between GNAT and C/C++ Compilation Models

The GNAT model of compilation is close to the C and C++ models. You can think of Ada specs as corresponding to header files in C. As in C, you don't need to compile specs; they are compiled when they are used. The Ada *with* is similar in effect to the `#include` of a C header.

One notable difference is that, in Ada, you may compile specs separately to check them for semantic and syntactic accuracy. This is not always possible with C headers because they are fragments of programs that have less specific syntactic or semantic rules.

The other major difference is the requirement for running the binder, which performs two important functions. First, it checks for consistency. In C or C++, the only defense against assembling inconsistent programs lies outside the compiler, in a makefile, for example. The binder satisfies the Ada requirement that it be impossible to construct an inconsistent program when the compiler is used in normal mode.

The other important function of the binder is to deal with elaboration issues. There are also elaboration issues in C++ that are handled automatically. This automatic handling has the advantage of being simpler to use, but the C++ programmer has no control over elaboration. Where `gnatbind` might complain there was no valid order of elaboration, a C++ compiler would simply construct a program that malfunctioned at run time.

3.12.2 Comparison between GNAT and Conventional Ada Library Models

This section is intended for Ada programmers who have used an Ada compiler implementing the traditional Ada library model, as described in the Ada Reference Manual.

In GNAT, there is no ‘library’ in the normal sense. Instead, the set of source files themselves acts as the library. Compiling Ada programs does not generate any centralized information, but rather an object file and a ALI file, which are of interest only to the binder and linker. In a traditional system, the compiler reads information not only from the source file being compiled, but also from the centralized library. This means that the effect of a compilation depends on what has been previously compiled. In particular:

- When a unit is *withed*, the unit seen by the compiler corresponds to the version of the unit most recently compiled into the library.
- Inlining is effective only if the necessary body has already been compiled into the library.
- Compiling a unit may obsolete other units in the library.

In GNAT, compiling one unit never affects the compilation of any other units because the compiler reads only source files. Only changes to source files can affect the results of a compilation. In particular:

- When a unit is *withed*, the unit seen by the compiler corresponds to the source version of the unit that is currently accessible to the compiler.
- Inlining requires the appropriate source files for the package or subprogram bodies to be available to the compiler. Inlining is always effective, independent of the order in which units are compiled.
- Compiling a unit never affects any other compilations. The editing of sources may cause previous compilations to be out of date if they depended on the source file being modified.

The most important result of these differences is that order of compilation is never significant in GNAT. There is no situation in which one is required to do one compilation before another. What shows up as order of compilation requirements in the traditional Ada library becomes, in GNAT, simple source dependencies; in other words, there is only a set of rules saying what source files must be present when a file is compiled.

3.13 Using GNAT Files with External Tools

This section explains how files that are produced by GNAT may be used with tools designed for other languages.

3.13.1 Using Other Utility Programs with GNAT

The object files generated by GNAT are in standard system format and in particular the debugging information uses this format. This means programs generated by GNAT can be used with existing utilities that depend on these formats.

In general, any utility program that works with C will also often work with Ada programs generated by GNAT. This includes software utilities such as gprof (a profiling program), gdb (the FSF debugger), and utilities such as Purify.

3.13.2 The External Symbol Naming Scheme of GNAT

In order to interpret the output from GNAT, when using tools that are originally intended for use with other languages, it is useful to understand the conventions used to generate link names from the Ada entity names.

All link names are in all lowercase letters. With the exception of library procedure names, the mechanism used is simply to use the full expanded Ada name with dots replaced by double underscores. For example, suppose we have the following package spec:

```
package QRS is
  MN : Integer;
end QRS;
```

The variable MN has a full expanded Ada name of QRS.MN, so the corresponding link name is qrs__mn. Of course if a pragma Export is used this may be overridden:

```
package Exports is
  Var1 : Integer;
  pragma Export (Var1, C, External_Name => "var1_name");
  Var2 : Integer;
  pragma Export (Var2, C, Link_Name => "var2_link_name");
end Exports;
```

In this case, the link name for Var1 is whatever link name the C compiler would assign for the C function var1_name. This typically would be either var1_name or _var1_name, depending on operating system conventions, but other possibilities exist. The link name for Var2 is var2_link_name, and this is not operating system dependent.

One exception occurs for library level procedures. A potential ambiguity arises between the required name _main for the C main program, and the name we would otherwise assign to an Ada library level procedure called Main (which might well not be the main program).

To avoid this ambiguity, we attach the prefix _ada_ to such names. So if we have a library level procedure such as:

```
procedure Hello (S : String);
```

the external name of this procedure will be _ada_hello.

BUILDING EXECUTABLE PROGRAMS WITH GNAT

This chapter describes first the `gnatmake` tool ([Building with `gnatmake`](#)), which automatically determines the set of sources needed by an Ada compilation unit and executes the necessary (re)compilations, binding and linking. It also explains how to use each tool individually: the compiler (`gcc`, see [Compiling with `gcc`](#)), binder (`gnatbind`, see [Binding with `gnatbind`](#)), and linker (`gnatlink`, see [Linking with `gnatlink`](#)) to build executable programs. Finally, this chapter provides examples of how to make use of the general GNU make mechanism in a GNAT context (see [Using the GNU make Utility](#)).

4.1 Building with `gnatmake`

A typical development cycle when working on an Ada program consists of the following steps:

1. Edit some sources to fix bugs;
2. Add enhancements;
3. Compile all sources affected;
4. Rebind and relink; and
5. Test.

The third step in particular can be tricky, because not only do the modified files have to be compiled, but any files depending on these files must also be recompiled. The dependency rules in Ada can be quite complex, especially in the presence of overloading, `use` clauses, generics and inlined subprograms.

`gnatmake` automatically takes care of the third and fourth steps of this process. It determines which sources need to be compiled, compiles them, and binds and links the resulting object files.

Unlike some other Ada make programs, the dependencies are always accurately recomputed from the new sources. The source based approach of the GNAT compilation model makes this possible. This means that if changes to the source program cause corresponding changes in dependencies, they will always be tracked exactly correctly by `gnatmake`.

Note that for advanced forms of project structure, we recommend creating a project file as explained in the *GNAT_Project_Manager* chapter in the *GPRbuild User's Guide*, and using the `gprbuild` tool which supports building with project files and works similarly to `gnatmake`.

4.1.1 Running gnatmake

The usual form of the `gnatmake` command is

```
$ gnatmake [<switches>] <file_name> [<file_names>] [<mode_switches>]
```

The only required argument is one `file_name`, which specifies a compilation unit that is a main program. Several `file_names` can be specified: this will result in several executables being built. If `switches` are present, they can be placed before the first `file_name`, between `file_names` or after the last `file_name`. If `mode_switches` are present, they must always be placed after the last `file_name` and all `switches`.

If you are using standard file extensions (`.adb` and `.ads`), then the extension may be omitted from the `file_name` arguments. However, if you are using non-standard extensions, then it is required that the extension be given. A relative or absolute directory path can be specified in a `file_name`, in which case, the input source file will be searched for in the specified directory only. Otherwise, the input source file will first be searched in the directory where `gnatmake` was invoked and if it is not found, it will be search on the source path of the compiler as described in [Search Paths and the Run-Time Library \(RTL\)](#).

All `gnatmake` output (except when you specify `-M`) is sent to `stderr`. The output produced by the `-M` switch is sent to `stdout`.

4.1.2 Switches for gnatmake

You may specify any of the following switches to `gnatmake`:

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

-Pproject

Build GNAT project file `project` using GPRbuild. When this switch is present, all other command-line switches are treated as GPRbuild switches and not `gnatmake` switches.

--GCC=compiler_name

Program used for compiling. The default is `gcc`. You need to use quotes around `compiler_name` if `compiler_name` contains spaces or other separator characters. As an example `--GCC="foo -x -y"` will instruct `gnatmake` to use `foo -x -y` as your compiler. A limitation of this syntax is that the name and path name of the executable itself must not include any embedded spaces. Note that switch `-c` is always inserted after your command name. Thus in the above example the compiler command that will be used by `gnatmake` will be `foo -c -x -y`. If several `--GCC=compiler_name` are used, only the last `compiler_name` is taken into account. However, all the additional switches are also taken into account. Thus, `--GCC="foo -x -y" --GCC="bar -z -t"` is equivalent to `--GCC="bar -x -y -z -t"`.

--GNATBIND=binder_name

Program used for binding. The default is `gnatbind`. You need to use quotes around `binder_name` if `binder_name` contains spaces or other separator characters. As an example

--GNATBIND="bar -x -y" will instruct `gnatmake` to use `bar -x -y` as your binder. Binder switches that are normally appended by `gnatmake` to `gnatbind` are now appended to the end of `bar -x -y`. A limitation of this syntax is that the name and path name of the executable itself must not include any embedded spaces.

--GNATLINK=*linker_name*

Program used for linking. The default is `gnatlink`. You need to use quotes around `linker_name` if `linker_name` contains spaces or other separator characters. As an example --GNATLINK="lan -x -y" will instruct `gnatmake` to use `lan -x -y` as your linker. Linker switches that are normally appended by `gnatmake` to `gnatlink` are now appended to the end of `lan -x -y`. A limitation of this syntax is that the name and path name of the executable itself must not include any embedded spaces.

--create-map-file

When linking an executable, create a map file. The name of the map file has the same name as the executable with extension ".map".

--create-map-file=*mapfile*

When linking an executable, create a map file with the specified name.

--create-missing-dirs

When using project files (*-Pproject*), automatically create missing object directories, library directories and exec directories.

--single-compile-per-obj-dir

Disallow simultaneous compilations in the same object directory when project files are used.

--subdirs=*subdir*

Actual object directory of each project file is the subdirectory `subdir` of the object directory specified or defaulted in the project file.

--unchecked-shared-lib-imports

By default, shared library projects are not allowed to import static library projects. When this switch is used on the command line, this restriction is relaxed.

--source-info=*source info file*

Specify a source info file. This switch is active only when project files are used. If the source info file is specified as a relative path, then it is relative to the object directory of the main project. If the source info file does not exist, then after the Project Manager has successfully parsed and processed the project files and found the sources, it creates the source info file. If the source info file already exists and can be read successfully, then the Project Manager will get all the needed information about the sources from the source info file and will not look for them. This reduces the time to process the project files, especially when looking for sources that take a long time. If the source info file exists but cannot be parsed successfully, the Project Manager will attempt to recreate it. If the Project Manager fails to create the source info file, a message is issued, but `gnatmake` does not fail. `gnatmake` "trusts" the source info file. This means that if the source files have changed (addition, deletion, moving to a different source directory), then the source info file need to be deleted and recreated.

-a

Consider all files in the make process, even the GNAT internal system files (for example, the predefined Ada library files), as well as any locked files. Locked files are files whose ALI file

is write-protected. By default, `gnatmake` does not check these files, because the assumption is that the GNAT internal files are properly up to date, and also that any write protected ALI files have been properly installed. Note that if there is an installation problem, such that one of these files is not up to date, it will be properly caught by the binder. You may have to specify this switch if you are working on GNAT itself. The switch `-a` is also useful in conjunction with `-f` if you need to recompile an entire application, including run-time files, using special configuration pragmas, such as a `Normalize_Scalars` pragma.

By default `gnatmake -a` compiles all GNAT internal files with `gcc -c -gnatpg` rather than `gcc -c`.

-b

Bind only. Can be combined with `-c` to do compilation and binding, but no link. Can be combined with `-l` to do binding and linking. When not combined with `-c` all the units in the closure of the main program must have been previously compiled and must be up to date. The root unit specified by `file_name` may be given without extension, with the source extension or, if no GNAT Project File is specified, with the ALI file extension.

-c

Compile only. Do not perform binding, except when `-b` is also specified. Do not perform linking, except if both `-b` and `-l` are also specified. If the root unit specified by `file_name` is not a main unit, this is the default. Otherwise `gnatmake` will attempt binding and linking unless all objects are up to date and the executable is more recent than the objects.

-C

Use a temporary mapping file. A mapping file is a way to communicate to the compiler two mappings: from unit names to file names (without any directory information) and from file names to path names (with full directory information). A mapping file can make the compiler's file searches faster, especially if there are many source directories, or the sources are read over a slow network connection. If `-P` is used, a mapping file is always used, so `-C` is unnecessary; in this case the mapping file is initially populated based on the project file. If `-C` is used without `-P`, the mapping file is initially empty. Each invocation of the compiler will add any newly accessed sources to the mapping file.

-C=file

Use a specific mapping file. The file, specified as a path name (absolute or relative) by this switch, should already exist, otherwise the switch is ineffective. The specified mapping file will be communicated to the compiler. This switch is not compatible with a project file (`-P'file'`) or with multiple compiling processes (`-jnnn`, when `nnn` is greater than 1).

-d

Display progress for each source, up to date or not, as a single line:

```
completed x out of y (zz%)
```

If the file needs to be compiled this is displayed after the invocation of the compiler. These lines are displayed even in quiet output mode.

-D dir

Put all object files and ALI file in directory `dir`. If the `-D` switch is not used, all object files and ALI files go in the current working directory.

This switch cannot be used when using a project file.

-eInnn

Indicates that the main source is a multi-unit source and the rank of the unit in the source file is *nnn*. *nnn* needs to be a positive number and a valid index in the source. This switch cannot be used when `gnatmake` is invoked for several mains.

-eL

Follow all symbolic links when processing project files. This should be used if your project uses symbolic links for files or directories, but is not needed in other cases.

This also assumes that no directory matches the naming scheme for files (for instance that you do not have a directory called “sources.ads” when using the default GNAT naming scheme).

When you do not have to use this switch (i.e., by default), `gnatmake` is able to save a lot of system calls (several per source file and object file), which can result in a significant speed up to load and manipulate a project file, especially when using source files from a remote system.

-eS

Output the commands for the compiler, the binder and the linker on standard output, instead of standard error.

-f

Force recompilations. Recompile all sources, even though some object files may be up to date, but don't recompile predefined or GNAT internal files or locked files (files with a write-protected ALI file), unless the `-a` switch is also specified.

-F

When using project files, if some errors or warnings are detected during parsing and verbose mode is not in effect (no use of switch `-v`), then error lines start with the full path name of the project file, rather than its simple file name.

-g

Enable debugging. This switch is simply passed to the compiler and to the linker.

-i

In normal mode, `gnatmake` compiles all object files and ALI files into the current directory. If the `-i` switch is used, then instead object files and ALI files that already exist are overwritten in place. This means that once a large project is organized into separate directories in the desired manner, then `gnatmake` will automatically maintain and update this organization. If no ALI files are found on the Ada object path (see [Search Paths and the Run-Time Library \(RTL\)](#)), the new object and ALI files are created in the directory containing the source being compiled. If another organization is desired, where objects and sources are kept in different directories, a useful technique is to create dummy ALI files in the desired directories. When detecting such a dummy file, `gnatmake` will be forced to recompile the corresponding source file, and it will be put the resulting object and ALI files in the directory where it found the dummy file.

-jn

Use *n* processes to carry out the (re)compilations. On a multiprocessor machine compilations will occur in parallel. If *n* is 0, then the maximum number of parallel compilations is the number of core processors on the platform. In the event of compilation errors, messages from various compilations might get interspersed (but `gnatmake` will give you the full ordered list

of failing compiles at the end). If this is problematic, rerun the make process with `n` set to 1 to get a clean list of messages.

-k

Keep going. Continue as much as possible after a compilation error. To ease the programmer's task in case of compilation errors, the list of sources for which the compile fails is given when `gnatmake` terminates.

If `gnatmake` is invoked with several `file_names` and with this switch, if there are compilation errors when building an executable, `gnatmake` will not attempt to build the following executables.

-l

Link only. Can be combined with `-b` to binding and linking. Linking will not be performed if combined with `-c` but not with `-b`. When not combined with `-b` all the units in the closure of the main program must have been previously compiled and must be up to date, and the main program needs to have been bound. The root unit specified by `file_name` may be given without extension, with the source extension or, if no GNAT Project File is specified, with the ALI file extension.

-m

Specify that the minimum necessary amount of recompilations be performed. In this mode `gnatmake` ignores time stamp differences when the only modifications to a source file consist in adding/removing comments, empty lines, spaces or tabs. This means that if you have changed the comments in a source file or have simply reformatted it, using this switch will tell `gnatmake` not to recompile files that depend on it (provided other sources on which these files depend have undergone no semantic modifications). Note that the debugging information may be out of date with respect to the sources if the `-m` switch causes a compilation to be switched, so the use of this switch represents a trade-off between compilation time and accurate debugging information.

-M

Check if all objects are up to date. If they are, output the object dependences to `stdout` in a form that can be directly exploited in a `Makefile`. By default, each source file is prefixed with its (relative or absolute) directory name. This name is whatever you specified in the various `-aI` and `-I` switches. If you use `gnatmake -M -q` (see below), only the source file names, without relative paths, are output. If you just specify the `-M` switch, dependencies of the GNAT internal system files are omitted. This is typically what you want. If you also specify the `-a` switch, dependencies of the GNAT internal files are also listed. Note that dependencies of the objects in external Ada libraries (see switch `-aLdir` in the following list) are never reported.

-n

Don't compile, bind, or link. Checks if all objects are up to date. If they are not, the full name of the first file that needs to be recompiled is printed. Repeated use of this option, followed by compiling the indicated source file, will eventually result in recompiling all required units.

-o *exec_name*

Output executable name. The name of the final executable program will be `exec_name`. If the `-o` switch is omitted the default name for the executable will be the name of the input file in appropriate form for an executable file on the host system.

This switch cannot be used when invoking `gnatmake` with several `file_names`.

-p

Same as `--create-missing-dirs`

-q

Quiet. When this flag is not set, the commands carried out by `gnatmake` are displayed.

-s

Recompile if compiler switches have changed since last compilation. All compiler switches but `-I` and `-o` are taken into account in the following way: orders between different 'first letter' switches are ignored, but orders between same switches are taken into account. For example, `-O -O2` is different than `-O2 -O`, but `-g -O` is equivalent to `-O -g`.

This switch is recommended when Integrated Preprocessing is used.

-u

Unique. Recompile at most the main files. It implies `-c`. Combined with `-f`, it is equivalent to calling the compiler directly. Note that using `-u` with a project file and no main has a special meaning.

-U

When used without a project file or with one or several mains on the command line, is equivalent to `-u`. When used with a project file and no main on the command line, all sources of all project files are checked and compiled if not up to date, and libraries are rebuilt, if necessary.

-v

Verbose. Display the reason for all recompilations `gnatmake` decides are necessary, with the highest verbosity level.

-v`l`

Verbosity level Low. Display fewer lines than in verbosity Medium.

-v`m`

Verbosity level Medium. Potentially display fewer lines than in verbosity High.

-v`h`

Verbosity level High. Equivalent to `-v`.

-v`Px`

Indicate the verbosity of the parsing of GNAT project files. See Switches Related to Project Files.

-x

Indicate that sources that are not part of any Project File may be compiled. Normally, when using Project Files, only sources that are part of a Project File may be compile. When this switch is used, a source outside of all Project Files may be compiled. The ALI file and the object file will be put in the object directory of the main Project. The compilation switches used will only be those specified on the command line. Even when `-x` is used, mains specified on the command line need to be sources of a project file.

-X`name=value`

Indicate that external variable `name` has the value `value`. The Project Manager will use this

value for occurrences of `external(name)` when parsing the project file. Switches Related to Project Files.

-z

No main subprogram. Bind and link the program even if the unit name given on the command line is a package name. The resulting executable will execute the elaboration routines of the package and its closure, then the finalization routines.

GCC switches

Any uppercase or multi-character switch that is not a `gnatmake` switch is passed to `gcc` (e.g., `-O`, `-gnato`, etc.)

Source and library search path switches

-aIdir

When looking for source files also look in directory `dir`. The order in which source files search is undertaken is described in [Search Paths](#) and the [Run-Time Library \(RTL\)](#).

-aLdir

Consider `dir` as being an externally provided Ada library. Instructs `gnatmake` to skip compilation units whose `.ALI` files have been located in directory `dir`. This allows you to have missing bodies for the units in `dir` and to ignore out of date bodies for the same units. You still need to specify the location of the specs for these units by using the switches `-aIdir` or `-Idir`. Note: this switch is provided for compatibility with previous versions of `gnatmake`. The easier method of causing standard libraries to be excluded from consideration is to write-protect the corresponding ALI files.

-aOdir

When searching for library and object files, look in directory `dir`. The order in which library files are searched is described in [Search Paths for gnatbind](#).

-Adir

Equivalent to `-aLdir -aIdir`.

-Idir

Equivalent to `-aOdir -aIdir`.

-I-

Do not look for source files in the directory containing the source file named in the command line. Do not look for ALI or object files in the directory where `gnatmake` was invoked.

-Ldir

Add directory `dir` to the list of directories in which the linker will search for libraries. This is equivalent to `-larg -Ldir`. Furthermore, under Windows, the sources pointed to by the `libraries` path set in the registry are not searched for.

-nostdinc

Do not look for source files in the system default directory.

-nostdlib

Do not look for library files in the system default directory.

--RTS=*rts-path*

Specifies the default location of the run-time library. GNAT looks for the run-time in the following directories, and stops as soon as a valid run-time is found (`adainclude` or `ada_source_path`, and `adalib` or `ada_object_path` present):

- *<current directory>/\$rts_path*
- *<default-search-dir>/\$rts_path*
- *<default-search-dir>/rts-\$rts_path*
- The selected path is handled like a normal RTS path.

4.1.3 Mode Switches for gnatmake

The mode switches (referred to as `mode_switches`) allow the inclusion of switches that are to be passed to the compiler itself, the binder or the linker. The effect of a mode switch is to cause all subsequent switches up to the end of the switch list, or up to the next mode switch, to be interpreted as switches to be passed on to the designated component of GNAT.

-cargs *switches*

Compiler switches. Here `switches` is a list of switches that are valid switches for `gcc`. They will be passed on to all compile steps performed by `gnatmake`.

-bargs *switches*

Binder switches. Here `switches` is a list of switches that are valid switches for `gnatbind`. They will be passed on to all bind steps performed by `gnatmake`.

-largs *switches*

Linker switches. Here `switches` is a list of switches that are valid switches for `gnatlink`. They will be passed on to all link steps performed by `gnatmake`.

-margs *switches*

Make switches. The switches are directly interpreted by `gnatmake`, regardless of any previous occurrence of `-cargs`, `-bargs` or `-largs`.

4.1.4 Notes on the Command Line

This section contains some additional useful notes on the operation of the `gnatmake` command.

- If `gnatmake` finds no ALI files, it recompiles the main program and all other units required by the main program. This means that `gnatmake` can be used for the initial compile, as well as during subsequent steps of the development cycle.
- If you enter `gnatmake foo.adb`, where `foo` is a subunit or body of a generic unit, `gnatmake` recompiles `foo.adb` (because it finds no ALI) and stops, issuing a warning.
- In `gnatmake` the switch `-I` is used to specify both source and library file paths. Use `-aI` instead if you just want to specify source paths only and `-a0` if you want to specify library paths only.

- `gnatmake` will ignore any files whose ALI file is write-protected. This may conveniently be used to exclude standard libraries from consideration and in particular it means that the use of the `-f` switch will not recompile these files unless `-a` is also specified.
- `gnatmake` has been designed to make the use of Ada libraries particularly convenient. Assume you have an Ada library organized as follows: `obj-dir` contains the objects and ALI files for of your Ada compilation units, whereas `include-dir` contains the specs of these units, but no bodies. Then to compile a unit stored in `main.adb`, which uses this Ada library you would just type:

```
$ gnatmake -aI`include-dir` -aL`obj-dir` main
```

- Using `gnatmake` along with the `-m` (minimal recompilation) switch provides a mechanism for avoiding unnecessary recompilations. Using this switch, you can update the comments/format of your source files without having to recompile everything. Note, however, that adding or deleting lines in a source files may render its debugging info obsolete. If the file in question is a spec, the impact is rather limited, as that debugging info will only be useful during the elaboration phase of your program. For bodies the impact can be more significant. In all events, your debugger will warn you if a source file is more recent than the corresponding object, and alert you to the fact that the debugging information may be out of date.

4.1.5 How gnatmake Works

Generally `gnatmake` automatically performs all necessary recompilations and you don't need to worry about how it works. However, it may be useful to have some basic understanding of the `gnatmake` approach and in particular to understand how it uses the results of previous compilations without incorrectly depending on them.

First a definition: an object file is considered *up to date* if the corresponding ALI file exists and if all the source files listed in the dependency section of this ALI file have time stamps matching those in the ALI file. This means that neither the source file itself nor any files that it depends on have been modified, and hence there is no need to recompile this file.

`gnatmake` works by first checking if the specified main unit is up to date. If so, no compilations are required for the main unit. If not, `gnatmake` compiles the main program to build a new ALI file that reflects the latest sources. Then the ALI file of the main unit is examined to find all the source files on which the main program depends, and `gnatmake` recursively applies the above procedure on all these files.

This process ensures that `gnatmake` only trusts the dependencies in an existing ALI file if they are known to be correct. Otherwise it always recompiles to determine a new, guaranteed accurate set of dependencies. As a result the program is compiled 'upside down' from what may be more familiar as the required order of compilation in some other Ada systems. In particular, clients are compiled before the units on which they depend. The ability of GNAT to compile in any order is critical in allowing an order of compilation to be chosen that guarantees that `gnatmake` will recompute a correct set of new dependencies if necessary.

When invoking `gnatmake` with several `file_names`, if a unit is imported by several of the executables, it will be recompiled at most once.

Note: when using non-standard naming conventions ([Using Other File Names](#)), changing through a configuration pragmas file the version of a source and invoking `gnatmake` to recompile may have no effect, if the previous version of the source is still accessible by `gnatmake`. It may be necessary to use the switch `-f`.

4.1.6 Examples of `gnatmake` Usage

gnatmake hello.adb

Compile all files necessary to bind and link the main program `hello.adb` (containing unit `Hello`) and bind and link the resulting object files to generate an executable file `hello`.

gnatmake main1 main2 main3

Compile all files necessary to bind and link the main programs `main1.adb` (containing unit `Main1`), `main2.adb` (containing unit `Main2`) and `main3.adb` (containing unit `Main3`) and bind and link the resulting object files to generate three executable files `main1`, `main2` and `main3`.

gnatmake -q Main_Unit -cargs -O2 -bargs -l

Compile all files necessary to bind and link the main program unit `Main_Unit` (from file `main_unit.adb`). All compilations will be done with optimization level 2 and the order of elaboration will be listed by the binder. `gnatmake` will operate in quiet mode, not displaying commands it is executing.

4.2 Compiling with `gcc`

This section discusses how to compile Ada programs using the `gcc` command. It also describes the set of switches that can be used to control the behavior of the compiler.

4.2.1 Compiling Programs

The first step in creating an executable program is to compile the units of the program using the `gcc` command. You must compile the following files:

- the body file (`.adb`) for a library level subprogram or generic subprogram
- the spec file (`.ads`) for a library level package or generic package that has no body
- the body file (`.adb`) for a library level package or generic package that has a body

You need *not* compile the following files

- the spec of a library unit which has a body
- subunits

because they are compiled as part of compiling related units. GNAT package specs when the corresponding body is compiled, and subunits when the parent is compiled.

If you attempt to compile any of these files, you will get one of the following error messages (where `fff` is the name of the file you compiled):

```
cannot generate code for file ``fff`` (package spec)
to check package spec, use -gnatc

cannot generate code for file ``fff`` (missing subunits)
to check parent unit, use -gnatc

cannot generate code for file ``fff`` (subprogram spec)
to check subprogram spec, use -gnatc

cannot generate code for file ``fff`` (subunit)
to check subunit, use -gnatc
```

As indicated by the above error messages, if you want to submit one of these files to the compiler to check for correct semantics without generating code, then use the `-gnatc` switch.

The basic command for compiling a file containing an Ada unit is:

```
$ gcc -c [switches] <file name>
```

where `file name` is the name of the Ada file (usually having an extension `.ads` for a spec or `.adb` for a body). You specify the `-c` switch to tell `gcc` to compile, but not link, the file. The result of a successful compilation is an object file, which has the same name as the source file but an extension of `.o` and an Ada Library Information (ALI) file, which also has the same name as the source file, but with `.ali` as the extension. GNAT creates these two output files in the current directory, but you may specify a source file in any directory using an absolute or relative path specification containing the directory information.

TESTING: the `--foobarNW` switch

`gcc` is actually a driver program that looks at the extensions of the file arguments and loads the appropriate compiler. For example, the GNU C compiler is `cc1`, and the Ada compiler is `gnat1`. These programs are in directories known to the driver program (in some configurations via environment variables you set), but need not be in your path. The `gcc` driver also calls the assembler and any other utilities needed to complete the generation of the required object files.

It is possible to supply several file names on the same `gcc` command. This causes `gcc` to call the appropriate compiler for each file. For example, the following command lists two separate files to be compiled:

```
$ gcc -c x.adb y.adb
```

calls `gnat1` (the Ada compiler) twice to compile `x.adb` and `y.adb`. The compiler generates two object files `x.o` and `y.o` and the two ALI files `x.ali` and `y.ali`.

Any switches apply to all the files listed, see [Compiler Switches](#) for a list of available `gcc` switches.

4.2.2 Search Paths and the Run-Time Library (RTL)

With the GNAT source-based library system, the compiler must be able to find source files for units that are needed by the unit being compiled. Search paths are used to guide this process.

The compiler compiles one source file whose name must be given explicitly on the command line. In other words, no searching is done for this file. To find all other source files that are needed (the most common being the specs of units), the compiler examines the following directories, in the following order:

- The directory containing the source file of the main unit being compiled (the file name on the command line).
- Each directory named by an `-I` switch given on the `gcc` command line, in the order given.
- Each of the directories listed in the text file whose name is given by the `ADA_PRJ_INCLUDE_FILE` environment variable. `ADA_PRJ_INCLUDE_FILE` is normally set by `gnatmake` or by the `gnat` driver when project files are used. It should not normally be set by other means.
- Each of the directories listed in the value of the `ADA_INCLUDE_PATH` environment variable. Construct this value exactly as the `PATH` environment variable: a list of directory names separated by colons (semicolons when working with the NT version).
- The content of the `ada_source_path` file which is part of the GNAT installation tree and is used to store standard libraries such as the GNAT Run Time Library (RTL) source files. [Installing a library](#)

Specifying the switch `-I-` inhibits the use of the directory containing the source file named in the command line. You can still have this directory on your search path, but in this case it must be explicitly requested with a `-I` switch.

Specifying the switch `-nostdinc` inhibits the search of the default location for the GNAT Run Time Library (RTL) source files.

The compiler outputs its object files and ALI files in the current working directory. Caution: The object file can be redirected with the `-o` switch; however, `gcc` and `gnat1` have not been coordinated on this so the ALI file will not go to the right place. Therefore, you should avoid using the `-o` switch.

The packages `Ada`, `System`, and `Interfaces` and their children make up the GNAT RTL, together with the simple `System.IO` package used in the "Hello World" example. The sources for these units are needed by the compiler and are kept together in one directory. Not all of the bodies are needed, but all of the sources are kept together anyway. In a normal installation, you need not specify these directory names when compiling or binding. Either the environment variables or the built-in defaults cause these files to be found.

In addition to the language-defined hierarchies (`System`, `Ada` and `Interfaces`), the GNAT distribution provides a fourth hierarchy, consisting of child units of `GNAT`. This is a collection of generally useful types, subprograms, etc. See the *GNAT Reference Manual* for further details.

Besides simplifying access to the RTL, a major use of search paths is in compiling sources from multiple directories. This can make development environments much more flexible.

4.2.3 Order of Compilation Issues

If, in our earlier example, there was a spec for the `hello` procedure, it would be contained in the file `hello.ads`; yet this file would not have to be explicitly compiled. This is the result of the model we chose to implement library management. Some of the consequences of this model are as follows:

- There is no point in compiling specs (except for package specs with no bodies) because these are compiled as needed by clients. If you attempt a useless compilation, you will receive an error message. It is also useless to compile subunits because they are compiled as needed by the parent.
- There are no order of compilation requirements: performing a compilation never obsoletes anything. The only way you can obsolete something and require recompilations is to modify one of the source files on which it depends.
- There is no library as such, apart from the ALI files ([The Ada Library Information Files](#), for information on the format of these files). For now we find it convenient to create separate ALI files, but eventually the information therein may be incorporated into the object file directly.
- When you compile a unit, the source files for the specs of all units that it *withs*, all its subunits, and the bodies of any generics it instantiates must be available (reachable by the search-paths mechanism described above), or you will receive a fatal error message.

4.2.4 Examples

The following are some typical Ada compilation command line examples:

```
$ gcc -c xyz.adb
```

Compile body in file `xyz.adb` with all default options.

```
$ gcc -c -O2 -gnata xyz-def.adb
```

Compile the child unit package in file `xyz-def.adb` with extensive optimizations, and pragma `Assert/Debug` statements enabled.

```
$ gcc -c -gnatc abc-def.adb
```

Compile the subunit in file `abc-def.adb` in semantic-checking-only mode.

4.3 Compiler Switches

The `gcc` command accepts switches that control the compilation process. These switches are fully described in this section: first an alphabetical listing of all switches with a brief description, and then functionally grouped sets of switches with more detailed information.

More switches exist for GCC than those documented here, especially for specific targets. However, their use is not recommended as they may change code generation in ways that are incompatible with the Ada run-time library, or can cause inconsistencies between compilation units.

4.3.1 Alphabetical List of All Switches

-b target

Compile your program to run on `target`, which is the name of a system configuration. You must have a GNAT cross-compiler built if `target` is not the same as your host system.

-Bdir

Load compiler executables (for example, `gnat1`, the Ada compiler) from `dir` instead of the default location. Only use this switch when multiple versions of the GNAT compiler are available. See the “Options for Directory Search” section in the *Using the GNU Compiler Collection (GCC)* manual for further details. You would normally use the `-b` or `-v` switch instead.

-c

Compile. Always use this switch when compiling Ada programs.

Note: for some other languages when using `gcc`, notably in the case of C and C++, it is possible to use `gcc` without a `-c` switch to compile and link in one step. In the case of GNAT, you cannot use this approach, because the binder must be run and `gcc` cannot be used to run the GNAT binder.

-fcallgraph-info[=su,da]

Makes the compiler output callgraph information for the program, on a per-file basis. The information is generated in the VCG format. It can be decorated with additional, per-node and/or per-edge information, if a list of comma-separated markers is additionally specified. When the `su` marker is specified, the callgraph is decorated with stack usage information; it is equivalent to `-fstack-usage`. When the `da` marker is specified, the callgraph is decorated with information about dynamically allocated objects.

-fdiagnostics-format=json

Makes GNAT emit warning and error messages as JSON. Inhibits printing of text warning and errors messages except if `-gnatv` or `-gnatl` are present. Uses absolute file paths when used along `-gnatf`.

-fdump-scos

Generates SCO (Source Coverage Obligation) information in the ALI file. This information is used by advanced coverage tools. See unit `SCOs` in the compiler sources for details in files `scos.ads` and `scos.adb`.

-fgnat-encodings=[all|gdb|minimal]

This switch controls the balance between GNAT encodings and standard DWARF emitted in the debug information.

-flto[=n]

Enables Link Time Optimization. This switch must be used in conjunction with the `-Ox` switches (but not with the `-gnatn` switch since it is a full replacement for the latter) and instructs the compiler to defer most optimizations until the link stage. The advantage of this approach is that the compiler can do a whole-program analysis and choose the best interprocedural optimization strategy based on a complete view of the program, instead of a fragmentary view with the usual approach. This can also speed up the compilation of big programs and reduce the size of the executable, compared with a traditional per-unit

compilation with inlining across units enabled by the `-gnatn` switch. The drawback of this approach is that it may require more memory and that the debugging information generated by `-g` with it might be hardly usable. The switch, as well as the accompanying `-Ox` switches, must be specified both for the compilation and the link phases. If the `n` parameter is specified, the optimization and final code generation at link time are executed using `n` parallel jobs by means of an installed `make` program.

-fno-inline

Suppresses all inlining, unless requested with `pragma Inline_Always`. The effect is enforced regardless of other optimization or inlining switches. Note that inlining can also be suppressed on a finer-grained basis with `pragma No_Inline`.

-fno-inline-functions

Suppresses automatic inlining of subprograms, which is enabled if `-O3` is used.

-fno-inline-small-functions

Suppresses automatic inlining of small subprograms, which is enabled if `-O2` is used.

-fno-inline-functions-called-once

Suppresses inlining of subprograms local to the unit and called once from within it, which is enabled if `-O1` is used.

-fno-ivopts

Suppresses high-level loop induction variable optimizations, which are enabled if `-O1` is used. These optimizations are generally profitable but, for some specific cases of loops with numerous uses of the iteration variable that follow a common pattern, they may end up destroying the regularity that could be exploited at a lower level and thus producing inferior code.

-fno-strict-aliasing

Causes the compiler to avoid assumptions regarding non-aliasing of objects of different types. See [Optimization and Strict Aliasing](#) for details.

-fno-strict-overflow

Causes the compiler to avoid assumptions regarding the rules of signed integer overflow. These rules specify that signed integer overflow will result in a `Constraint_Error` exception at run time and are enforced in default mode by the compiler, so this switch should not be necessary in normal operating mode. It might be useful in conjunction with `-gnatO0` for very peculiar cases of low-level programming.

-fstack-check

Activates stack checking. See [Stack Overflow Checking](#) for details.

-fstack-usage

Makes the compiler output stack usage information for the program, on a per-subprogram basis. See [Static Stack Usage Analysis](#) for details.

-g

Generate debugging information. This information is stored in the object file and copied from there to the final executable file by the linker, where it can be read by the debugger. You must use the `-g` switch if you plan on using the debugger.

-gnat05

Allow full Ada 2005 features.

-gnat12

Allow full Ada 2012 features.

-gnat2005

Allow full Ada 2005 features (same as `-gnat05`)

-gnat2012

Allow full Ada 2012 features (same as `-gnat12`)

-gnat2022

Allow full Ada 2022 features

-gnat83

Enforce Ada 83 restrictions.

-gnat95

Enforce Ada 95 restrictions.

Note: for compatibility with some Ada 95 compilers which support only the `overriding` keyword of Ada 2005, the `-gnatd.D` switch can be used along with `-gnat95` to achieve a similar effect with GNAT.

`-gnatd.D` instructs GNAT to consider `overriding` as a keyword and handle its associated semantic checks, even in Ada 95 mode.

-gnata

Assertions enabled. `Pragma Assert` and `pragma Debug` to be activated. Note that these pragmas can also be controlled using the configuration pragmas `Assertion_Policy` and `Debug_Policy`. It also activates pragmas `Check`, `Precondition`, and `Postcondition`. Note that these pragmas can also be controlled using the configuration pragma `Check_Policy`. In Ada 2012, it also activates all assertions defined in the RM as aspects: preconditions, postconditions, type invariants and (sub)type predicates. In all Ada modes, corresponding pragmas for type invariants and (sub)type predicates are also activated. The default is that all these assertions are disabled, and have no effect, other than being checked for syntactic validity, and in the case of subtype predicates, constructions such as membership tests still test predicates even if assertions are turned off.

-gnatA

Avoid processing `gnat.adc`. If a `gnat.adc` file is present, it will be ignored.

-gnatb

Generate brief messages to `stderr` even if verbose mode set.

-gnatB

Assume no invalid (bad) values except for 'Valid attribute use ([Validity Checking](#)).

-gnatc

Check syntax and semantics only (no code generation attempted). When the compiler is invoked by `gnatmake`, if the switch `-gnatc` is only given to the compiler (after `-cargs` or in package `Compiler` of the project file), `gnatmake` will fail because it will not find the object file after compilation. If `gnatmake` is called with `-gnatc` as a builder switch (before `-cargs` or in

package Builder of the project file) then `gnatmake` will not fail because it will not look for the object files after compilation, and it will not try to build and link.

-gnatC

Generate CodePeer intermediate format (no code generation attempted). This switch will generate an intermediate representation suitable for use by CodePeer (`.scil` files). This switch is not compatible with code generation (it will, among other things, disable some switches such as `-gnatn`, and enable others such as `-gnata`).

-gnatd

Specify debug options for the compiler. The string of characters after the `-gnatd` specifies the specific debug options. The possible characters are 0-9, a-z, A-Z, optionally preceded by a dot or underscore. See compiler source file `debug.adb` for details of the implemented debug options. Certain debug options are relevant to applications programmers, and these are documented at appropriate points in this users guide.

-gnatD

Create expanded source files for source level debugging. This switch also suppresses generation of cross-reference information (see `-gnatx`). Note that this switch is not allowed if a previous `-gnatR` switch has been given, since these two switches are not compatible.

-gnateA

Check that the actual parameters of a subprogram call are not aliases of one another. To qualify as aliasing, their memory locations must be identical or overlapping, at least one of the corresponding formal parameters must be of mode OUT or IN OUT, and at least one of the corresponding formal parameters must have its parameter passing mechanism not specified.

```
type Rec_Typ is record
  Data : Integer := 0;
end record;

function Self (Val : Rec_Typ) return Rec_Typ is
begin
  return Val;
end Self;

procedure Detect_Aliasing (Val_1 : in out Rec_Typ; Val_2 : Rec_Typ) is
begin
  null;
end Detect_Aliasing;

Obj : Rec_Typ;

Detect_Aliasing (Obj, Obj);
Detect_Aliasing (Obj, Self (Obj));
```

In the example above, the first call to `Detect_Aliasing` fails with a `Program_Error` at run time because the actuals for `Val_1` and `Val_2` denote the same object. The second call executes without raising an exception because `Self(Obj)` produces an anonymous object which does not share the memory location of `Obj`.

-gnateb

Store configuration files by their basename in ALI files. This switch is used for instance by gprbuild for distributed builds in order to prevent issues where machine-specific absolute paths could end up being stored in ALI files.

-gnatec=*path*

Specify a configuration pragma file (the equal sign is optional) ([The Configuration Pragma Files](#)).

-gnateC

Generate CodePeer messages in a compiler-like format. This switch is only effective if `-gnatcC` is also specified and requires an installation of CodePeer.

-gnated

Disable atomic synchronization

-gnateDsymbol[=*value***]**

Defines a symbol, associated with *value*, for preprocessing. ([Integrated Preprocessing](#)).

-gnateE

Generate extra information in exception messages. In particular, display extra column information and the value and range associated with index and range check failures, and extra column information for access checks. In cases where the compiler is able to determine at compile time that a check will fail, it gives a warning, and the extra information is not produced at run time.

-gnatef

Display full source path name in brief error messages and absolute paths in `-fdiagnostics-format=json`'s output.

-gnateF

Check for overflow on all floating-point operations, including those for unconstrained predefined types. See description of pragma `Check_Float_Overflow` in GNAT RM.

-gnateg -gnatceg

The `-gnatc` switch must always be specified before this switch, e.g. `-gnatceg`. Generate a C header from the Ada input file. See [Generating C Headers for Ada Specifications](#) for more information.

-gnateG

Save result of preprocessing in a text file.

-gnateinnn

Set maximum number of instantiations during compilation of a single unit to *nnn*. This may be useful in increasing the default maximum of 8000 for the rare case when a single unit legitimately exceeds this limit.

-gnateInnn

Indicates that the source is a multi-unit source and that the index of the unit to compile is *nnn*. *nnn* needs to be a positive number and need to be a valid index in the multi-unit source.

-gnatel

This switch can be used with the static elaboration model to issue info messages showing where implicit pragma `Elaborate` and pragma `Elaborate_All` are generated. This is useful in

diagnosing elaboration circularities caused by these implicit pragmas when using the static elaboration model. See the section in this guide on elaboration checking for further details. These messages are not generated by default, and are intended only for temporary use when debugging circularity problems.

-gnateL

This switch turns off the info messages about implicit elaboration pragmas.

-gnatem=*path*

Specify a mapping file (the equal sign is optional) ([Units to Sources Mapping Files](#)).

-gnatep=*file*

Specify a preprocessing data file (the equal sign is optional) ([Integrated Preprocessing](#)).

-gnateP

Turn categorization dependency errors into warnings. Ada requires that units that WITH one another have compatible categories, for example a Pure unit cannot WITH a Preelaborate unit. If this switch is used, these errors become warnings (which can be ignored, or suppressed in the usual manner). This can be useful in some specialized circumstances such as the temporary use of special test software.

-gnateS

Synonym of `-fdump-scos`, kept for backwards compatibility.

-gnatet=*path*

Generate target dependent information. The format of the output file is described in the section about switch `-gnateT`.

-gnateT=*path*

Read target dependent information, such as endianness or sizes and alignments of base type. If this switch is passed, the default target dependent information of the compiler is replaced by the one read from the input file. This is used by tools other than the compiler, e.g. to do semantic analysis of programs that will run on some other target than the machine on which the tool is run.

The following target dependent values should be defined, where `Nat` denotes a natural integer value, `Pos` denotes a positive integer value, and fields marked with a question mark are boolean fields, where a value of 0 is False, and a value of 1 is True:

<code>Bits_BE</code>	: Nat; -- Bits stored big-endian?
<code>Bits_Per_Unit</code>	: Pos; -- Bits in a storage unit
<code>Bits_Per_Word</code>	: Pos; -- Bits in a word
<code>Bytes_BE</code>	: Nat; -- Bytes stored big-endian?
<code>Char_Size</code>	: Pos; -- Standard.Character'Size
<code>Double_Float_Alignment</code>	: Nat; -- Alignment of double float
<code>Double_Scalar_Alignment</code>	: Nat; -- Alignment of double length scalar
<code>Double_Size</code>	: Pos; -- Standard.Long_Float'Size
<code>Float_Size</code>	: Pos; -- Standard.Float'Size
<code>Float_Words_BE</code>	: Nat; -- Float words stored big-endian?
<code>Int_Size</code>	: Pos; -- Standard.Integer'Size
<code>Long_Double_Size</code>	: Pos; -- Standard.Long_Long_Float'Size
<code>Long_Long_Long_Size</code>	: Pos; -- Standard.Long_Long_Long_Integer'Size
<code>Long_Long_Size</code>	: Pos; -- Standard.Long_Long_Integer'Size

(continues on next page)

(continued from previous page)

Long_Size	: Pos; -- Standard.Long_Integer'Size
Maximum_Alignment	: Pos; -- Maximum permitted alignment
Max_Unaligned_Field	: Pos; -- Maximum size for unaligned bit field
Pointer_Size	: Pos; -- System.Address'Size
Short_Enums	: Nat; -- Foreign enums use short size?
Short_Size	: Pos; -- Standard.Short_Integer'Size
Strict_Alignment	: Nat; -- Strict alignment?
System_Allocator_Alignment	: Nat; -- Alignment for malloc calls
Wchar_T_Size	: Pos; -- Interfaces.C.wchar_t'Size
Words_BE	: Nat; -- Words stored big-endian?

`Bits_Per_Unit` is the number of bits in a storage unit, the equivalent of GCC macro `BITS_PER_UNIT` documented as follows: *Define this macro to be the number of bits in an addressable storage unit (byte); normally 8.*

`Bits_Per_Word` is the number of bits in a machine word, the equivalent of GCC macro `BITS_PER_WORD` documented as follows: *Number of bits in a word; normally 32.*

`Double_Float_Alignment`, if not zero, is the maximum alignment that the compiler can choose by default for a 64-bit floating-point type or object.

`Double_Scalar_Alignment`, if not zero, is the maximum alignment that the compiler can choose by default for a 64-bit or larger scalar type or object.

`Maximum_Alignment` is the maximum alignment that the compiler can choose by default for a type or object, which is also the maximum alignment that can be specified in GNAT. It is computed for GCC backends as `BIGGEST_ALIGNMENT / BITS_PER_UNIT` where GCC macro `BIGGEST_ALIGNMENT` is documented as follows: *Biggest alignment that any data type can require on this machine, in bits.*

`Max_Unaligned_Field` is the maximum size for unaligned bit field, which is 64 for the majority of GCC targets (but can be different on some targets).

`Strict_Alignment` is the equivalent of GCC macro `STRICT_ALIGNMENT` documented as follows: *Define this macro to be the value 1 if instructions will fail to work if given data not on the nominal alignment. If instructions will merely go slower in that case, define this macro as 0.*

`System_Allocator_Alignment` is the guaranteed alignment of data returned by calls to `malloc`.

The format of the input file is as follows. First come the values of the variables defined above, with one line per value:

name	value
------	-------

where `name` is the name of the parameter, spelled out in full, and cased as in the above list, and `value` is an unsigned decimal integer. Two or more blanks separates the name from the value.

All the variables must be present, in alphabetical order (i.e. the same order as the list above).

Then there is a blank line to separate the two parts of the file. Then come the lines showing the floating-point types to be registered, with one line per registered mode:

```
name  digs float_rep size alignment
```

where `name` is the string name of the type (which can have single spaces embedded in the name, e.g. `long double`), `digs` is the number of digits for the floating-point type, `float_rep` is the float representation (I for IEEE-754-Binary, which is the only one supported at this time), `size` is the size in bits, `alignment` is the alignment in bits. The name is followed by at least two blanks, fields are separated by at least one blank, and a LF character immediately follows the alignment field.

Here is an example of a target parameterization file:

```
Bits_BE           0
Bits_Per_Unit     8
Bits_Per_Word     64
Bytes_BE          0
Char_Size         8
Double_Float_Alignment  0
Double_Scalar_Alignment  0
Double_Size       64
Float_Size        32
Float_Words_BE    0
Int_Size          64
Long_Double_Size  128
Long_Long_Long_Size  128
Long_Long_Size    64
Long_Size         64
Maximum_Alignment  16
Max_Unaligned_Field  64
Pointer_Size      64
Short_Size        16
Strict_Alignment  0
System_Allocator_Alignment  16
Wchar_T_Size      32
Words_BE          0

float           15  I  64  64
double          15  I  64  64
long double     18  I  80  128
TF              33  I 128  128
```

-gnateu

Ignore unrecognized validity, warning, and style switches that appear after this switch is given. This may be useful when compiling sources developed on a later version of the compiler with an earlier version. Of course the earlier version must support this switch.

-gnateV

Check that all actual parameters of a subprogram call are valid according to the rules of validity checking ([Validity Checking](#)).

-gnateY

Ignore all `STYLE_CHECKS` pragmas. Full legality checks are still carried out, but the pragmas have no effect on what style checks are active. This allows all style checking options

to be controlled from the command line.

-gnatE

Dynamic elaboration checking mode enabled. For further details see [Elaboration Order Handling in GNAT](#).

-gnatf

Full errors. Multiple errors per line, all undefined references, do not attempt to suppress cascaded errors.

-gnatF

Externals names are folded to all uppercase.

-gnatg

Internal GNAT implementation mode. This should not be used for applications programs, it is intended only for use by the compiler and its run-time library. For documentation, see the GNAT sources. Note that **-gnatg** implies **-gnatw.ge** and **-gnatyg** so that all standard warnings and all standard style options are turned on. All warnings and style messages are treated as errors.

-gnatG=nn

List generated expanded code in source form.

-gnath

Output usage information. The output is written to `stdout`.

-gnathH

Legacy elaboration-checking mode enabled. When this switch is in effect, the pre-18.x access-before-elaboration model becomes the de facto model. For further details see [Elaboration Order Handling in GNAT](#).

-gnatic

Identifier character set ($c = 1/2/3/4/5/9/p/8/f/n/w$). For details of the possible selections for `c`, see [Character Set Control](#).

-gnatI

Ignore representation clauses. When this switch is used, representation clauses are treated as comments. This is useful when initially porting code where you want to ignore rep clause problems, and also for compiling foreign code (particularly for use with ASIS). The representation clauses that are ignored are: `enumeration_representation_clause`, `record_representation_clause`, and `attribute_definition_clause` for the following attributes: `Address`, `Alignment`, `Bit_Order`, `Component_Size`, `Machine_Radix`, `Object_Size`, `Scalar_Storage_Order`, `Size`, `Small`, `Stream_Size`, and `Value_Size`. `Pragma_Default_Scalar_Storage_Order` is also ignored. Note that this option should be used only for compiling – the code is likely to malfunction at run time.

-gnatjnn

Reformat error messages to fit on `nn` character lines

-gnatJ

Permissive elaboration-checking mode enabled. When this switch is in effect, the post-18.x access-before-elaboration model ignores potential issues with:

- Accept statements

- Activations of tasks defined in instances
- Assertion pragmas
- Calls from within an instance to its enclosing context
- Calls through generic formal parameters
- Calls to subprograms defined in instances
- Entry calls
- Indirect calls using 'Access
- Requeue statements
- Select statements
- Synchronous task suspension

and does not emit compile-time diagnostics or run-time checks. For further details see [Elaboration Order Handling in GNAT](#).

-gnatk=*n*

Limit file names to *n* (1-999) characters (*k* = crunch).

-gnatl

Output full source listing with embedded error messages.

-gnatL

Used in conjunction with -gnatG or -gnatD to intersperse original source lines (as comment lines with line numbers) in the expanded source output.

-gnatm=*n*

Limit number of detected error or warning messages to *n* where *n* is in the range 1..999999. The default setting if no switch is given is 9999. If the number of warnings reaches this limit, then a message is output and further warnings are suppressed, but the compilation is continued. If the number of error messages reaches this limit, then a message is output and the compilation is abandoned. The equal sign here is optional. A value of zero means that no limit applies.

-gnatn[12]

Activate inlining across units for subprograms for which pragma `InLine` is specified. This inlining is performed by the GCC back-end. An optional digit sets the inlining level: 1 for moderate inlining across units or 2 for full inlining across units. If no inlining level is specified, the compiler will pick it based on the optimization level.

-gnatN

Activate front end inlining for subprograms for which pragma `InLine` is specified. This inlining is performed by the front end and will be visible in the -gnatG output.

When using a gcc-based back end, then the use of -gnatN is deprecated, and the use of -gnatn is preferred. Historically front end inlining was more extensive than the gcc back end inlining, but that is no longer the case.

-gnato0

Suppresses overflow checking. This causes the behavior of the compiler to match the default

for older versions where overflow checking was suppressed by default. This is equivalent to having `pragma Suppress (Overflow_Check)` in a configuration pragma file.

-gnato??

Set default mode for handling generation of code to avoid intermediate arithmetic overflow. Here ?? is two digits, a single digit, or nothing. Each digit is one of the digits 1 through 3:

Digit	Interpretation
1	All intermediate overflows checked against base type (STRICT)
2	Minimize intermediate overflows (MINIMIZED)
3	Eliminate intermediate overflows (ELIMINATED)

If only one digit appears, then it applies to all cases; if two digits are given, then the first applies outside assertions, pre/postconditions, and type invariants, and the second applies within assertions, pre/postconditions, and type invariants.

If no digits follow the `-gnato`, then it is equivalent to `-gnato11`, causing all intermediate overflows to be handled in strict mode.

This switch also causes arithmetic overflow checking to be performed (as though `pragma Unsuppress (Overflow_Check)` had been specified).

The default if no option `-gnato` is given is that overflow handling is in **STRICT** mode (computations done using the base type), and that overflow checking is enabled.

Note that division by zero is a separate check that is not controlled by this switch (divide-by-zero checking is on by default).

See also [Specifying the Desired Mode](#).

-gnatp

Suppress all checks. See [Run-Time Checks](#) for details. This switch has no effect if cancelled by a subsequent `-gnat-p` switch.

-gnat-p

Cancel effect of previous `-gnatp` switch.

-gnatq

Don't quit. Try semantics, even if parse errors.

-gnatQ

Don't quit. Generate ALI and tree files even if illegalities. Note that code generation is still suppressed in the presence of any errors, so even with `-gnatQ` no object file is generated.

-gnatr

Treat pragma Restrictions as `Restriction_Warnings`.

-gnatR[0|1|2|3|4][e][j][m][s]

Output representation information for declared types, objects and subprograms. Note that this switch is not allowed if a previous `-gnatD` switch has been given, since these two switches are not compatible.

-gnats

Syntax check only.

- gnatS**
Print package Standard.
- gnatTnnn**
All compiler tables start at nnn times usual starting size.
- gnatu**
List units for this compilation.
- gnatU**
Tag all error messages with the unique string 'error:'
- gnatv**
Verbose mode. Full error output with source lines to `stdout`.
- gnatV**
Control level of validity checking ([Validity Checking](#)).
- gnatwxxx**
Warning mode where xxx is a string of option letters that denotes the exact warnings that are enabled or disabled ([Warning Message Control](#)).
- gnatWe**
Wide character encoding method (e=n/h/u/s/e/8).
- gnatx**
Suppress generation of cross-reference information.
- gnatX**
Enable core GNAT implementation extensions and latest Ada version.
- gnatX0**
Enable all GNAT implementation extensions and latest Ada version.
- gnaty**
Enable built-in style checks ([Style Checking](#)).
- gnatzm**
Distribution stub generation and compilation (m=r/c for receiver/caller stubs).
- Idir**
Direct GNAT to search the `dir` directory for source files needed by the current compilation (see [Search Paths](#) and the [Run-Time Library \(RTL\)](#)).
- I-**
Except for the source file named in the command line, do not look for source files in the directory containing the source file named in the command line (see [Search Paths](#) and the [Run-Time Library \(RTL\)](#)).
- o file**
This switch is used in `gcc` to redirect the generated object file and its associated ALI file. Beware of this switch with GNAT, because it may cause the object file and ALI file to have different names which in turn may confuse the binder and the linker.
- nostdinc**
Inhibit the search of the default location for the GNAT Run Time Library (RTL) source files.

-nostdlib

Inhibit the search of the default location for the GNAT Run Time Library (RTL) ALI files.

-O[n]

n controls the optimization level:

<i>n</i>	Effect
<i>0</i>	No optimization, the default setting if no <code>-O</code> appears
<i>1</i>	Normal optimization, the default if you specify <code>-O</code> without an operand. A good compromise between code quality and compilation time.
<i>2</i>	Extensive optimization, may improve execution time, possibly at the cost of substantially increased compilation time.
<i>3</i>	Same as <code>-O2</code> , and also includes inline expansion for small subprograms in the same unit.
<i>s</i>	Optimize space usage

See also [Optimization Levels](#).

-pass-exit-codes

Catch exit codes from the compiler and use the most meaningful as exit status.

--RTS=*rts-path*

Specifies the default location of the run-time library. Same meaning as the equivalent `gnatmake` flag ([Switches for gnatmake](#)).

-S

Used in place of `-c` to cause the assembler source file to be generated, using `.s` as the extension, instead of the object file. This may be useful if you need to examine the generated assembly code.

-fverbose-asm

Used in conjunction with `-S` to cause the generated assembly code file to be annotated with variable names, making it significantly easier to follow.

-v

Show commands generated by the `gcc` driver. Normally used only for debugging purposes or if you need to be sure what version of the compiler you are executing.

-V *ver*

Execute `ver` version of the compiler. This is the `gcc` version, not the GNAT version.

-w

Turn off warnings generated by the back end of the compiler. Use of this switch also causes the default for front end warnings to be set to suppress (as though `-gnatws` had appeared at the start of the options).

You may combine a sequence of GNAT switches into a single switch. For example, the combined switch

```
-gnatofi3
```

is equivalent to specifying the following sequence of switches:

```
-gnato -gnatf -gnati3
```

The following restrictions apply to the combination of switches in this manner:

- The switch `-gnatc` if combined with other switches must come first in the string.
- The switch `-gnats` if combined with other switches must come first in the string.
- The switches `-gnatzc` and `-gnatzr` may not be combined with any other switches, and only one of them may appear in the command line.
- The switch `-gnat-p` may not be combined with any other switch.
- Once a 'y' appears in the string (that is a use of the `-gnaty` switch), then all further characters in the switch are interpreted as style modifiers (see description of `-gnaty`).
- Once a 'd' appears in the string (that is a use of the `-gnatd` switch), then all further characters in the switch are interpreted as debug flags (see description of `-gnatd`).
- Once a 'w' appears in the string (that is a use of the `-gnatw` switch), then all further characters in the switch are interpreted as warning mode modifiers (see description of `-gnatw`).
- Once a 'V' appears in the string (that is a use of the `-gnatV` switch), then all further characters in the switch are interpreted as validity checking options ([Validity Checking](#)).
- Option 'em', 'ec', 'ep', 'l=' and 'R' must be the last options in a combined list of options.

4.3.2 Output and Error Message Control

The standard default format for error messages is called 'brief format'. Brief format messages are written to `stderr` (the standard error file) and have the following form:

```
e.adb:3:04: Incorrect spelling of keyword "function"  
e.adb:4:20: ";" should be "is"
```

The first integer after the file name is the line number in the file, and the second integer is the column number within the line. `GNAT Studio` can parse the error messages and point to the referenced character. The following switches provide control over the error message format:

-gnatv

The `v` stands for verbose. The effect of this setting is to write long-format error messages to `stdout` (the standard output file). The same program compiled with the `-gnatv` switch would generate:

```
3. function X (Q : Integer)  
   |  
>>> Incorrect spelling of keyword "function"  
4. return Integer;  
   |  
>>> ";" should be "is"
```

The vertical bar indicates the location of the error, and the >>> prefix can be used to search for error messages. When this switch is used the only source lines output are those with errors.

-gnatl

The `l` stands for list. This switch causes a full listing of the file to be generated. In the case where a body is compiled, the corresponding spec is also listed, along with any subunits. Typical output from compiling a package body `p.adb` might look like:

```
Compiling: p.adb

  1. package body p is
  2.   procedure a;
  3.   procedure a is separate;
  4. begin
  5.   null
      |
      >>> missing ";"

  6. end;

Compiling: p.ads

  1. package p is
  2.   pragma Elaborate_Body
      |
      >>> missing ";"

  3. end p;

Compiling: p-a.adb

  1. separate p
      |
      >>> missing "("

  2. procedure a is
  3. begin
  4.   null
      |
      >>> missing ";"

  5. end;
```

When you specify the `-gnatv` or `-gnatl` switches and standard output is redirected, a brief summary is written to `stderr` (standard error) giving the number of error messages and warning messages generated.

-gnatl=fname

This has the same effect as `-gnatl` except that the output is written to a file instead of to standard output. If the given name `fname` does not start with a period, then it is the full name of the file to be written. If `fname` is an extension, it is appended to the name of the file being compiled. For example, if file `xyz.adb` is compiled with `-gnatl=.lst`, then the output

is written to file xyz.adb.lst.

-gnatU

This switch forces all error messages to be preceded by the unique string 'error:'. This means that error messages take a few more characters in space, but allows easy searching for and identification of error messages.

-gnatb

The **b** stands for brief. This switch causes GNAT to generate the brief format error messages to `stderr` (the standard error file) as well as the verbose format message or full listing (which as usual is written to `stdout`, the standard output file).

-gnatm=*n*

The **m** stands for maximum. *n* is a decimal integer in the range of 1 to 999999 and limits the number of error or warning messages to be generated. For example, using `-gnatm2` might yield

```
e.adb:3:04: Incorrect spelling of keyword "function"
e.adb:5:35: missing ".."
fatal error: maximum number of errors detected
compilation abandoned
```

The default setting if no switch is given is 9999. If the number of warnings reaches this limit, then a message is output and further warnings are suppressed, but the compilation is continued. If the number of error messages reaches this limit, then a message is output and the compilation is abandoned. A value of zero means that no limit applies.

Note that the equal sign is optional, so the switches `-gnatm2` and `-gnatm=2` are equivalent.

-gnatf

The **f** stands for full. Normally, the compiler suppresses error messages that are likely to be redundant. This switch causes all error messages to be generated. In particular, in the case of references to undefined variables. If a given variable is referenced several times, the normal format of messages is

```
e.adb:7:07: "V" is undefined (more references follow)
```

where the parenthetical comment warns that there are additional references to the variable *V*. Compiling the same program with the `-gnatf` switch yields

```
e.adb:7:07: "V" is undefined
e.adb:8:07: "V" is undefined
e.adb:8:12: "V" is undefined
e.adb:8:16: "V" is undefined
e.adb:9:07: "V" is undefined
e.adb:9:12: "V" is undefined
```

The `-gnatf` switch also generates additional information for some error messages. Some examples are:

- Details on possibly non-portable unchecked conversion
- List possible interpretations for ambiguous calls

- Additional details on incorrect parameters

-gnatjnn

In normal operation mode (or if `-gnatj0` is used), then error messages with continuation lines are treated as though the continuation lines were separate messages (and so a warning with two continuation lines counts as three warnings, and is listed as three separate messages).

If the `-gnatjnn` switch is used with a positive value for `nn`, then messages are output in a different manner. A message and all its continuation lines are treated as a unit, and count as only one warning or message in the statistics totals. Furthermore, the message is reformatted so that no line is longer than `nn` characters.

-gnatq

The `q` stands for quit (really ‘don’t quit’). In normal operation mode, the compiler first parses the program and determines if there are any syntax errors. If there are, appropriate error messages are generated and compilation is immediately terminated. This switch tells GNAT to continue with semantic analysis even if syntax errors have been found. This may enable the detection of more errors in a single run. On the other hand, the semantic analyzer is more likely to encounter some internal fatal error when given a syntactically invalid tree.

-gnatQ

In normal operation mode, the ALI file is not generated if any illegalities are detected in the program. The use of `-gnatQ` forces generation of the ALI file. This file is marked as being in error, so it cannot be used for binding purposes, but it does contain reasonably complete cross-reference information, and thus may be useful for use by tools (e.g., semantic browsing tools or integrated development environments) that are driven from the ALI file. This switch implies `-gnatq`, since the semantic phase must be run to get a meaningful ALI file.

When `-gnatQ` is used and the generated ALI file is marked as being in error, `gnatmake` will attempt to recompile the source when it finds such an ALI file, including with switch `-gnatc`.

Note that `-gnatQ` has no effect if `-gnats` is specified, since ALI files are never generated if `-gnats` is set.

4.3.3 Warning Message Control

In addition to error messages, which correspond to illegalities as defined in the Ada Reference Manual, the compiler detects two kinds of warning situations.

First, the compiler considers some constructs suspicious and generates a warning message to alert you to a possible error. Second, if the compiler detects a situation that is sure to raise an exception at run time, it generates a warning message. The following shows an example of warning messages:

```
e.adb:4:24: warning: creation of object may raise Storage_Error
e.adb:10:17: warning: static value out of range
e.adb:10:17: warning: "Constraint_Error" will be raised at run time
```

GNAT considers a large number of situations as appropriate for the generation of warning messages. As always, warnings are not definite indications of errors. For example, if you do an out-of-range assignment with the deliberate intention of raising a `Constraint_Error` exception, then the warning that may be issued does not indicate an error. Some of the situations for which GNAT issues

warnings (at least some of the time) are given in the following list. This list is not complete, and new warnings are often added to subsequent versions of GNAT. The list is intended to give a general idea of the kinds of warnings that are generated.

- Possible infinitely recursive calls
- Out-of-range values being assigned
- Possible order of elaboration problems
- Size not a multiple of alignment for a record type
- Assertions (pragma Assert) that are sure to fail
- Unreachable code
- Address clauses with possibly unaligned values, or where an attempt is made to overlay a smaller variable with a larger one.
- Fixed-point type declarations with a null range
- Direct_IO or Sequential_IO instantiated with a type that has access values
- Variables that are never assigned a value
- Variables that are referenced before being initialized
- Task entries with no corresponding `accept` statement
- Duplicate accepts for the same task entry in a `select`
- Objects that take too much storage
- Unchecked conversion between types of differing sizes
- Missing `return` statement along some execution path in a function
- Incorrect (unrecognized) pragmas
- Incorrect external names
- Allocation from empty storage pool
- Potentially blocking operation in protected type
- Suspicious parenthesization of expressions
- Mismatching bounds in an aggregate
- Attempt to return local value by reference
- Premature instantiation of a generic body
- Attempt to pack aliased components
- Out of bounds array subscripts
- Wrong length on string assignment
- Violations of style rules if style checking is enabled
- Unused *with* clauses

- `Bit_Order` usage that does not have any effect
- `Standard.Duration` used to resolve universal fixed expression
- Dereference of possibly null value
- Declaration that is likely to cause storage error
- Internal GNAT unit *withed* by application unit
- Values known to be out of range at compile time
- Unreferenced or unmodified variables. Note that a special exemption applies to variables which contain any of the substrings `DISCARD`, `DUMMY`, `IGNORE`, `JUNK`, `UNUSED`, in any casing. Such variables are considered likely to be intentionally used in a situation where otherwise a warning would be given, so warnings of this kind are always suppressed for such variables.
- Address overlays that could clobber memory
- Unexpected initialization when address clause present
- Bad alignment for address clause
- Useless type conversions
- Redundant assignment statements and other redundant constructs
- Useless exception handlers
- Accidental hiding of name by child unit
- Access before elaboration detected at compile time
- A range in a `for` loop that is known to be null or might be null

The following section lists compiler switches that are available to control the handling of warning messages. It is also possible to exercise much finer control over what warnings are issued and suppressed using the GNAT pragma `Warnings` (see the description of the pragma in the *GNAT_Reference_manual*).

-gnatwa

Activate most optional warnings.

This switch activates most optional warning messages. See the remaining list in this section for details on optional warning messages that can be individually controlled. The warnings that are not turned on by this switch are:

- `-gnatwd` (implicit dereferencing)
- `-gnatw.d` (tag warnings with `-gnatw` switch)
- `-gnatwh` (hiding)
- `-gnatw.h` (holes in record layouts)
- `-gnatw.j` (late primitives of tagged types)
- `-gnatw.k` (redefinition of names in standard)
- `-gnatwl` (elaboration warnings)

- `-gnatw.l` (inherited aspects)
- `-gnatw.n` (atomic synchronization)
- `-gnatwo` (address clause overlay)
- `-gnatw.o` (values set by out parameters ignored)
- `-gnatw.q` (questionable layout of record types)
- `-gnatw_q` (ignored equality)
- `-gnatw_r` (out-of-order record representation clauses)
- `-gnatw.s` (overridden size clause)
- `-gnatwt` (tracking of deleted conditional code)
- `-gnatw.u` (unordered enumeration)
- `-gnatw.w` (use of Warnings Off)
- `-gnatw.y` (reasons for package needing body)

All other optional warnings are turned on.

-gnatwA

Suppress all optional errors.

This switch suppresses all optional warning messages, see remaining list in this section for details on optional warning messages that can be individually controlled. Note that unlike switch `-gnatws`, the use of switch `-gnatwA` does not suppress warnings that are normally given unconditionally and cannot be individually controlled (for example, the warning about a missing exit path in a function). Also, again unlike switch `-gnatws`, warnings suppressed by the use of switch `-gnatwA` can be individually turned back on. For example the use of switch `-gnatwA` followed by switch `-gnatwd` will suppress all optional warnings except the warnings for implicit dereferencing.

-gnatw.a

Activate warnings on failing assertions.

This switch activates warnings for assertions where the compiler can tell at compile time that the assertion will fail. Note that this warning is given even if assertions are disabled. The default is that such warnings are generated.

-gnatw.A

Suppress warnings on failing assertions.

This switch suppresses warnings for assertions where the compiler can tell at compile time that the assertion will fail.

-gnatw_a

Activate warnings on anonymous allocators.

This switch activates warnings for allocators of anonymous access types, which can involve run-time accessibility checks and lead to unexpected accessibility violations. For more details on the rules involved, see RM 3.10.2 (14).

-gnatw_A

Suppress warnings on anonymous allocators.

This switch suppresses warnings for anonymous access type allocators.

-gnatwb

Activate warnings on bad fixed values.

This switch activates warnings for static fixed-point expressions whose value is not an exact multiple of Small. Such values are implementation dependent, since an implementation is free to choose either of the multiples that surround the value. GNAT always chooses the closer one, but this is not required behavior, and it is better to specify a value that is an exact multiple, ensuring predictable execution. The default is that such warnings are not generated.

-gnatwB

Suppress warnings on bad fixed values.

This switch suppresses warnings for static fixed-point expressions whose value is not an exact multiple of Small.

-gnatw.b

Activate warnings on biased representation.

This switch activates warnings when a size clause, value size clause, component clause, or component size clause forces the use of biased representation for an integer type (e.g. representing a range of 10..11 in a single bit by using 0/1 to represent 10/11). The default is that such warnings are generated.

-gnatw.B

Suppress warnings on biased representation.

This switch suppresses warnings for representation clauses that force the use of biased representation.

-gnatwc

Activate warnings on conditionals.

This switch activates warnings for conditional expressions used in tests that are known to be True or False at compile time. The default is that such warnings are not generated. Note that this warning does not get issued for the use of boolean constants whose values are known at compile time, since this is a standard technique for conditional compilation in Ada, and this would generate too many false positive warnings.

This warning option also activates a special test for comparisons using the operators '>=' and '<='. If the compiler can tell that only the equality condition is possible, then it will warn that the '>' or '<' part of the test is useless and that the operator could be replaced by '='. An example would be comparing a `Natural` variable `<= 0`.

This warning option also generates warnings if one or both tests is optimized away in a membership test for integer values if the result can be determined at compile time. Range tests on enumeration types are not included, since it is common for such tests to include an end point.

This warning can also be turned on using `-gnatwa`.

-gnatwC

Suppress warnings on conditionals.

This switch suppresses warnings for conditional expressions used in tests that are known to be True or False at compile time.

-gnatw.c

Activate warnings on missing component clauses.

This switch activates warnings for record components where a record representation clause is present and has component clauses for the majority, but not all, of the components. A warning is given for each component for which no component clause is present.

-gnatw.C

Suppress warnings on missing component clauses.

This switch suppresses warnings for record components that are missing a component clause in the situation described above.

-gnatw_c

Activate warnings on unknown condition in Compile_Time_Warning.

This switch activates warnings on a pragma `Compile_Time_Warning` or `Compile_Time_Error` whose condition has a value that is not known at compile time. The default is that such warnings are generated.

-gnatw_C

Suppress warnings on unknown condition in Compile_Time_Warning.

This switch suppresses warnings on a pragma `Compile_Time_Warning` or `Compile_Time_Error` whose condition has a value that is not known at compile time.

-gnatwd

Activate warnings on implicit dereferencing.

If this switch is set, then the use of a prefix of an access type in an indexed component, slice, or selected component without an explicit `.all` will generate a warning. With this warning enabled, access checks occur only at points where an explicit `.all` appears in the source code (assuming no warnings are generated as a result of this switch). The default is that such warnings are not generated.

-gnatwD

Suppress warnings on implicit dereferencing.

This switch suppresses warnings for implicit dereferences in indexed components, slices, and selected components.

-gnatw.d

Activate tagging of warning and info messages.

If this switch is set, then warning messages are tagged, with one of the following strings:

- `[-gnatw?]` Used to tag warnings controlled by the switch `-gnatwx` where `x` is a letter a-z.
- `[-gnatw.?)]` Used to tag warnings controlled by the switch `-gnatw.x` where `x` is a letter a-z.

- *[-gnatel]* Used to tag elaboration information (info) messages generated when the static model of elaboration is used and the `-gnatel` switch is set.
- *[restriction warning]* Used to tag warning messages for restriction violations, activated by use of the pragma `Restriction_Warnings`.
- *[warning-as-error]* Used to tag warning messages that have been converted to error messages by use of the pragma `Warning_As_Error`. Note that such warnings are prefixed by the string “error: ” rather than “warning: “.
- *[enabled by default]* Used to tag all other warnings that are always given by default, unless warnings are completely suppressed using pragma `Warnings(Off)` or the switch `-gnatws`.

-gnatw.D

Deactivate tagging of warning and info messages messages.

If this switch is set, then warning messages return to the default mode in which warnings and info messages are not tagged as described above for `-gnatw.d`.

-gnatwe

Treat warnings and style checks as errors.

This switch causes warning messages and style check messages to be treated as errors. The warning string still appears, but the warning messages are counted as errors, and prevent the generation of an object file. Note that this is the only `-gnatw` switch that affects the handling of style check messages. Note also that this switch has no effect on info (information) messages, which are not treated as errors if this switch is present.

-gnatw.e

Activate every optional warning.

This switch activates all optional warnings, including those which are not activated by `-gnatwa`. The use of this switch is not recommended for normal use. If you turn this switch on, it is almost certain that you will get large numbers of useless warnings. The warnings that are excluded from `-gnatwa` are typically highly specialized warnings that are suitable for use only in code that has been specifically designed according to specialized coding rules.

-gnatwE

Treat all run-time exception warnings as errors.

This switch causes warning messages regarding errors that will be raised during run-time execution to be treated as errors.

-gnatwf

Activate warnings on unreferenced formals.

This switch causes a warning to be generated if a formal parameter is not referenced in the body of the subprogram. This warning can also be turned on using `-gnatwu`. The default is that these warnings are not generated.

-gnatwF

Suppress warnings on unreferenced formals.

This switch suppresses warnings for unreferenced formal parameters. Note that the combination `-gnatwu` followed by `-gnatwF` has the effect of warning on unreferenced entities other than subprogram formals.

-gnatwg

Activate warnings on unrecognized pragmas.

This switch causes a warning to be generated if an unrecognized pragma is encountered. Apart from issuing this warning, the pragma is ignored and has no effect. The default is that such warnings are issued (satisfying the Ada Reference Manual requirement that such warnings appear).

-gnatwG

Suppress warnings on unrecognized pragmas.

This switch suppresses warnings for unrecognized pragmas.

-gnatw.g

Warnings used for GNAT sources.

This switch sets the warning categories that are used by the standard GNAT style. Currently this is equivalent to `-gnatwAao.q.s.CI.V.X.Z` but more warnings may be added in the future without advanced notice.

-gnatwh

Activate warnings on hiding.

This switch activates warnings on hiding declarations that are considered potentially confusing. Not all cases of hiding cause warnings; for example an overriding declaration hides an implicit declaration, which is just normal code. The default is that warnings on hiding are not generated.

-gnatwH

Suppress warnings on hiding.

This switch suppresses warnings on hiding declarations.

-gnatw.h

Activate warnings on holes/gaps in records.

This switch activates warnings on component clauses in record representation clauses that leave holes (gaps) in the record layout. If this warning option is active, then record representation clauses should specify a contiguous layout, adding unused fill fields if needed.

-gnatw.H

Suppress warnings on holes/gaps in records.

This switch suppresses warnings on component clauses in record representation clauses that leave holes (haps) in the record layout.

-gnatwi

Activate warnings on implementation units.

This switch activates warnings for a *with* of an internal GNAT implementation unit, defined as any unit from the `Ada`, `Interfaces`, `GNAT`, or `System` hierarchies that is not documented in either the Ada Reference Manual or the GNAT Programmer's Reference Manual. Such units

are intended only for internal implementation purposes and should not be *withed* by user programs. The default is that such warnings are generated

-gnatwI

Disable warnings on implementation units.

This switch disables warnings for a *with* of an internal GNAT implementation unit.

-gnatw.i

Activate warnings on overlapping actuals.

This switch enables a warning on statically detectable overlapping actuals in a subprogram call, when one of the actuals is an in-out parameter, and the types of the actuals are not by-copy types. This warning is off by default.

-gnatw.I

Disable warnings on overlapping actuals.

This switch disables warnings on overlapping actuals in a call.

-gnatwj

Activate warnings on obsolescent features (Annex J).

If this warning option is activated, then warnings are generated for calls to subprograms marked with `pragma Obsolescent` and for use of features in Annex J of the Ada Reference Manual. In the case of Annex J, not all features are flagged. In particular, uses of package ASCII are not flagged, since these are very common and would generate many annoying positive warnings. The default is that such warnings are not generated.

In addition to the above cases, warnings are also generated for GNAT features that have been provided in past versions but which have been superseded (typically by features in the new Ada standard). For example, `pragma Ravenscar` will be flagged since its function is replaced by `pragma Profile(Ravenscar)`, and `pragma Interface_Name` will be flagged since its function is replaced by `pragma Import`.

Note that this warning option functions differently from the restriction `No_Obsolescent_Features` in two respects. First, the restriction applies only to annex J features. Second, the restriction does flag uses of package ASCII.

-gnatwJ

Suppress warnings on obsolescent features (Annex J).

This switch disables warnings on use of obsolescent features.

-gnatw.j

Activate warnings on late declarations of tagged type primitives.

This switch activates warnings on visible primitives added to a tagged type after deriving a private extension from it.

-gnatw.J

Suppress warnings on late declarations of tagged type primitives.

This switch suppresses warnings on visible primitives added to a tagged type after deriving a private extension from it.

-gnatwk

Activate warnings on variables that could be constants.

This switch activates warnings for variables that are initialized but never modified, and then could be declared constants. The default is that such warnings are not given.

-gnatwK

Suppress warnings on variables that could be constants.

This switch disables warnings on variables that could be declared constants.

-gnatw.k

Activate warnings on redefinition of names in standard.

This switch activates warnings for declarations that declare a name that is defined in package Standard. Such declarations can be confusing, especially since the names in package Standard continue to be directly visible, meaning that use visibility on such redeclared names does not work as expected. Names of discriminants and components in records are not included in this check.

-gnatw.K

Suppress warnings on redefinition of names in standard.

This switch disables warnings for declarations that declare a name that is defined in package Standard.

-gnatwl

Activate warnings for elaboration pragmas.

This switch activates warnings for possible elaboration problems, including suspicious use of Elaborate pragmas, when using the static elaboration model, and possible situations that may raise Program_Error when using the dynamic elaboration model. See the section in this guide on elaboration checking for further details. The default is that such warnings are not generated.

-gnatwL

Suppress warnings for elaboration pragmas.

This switch suppresses warnings for possible elaboration problems.

-gnatw.l

List inherited aspects.

This switch causes the compiler to list inherited invariants, preconditions, and postconditions from Type_Invariant'Class, Invariant'Class, Pre'Class, and Post'Class aspects. Also list inherited subtype predicates.

-gnatw.L

Suppress listing of inherited aspects.

This switch suppresses listing of inherited aspects.

-gnatwm

Activate warnings on modified but unreferenced variables.

This switch activates warnings for variables that are assigned (using an initialization value or with one or more assignment statements) but whose value is never read. The warning is suppressed for volatile variables and also for variables that are renamings of other variables or for which an address clause is given. The default is that these warnings are not given.

-gnatwM

Disable warnings on modified but unreferenced variables.

This switch disables warnings for variables that are assigned or initialized, but never read.

-gnatw.m

Activate warnings on suspicious modulus values.

This switch activates warnings for modulus values that seem suspicious. The cases caught are where the size is the same as the modulus (e.g. a modulus of 7 with a size of 7 bits), and modulus values of 32 or 64 with no size clause. The guess in both cases is that 2^*x was intended rather than x . In addition expressions of the form 2^*x for small x generate a warning (the almost certainly accurate guess being that 2^*x was intended). This switch also activates warnings for negative literal values of a modular type, which are interpreted as large positive integers after wrap-around. The default is that these warnings are given.

-gnatw.M

Disable warnings on suspicious modulus values.

This switch disables warnings for suspicious modulus values.

-gnatwn

Set normal warnings mode.

This switch sets normal warning mode, in which enabled warnings are issued and treated as warnings rather than errors. This is the default mode. the switch `-gnatwn` can be used to cancel the effect of an explicit `-gnatws` or `-gnatwe`. It also cancels the effect of the implicit `-gnatwe` that is activated by the use of `-gnatg`.

-gnatw.n

Activate warnings on atomic synchronization.

This switch activates warnings when an access to an atomic variable requires the generation of atomic synchronization code. These warnings are off by default.

-gnatw.N

Suppress warnings on atomic synchronization.

This switch suppresses warnings when an access to an atomic variable requires the generation of atomic synchronization code.

-gnatwo

Activate warnings on address clause overlays.

This switch activates warnings for possibly unintended initialization effects of defining address clauses that cause one variable to overlap another. The default is that such warnings are generated.

-gnatw0

Suppress warnings on address clause overlays.

This switch suppresses warnings on possibly unintended initialization effects of defining address clauses that cause one variable to overlap another.

-gnatw.o

Activate warnings on modified but unreferenced out parameters.

This switch activates warnings for variables that are modified by using them as actuals for a call to a procedure with an out mode formal, where the resulting assigned value is never read. It is applicable in the case where there is more than one out mode formal. If there is only one out mode formal, the warning is issued by default (controlled by -gnatwu). The warning is suppressed for volatile variables and also for variables that are renamings of other variables or for which an address clause is given. The default is that these warnings are not given.

-gnatw.O

Disable warnings on modified but unreferenced out parameters.

This switch suppresses warnings for variables that are modified by using them as actuals for a call to a procedure with an out mode formal, where the resulting assigned value is never read.

-gnatwp

Activate warnings on ineffective pragma Inlines.

This switch activates warnings for failure of front end inlining (activated by -gnatN) to inline a particular call. There are many reasons for not being able to inline a call, including most commonly that the call is too complex to inline. The default is that such warnings are not given. Warnings on ineffective inlining by the gcc back-end can be activated separately, using the gcc switch -Winline.

-gnatwP

Suppress warnings on ineffective pragma Inlines.

This switch suppresses warnings on ineffective pragma Inlines. If the inlining mechanism cannot inline a call, it will simply ignore the request silently.

-gnatw.p

Activate warnings on parameter ordering.

This switch activates warnings for cases of suspicious parameter ordering when the list of arguments are all simple identifiers that match the names of the formals, but are in a different order. The warning is suppressed if any use of named parameter notation is used, so this is the appropriate way to suppress a false positive (and serves to emphasize that the “misordering” is deliberate). The default is that such warnings are not given.

-gnatw.P

Suppress warnings on parameter ordering.

This switch suppresses warnings on cases of suspicious parameter ordering.

-gnatw_p

Activate warnings for pedantic checks.

This switch activates warnings for the failure of certain pedantic checks. The only case currently supported is a check that the subtype_marks given for corresponding formal

parameter and function results in a subprogram declaration and its body denote the same subtype declaration. The default is that such warnings are not given.

-gnatw_P

Suppress warnings for pedantic checks.

This switch suppresses warnings on violations of pedantic checks.

-gnatwq

Activate warnings on questionable missing parentheses.

This switch activates warnings for cases where parentheses are not used and the result is potential ambiguity from a readers point of view. For example `(not a > b)` when `a` and `b` are modular means `((not a) > b)` and very likely the programmer intended `(not (a > b))`. Similarly `(-x mod 5)` means `((-x mod 5))` and quite likely `((-x) mod 5)` was intended. In such situations it seems best to follow the rule of always parenthesizing to make the association clear, and this warning switch warns if such parentheses are not present. The default is that these warnings are given.

-gnatwQ

Suppress warnings on questionable missing parentheses.

This switch suppresses warnings for cases where the association is not clear and the use of parentheses is preferred.

-gnatw.q

Activate warnings on questionable layout of record types.

This switch activates warnings for cases where the default layout of a record type, that is to say the layout of its components in textual order of the source code, would very likely cause inefficiencies in the code generated by the compiler, both in terms of space and speed during execution. One warning is issued for each problematic component without representation clause in the nonvariant part and then in each variant recursively, if any.

The purpose of these warnings is neither to prescribe an optimal layout nor to force the use of representation clauses, but rather to get rid of the most blatant inefficiencies in the layout. Therefore, the default layout is matched against the following synthetic ordered layout and the deviations are flagged on a component-by-component basis:

- first all components or groups of components whose length is fixed and a multiple of the storage unit,
- then the remaining components whose length is fixed and not a multiple of the storage unit,
- then the remaining components whose length doesn't depend on discriminants (that is to say, with variable but uniform length for all objects),
- then all components whose length depends on discriminants,
- finally the variant part (if any),

for the nonvariant part and for each variant recursively, if any.

The exact wording of the warning depends on whether the compiler is allowed to reorder the components in the record type or precluded from doing it by means of pragma `No_Component_Reordering`.

The default is that these warnings are not given.

-gnatw.Q

Suppress warnings on questionable layout of record types.

This switch suppresses warnings for cases where the default layout of a record type would very likely cause inefficiencies.

-gnatw.q

Activate warnings for ignored equality operators.

This switch activates warnings for a user-defined “=” function that does not compose (i.e. is ignored for a predefined “=” for a composite type containing a component whose type has the user-defined “=” as primitive). Note that the user-defined “=” must be a primitive operator in order to trigger the warning.

The default is that these warnings are not given.

-gnatw_Q

Suppress warnings for ignored equality operators.

-gnatwr

Activate warnings on redundant constructs.

This switch activates warnings for redundant constructs. The following is the current list of constructs regarded as redundant:

- Assignment of an item to itself.
- Type conversion that converts an expression to its own type.
- Use of the attribute `Base` where `typ'Base` is the same as `typ`.
- Use of pragma `Pack` when all components are placed by a record representation clause.
- Exception handler containing only a `raise` statement (`raise` with no operand) which has no effect.
- Use of the operator `abs` on an operand that is known at compile time to be non-negative.
- Comparison of an object or (unary or binary) operation of boolean type to an explicit `True` value.
- Import of parent package.

The default is that warnings for redundant constructs are not given.

-gnatwR

Suppress warnings on redundant constructs.

This switch suppresses warnings for redundant constructs.

-gnatw.r

Activate warnings for object renaming function.

This switch activates warnings for an object renaming that renames a function call, which is equivalent to a constant declaration (as opposed to renaming the function itself). The default is that these warnings are given.

-gnatw.R

Suppress warnings for object renaming function.

This switch suppresses warnings for object renaming function.

-gnatw_r

Activate warnings for out-of-order record representation clauses.

This switch activates warnings for record representation clauses, if the order of component declarations, component clauses, and bit-level layout do not all agree. The default is that these warnings are not given.

-gnatw_R

Suppress warnings for out-of-order record representation clauses.

-gnatws

Suppress all warnings.

This switch completely suppresses the output of all warning messages from the GNAT front end, including both warnings that can be controlled by switches described in this section, and those that are normally given unconditionally. The effect of this suppress action can only be cancelled by a subsequent use of the switch `-gnatwn`.

Note that switch `-gnatws` does not suppress warnings from the `gcc` back end. To suppress these back end warnings as well, use the switch `-w` in addition to `-gnatws`. Also this switch has no effect on the handling of style check messages.

-gnatw.s

Activate warnings on overridden size clauses.

This switch activates warnings on component clauses in record representation clauses where the length given overrides that specified by an explicit size clause for the component type. A warning is similarly given in the array case if a specified component size overrides an explicit size clause for the array component type.

-gnatw.S

Suppress warnings on overridden size clauses.

This switch suppresses warnings on component clauses in record representation clauses that override size clauses, and similar warnings when an array component size overrides a size clause.

-gnatwt

Activate warnings for tracking of deleted conditional code.

This switch activates warnings for tracking of code in conditionals (IF and CASE statements) that is detected to be dead code which cannot be executed, and which is removed by the front end. This warning is off by default. This may be useful for detecting deactivated code in certified applications.

-gnatwT

Suppress warnings for tracking of deleted conditional code.

This switch suppresses warnings for tracking of deleted conditional code.

-gnatw.t

Activate warnings on suspicious contracts.

This switch activates warnings on suspicious contracts. This includes warnings on suspicious postconditions (whether a pragma `Postcondition` or a `Post` aspect in Ada 2012) and suspicious contract cases (pragma or aspect `Contract_Cases`). A function postcondition or contract case is suspicious when no postcondition or contract case for this function mentions the result of the function. A procedure postcondition or contract case is suspicious when it only refers to the pre-state of the procedure, because in that case it should rather be expressed as a precondition. This switch also controls warnings on suspicious cases of expressions typically found in contracts like quantified expressions and uses of `Update` attribute. The default is that such warnings are generated.

-gnatw.T

Suppress warnings on suspicious contracts.

This switch suppresses warnings on suspicious contracts.

-gnatwu

Activate warnings on unused entities.

This switch activates warnings to be generated for entities that are declared but not referenced, and for units that are *withed* and not referenced. In the case of packages, a warning is also generated if no entities in the package are referenced. This means that if a *with*'ed package is referenced but the only references are in `use` clauses or `renames` declarations, a warning is still generated. A warning is also generated for a generic package that is *withed* but never instantiated. In the case where a package or subprogram body is compiled, and there is a *with* on the corresponding spec that is only referenced in the body, a warning is also generated, noting that the *with* can be moved to the body. The default is that such warnings are not generated. This switch also activates warnings on unreferenced formals (it includes the effect of `-gnatwf`).

-gnatwU

Suppress warnings on unused entities.

This switch suppresses warnings for unused entities and packages. It also turns off warnings on unreferenced formals (and thus includes the effect of `-gnatwF`).

-gnatw.u

Activate warnings on unordered enumeration types.

This switch causes enumeration types to be considered as conceptually unordered, unless an explicit pragma `Ordered` is given for the type. The effect is to generate warnings in clients that use explicit comparisons or subranges, since these constructs both treat objects of the type as ordered. (A *client* is defined as a unit that is other than the unit in which the type is declared, or its body or subunits.) Please refer to the description of pragma `Ordered` in the *GNAT Reference Manual* for further details. The default is that such warnings are not generated.

-gnatw.U

Deactivate warnings on unordered enumeration types.

This switch causes all enumeration types to be considered as ordered, so that no warnings are given for comparisons or subranges for any type.

-gnatwv

Activate warnings on unassigned variables.

This switch activates warnings for access to variables which may not be properly initialized. The default is that such warnings are generated. This switch will also be emitted when initializing an array or record object via the following aggregate:

```
Array_Or_Record : XXX := (others => <>);
```

unless the relevant type fully initializes all components.

-gnatwV

Suppress warnings on unassigned variables.

This switch suppresses warnings for access to variables which may not be properly initialized.

-gnatw.v

Activate info messages for non-default bit order.

This switch activates messages (labeled “info”, they are not warnings, just informational messages) about the effects of non-default bit-order on records to which a component clause is applied. The effect of specifying non-default bit ordering is a bit subtle (and changed with Ada 2005), so these messages, which are given by default, are useful in understanding the exact consequences of using this feature.

-gnatw.V

Suppress info messages for non-default bit order.

This switch suppresses information messages for the effects of specifying non-default bit order on record components with component clauses.

-gnatww

Activate warnings on wrong low bound assumption.

This switch activates warnings for indexing an unconstrained string parameter with a literal or S'Length. This is a case where the code is assuming that the low bound is one, which is in general not true (for example when a slice is passed). The default is that such warnings are generated.

-gnatwW

Suppress warnings on wrong low bound assumption.

This switch suppresses warnings for indexing an unconstrained string parameter with a literal or S'Length. Note that this warning can also be suppressed in a particular case by adding an assertion that the lower bound is 1, as shown in the following example:

```
procedure K (S : String) is
  pragma Assert (S'First = 1);
  ...
```

-gnatw.w

Activate warnings on Warnings Off pragmas.

This switch activates warnings for use of `pragma Warnings (Off, entity)` where either the pragma is entirely useless (because it suppresses no warnings), or it could be replaced by `pragma Unreferenced` or `pragma Unmodified`. Also activates warnings for the case of `Warnings (Off, String)`, where either there is no matching `Warnings (On, String)`, or the `Warnings (Off)` did not suppress any warning. The default is that these warnings are not given.

-gnatw.W

Suppress warnings on unnecessary Warnings Off pragmas.

This switch suppresses warnings for use of `pragma Warnings (Off, ...)`.

-gnatwx

Activate warnings on Export/Import pragmas.

This switch activates warnings on `Export/Import` pragmas when the compiler detects a possible conflict between the Ada and foreign language calling sequences. For example, the use of default parameters in a `convention C` procedure is dubious because the C compiler cannot supply the proper default, so a warning is issued. The default is that such warnings are generated.

-gnatwX

Suppress warnings on Export/Import pragmas.

This switch suppresses warnings on `Export/Import` pragmas. The sense of this is that you are telling the compiler that you know what you are doing in writing the pragma, and it should not complain at you.

-gnatw.x

Activate warnings for No_Exception_Propagation mode.

This switch activates warnings for exception usage when `pragma Restrictions (No_Exception_Propagation)` is in effect. Warnings are given for implicit or explicit exception raises which are not covered by a local handler, and for exception handlers which do not cover a local raise. The default is that these warnings are given for units that contain exception handlers.

-gnatw.X

Disable warnings for No_Exception_Propagation mode.

This switch disables warnings for exception usage when `pragma Restrictions (No_Exception_Propagation)` is in effect.

-gnatwy

Activate warnings for Ada compatibility issues.

For the most part, newer versions of Ada are upwards compatible with older versions. For example, Ada 2005 programs will almost always work when compiled as Ada 2012. However there are some exceptions (for example the fact that `some` is now a reserved word in Ada 2012). This switch activates several warnings to help in identifying and correcting such incompatibilities. The default is that these warnings are generated. Note that at one point Ada 2005 was called Ada 0Y, hence the choice of character.

-gnatwY

Disable warnings for Ada compatibility issues.

This switch suppresses the warnings intended to help in identifying incompatibilities between Ada language versions.

-gnatw.y

Activate information messages for why package spec needs body.

There are a number of cases in which a package spec needs a body. For example, the use of pragma Elaborate_Body, or the declaration of a procedure specification requiring a completion. This switch causes information messages to be output showing why a package specification requires a body. This can be useful in the case of a large package specification which is unexpectedly requiring a body. The default is that such information messages are not output.

-gnatw.Y

Disable information messages for why package spec needs body.

This switch suppresses the output of information messages showing why a package specification needs a body.

-gnatwz

Activate warnings on unchecked conversions.

This switch activates warnings for unchecked conversions where the types are known at compile time to have different sizes. The default is that such warnings are generated. Warnings are also generated for subprogram pointers with different conventions.

-gnatwZ

Suppress warnings on unchecked conversions.

This switch suppresses warnings for unchecked conversions where the types are known at compile time to have different sizes or conventions.

-gnatw.z

Activate warnings for size not a multiple of alignment.

This switch activates warnings for cases of array and record types with specified Size and Alignment attributes where the size is not a multiple of the alignment, resulting in an object size that is greater than the specified size. The default is that such warnings are generated.

-gnatw.Z

Suppress warnings for size not a multiple of alignment.

This switch suppresses warnings for cases of array and record types with specified Size and Alignment attributes where the size is not a multiple of the alignment, resulting in an object size that is greater than the specified size. The warning can also be suppressed by giving an explicit Object_Size value.

-Wunused

The warnings controlled by the -gnatw switch are generated by the front end of the compiler. The GCC back end can provide additional warnings and they are controlled by the -W switch. For example, -Wunused activates back end warnings for entities that are declared but not referenced.

-Wuninitialized

Similarly, `-Wuninitialized` activates the back end warning for uninitialized variables. This switch must be used in conjunction with an optimization level greater than zero.

-Wstack-usage=len

Warn if the stack usage of a subprogram might be larger than `len` bytes. See [Static Stack Usage Analysis](#) for details.

-Wall

This switch enables most warnings from the GCC back end. The code generator detects a number of warning situations that are missed by the GNAT front end, and this switch can be used to activate them. The use of this switch also sets the default front-end warning mode to `-gnatwa`, that is, most front-end warnings are activated as well.

-w

Conversely, this switch suppresses warnings from the GCC back end. The use of this switch also sets the default front-end warning mode to `-gnatws`, that is, front-end warnings are suppressed as well.

-Werror

This switch causes warnings from the GCC back end to be treated as errors. The warning string still appears, but the warning messages are counted as errors, and prevent the generation of an object file. The use of this switch also sets the default front-end warning mode to `-gnatwe`, that is, front-end warning messages and style check messages are treated as errors as well.

A string of warning parameters can be used in the same parameter. For example:

```
-gnatwaGe
```

will turn on all optional warnings except for unrecognized pragma warnings, and also specify that warnings should be treated as errors.

When no switch `-gnatw` is used, this is equivalent to:

- `-gnatw.a`
- `-gnatwB`
- `-gnatw.b`
- `-gnatwC`
- `-gnatw.C`
- `-gnatwD`
- `-gnatw.D`
- `-gnatwF`
- `-gnatw.F`
- `-gnatwg`
- `-gnatwH`

- -gnatw.H
- -gnatwi
- -gnatwJ
- -gnatw.J
- -gnatwK
- -gnatw.K
- -gnatwL
- -gnatw.L
- -gnatwM
- -gnatw.m
- -gnatwn
- -gnatw.N
- -gnatwo
- -gnatw.O
- -gnatwP
- -gnatw.P
- -gnatwq
- -gnatw.Q
- -gnatwR
- -gnatw.R
- -gnatw.S
- -gnatwT
- -gnatw.t
- -gnatwU
- -gnatw.U
- -gnatwv
- -gnatw.v
- -gnatww
- -gnatw.W
- -gnatwx
- -gnatw.X
- -gnatwy

- -gnatw.Y
- -gnatwz
- -gnatw.z

4.3.4 Debugging and Assertion Control

-gnata

The -gnata option is equivalent to the following `Assertion_Policy` pragma:

```
pragma Assertion_Policy (Check);
```

Which is a shorthand for:

```
pragma Assertion_Policy
-- Ada RM assertion pragmas
(Assert           => Check,
 Static_Predicate => Check,
 Dynamic_Predicate => Check,
 Pre              => Check,
 Pre'Class        => Check,
 Post             => Check,
 Post'Class       => Check,
 Type_Invariant   => Check,
 Type_Invariant'Class => Check,
 Default_Initial_Condition => Check,
-- GNAT specific assertion pragmas
Assert_And_Cut    => Check,
Assume            => Check,
Contract_Cases    => Check,
Debug             => Check,
Ghost             => Check,
Initial_Condition => Check,
Loop_Invariant    => Check,
Loop_Variant      => Check,
Postcondition     => Check,
Precondition      => Check,
Predicate         => Check,
Refined_Post      => Check,
Subprogram_Variant => Check);
```

The pragmas `Assert` and `Debug` normally have no effect and are ignored. This switch, where `a` stands for ‘assert’, causes pragmas `Assert` and `Debug` to be activated. This switch also causes preconditions, postconditions, subtype predicates, and type invariants to be activated.

The pragmas have the form:

```
pragma Assert (<Boolean-expression> [, <static-string-expression>])
pragma Debug (<procedure call>)
pragma Type_Invariant (<type-local-name>, <Boolean-expression>)
pragma Predicate (<type-local-name>, <Boolean-expression>)
```

(continues on next page)

(continued from previous page)

```
pragma Precondition (<Boolean-expression>, <string-expression>)
pragma Postcondition (<Boolean-expression>, <string-expression>)
```

The aspects have the form:

```
with [Pre|Post|Type_Invariant|Dynamic_Predicate|Static_Predicate]
=> <Boolean-expression>;
```

The **Assert** pragma causes **Boolean-expression** to be tested. If the result is **True**, the pragma has no effect (other than possible side effects from evaluating the expression). If the result is **False**, the exception **Assert_Failure** declared in the package **System.Assertions** is raised (passing **static-string-expression**, if present, as the message associated with the exception). If no string expression is given, the default is a string containing the file name and line number of the pragma.

The **Debug** pragma causes **procedure** to be called. Note that **pragma Debug** may appear within a declaration sequence, allowing debugging procedures to be called between declarations.

For the aspect specification, the **Boolean-expression** is evaluated. If the result is **True**, the aspect has no effect. If the result is **False**, the exception **Assert_Failure** is raised.

4.3.5 Validity Checking

The Ada Reference Manual defines the concept of invalid values (see RM 13.9.1). The primary source of invalid values is uninitialized variables. A scalar variable that is left uninitialized may contain an invalid value; the concept of invalid does not apply to access or composite types.

It is an error to read an invalid value, but the RM does not require run-time checks to detect such errors, except for some minimal checking to prevent erroneous execution (i.e. unpredictable behavior). This corresponds to the **-gnatVd** switch below, which is the default. For example, by default, if the expression of a case statement is invalid, it will raise **Constraint_Error** rather than causing a wild jump, and if an array index on the left-hand side of an assignment is invalid, it will raise **Constraint_Error** rather than overwriting an arbitrary memory location.

The **-gnatVa** may be used to enable additional validity checks, which are not required by the RM. These checks are often very expensive (which is why the RM does not require them). These checks are useful in tracking down uninitialized variables, but they are not usually recommended for production builds, and in particular we do not recommend using these extra validity checking options in combination with optimization, since this can confuse the optimizer. If performance is a consideration, leading to the need to optimize, then the validity checking options should not be used.

The other **-gnatVx** switches below allow finer-grained control; you can enable whichever validity checks you desire. However, for most debugging purposes, **-gnatVa** is sufficient, and the default **-gnatVd** (i.e. standard Ada behavior) is usually sufficient for non-debugging use.

The **-gnatB** switch tells the compiler to assume that all values are valid (that is, within their declared subtype range) except in the context of a use of the **Valid** attribute. This means the compiler can

generate more efficient code, since the range of values is better known at compile time. However, an uninitialized variable can cause wild jumps and memory corruption in this mode.

The `-gnatVx` switch allows control over the validity checking mode as described below. The `x` argument is a string of letters that indicate validity checks that are performed or not performed in addition to the default checks required by Ada as described above.

-gnatVa

All validity checks.

All validity checks are turned on. That is, `-gnatVa` is equivalent to `gnatVcdefimoprst`.

-gnatVc

Validity checks for copies.

The right-hand side of assignments, and the (explicit) initializing values of object declarations are validity checked.

-gnatVd

Default (RM) validity checks.

Some validity checks are required by Ada (see RM 13.9.1 (9-11)); these (and only these) validity checks are enabled by default. For case statements (and case expressions) that lack a “when others =>” choice, a check is made that the value of the selector expression belongs to its nominal subtype. If it does not, `Constraint_Error` is raised. For assignments to array components (and for indexed components in some other contexts), a check is made that each index expression belongs to the corresponding index subtype. If it does not, `Constraint_Error` is raised. Both these validity checks may be turned off using switch `-gnatVD`. They are turned on by default. If `-gnatVD` is specified, a subsequent switch `-gnatVd` will leave the checks turned on. Switch `-gnatVD` should be used only if you are sure that all such expressions have valid values. If you use this switch and invalid values are present, then the program is erroneous, and wild jumps or memory overwriting may occur.

-gnatVe

Validity checks for scalar components.

In the absence of this switch, assignments to scalar components of enclosing record or array objects are not validity checked, even if validity checks for assignments generally (`-gnatVc`) are turned on. Specifying this switch enables such checks. This switch has no effect if the `-gnatVc` switch is not specified.

-gnatVf

Validity checks for floating-point values.

Specifying this switch enables validity checking for floating-point values in the same contexts where validity checking is enabled for other scalar values. In the absence of this switch, validity checking is not performed for floating-point values. This takes precedence over other statements about performing validity checking for scalar objects in various scenarios. One way to look at it is that if this switch is not set, then whenever any of the other rules in this section use the word “scalar” they really mean “scalar and not floating-point”. If `-gnatVf` is specified, then validity checking also applies for floating-point values, and NaNs and infinities are considered invalid, as well as out-of-range values for constrained types. The exact contexts in which floating-point values are checked depends on the setting of

other options. For example, `-gnatVif` or `-gnatVfi` (the order does not matter) specifies that floating-point parameters of mode `in` should be validity checked.

-gnatVi

Validity checks for “in” mode parameters.

Arguments for parameters of mode `in` are validity checked in function and procedure calls at the point of call.

-gnatVm

Validity checks for “in out” mode parameters.

Arguments for parameters of mode `in out` are validity checked in procedure calls at the point of call. The 'm' here stands for modify, since this concerns parameters that can be modified by the call. Note that there is no specific option to test `out` parameters, but any reference within the subprogram will be tested in the usual manner, and if an invalid value is copied back, any reference to it will be subject to validity checking.

-gnatVn

No validity checks.

This switch turns off all validity checking, including the default checking for case statements and left hand side subscripts. Note that the use of the switch `-gnatp` suppresses all run-time checks, including validity checks, and thus implies `-gnatVn`. When this switch is used, it cancels any other `-gnatV` previously issued.

-gnatVo

Validity checks for operator and attribute operands.

Scalar arguments for predefined operators and for attributes are validity checked. This includes all operators in package `Standard`, the shift operators defined as intrinsic in package `Interfaces` and operands for attributes such as `Pos`. Checks are also made on individual component values for composite comparisons, and on the expressions in type conversions and qualified expressions. Checks are also made on explicit ranges using `..` (e.g., slices, loops etc).

-gnatVp

Validity checks for parameters.

This controls the treatment of formal parameters within a subprogram (as opposed to `-gnatVi` and `-gnatVm`, which control validity testing of actual parameters of a call). If either of these call options is specified, then normally an assumption is made within a subprogram that the validity of any incoming formal parameters of the corresponding mode(s) has already been checked at the point of call and does not need rechecking. If `-gnatVp` is set, then this assumption is not made and so their validity may be checked (or rechecked) within the subprogram. If neither of the two call-related options is specified, then this switch has no effect.

-gnatVr

Validity checks for function returns.

The expression in simple return statements in functions is validity checked.

-gnatVs

Validity checks for subscripts.

All subscript expressions are checked for validity, whatever context they occur in (in default mode some subscripts are not validity checked; for example, validity checking may be omitted in some cases involving a read of a component of an array).

-gnatVt

Validity checks for tests.

Expressions used as conditions in `if`, `while` or `exit` statements are checked, as well as guard expressions in entry calls.

The `-gnatV` switch may be followed by a string of letters to turn on a series of validity checking options. For example, `-gnatVcr` specifies that in addition to the default validity checking, copies and function return expressions are to be validity checked. In order to make it easier to specify the desired combination of effects, the upper case letters `CDFIMORST` may be used to turn off the corresponding lower case option. Thus `-gnatVaM` turns on all validity checking options except for checking of `in out` parameters.

The specification of additional validity checking generates extra code (and in the case of `-gnatVa` the code expansion can be substantial). However, these additional checks can be very useful in detecting uninitialized variables, incorrect use of unchecked conversion, and other errors leading to invalid values. The use of `pragma Initialize_Scalars` is useful in conjunction with the extra validity checking, since this ensures that wherever possible uninitialized variables have invalid values.

See also the `pragma Validity_Checks` which allows modification of the validity checking mode at the program source level, and also allows for temporary disabling of validity checks.

4.3.6 Style Checking

The `-gnatyx` switch causes the compiler to enforce specified style rules. A limited set of style rules has been used in writing the GNAT sources themselves. This switch allows user programs to activate all or some of these checks. If the source program fails a specified style check, an appropriate message is given, preceded by the character sequence '(style)'. This message does not prevent successful compilation (unless the `-gnatwe` switch is used).

Note that this is by no means intended to be a general facility for checking arbitrary coding standards. It is simply an embedding of the style rules we have chosen for the GNAT sources. If you are starting a project which does not have established style standards, you may find it useful to adopt the entire set of GNAT coding standards, or some subset of them.

The string `x` is a sequence of letters or digits indicating the particular style checks to be performed. The following checks are defined:

-gnaty0

Specify indentation level.

If a digit from 1-9 appears in the string after `-gnaty` then proper indentation is checked, with the digit indicating the indentation level required. A value of zero turns off this style check. The rule checks that the following constructs start on a column that is a multiple of the alignment level:

- beginnings of declarations (except record component declarations) and statements;

- beginnings of the structural components of compound statements;
- `end` keyword that completes the declaration of a program unit declaration or body or that completes a compound statement.

Full line comments must be aligned with the `--` starting on a column that is a multiple of the alignment level, or they may be aligned the same way as the following non-blank line (this is useful when full line comments appear in the middle of a statement, or they may be aligned with the source line on the previous non-blank line).

-gnatya

Check attribute casing.

Attribute names, including the case of keywords such as `digits` used as attributes names, must be written in mixed case, that is, the initial letter and any letter following an underscore must be uppercase. All other letters must be lowercase.

-gnatyA

Use of array index numbers in array attributes.

When using the array attributes `First`, `Last`, `Range`, or `Length`, the index number must be omitted for one-dimensional arrays and is required for multi-dimensional arrays.

-gnatyb

Blanks not allowed at statement end.

Trailing blanks are not allowed at the end of statements. The purpose of this rule, together with `h` (no horizontal tabs), is to enforce a canonical format for the use of blanks to separate source tokens.

-gnatyB

Check Boolean operators.

The use of `AND/OR` operators is not permitted except in the cases of modular operands, array operands, and simple stand-alone boolean variables or boolean constants. In all other cases `and` `then/or else` are required.

-gnatyc

Check comments, double space.

Comments must meet the following set of rules:

- The `--` that starts the column must either start in column one, or else at least one blank must precede this sequence.
- Comments that follow other tokens on a line must have at least one blank following the `--` at the start of the comment.
- Full line comments must have at least two blanks following the `--` that starts the comment, with the following exceptions.
- A line consisting only of the `--` characters, possibly preceded by blanks is permitted.
- A comment starting with `--x` where `x` is a special character is permitted. This allows proper processing of the output from specialized tools such as `gnatprep` (where `--!` is used) and in earlier versions of the SPARK annotation language (where `--#` is used).

For the purposes of this rule, a special character is defined as being in one of the ASCII ranges 16#21#...16#2F# or 16#3A#...16#3F#. Note that this usage is not permitted in GNAT implementation units (i.e., when `-gnatg` is used).

- A line consisting entirely of minus signs, possibly preceded by blanks, is permitted. This allows the construction of box comments where lines of minus signs are used to form the top and bottom of the box.
- A comment that starts and ends with `--` is permitted as long as at least one blank follows the initial `--`. Together with the preceding rule, this allows the construction of box comments, as shown in the following example:

```
-----  
-- This is a box comment --  
-- with two text lines. --  
-----
```

-gnatyc

Check comments, single space.

This is identical to `c` except that only one space is required following the `--` of a comment instead of two.

-gnatyd

Check no DOS line terminators present.

All lines must be terminated by a single ASCII.LF character (in particular the DOS line terminator sequence CR/LF is not allowed).

-gnatyD

Check declared identifiers in mixed case.

Declared identifiers must be in mixed case, as in `This_Is_An_Identifier`. Use `-gnatyr` in addition to ensure that references match declarations.

-gnatye

Check end/exit labels.

Optional labels on `end` statements ending subprograms and on `exit` statements exiting named loops, are required to be present.

-gnatyf

No form feeds or vertical tabs.

Neither form feeds nor vertical tab characters are permitted in the source text.

-gnatyG

GNAT style mode.

The set of style check switches is set to match that used by the GNAT sources. This may be useful when developing code that is eventually intended to be incorporated into GNAT. Currently this is equivalent to `-gnatyydISux`) but additional style switches may be added to this set in the future without advance notice.

-gnatyh

No horizontal tabs.

Horizontal tab characters are not permitted in the source text. Together with the `b` (no blanks at end of line) check, this enforces a canonical form for the use of blanks to separate source tokens.

-gnatyi

Check if-then layout.

The keyword `then` must appear either on the same line as corresponding `if`, or on a line on its own, lined up under the `if`.

-gnatyI

check mode IN keywords.

Mode `in` (the default mode) is not allowed to be given explicitly. `in out` is fine, but not `in` on its own.

-gnatyk

Check keyword casing.

All keywords must be in lower case (with the exception of keywords such as `digits` used as attribute names to which this check does not apply). A single error is reported for each line breaking this rule even if multiple casing issues exist on a same line.

-gnatyl

Check layout.

Layout of statement and declaration constructs must follow the recommendations in the Ada Reference Manual, as indicated by the form of the syntax rules. For example an `else` keyword must be lined up with the corresponding `if` keyword.

There are two respects in which the style rule enforced by this check option are more liberal than those in the Ada Reference Manual. First in the case of record declarations, it is permissible to put the `record` keyword on the same line as the `type` keyword, and then the `end` in `end record` must line up under `type`. This is also permitted when the type declaration is split on two lines. For example, any of the following three layouts is acceptable:

```

type q is record
  a : integer;
  b : integer;
end record;

type q is
  record
    a : integer;
    b : integer;
  end record;

type q is
  record
    a : integer;
```

(continues on next page)

(continued from previous page)

```
    b : integer;  
end record;
```

Second, in the case of a block statement, a permitted alternative is to put the block label on the same line as the `declare` or `begin` keyword, and then line the `end` keyword up under the block label. For example both the following are permitted:

```
Block : declare  
    A : Integer := 3;  
begin  
    Proc (A, A);  
end Block;  
  
Block :  
    declare  
        A : Integer := 3;  
    begin  
        Proc (A, A);  
    end Block;
```

The same alternative format is allowed for loops. For example, both of the following are permitted:

```
Clear : while J < 10 loop  
    A (J) := 0;  
end loop Clear;  
  
Clear :  
    while J < 10 loop  
        A (J) := 0;  
    end loop Clear;
```

-gnatyL

Set maximum nesting level.

The maximum level of nesting of constructs (including subprograms, loops, blocks, packages, and conditionals) may not exceed the given value *nnn*. A value of zero disconnects this style check.

-gnatym

Check maximum line length.

The length of source lines must not exceed 79 characters, including any trailing blanks. The value of 79 allows convenient display on an 80 character wide device or window, allowing for possible special treatment of 80 character lines. Note that this count is of characters in the source text. This means that a tab character counts as one character in this count and a wide character sequence counts as a single character (however many bytes are needed in the encoding).

-gnatyM

Set maximum line length.

The length of lines must not exceed the given value *nnn*. The maximum value that can be specified is 32767. If neither style option for setting the line length is used, then the default is 255. This also controls the maximum length of lexical elements, where the only restriction is that they must fit on a single line.

-gnatyn

Check casing of entities in Standard.

Any identifier from Standard must be cased to match the presentation in the Ada Reference Manual (for example, `Integer` and `ASCII.NUL`).

-gnatyN

Turn off all style checks.

All style check options are turned off.

-gnatyo

Check order of subprogram bodies.

All subprogram bodies in a given scope (e.g., a package body) must be in alphabetical order. The ordering rule uses normal Ada rules for comparing strings, ignoring casing of letters, except that if there is a trailing numeric suffix, then the value of this suffix is used in the ordering (e.g., `Junk2` comes before `Junk10`).

-gnaty0

Check that overriding subprograms are explicitly marked as such.

This applies to all subprograms of a derived type that override a primitive operation of the type, for both tagged and untagged types. In particular, the declaration of a primitive operation of a type extension that overrides an inherited operation must carry an overriding indicator. Another case is the declaration of a function that overrides a predefined operator (such as an equality operator).

-gnatyp

Check pragma casing.

Pragma names must be written in mixed case, that is, the initial letter and any letter following an underscore must be uppercase. All other letters must be lowercase. An exception is that `SPARK_Mode` is allowed as an alternative for `Spark_Mode`.

-gnatyr

Check references.

All identifier references must be cased in the same way as the corresponding declaration. No specific casing style is imposed on identifiers. The only requirement is for consistency of references with declarations.

-gnatys

Check separate specs.

Separate declarations ('specs') are required for subprograms (a body is not allowed to serve as its own declaration). The only exception is that parameterless library level procedures are not required to have a separate declaration. This exception covers the most frequent form of main program procedures.

-gnatyS

Check no statements after then/else.

No statements are allowed on the same line as a `then` or `else` keyword following the keyword in an `if` statement. `or else` and `and then` are not affected, and a special exception allows a pragma to appear after `else`.

-gnatyt

Check token spacing.

The following token spacing rules are enforced:

- The keywords `abs` and `not` must be followed by a space.
- The token `=>` must be surrounded by spaces.
- The token `<>` must be preceded by a space or a left parenthesis.
- Binary operators other than `**` must be surrounded by spaces. There is no restriction on the layout of the `**` binary operator.
- Colon must be surrounded by spaces.
- Colon-equal (assignment, initialization) must be surrounded by spaces.
- Comma must be the first non-blank character on the line, or be immediately preceded by a non-blank character, and must be followed by a space.
- If the token preceding a left parenthesis ends with a letter or digit, then a space must separate the two tokens.
- If the token following a right parenthesis starts with a letter or digit, then a space must separate the two tokens.
- A right parenthesis must either be the first non-blank character on a line, or it must be preceded by a non-blank character.
- A semicolon must not be preceded by a space, and must not be followed by a non-blank character.
- A unary plus or minus may not be followed by a space.
- A vertical bar must be surrounded by spaces.

Exactly one blank (and no other white space) must appear between a `not` token and a following `in` token.

-gnatyu

Check unnecessary blank lines.

Unnecessary blank lines are not allowed. A blank line is considered unnecessary if it appears at the end of the file, or if more than one blank line occurs in sequence.

-gnatyx

Check extra parentheses.

Unnecessary extra level of parentheses (C-style) are not allowed around conditions in `if` statements, `while` statements and `exit` statements.

-gnatyy

Set all standard style check options.

This is equivalent to `gnaty3aAbcefghiklmnrst`, that is all checking options enabled with the exception of `-gnatyB`, `-gnatyD`, `-gnatyI`, `-gnatyLnnn`, `-gnatyO`, `-gnaty0`, `-gnatyS`, `-gnatyu`, and `-gnatyx`.

-gnaty-

Remove style check options.

This causes any subsequent options in the string to act as canceling the corresponding style check option. To cancel maximum nesting level control, use the `L` parameter without any integer value after that, because any digit following `-` in the parameter string of the `-gnaty` option will be treated as canceling the indentation check. The same is true for the `M` parameter. `y` and `N` parameters are not allowed after `-`.

-gnaty+

Enable style check options.

This causes any subsequent options in the string to enable the corresponding style check option. That is, it cancels the effect of a previous `-`, if any.

In the above rules, appearing in column one is always permitted, that is, counts as meeting either a requirement for a required preceding space, or as meeting a requirement for no preceding space.

Appearing at the end of a line is also always permitted, that is, counts as meeting either a requirement for a following space, or as meeting a requirement for no following space.

If any of these style rules is violated, a message is generated giving details on the violation. The initial characters of such messages are always *'(style)'*. Note that these messages are treated as warning messages, so they normally do not prevent the generation of an object file. The `-gnatwe` switch can be used to treat warning messages, including style messages, as fatal errors.

The switch `-gnaty` on its own (that is not followed by any letters or digits) is equivalent to the use of `-gnatyy` as described above, that is all built-in standard style check options are enabled.

The switch `-gnatyN` clears any previously set style checks.

4.3.7 Run-Time Checks

By default, the following checks are suppressed: stack overflow checks, and checks for access before elaboration on subprogram calls. All other checks, including overflow checks, range checks and array bounds checks, are turned on by default. The following `gcc` switches refine this default behavior.

-gnatp

This switch causes the unit to be compiled as though `pragma Suppress (All_checks)` had been present in the source. Validity checks are also eliminated (in other words `-gnatp` also implies `-gnatVn`). Use this switch to improve the performance of the code at the expense of safety in the presence of invalid data or program bugs.

Note that when checks are suppressed, the compiler is allowed, but not required, to omit the checking code. If the run-time cost of the checking code is zero or near-zero, the compiler will generate it even if checks are suppressed. In particular, if the compiler can prove that

a certain check will necessarily fail, it will generate code to do an unconditional 'raise', even if checks are suppressed. The compiler warns in this case. Another case in which checks may not be eliminated is when they are embedded in certain run-time routines such as math library routines.

Of course, run-time checks are omitted whenever the compiler can prove that they will not fail, whether or not checks are suppressed.

Note that if you suppress a check that would have failed, program execution is erroneous, which means the behavior is totally unpredictable. The program might crash, or print wrong answers, or do anything else. It might even do exactly what you wanted it to do (and then it might start failing mysteriously next week or next year). The compiler will generate code based on the assumption that the condition being checked is true, which can result in erroneous execution if that assumption is wrong.

The checks subject to suppression include all the checks defined by the Ada standard, the additional implementation defined checks `Alignment_Check`, `Duplicated_Tag_Check`, `Predicate_Check`, `Container_Checks`, `Tampering_Check`, and `Validity_Check`, as well as any checks introduced using `pragma Check_Name`. Note that `Atomic_Synchronization` is not automatically suppressed by use of this option.

If the code depends on certain checks being active, you can use `pragma Unsuppress` either as a configuration pragma or as a local pragma to make sure that a specified check is performed even if `gnatp` is specified.

The `-gnatp` switch has no effect if a subsequent `-gnat-p` switch appears.

-gnat-p

This switch cancels the effect of a previous `gnatp` switch.

-gnato??

This switch controls the mode used for computing intermediate arithmetic integer operations, and also enables overflow checking. For a full description of overflow mode and checking control, see the 'Overflow Check Handling in GNAT' appendix in this User's Guide.

Overflow checks are always enabled by this switch. The argument controls the mode, using the codes

1 = *STRICT*

In *STRICT* mode, intermediate operations are always done using the base type, and overflow checking ensures that the result is within the base type range.

2 = *MINIMIZED*

In *MINIMIZED* mode, overflows in intermediate operations are avoided where possible by using a larger integer type for the computation (typically `Long_Long_Integer`). Overflow checking ensures that the result fits in this larger integer type.

3 = *ELIMINATED*

In *ELIMINATED* mode, overflows in intermediate operations are avoided by using multi-precision arithmetic. In this case, overflow checking has no effect on intermediate operations (since overflow is impossible).

If two digits are present after `-gnato` then the first digit sets the mode for expressions outside assertions, and the second digit sets the mode for expressions within assertions.

Here assertions is used in the technical sense (which includes for example precondition and postcondition expressions).

If one digit is present, the corresponding mode is applicable to both expressions within and outside assertion expressions.

If no digits are present, the default is to enable overflow checks and set STRICT mode for both kinds of expressions. This is compatible with the use of `-gnato` in previous versions of GNAT.

Note that the `-gnato??` switch does not affect the code generated for any floating-point operations; it applies only to integer semantics. For floating-point, GNAT has the `Machine_Overflows` attribute set to `False` and the normal mode of operation is to generate IEEE NaN and infinite values on overflow or invalid operations (such as dividing 0.0 by 0.0).

The reason that we distinguish overflow checking from other kinds of range constraint checking is that a failure of an overflow check, unlike for example the failure of a range check, can result in an incorrect value, but cannot cause random memory destruction (like an out of range subscript), or a wild jump (from an out of range case value). Overflow checking is also quite expensive in time and space, since in general it requires the use of double length arithmetic.

Note again that the default is `-gnato11` (equivalent to `-gnato1`), so overflow checking is performed in STRICT mode by default.

-gnatE

Enables dynamic checks for access-before-elaboration on subprogram calls and generic instantiations. Note that `-gnatE` is not necessary for safety, because in the default mode, GNAT ensures statically that the checks would not fail. For full details of the effect and use of this switch, [Compiling with gcc](#).

-fstack-check

Activates stack overflow checking. For full details of the effect and use of this switch see [Stack Overflow Checking](#).

The setting of these switches only controls the default setting of the checks. You may modify them using either `Suppress` (to remove checks) or `Unsuppress` (to add back suppressed checks) pragmas in the program source.

4.3.8 Using gcc for Syntax Checking

-gnats

The `s` stands for 'syntax'.

Run GNAT in syntax checking only mode. For example, the command

```
$ gcc -c -gnats x.adb
```

compiles file `x.adb` in syntax-check-only mode. You can check a series of files in a single command, and can use wildcards to specify such a group of files. Note that you must specify the `-c` (compile only) flag in addition to the `-gnats` flag.

You may use other switches in conjunction with `-gnats`. In particular, `-gnatl` and `-gnatv` are useful to control the format of any generated error messages.

When the source file is empty or contains only empty lines and/or comments, the output is a warning:

```
$ gcc -c -gnats -x ada toto.txt
toto.txt:1:01: warning: empty file, contains no compilation units
$
```

Otherwise, the output is simply the error messages, if any. No object file or ALI file is generated by a syntax-only compilation. Also, no units other than the one specified are accessed. For example, if a unit *X* *withs* a unit *Y*, compiling unit *X* in syntax check only mode does not access the source file containing unit *Y*.

Normally, GNAT allows only a single unit in a source file. However, this restriction does not apply in syntax-check-only mode, and it is possible to check a file containing multiple compilation units concatenated together. This is primarily used by the `gnatchop` utility ([Renaming Files with gnatchop](#)).

4.3.9 Using gcc for Semantic Checking

-gnatc

The `c` stands for ‘check’. Causes the compiler to operate in semantic check mode, with full checking for all illegalities specified in the Ada Reference Manual, but without generation of any object code (no object file is generated).

Because dependent files must be accessed, you must follow the GNAT semantic restrictions on file structuring to operate in this mode:

- The needed source files must be accessible (see [Search Paths and the Run-Time Library \(RTL\)](#)).
- Each file must contain only one compilation unit.
- The file name and unit name must match ([File Naming Rules](#)).

The output consists of error messages as appropriate. No object file is generated. An ALI file is generated for use in the context of cross-reference tools, but this file is marked as not being suitable for binding (since no object file is generated). The checking corresponds exactly to the notion of legality in the Ada Reference Manual.

Any unit can be compiled in semantics-checking-only mode, including units that would not normally be compiled (subunits, and specifications where a separate body is present).

4.3.10 Compiling Different Versions of Ada

The switches described in this section allow you to explicitly specify the version of the Ada language that your programs are written in. The default mode is Ada 2012, but you can also specify Ada 95, Ada 2005 mode, or indicate Ada 83 compatibility mode.

-gnat83 (Ada 83 Compatibility Mode)

Although GNAT is primarily an Ada 95 / Ada 2005 compiler, this switch specifies that the program is to be compiled in Ada 83 mode. With `-gnat83`, GNAT rejects most post-Ada 83 extensions and applies Ada 83 semantics where this can be done easily. It is not possible to guarantee this switch does a perfect job; some subtle tests, such as are found in earlier ACVC tests (and that have been removed from the ACATS suite for Ada 95), might not compile correctly. Nevertheless, this switch may be useful in some circumstances, for example where, due to contractual reasons, existing code needs to be maintained using only Ada 83 features.

With few exceptions (most notably the need to use `<>` on unconstrained generic formal parameters, the use of the new Ada 95 / Ada 2005 reserved words, and the use of packages with optional bodies), it is not necessary to specify the `-gnat83` switch when compiling Ada 83 programs, because, with rare exceptions, Ada 95 and Ada 2005 are upwardly compatible with Ada 83. Thus a correct Ada 83 program is usually also a correct program in these later versions of the language standard. For further information please refer to the *Compatibility and Porting Guide* chapter in the *GNAT Reference Manual*.

-gnat95 (Ada 95 mode)

This switch directs the compiler to implement the Ada 95 version of the language. Since Ada 95 is almost completely upwards compatible with Ada 83, Ada 83 programs may generally be compiled using this switch (see the description of the `-gnat83` switch for further information about Ada 83 mode). If an Ada 2005 program is compiled in Ada 95 mode, uses of the new Ada 2005 features will cause error messages or warnings.

This switch also can be used to cancel the effect of a previous `-gnat83`, `-gnat05/2005`, or `-gnat12/2012` switch earlier in the command line.

-gnat05 or -gnat2005 (Ada 2005 mode)

This switch directs the compiler to implement the Ada 2005 version of the language, as documented in the official Ada standards document. Since Ada 2005 is almost completely upwards compatible with Ada 95 (and thus also with Ada 83), Ada 83 and Ada 95 programs may generally be compiled using this switch (see the description of the `-gnat83` and `-gnat95` switches for further information).

-gnat12 or -gnat2012 (Ada 2012 mode)

This switch directs the compiler to implement the Ada 2012 version of the language (also the default). Since Ada 2012 is almost completely upwards compatible with Ada 2005 (and thus also with Ada 83, and Ada 95), Ada 83 and Ada 95 programs may generally be compiled using this switch (see the description of the `-gnat83`, `-gnat95`, and `-gnat05/2005` switches for further information).

-gnat2022 (Ada 2022 mode)

This switch directs the compiler to implement the Ada 2022 version of the language.

-gnatX0 (Enable GNAT Extensions)

This switch directs the compiler to implement the latest version of the language (currently Ada 2022) and also to enable certain GNAT implementation extensions that are not part of any Ada standard. For a full list of these extensions, see the GNAT reference manual, `Pragma Extensions_Allowed`.

-gnatX (Enable core GNAT Extensions)

This switch is similar to `-gnatX0` except that only some, not all, of the GNAT-defined language extensions are enabled. For a list of the extensions enabled by this switch, see the GNAT reference manual `Pragma Extensions_Allowed` and the description of that pragma's "On" (as opposed to "All") argument.

4.3.11 Character Set Control

-gnatic

Normally GNAT recognizes the Latin-1 character set in source program identifiers, as described in the Ada Reference Manual. This switch causes GNAT to recognize alternate character sets in identifiers. `c` is a single character indicating the character set, as follows:

<i>1</i>	ISO 8859-1 (Latin-1) identifiers
<i>2</i>	ISO 8859-2 (Latin-2) letters allowed in identifiers
<i>3</i>	ISO 8859-3 (Latin-3) letters allowed in identifiers
<i>4</i>	ISO 8859-4 (Latin-4) letters allowed in identifiers
<i>5</i>	ISO 8859-5 (Cyrillic) letters allowed in identifiers
<i>9</i>	ISO 8859-15 (Latin-9) letters allowed in identifiers
<i>p</i>	IBM PC letters (code page 437) allowed in identifiers
<i>8</i>	IBM PC letters (code page 850) allowed in identifiers
<i>f</i>	Full upper-half codes allowed in identifiers
<i>n</i>	No upper-half codes allowed in identifiers
<i>w</i>	Wide-character codes (that is, codes greater than 255) allowed in identifiers

See [Foreign Language Representation](#) for full details on the implementation of these character sets.

-gnatWe

Specify the method of encoding for wide characters. `e` is one of the following:

<i>h</i>	Hex encoding (brackets coding also recognized)
<i>u</i>	Upper half encoding (brackets encoding also recognized)
<i>s</i>	Shift/JIS encoding (brackets encoding also recognized)
<i>e</i>	EUC encoding (brackets encoding also recognized)
<i>8</i>	UTF-8 encoding (brackets encoding also recognized)
<i>b</i>	Brackets encoding only (default value)

For full details on these encoding methods see [Wide_Character Encodings](#). Note that brackets coding is always accepted, even if one of the other options is specified, so for example `-gnatW8` specifies that both brackets and UTF-8 encodings will be recognized. The units that are with'ed directly or indirectly will be scanned using the specified representation scheme, and

so if one of the non-brackets scheme is used, it must be used consistently throughout the program. However, since brackets encoding is always recognized, it may be conveniently used in standard libraries, allowing these libraries to be used with any of the available coding schemes.

Note that brackets encoding only applies to program text. Within comments, brackets are considered to be normal graphic characters, and bracket sequences are never recognized as wide characters.

If no `-gnatw?` parameter is present, then the default representation is normally Brackets encoding only. However, if the first three characters of the file are `16#EF# 16#BB# 16#BF#` (the standard byte order mark or BOM for UTF-8), then these three characters are skipped and the default representation for the file is set to UTF-8.

Note that the wide character representation that is specified (explicitly or by default) for the main program also acts as the default encoding used for `Wide_Text_IO` files if not specifically overridden by a `WCEM` form parameter.

When no `-gnatw?` is specified, then characters (other than wide characters represented using brackets notation) are treated as 8-bit Latin-1 codes. The codes recognized are the Latin-1 graphic characters, and ASCII format effectors (CR, LF, HT, VT). Other lower half control characters in the range `16#00#..16#1F#` are not accepted in program text or in comments. Upper half control characters (`16#80#..16#9F#`) are rejected in program text, but allowed and ignored in comments. Note in particular that the Next Line (NEL) character whose encoding is `16#85#` is not recognized as an end of line in this default mode. If your source program contains instances of the NEL character used as a line terminator, you must use UTF-8 encoding for the whole source program. In default mode, all lines must be ended by a standard end of line sequence (CR, CR/LF, or LF).

Note that the convention of simply accepting all upper half characters in comments means that programs that use standard ASCII for program text, but UTF-8 encoding for comments are accepted in default mode, providing that the comments are ended by an appropriate (CR, or CR/LF, or LF) line terminator. This is a common mode for many programs with foreign language comments.

4.3.12 File Naming Control

-gnatkn

Activates file name 'krunching'. `n`, a decimal integer in the range 1-999, indicates the maximum allowable length of a file name (not including the `.ads` or `.adb` extension). The default is not to enable file name krunching.

For the source file naming rules, [File Naming Rules](#).

4.3.13 Subprogram Inlining Control

-gnatn[12]

The `n` here is intended to suggest the first syllable of the word ‘inline’. GNAT recognizes and processes `InLine` pragmas. However, for inlining to actually occur, optimization must be enabled and, by default, inlining of subprograms across units is not performed. If you want to additionally enable inlining of subprograms specified by pragma `InLine` across units, you must also specify this switch.

In the absence of this switch, GNAT does not attempt inlining across units and does not access the bodies of subprograms for which `pragma InLine` is specified if they are not in the current unit.

You can optionally specify the inlining level: 1 for moderate inlining across units, which is a good compromise between compilation times and performances at run time, or 2 for full inlining across units, which may bring about longer compilation times. If no inlining level is specified, the compiler will pick it based on the optimization level: 1 for `-O1`, `-O2` or `-Os` and 2 for `-O3`.

If you specify this switch the compiler will access these bodies, creating an extra source dependency for the resulting object file, and where possible, the call will be inlined. For further details on when inlining is possible see [Inlining of Subprograms](#).

-gnatN

This switch activates front-end inlining which also generates additional dependencies.

When using a gcc-based back end, then the use of `-gnatN` is deprecated, and the use of `-gnatn` is preferred. Historically front end inlining was more extensive than the gcc back end inlining, but that is no longer the case.

4.3.14 Auxiliary Output Control

-gnatu

Print a list of units required by this compilation on `stdout`. The listing includes all units on which the unit being compiled depends either directly or indirectly.

-pass-exit-codes

If this switch is not used, the exit code returned by `gcc` when compiling multiple files indicates whether all source files have been successfully used to generate object files or not.

When `-pass-exit-codes` is used, `gcc` exits with an extended exit status and allows an integrated development environment to better react to a compilation failure. Those exit status are:

5	There was an error in at least one source file.
3	At least one source file did not generate an object file.
2	The compiler died unexpectedly (internal error for example).
0	An object file has been generated for every source file.

4.3.15 Debugging Control

-gnatdx

Activate internal debugging switches. *x* is a letter or digit, or string of letters or digits, which specifies the type of debugging outputs desired. Normally these are used only for internal development or system debugging purposes. You can find full documentation for these switches in the body of the `Debug` unit in the compiler source file `debug.adb`.

-gnatG[=nn]

This switch causes the compiler to generate auxiliary output containing a pseudo-source listing of the generated expanded code. Like most Ada compilers, GNAT works by first transforming the high level Ada code into lower level constructs. For example, tasking operations are transformed into calls to the tasking run-time routines. A unique capability of GNAT is to list this expanded code in a form very close to normal Ada source. This is very useful in understanding the implications of various Ada usage on the efficiency of the generated code. There are many cases in Ada (e.g., the use of controlled types), where simple Ada statements can generate a lot of run-time code. By using `-gnatG` you can identify these cases, and consider whether it may be desirable to modify the coding approach to improve efficiency.

The optional parameter *nn* if present after `-gnatG` specifies an alternative maximum line length that overrides the normal default of 72. This value is in the range 40-999999, values less than 40 being silently reset to 40. The equal sign is optional.

The format of the output is very similar to standard Ada source, and is easily understood by an Ada programmer. The following special syntactic additions correspond to low level features used in the generated code that do not have any exact analogies in pure Ada source form. The following is a partial list of these special constructions. See the spec of package `Sprint` in file `sprint.ads` for a full list.

If the switch `-gnatL` is used in conjunction with `-gnatG`, then the original source lines are interspersed in the expanded source (as comment lines with the original line number).

new xxx [storage_pool = yyy]

Shows the storage pool being used for an allocator.

at end procedure-name;

Shows the finalization (cleanup) procedure for a scope.

(if expr then expr else expr)

Conditional expression equivalent to the `x?y:z` construction in C.

target^(source)

A conversion with floating-point truncation instead of rounding.

target?(source)

A conversion that bypasses normal Ada semantic checking. In particular enumeration types and fixed-point types are treated simply as integers.

target?^(source)

Combines the above two cases.

x #/ y

x #mod y

x # y

x #rem y

A division or multiplication of fixed-point values which are treated as integers without any kind of scaling.

free *expr* [storage_pool = *xxx*]

Shows the storage pool associated with a `free` statement.

[subtype or type declaration]

Used to list an equivalent declaration for an internally generated type that is referenced elsewhere in the listing.

freeze *type-name* [*actions*]

Shows the point at which *type-name* is frozen, with possible associated actions to be performed at the freeze point.

reference *itype*

Reference (and hence definition) to internal type *itype*.

***function-name!* (*arg*, *arg*, *arg*)**

Intrinsic function call.

***label-name* : label**

Declaration of label *labelname*.

#\$ *subprogram-name*

An implicit call to a run-time support routine (to meet the requirement of H.3.1(9) in a convenient manner).

expr* && *expr* && *expr* ... && *expr

A multiple concatenation (same effect as *expr* & *expr* & *expr*, but handled more efficiently).

[constraint_error]

Raise the `Constraint_Error` exception.

***expression*'reference**

A pointer to the result of evaluating {*expression*}.

***target-type!*(*source-expression*)**

An unchecked conversion of *source-expression* to *target-type*.

[*numerator/denominator*]

Used to represent internal real literals (that) have no exact representation in base 2-16 (for example, the result of compile time evaluation of the expression 1.0/27.0).

-gnatD[=nn]

When used in conjunction with `-gnatG`, this switch causes the expanded source, as described above for `-gnatG` to be written to files with names *xxx.dg*, where *xxx* is the normal file name, instead of to the standard output file. For example, if the source file name is `hello.adb`, then a file `hello.adb.dg` will be written. The debugging information generated by the `gcc -g`

switch will refer to the generated `xxx.dg` file. This allows you to do source level debugging using the generated code which is sometimes useful for complex code, for example to find out exactly which part of a complex construction raised an exception. This switch also suppresses generation of cross-reference information (see `-gnatx`) since otherwise the cross-reference information would refer to the `.dg` file, which would cause confusion since this is not the original source file.

Note that `-gnatD` actually implies `-gnatG` automatically, so it is not necessary to give both options. In other words `-gnatD` is equivalent to `-gnatDG`.

If the switch `-gnatL` is used in conjunction with `-gnatDG`, then the original source lines are interspersed in the expanded source (as comment lines with the original line number).

The optional parameter `nn` if present after `-gnatD` specifies an alternative maximum line length that overrides the normal default of 72. This value is in the range 40-999999, values less than 40 being silently reset to 40. The equal sign is optional.

-gnatr

This switch causes pragma Restrictions to be treated as `Restriction_Warnings` so that violation of restrictions causes warnings rather than illegalities. This is useful during the development process when new restrictions are added or investigated. The switch also causes pragma Profile to be treated as `Profile_Warnings`, and pragma Restricted_Run_Time and pragma Ravenscar set restriction warnings rather than restrictions.

-gnatR[0|1|2|3|4][e][j][m][s]

This switch controls output from the compiler of a listing showing representation information for declared types, objects and subprograms. For `-gnatR0`, no information is output (equivalent to omitting the `-gnatR` switch). For `-gnatR1` (which is the default, so `-gnatR` with no parameter has the same effect), size and alignment information is listed for declared array and record types.

For `-gnatR2`, size and alignment information is listed for all declared types and objects. The `Linker_Section` is also listed for any entity for which the `Linker_Section` is set explicitly or implicitly (the latter case occurs for objects of a type for which a `Linker_Section` is set).

For `-gnatR3`, symbolic expressions for values that are computed at run time for records are included. These symbolic expressions have a mostly obvious format with `#n` being used to represent the value of the `n`'th discriminant. See source files `repinfo.ads/adb` in the GNAT sources for full details on the format of `-gnatR3` output.

For `-gnatR4`, information for relevant compiler-generated types is also listed, i.e. when they are structurally part of other declared types and objects.

If the switch is followed by an `e` (e.g. `-gnatR2e`), then extended representation information for record sub-components of records is included.

If the switch is followed by an `m` (e.g. `-gnatRm`), then subprogram conventions and parameter passing mechanisms for all the subprograms are included.

If the switch is followed by a `j` (e.g., `-gnatRj`), then the output is in the JSON data interchange format specified by the ECMA-404 standard. The semantic description of this JSON output is available in the specification of the `Repinfo` unit present in the compiler sources.

If the switch is followed by an `s` (e.g., `-gnatR3s`), then the output is to a file with the name `file.rep` where `file` is the name of the corresponding source file, except if `j` is also specified, in which case the file name is `file.json`.

Note that it is possible for record components to have zero size. In this case, the component clause uses an obvious extension of permitted Ada syntax, for example `at 0 range 0 .. -1`.

-gnatS

The use of the switch `-gnatS` for an Ada compilation will cause the compiler to output a representation of package Standard in a form very close to standard Ada. It is not quite possible to do this entirely in standard Ada (since new numeric base types cannot be created in standard Ada), but the output is easily readable to any Ada programmer, and is useful to determine the characteristics of target dependent types in package Standard.

-gnatx

Normally the compiler generates full cross-referencing information in the ALI file. This information is used by a number of tools. The `-gnatx` switch suppresses this information. This saves some space and may slightly speed up compilation, but means that tools depending on this information cannot be used.

-fgnat-encodings=[all|gdb|minimal]

This switch controls the balance between GNAT encodings and standard DWARF emitted in the debug information.

Historically, old debug formats like stabs were not powerful enough to express some Ada types (for instance, variant records or fixed-point types). To work around this, GNAT introduced proprietary encodings that embed the missing information (“GNAT encodings”).

Recent versions of the DWARF debug information format are now able to correctly describe most of these Ada constructs (“standard DWARF”). As third-party tools started to use this format, GNAT has been enhanced to generate it. However, most tools (including GDB) are still relying on GNAT encodings.

To support all tools, GNAT needs to be versatile about the balance between generation of GNAT encodings and standard DWARF. This is what `-fgnat-encodings` is about.

- `=all`: Emit all GNAT encodings, and then emit as much standard DWARF as possible so it does not conflict with GNAT encodings.
- `=gdb`: Emit as much standard DWARF as possible as long as the current GDB handles it. Emit GNAT encodings for the rest.
- `=minimal`: Emit as much standard DWARF as possible and emit GNAT encodings for the rest.

4.3.16 Exception Handling Control

GNAT uses two methods for handling exceptions at run time. The `setjmp/longjmp` method saves the context when entering a frame with an exception handler. Then when an exception is raised, the context can be restored immediately, without the need for tracing stack frames. This method provides very fast exception propagation, but introduces significant overhead for the use of exception handlers, even if no exception is raised.

The other approach is called 'zero cost' exception handling. With this method, the compiler builds static tables to describe the exception ranges. No dynamic code is required when entering a frame containing an exception handler. When an exception is raised, the tables are used to control a back trace of the subprogram invocation stack to locate the required exception handler. This method has considerably poorer performance for the propagation of exceptions, but there is no overhead for exception handlers if no exception is raised. Note that in this mode and in the context of mixed Ada and C/C++ programming, to propagate an exception through a C/C++ code, the C/C++ code must be compiled with the `-funwind-tables` GCC's option.

The following switches may be used to control which of the two exception handling methods is used.

--RTS=sjlj

This switch causes the `setjmp/longjmp` run-time (when available) to be used for exception handling. If the default mechanism for the target is zero cost exceptions, then this switch can be used to modify this default, and must be used for all units in the partition. This option is rarely used. One case in which it may be advantageous is if you have an application where exception raising is common and the overall performance of the application is improved by favoring exception propagation.

--RTS=zcx

This switch causes the zero cost approach to be used for exception handling. If this is the default mechanism for the target (see below), then this switch is unneeded. If the default mechanism for the target is `setjmp/longjmp` exceptions, then this switch can be used to modify this default, and must be used for all units in the partition. This option can only be used if the zero cost approach is available for the target in use, otherwise it will generate an error.

The same option `--RTS` must be used both for `gcc` and `gnatbind`. Passing this option to `gnatmake` ([Switches for gnatmake](#)) will ensure the required consistency through the compilation and binding steps.

4.3.17 Units to Sources Mapping Files

-gnatem=path

A mapping file is a way to communicate to the compiler two mappings: from unit names to file names (without any directory information) and from file names to path names (with full directory information). These mappings are used by the compiler to short-circuit the path search.

The use of mapping files is not required for correct operation of the compiler, but mapping files can improve efficiency, particularly when sources are read over a slow network connection. In

normal operation, you need not be concerned with the format or use of mapping files, and the `-gnatem` switch is not a switch that you would use explicitly. It is intended primarily for use by automatic tools such as `gnatmake` running under the project file facility. The description here of the format of mapping files is provided for completeness and for possible use by other tools.

A mapping file is a sequence of sets of three lines. In each set, the first line is the unit name, in lower case, with `%s` appended for specs and `%b` appended for bodies; the second line is the file name; and the third line is the path name.

Example:

```
main%b
main.2.ada
/gnat/project1/sources/main.2.ada
```

When the switch `-gnatem` is specified, the compiler will create in memory the two mappings from the specified file. If there is any problem (nonexistent file, truncated file or duplicate entries), no mapping will be created.

Several `-gnatem` switches may be specified; however, only the last one on the command line will be taken into account.

When using a project file, `gnatmake` creates a temporary mapping file and communicates it to the compiler using this switch.

4.3.18 Code Generation Control

The GCC technology provides a wide range of target dependent `-m` switches for controlling details of code generation with respect to different versions of architectures. This includes variations in instruction sets (e.g., different members of the power pc family), and different requirements for optimal arrangement of instructions (e.g., different members of the x86 family). The list of available `-m` switches may be found in the GCC documentation.

Use of these `-m` switches may in some cases result in improved code performance.

The GNAT technology is tested and qualified without any `-m` switches, so generally the most reliable approach is to avoid the use of these switches. However, we generally expect most of these switches to work successfully with GNAT, and many customers have reported successful use of these options.

Our general advice is to avoid the use of `-m` switches unless special needs lead to requirements in this area. In particular, there is no point in using `-m` switches to improve performance unless you actually see a performance improvement.

4.4 Linker Switches

Linker switches can be specified after `-larg`s builder switch.

`-fuse-ld=name`

Linker to be used. The default is `bfd` for `ld.bfd`; `gold` (for `ld.gold`) and `gold` (for `ld.gold`) are more recent and faster alternatives, but only available on GNU/Linux platforms.

4.5 Binding with `gnatbind`

This chapter describes the GNAT binder, `gnatbind`, which is used to bind compiled GNAT objects.

The `gnatbind` program performs four separate functions:

- Checks that a program is consistent, in accordance with the rules in Chapter 10 of the Ada Reference Manual. In particular, error messages are generated if a program uses inconsistent versions of a given unit.
- Checks that an acceptable order of elaboration exists for the program and issues an error message if it cannot find an order of elaboration that satisfies the rules in Chapter 10 of the Ada Language Manual.
- Generates a main program incorporating the given elaboration order. This program is a small Ada package (body and spec) that must be subsequently compiled using the GNAT compiler. The necessary compilation step is usually performed automatically by `gnatlink`. The two most important functions of this program are to call the elaboration routines of units in an appropriate order and to call the main program.
- Determines the set of object files required by the given main program. This information is output in the forms of comments in the generated program, to be read by the `gnatlink` utility used to link the Ada application.

4.5.1 Running `gnatbind`

The form of the `gnatbind` command is

```
$ gnatbind [ switches ] mainprog[.ali] [ switches ]
```

where `mainprog.adb` is the Ada file containing the main program unit body. `gnatbind` constructs an Ada package in two files whose names are `b-mainprog.ads`, and `b-mainprog.adb`. For example, if given the parameter `hello.ali`, for a main program contained in file `hello.adb`, the binder output files would be `b-hello.ads` and `b-hello.adb`.

When doing consistency checking, the binder takes into consideration any source files it can locate. For example, if the binder determines that the given main program requires the package `Pack`, whose `.ALI` file is `pack.ali` and whose corresponding source spec file is `pack.ads`, it attempts to locate the source file `pack.ads` (using the same search path conventions as previously described for the `gcc` command). If it can locate this source file, it checks that the time stamps or source checksums of the source and its references to in ALI files match. In other words, any ALI files that

mentions this spec must have resulted from compiling this version of the source file (or in the case where the source checksums match, a version close enough that the difference does not matter).

The effect of this consistency checking, which includes source files, is that the binder ensures that the program is consistent with the latest version of the source files that can be located at bind time. Editing a source file without compiling files that depend on the source file cause error messages to be generated by the binder.

For example, suppose you have a main program `hello.adb` and a package `P`, from file `p.ads` and you perform the following steps:

- Enter `gcc -c hello.adb` to compile the main program.
- Enter `gcc -c p.ads` to compile package `P`.
- Edit file `p.ads`.
- Enter `gnatbind hello`.

At this point, the file `p.ali` contains an out-of-date time stamp because the file `p.ads` has been edited. The attempt at binding fails, and the binder generates the following error messages:

```
error: "hello.adb" must be recompiled ("p.ads" has been modified)
error: "p.ads" has been modified and must be recompiled
```

Now both files must be recompiled as indicated, and then the bind can succeed, generating a main program. You need not normally be concerned with the contents of this file, but for reference purposes a sample binder output file is given in [Example of Binder Output File](#).

In most normal usage, the default mode of `gnatbind` which is to generate the main package in Ada, as described in the previous section. In particular, this means that any Ada programmer can read and understand the generated main program. It can also be debugged just like any other Ada code provided the `-g` switch is used for `gnatbind` and `gnatlink`.

4.5.2 Switches for `gnatbind`

The following switches are available with `gnatbind`; details will be presented in subsequent sections.

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

-a

Indicates that, if supported by the platform, the `adainit` procedure should be treated as an initialisation routine by the linker (a constructor). This is intended to be used by the Project Manager to automatically initialize shared Stand-Alone Libraries.

-a0

Specify directory to be searched for ALI files.

-aI

Specify directory to be searched for source file.

-A[=*filename*]

Output ALI list (to standard output or to the named file).

-b

Generate brief messages to `stderr` even if verbose mode set.

-c

Check only, no generation of binder output file.

-dnn[k|m]

This switch can be used to change the default task stack size value to a specified size `nn`, which is expressed in bytes by default, or in kilobytes when suffixed with `k` or in megabytes when suffixed with `m`. In the absence of a `[k|m]` suffix, this switch is equivalent, in effect, to completing all task specs with

```
pragma Storage_Size (nn);
```

When they do not already have such a pragma.

-Dnn[k|m]

Set the default secondary stack size to `nn`. The suffix indicates whether the size is in bytes (no suffix), kilobytes (`k` suffix) or megabytes (`m` suffix).

The secondary stack holds objects of unconstrained types that are returned by functions, for example unconstrained Strings. The size of the secondary stack can be dynamic or fixed depending on the target.

For most targets, the secondary stack grows on demand and is implemented as a chain of blocks in the heap. In this case, the default secondary stack size determines the initial size of the secondary stack for each task and the smallest amount the secondary stack can grow by.

For Ravenscar, ZFP, and Cert run-times the size of the secondary stack is fixed. This switch can be used to change the default size of these stacks. The default secondary stack size can be overridden on a per-task basis if individual tasks have different secondary stack requirements. This is achieved through the `Secondary_Stack_Size` aspect that takes the size of the secondary stack in bytes.

-e

Output complete list of elaboration-order dependencies.

-Ea

Store tracebacks in exception occurrences when the target supports it. The “a” is for “address”; tracebacks will contain hexadecimal addresses, unless symbolic tracebacks are enabled.

See also the packages `GNAT.Traceback` and `GNAT.Traceback.Symbolic` for more information. Note that on x86 ports, you must not use `-fomit-frame-pointer` gcc option.

-Es

Store tracebacks in exception occurrences when the target supports it. The “s” is for “symbolic”; symbolic tracebacks are enabled.

-E

Currently the same as `-Ea`.

-feLab-order

Force elaboration order. For further details see [Elaboration Control](#) and [Elaboration Order Handling in GNAT](#).

-F

Force the checks of elaboration flags. `gnatbind` does not normally generate checks of elaboration flags for the main executable, except when a Stand-Alone Library is used. However, there are cases when this cannot be detected by `gnatbind`. An example is importing an interface of a Stand-Alone Library through a pragma `Import` and only specifying through a linker switch this Stand-Alone Library. This switch is used to guarantee that elaboration flag checks are generated.

-h

Output usage (help) information.

-H

Legacy elaboration order model enabled. For further details see [Elaboration Order Handling in GNAT](#).

-H32

Use 32-bit allocations for `__gnat_malloc` (and thus for access types). For further details see [Dynamic Allocation Control](#).

-H64

Use 64-bit allocations for `__gnat_malloc` (and thus for access types). For further details see [Dynamic Allocation Control](#).

-I

Specify directory to be searched for source and ALI files.

-I-

Do not look for sources in the current directory where `gnatbind` was invoked, and do not look for ALI files in the directory containing the ALI file named in the `gnatbind` command line.

-k

Disable checking of elaboration flags. When using `-n` either explicitly or implicitly, `-F` is also implied, unless `-k` is used. This switch should be used with care and you should ensure manually that elaboration routines are not called twice unintentionally.

-K

Give list of linker options specified for link.

-l

Output chosen elaboration order.

-Lxxx

Bind the units for library building. In this case the `adainit` and `adafinal` procedures ([Binding with Non-Ada Main Programs](#)) are renamed to `xxxinit` and `xxxfinal`. Implies `-n`. ([GNAT and Libraries](#), for more details.)

-Mxyz

Rename generated main program from `main` to `xyz`. This option is supported on cross environments only.

-m*n*

Limit number of detected errors or warnings to *n*, where *n* is in the range 1..999999. The default value if no switch is given is 9999. If the number of warnings reaches this limit, then a message is output and further warnings are suppressed, the bind continues in this case. If the number of errors reaches this limit, then a message is output and the bind is abandoned. A value of zero means that no limit is enforced. The equal sign is optional.

-minimal

Generate a binder file suitable for space-constrained applications. When active, binder-generated objects not required for program operation are no longer generated. **Warning:** this option comes with the following limitations:

- Starting the program's execution in the debugger will cause it to stop at the start of the `main` function instead of the main subprogram. This can be worked around by manually inserting a breakpoint on that subprogram and resuming the program's execution until reaching that breakpoint.
- Programs using `GNAT.Compiler_Version` will not link.

-n

No main program.

-nostdinc

Do not look for sources in the system default directory.

-nostdlib

Do not look for library files in the system default directory.

--RTS=*rts-path*

Specifies the default location of the run-time library. Same meaning as the equivalent `gnatmake` flag ([Switches for gnatmake](#)).

-o *file*

Name the output file *file* (default is `b~`xxx.adb'`). Note that if this option is used, then linking must be done manually, `gnatlink` cannot be used.

-O[=*filename*]

Output object list (to standard output or to the named file).

-p

Pessimistic (worst-case) elaboration order.

-P

Generate binder file suitable for CodePeer.

-R

Output closure source list, which includes all non-run-time units that are included in the bind.

-Ra

Like `-R` but the list includes run-time units.

-s

Require all source files to be present.

-Sxxx

Specifies the value to be used when detecting uninitialized scalar objects with pragma `Initialize_Scalars`. The xxx string specified with the switch is one of:

- in for an invalid value.

If zero is invalid for the discrete type in question, then the scalar value is set to all zero bits. For signed discrete types, the largest possible negative value of the underlying scalar is set (i.e. a one bit followed by all zero bits). For unsigned discrete types, the underlying scalar value is set to all one bits. For floating-point types, a NaN value is set (see body of package `System.Scalar_Values` for exact values).

- lo for low value.

If zero is invalid for the discrete type in question, then the scalar value is set to all zero bits. For signed discrete types, the largest possible negative value of the underlying scalar is set (i.e. a one bit followed by all zero bits). For unsigned discrete types, the underlying scalar value is set to all zero bits. For floating-point, a small value is set (see body of package `System.Scalar_Values` for exact values).

- hi for high value.

If zero is invalid for the discrete type in question, then the scalar value is set to all one bits. For signed discrete types, the largest possible positive value of the underlying scalar is set (i.e. a zero bit followed by all one bits). For unsigned discrete types, the underlying scalar value is set to all one bits. For floating-point, a large value is set (see body of package `System.Scalar_Values` for exact values).

- xx for hex value (two hex digits).

The underlying scalar is set to a value consisting of repeated bytes, whose value corresponds to the given value. For example if `BF` is given, then a 32-bit scalar value will be set to the bit pattern `16#BFBFBFBF#`.

In addition, you can specify `-Sev` to indicate that the value is to be set at run time. In this case, the program will look for an environment variable of the form `GNAT_INIT_SCALARS=yy`, where `yy` is one of `in/lo/hi/xx` with the same meanings as above. If no environment variable is found, or if it does not have a valid value, then the default is `in` (invalid values).

-static

Link against a static GNAT run-time.

-shared

Link against a shared GNAT run-time when available.

-t

Tolerate time stamp and other consistency errors.

-Tn

Set the time slice value to `n` milliseconds. If the system supports the specification of a specific time slice value, then the indicated value is used. If the system does not support specific time slice values, but does support some general notion of round-robin scheduling, then any nonzero value will activate round-robin scheduling.

A value of zero is treated specially. It turns off time slicing, and in addition, indicates to the tasking run-time that the semantics should match as closely as possible the Annex D requirements of the Ada RM, and in particular sets the default scheduling policy to `FIFO_Within_Priorities`.

-un

Enable dynamic stack usage, with `n` results stored and displayed at program termination. A result is generated when a task terminates. Results that can't be stored are displayed on the fly, at task termination. This option is currently not supported on Itanium platforms. (See [Dynamic Stack Usage Analysis](#) for details.)

-v

Verbose mode. Write error messages, header, summary output to `stdout`.

-Vkey=value

Store the given association of `key` to `value` in the bind environment. Values stored this way can be retrieved at run time using `GNAT.Bind_Environment`.

-wx

Warning mode; `x = s/e` for suppress/treat as error.

-Wxe

Override default wide character encoding for standard `Text_IO` files.

-x

Exclude source files (check object consistency only).

-xdr

Use the target-independent XDR protocol for stream oriented attributes instead of the default implementation which is based on direct binary representations and is therefore target-and-endianness-dependent. However it does not support 128-bit integer types and the exception `Ada.IO_Exceptions.Device_Error` is raised if any attempt is made at streaming 128-bit integer types with it.

-Xnnn

Set default exit status value, normally 0 for POSIX compliance.

-y

Enable leap seconds support in `Ada.Calendar` and its children.

-z

No main subprogram.

You may obtain this listing of switches by running `gnatbind` with no arguments.

Consistency-Checking Modes

As described earlier, by default `gnatbind` checks that object files are consistent with one another and are consistent with any source files it can locate. The following switches control binder access to sources.

-s

Require source files to be present. In this mode, the binder must be able to locate all source files that are referenced, in order to check their consistency. In normal mode, if a source file cannot be located it is simply ignored. If you specify this switch, a missing source file is an error.

-Wxe

Override default wide character encoding for standard `Text_IO` files. Normally the default wide character encoding method used for standard `[Wide_[Wide_]]Text_IO` files is taken from the encoding specified for the main source input (see description of switch `-gnatWx` for the compiler). The use of this switch for the binder (which has the same set of possible arguments) overrides this default as specified.

-x

Exclude source files. In this mode, the binder only checks that ALI files are consistent with one another. Source files are not accessed. The binder runs faster in this mode, and there is still a guarantee that the resulting program is self-consistent. If a source file has been edited since it was last compiled, and you specify this switch, the binder will not detect that the object file is out of date with respect to the source file. Note that this is the mode that is automatically used by `gnatmake` because in this case the checking against sources has already been performed by `gnatmake` in the course of compilation (i.e., before binding).

Binder Error Message Control

The following switches provide control over the generation of error messages from the binder:

-v

Verbose mode. In the normal mode, brief error messages are generated to `stderr`. If this switch is present, a header is written to `stdout` and any error messages are directed to `stdout`. All that is written to `stderr` is a brief summary message.

-b

Generate brief error messages to `stderr` even if verbose mode is specified. This is relevant only when used with the `-v` switch.

-mn

Limits the number of error messages to `n`, a decimal integer in the range 1-999. The binder terminates immediately if this limit is reached.

-Mxxx

Renames the generated main program from `main` to `xxx`. This is useful in the case of

some cross-building environments, where the actual main program is separate from the one generated by `gnatbind`.

-ws

Suppress all warning messages.

-we

Treat any warning messages as fatal errors.

-t

The binder performs a number of consistency checks including:

- Check that time stamps of a given source unit are consistent
- Check that checksums of a given source unit are consistent
- Check that consistent versions of GNAT were used for compilation
- Check consistency of configuration pragmas as required

Normally failure of such checks, in accordance with the consistency requirements of the Ada Reference Manual, causes error messages to be generated which abort the binder and prevent the output of a binder file and subsequent link to obtain an executable.

The `-t` switch converts these error messages into warnings, so that binding and linking can continue to completion even in the presence of such errors. The result may be a failed link (due to missing symbols), or a non-functional executable which has undefined semantics.

Note: This means that `-t` should be used only in unusual situations, with extreme care.

Elaboration Control

The following switches provide additional control over the elaboration order. For further details see [Elaboration Order Handling in GNAT](#).

-felab-order

Force elaboration order.

`elab-order` should be the name of a “forced elaboration order file”, that is, a text file containing library item names, one per line. A name of the form “`some.unit%s`” or “`some.unit (spec)`” denotes the spec of `Some.Unit`. A name of the form “`some.unit%b`” or “`some.unit (body)`” denotes the body of `Some.Unit`. Each pair of lines is taken to mean that there is an elaboration dependence of the second line on the first. For example, if the file contains:

```
this (spec)
this (body)
that (spec)
that (body)
```

then the spec of `This` will be elaborated before the body of `This`, and the body of `This` will be elaborated before the spec of `That`, and the spec of `That` will be elaborated before the body

of That. The first and last of these three dependences are already required by Ada rules, so this file is really just forcing the body of This to be elaborated before the spec of That.

The given order must be consistent with Ada rules, or else `gnatbind` will give elaboration cycle errors. For example, if you say `x (body)` should be elaborated before `x (spec)`, there will be a cycle, because Ada rules require `x (spec)` to be elaborated before `x (body)`; you can't have the spec and body both elaborated before each other.

If you later add “with That;” to the body of This, there will be a cycle, in which case you should erase either “this (body)” or “that (spec)” from the above forced elaboration order file.

Blank lines and Ada-style comments are ignored. Unit names that do not exist in the program are ignored. Units in the GNAT predefined library are also ignored.

-p

Pessimistic elaboration order

This switch is only applicable to the pre-20.x legacy elaboration models. The post-20.x elaboration model uses a more informed approach of ordering the units.

Normally the binder attempts to choose an elaboration order that is likely to minimize the likelihood of an elaboration order error resulting in raising a `Program_Error` exception. This switch reverses the action of the binder, and requests that it deliberately choose an order that is likely to maximize the likelihood of an elaboration error. This is useful in ensuring portability and avoiding dependence on accidental fortuitous elaboration ordering.

Normally it only makes sense to use the `-p` switch if dynamic elaboration checking is used (`-gnatE` switch used for compilation). This is because in the default static elaboration mode, all necessary `Elaborate` and `Elaborate_All` pragmas are implicitly inserted. These implicit pragmas are still respected by the binder in `-p` mode, so a safe elaboration order is assured.

Note that `-p` is not intended for production use; it is more for debugging/experimental use.

Output Control

The following switches allow additional control over the output generated by the binder.

-c

Check only. Do not generate the binder output file. In this mode the binder performs all error checks but does not generate an output file.

-e

Output complete list of elaboration-order dependencies, showing the reason for each dependency. This output can be rather extensive but may be useful in diagnosing problems with elaboration order. The output is written to `stdout`.

-h

Output usage information. The output is written to `stdout`.

- K**
Output linker options to `stdout`. Includes library search paths, contents of pragmas `Ident` and `Linker_Options`, and libraries added by `gnatbind`.
- l**
Output chosen elaboration order. The output is written to `stdout`.
- O**
Output full names of all the object files that must be linked to provide the Ada component of the program. The output is written to `stdout`. This list includes the files explicitly supplied and referenced by the user as well as implicitly referenced run-time unit files. The latter are omitted if the corresponding units reside in shared libraries. The directory names for the run-time units depend on the system configuration.
- o file**
Set name of output file to `file` instead of the normal `b-`mainprog.adb'` default. Note that `file` denote the Ada binder generated body filename. Note that if this option is used, then linking must be done manually. It is not possible to use `gnatlink` in this case, since it cannot locate the binder file.
- r**
Generate list of `pragma Restrictions` that could be applied to the current unit. This is useful for code audit purposes, and also may be used to improve code generation in some cases.

Dynamic Allocation Control

The heap control switches – `-H32` and `-H64` – determine whether dynamic allocation uses 32-bit or 64-bit memory. They only affect compiler-generated allocations via `__gnat_malloc`; explicit calls to `malloc` and related functions from the C run-time library are unaffected.

-H32
Allocate memory on 32-bit heap

-H64
Allocate memory on 64-bit heap. This is the default unless explicitly overridden by a `'Size` clause on the access type.

These switches are only effective on VMS platforms.

Binding with Non-Ada Main Programs

The description so far has assumed that the main program is in Ada, and that the task of the binder is to generate a corresponding function `main` that invokes this Ada main program. GNAT also supports the building of executable programs where the main program is not in Ada, but some of the called routines are written in Ada and compiled using GNAT ([Mixed Language Programming](#)). The following switch is used in this situation:

-n
No main program. The main program is not in Ada.

In this case, most of the functions of the binder are still required, but instead of generating a main program, the binder generates a file containing the following callable routines:

adainit

You must call this routine to initialize the Ada part of the program by calling the necessary elaboration routines. A call to `adainit` is required before the first call to an Ada subprogram.

Note that it is assumed that the basic execution environment must be setup to be appropriate for Ada execution at the point where the first Ada subprogram is called. In particular, if the Ada code will do any floating-point operations, then the FPU must be setup in an appropriate manner. For the case of the x86, for example, full precision mode is required. The procedure `GNAT.Float_Control.Reset` may be used to ensure that the FPU is in the right state.

adafinal

You must call this routine to perform any library-level finalization required by the Ada subprograms. A call to `adafinal` is required after the last call to an Ada subprogram, and before the program terminates.

If the `-n` switch is given, more than one ALI file may appear on the command line for `gnatbind`. The normal `closure` calculation is performed for each of the specified units. Calculating the closure means finding out the set of units involved by tracing *with* references. The reason it is necessary to be able to specify more than one ALI file is that a given program may invoke two or more quite separate groups of Ada units.

The binder takes the name of its output file from the last specified ALI file, unless overridden by the use of the `-o file`.

The output is an Ada unit in source form that can be compiled with GNAT. This compilation occurs automatically as part of the `gnatlink` processing.

Currently the GNAT run-time requires a FPU using 80 bits mode precision. Under targets where this is not the default it is required to call `GNAT.Float_Control.Reset` before using floating point numbers (this include float computation, float input and output) in the Ada code. A side effect is that this could be the wrong mode for the foreign code where floating point computation could be broken after this call.

Binding Programs with No Main Subprogram

It is possible to have an Ada program which does not have a main subprogram. This program will call the elaboration routines of all the packages, then the finalization routines.

The following switch is used to bind programs organized in this manner:

-z

Normally the binder checks that the unit name given on the command line corresponds to a suitable main subprogram. When this switch is used, a list of ALI files can be given, and the execution of the program consists of elaboration of these units in an appropriate order. Note

that the default wide character encoding method for standard Text_IO files is always set to Brackets if this switch is set (you can use the binder switch `-Wx` to override this default).

4.5.3 Command-Line Access

The package `Ada.Command_Line` provides access to the command-line arguments and program name. In order for this interface to operate correctly, the two variables

```
int gnat_argc;
char **gnat_argv;
```

are declared in one of the GNAT library routines. These variables must be set from the actual `argc` and `argv` values passed to the main program. With no `n` present, `gnatbind` generates the C main program to automatically set these variables. If the `n` switch is used, there is no automatic way to set these variables. If they are not set, the procedures in `Ada.Command_Line` will not be available, and any attempt to use them will raise `Constraint_Error`. If command line access is required, your main program must set `gnat_argc` and `gnat_argv` from the `argc` and `argv` values passed to it.

4.5.4 Search Paths for `gnatbind`

The binder takes the name of an ALI file as its argument and needs to locate source files as well as other ALI files to verify object consistency.

For source files, it follows exactly the same search rules as `gcc` (see [Search Paths and the Run-Time Library \(RTL\)](#)). For ALI files the directories searched are:

- The directory containing the ALI file named in the command line, unless the switch `-I-` is specified.
- All directories specified by `-I` switches on the `gnatbind` command line, in the order given.
- Each of the directories listed in the text file whose name is given by the `ADA_PRJ_OBJECTS_FILE` environment variable.

`ADA_PRJ_OBJECTS_FILE` is normally set by `gnatmake` or by the `gnat` driver when project files are used. It should not normally be set by other means.

- Each of the directories listed in the value of the `ADA_OBJECTS_PATH` environment variable. Construct this value exactly as the `PATH` environment variable: a list of directory names separated by colons (semicolons when working with the NT version of GNAT).
- The content of the `ada_object_path` file which is part of the GNAT installation tree and is used to store standard libraries such as the GNAT Run-Time Library (RTL) unless the switch `-nostdlib` is specified. See [Installing a library](#)

In the binder the switch `-I` is used to specify both source and library file paths. Use `-aI` instead if you want to specify source paths only, and `-a0` if you want to specify library paths only. This means that for the binder `-I dir` is equivalent to `-aI dir -a0` dir`. The binder generates the bind file (a C language source file) in the current working directory.

The packages `Ada`, `System`, and `Interfaces` and their children make up the GNAT Run-Time Library, together with the package `GNAT` and its children, which contain a set of useful additional library functions provided by GNAT. The sources for these units are needed by the compiler and are kept together in one directory. The ALI files and object files generated by compiling the RTL are needed by the binder and the linker and are kept together in one directory, typically different from the directory containing the sources. In a normal installation, you need not specify these directory names when compiling or binding. Either the environment variables or the built-in defaults cause these files to be found.

Besides simplifying access to the RTL, a major use of search paths is in compiling sources from multiple directories. This can make development environments much more flexible.

4.5.5 Examples of `gnatbind` Usage

Here are some examples of `gnatbind` invocations:

```
gnatbind hello
```

The main program `Hello` (source program in `hello.adb`) is bound using the standard switch settings. The generated main program is `b-hello.adb`. This is the normal, default use of the binder.

```
gnatbind hello -o mainprog.adb
```

The main program `Hello` (source program in `hello.adb`) is bound using the standard switch settings. The generated main program is `mainprog.adb` with the associated spec in `mainprog.ads`. Note that you must specify the body here not the spec. Note that if this option is used, then linking must be done manually, since `gnatlink` will not be able to find the generated file.

4.6 Linking with `gnatlink`

This chapter discusses `gnatlink`, a tool that links an Ada program and builds an executable file. This utility invokes the system linker (via the `gcc` command) with a correct list of object files and library references. `gnatlink` automatically determines the list of files and references for the Ada part of a program. It uses the binder file generated by the `gnatbind` to determine this list.

4.6.1 Running `gnatlink`

The form of the `gnatlink` command is

```
$ gnatlink [ switches ] mainprog [.ali]
           [ non-Ada objects ] [ linker options ]
```

The arguments of `gnatlink` (switches, main ALI file, non-Ada objects or linker options) may be in any order, provided that no non-Ada object may be mistaken for a main ALI file. Any file name

F without the `.ali` extension will be taken as the main ALI file if a file exists whose name is the concatenation of F and `.ali`.

`mainprog.ali` references the ALI file of the main program. The `.ali` extension of this file can be omitted. From this reference, `gnatlink` locates the corresponding binder file `b-mainprog.adb` and, using the information in this file along with the list of non-Ada objects and linker options, constructs a linker command file to create the executable.

The arguments other than the `gnatlink` switches and the main ALI file are passed to the linker uninterpreted. They typically include the names of object files for units written in other languages than Ada and any library references required to resolve references in any of these foreign language units, or in `Import` pragmas in any Ada units.

`linker options` is an optional list of linker specific switches. The default linker called by `gnatlink` is `gcc` which in turn calls the appropriate system linker.

One useful option for the linker is `-s`: it reduces the size of the executable by removing all symbol table and relocation information from the executable.

Standard options for the linker such as `-lmy_lib` or `-Ldir` can be added as is. For options that are not recognized by `gcc` as linker options, use the `gcc` switches `-Xlinker` or `-Wl,.`

Refer to the GCC documentation for details.

Here is an example showing how to generate a linker map:

```
$ gnatlink my_prog -Wl,-Map,MAPFILE
```

Using `linker options` it is possible to set the program stack and heap size. See [Setting Stack Size from gnatlink](#) and [Setting Heap Size from gnatlink](#).

`gnatlink` determines the list of objects required by the Ada program and prepends them to the list of objects passed to the linker. `gnatlink` also gathers any arguments set by the use of `pragma Linker_Options` and adds them to the list of arguments presented to the linker.

4.6.2 Switches for gnatlink

The following switches are available with the `gnatlink` utility:

--version

Display Copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

-f

On some targets, the command line length is limited, and `gnatlink` will generate a separate file for the linker if the list of object files is too long. The `-f` switch forces this file to be generated even if the limit is not exceeded. This is useful in some cases to deal with special situations where the command line length is exceeded.

-g

The option to include debugging information causes the Ada bind file (in other words,

b~mainprog.adb) to be compiled with `-g`. In addition, the binder does not delete the `b~mainprog.adb`, `b~mainprog.o` and `b~mainprog.ali` files. Without `-g`, the binder removes these files by default.

-n

Do not compile the file generated by the binder. This may be used when a link is rerun with different options, but there is no need to recompile the binder file.

-v

Verbose mode. Causes additional information to be output, including a full list of the included object files. This switch option is most useful when you want to see what set of object files are being used in the link step.

-v -v

Very verbose mode. Requests that the compiler operate in verbose mode when it compiles the binder file, and that the system linker run in verbose mode.

-o *exec-name*

exec-name specifies an alternate name for the generated executable program. If this switch is omitted, the executable has the same name as the main unit. For example, `gnatlink try.ali` creates an executable called `try`.

-Bdir

Load compiler executables (for example, `gnat1`, the Ada compiler) from `dir` instead of the default location. Only use this switch when multiple versions of the GNAT compiler are available. See the **Directory Options** section in *The GNU Compiler Collection* for further details. You would normally use the `-b` or `-V` switch instead.

-M

When linking an executable, create a map file. The name of the map file has the same name as the executable with extension `“.map”`.

-M=*mapfile*

When linking an executable, create a map file. The name of the map file is `mapfile`.

--GCC=*compiler_name*

Program used for compiling the binder file. The default is `gcc`. You need to use quotes around `compiler_name` if `compiler_name` contains spaces or other separator characters. As an example `--GCC="foo -x -y"` will instruct `gnatlink` to use `foo -x -y` as your compiler. Note that switch `-c` is always inserted after your command name. Thus in the above example the compiler command that will be used by `gnatlink` will be `foo -c -x -y`. A limitation of this syntax is that the name and path name of the executable itself must not include any embedded spaces. If the compiler executable is different from the default one (`gcc` or `<prefix>-gcc`), then the back-end switches in the ALI file are not used to compile the binder generated source. For example, this is the case with `--GCC="foo -x -y"`. But the back end switches will be used for `--GCC="gcc -gnatv"`. If several `--GCC=compiler_name` are used, only the last `compiler_name` is taken into account. However, all the additional switches are also taken into account. Thus, `--GCC="foo -x -y" --GCC="bar -z -t"` is equivalent to `--GCC="bar -x -y -z -t"`.

--LINK=*name*

name is the name of the linker to be invoked. This is especially useful in mixed language programs since languages such as C++ require their own linker to be used. When this switch

is omitted, the default name for the linker is `gcc`. When this switch is used, the specified linker is called instead of `gcc` with exactly the same parameters that would have been passed to `gcc` so if the desired linker requires different parameters it is necessary to use a wrapper script that massages the parameters before invoking the real linker. It may be useful to control the exact invocation by using the verbose switch.

4.7 Using the GNU make Utility

This chapter offers some examples of makefiles that solve specific problems. It does not explain how to write a makefile, nor does it try to replace the `gnatmake` utility ([Building with gnatmake](#)).

All the examples in this section are specific to the GNU version of `make`. Although `make` is a standard utility, and the basic language is the same, these examples use some advanced features found only in GNU `make`.

4.7.1 Using gnatmake in a Makefile

Complex project organizations can be handled in a very powerful way by using GNU `make` combined with `gnatmake`. For instance, here is a Makefile which allows you to build each subsystem of a big project into a separate shared library. Such a makefile allows you to significantly reduce the link time of very big applications while maintaining full coherence at each step of the build process.

The list of dependencies are handled automatically by `gnatmake`. The Makefile is simply used to call `gnatmake` in each of the appropriate directories.

Note that you should also read the example on how to automatically create the list of directories ([Automatically Creating a List of Directories](#)) which might help you in case your project has a lot of subdirectories.

```
## This Makefile is intended to be used with the following directory
## configuration:
## - The sources are split into a series of csc (computer software components)
##   Each of these csc is put in its own directory.
##   Their name are referenced by the directory names.
##   They will be compiled into shared library (although this would also work
##   with static libraries)
## - The main program (and possibly other packages that do not belong to any
##   csc) is put in the top level directory (where the Makefile is).
##   toplevel_dir __ first_csc (sources) __ lib (will contain the library)
##           \_ second_csc (sources) __ lib (will contain the library)
##           \_ ...
## Although this Makefile is build for shared library, it is easy to modify
## to build partial link objects instead (modify the lines with -shared and
## gnatlink below)
##
## With this makefile, you can change any file in the system or add any new
## file, and everything will be recompiled correctly (only the relevant shared
## objects will be recompiled, and the main program will be re-linked).
```

(continues on next page)

(continued from previous page)

```

# The list of computer software component for your project. This might be
# generated automatically.
CSC_LIST=aa bb cc

# Name of the main program (no extension)
MAIN=main

# If we need to build objects with -fPIC, uncomment the following line
#NEED_FPIC=-fPIC

# The following variable should give the directory containing libgnat.so
# You can get this directory through 'gnatls -v'. This is usually the last
# directory in the Object_Path.
GLIB=...

# The directories for the libraries
# (This macro expands the list of CSC to the list of shared libraries, you
# could simply use the expanded form:
# LIB_DIR=aa/lib/libaa.so bb/lib/libbb.so cc/lib/libcc.so
LIB_DIR=${foreach dir,${CSC_LIST},${dir}/lib/lib${dir}.so}

${MAIN}: objects ${LIB_DIR}
    gnatbind ${MAIN} ${CSC_LIST:%=-a0%/lib} -shared
    gnatlink ${MAIN} ${CSC_LIST:%=-l%}

objects::
    # recompile the sources
    gnatmake -c -i ${MAIN}.adb ${NEED_FPIC} ${CSC_LIST:%=-I%}

# Note: In a future version of GNAT, the following commands will be simplified
# by a new tool, gnatlib
${LIB_DIR}:
    mkdir -p ${dir} ${@}
    cd ${dir} ${@} && gcc -shared -o ${notdir} ${@} ../*.o -L${GLIB} -lgnat
    cd ${dir} ${@} && cp -f ../*.ali .

# The dependencies for the modules
# Note that we have to force the expansion of *.o, since in some cases
# make won't be able to do it itself.
aa/lib/libaa.so: ${wildcard aa/*.o}
bb/lib/libbb.so: ${wildcard bb/*.o}
cc/lib/libcc.so: ${wildcard cc/*.o}

# Make sure all of the shared libraries are in the path before starting the
# program
run::
    LD_LIBRARY_PATH=`pwd`/aa/lib:`pwd`/bb/lib:`pwd`/cc/lib ./${MAIN}

clean::
    ${RM} -rf ${CSC_LIST:%=%/lib}
    ${RM} ${CSC_LIST:%=%/*.ali}

```

(continues on next page)

(continued from previous page)

```

${RM} ${CSC_LIST:%=/*/*.o}
${RM} *.o *.ali ${MAIN}

```

4.7.2 Automatically Creating a List of Directories

In most makefiles, you will have to specify a list of directories, and store it in a variable. For small projects, it is often easier to specify each of them by hand, since you then have full control over what is the proper order for these directories, which ones should be included.

However, in larger projects, which might involve hundreds of subdirectories, it might be more convenient to generate this list automatically.

The example below presents two methods. The first one, although less general, gives you more control over the list. It involves wildcard characters, that are automatically expanded by `make`. Its shortcoming is that you need to explicitly specify some of the organization of your project, such as for instance the directory tree depth, whether some directories are found in a separate tree, etc.

The second method is the most general one. It requires an external program, called `find`, which is standard on all Unix systems. All the directories found under a given root directory will be added to the list.

```

# The examples below are based on the following directory hierarchy:
# All the directories can contain any number of files
# ROOT_DIRECTORY -> a -> aa -> aaa
#                   -> ab
#                   -> ac
#                   -> b -> ba -> baa
#                   -> bb
#                   -> bc
# This Makefile creates a variable called DIRS, that can be reused any time
# you need this list (see the other examples in this section)

# The root of your project's directory hierarchy
ROOT_DIRECTORY=.

####
# First method: specify explicitly the list of directories
# This allows you to specify any subset of all the directories you need.
####
DIRS := a/aa/ a/ab/ b/ba/

####
# Second method: use wildcards
# Note that the argument(s) to wildcard below should end with a '/'.
# Since wildcards also return file names, we have to filter them out
# to avoid duplicate directory names.
# We thus use make's ``dir`` and ``sort`` functions.
# It sets DIRS to the following value (note that the directories aaa and baa
# are not given, unless you change the arguments to wildcard).

```

(continues on next page)

(continued from previous page)

```

# DIRS= ./a/a/ ./b/ ./a/aa/ ./a/ab/ ./a/ac/ ./b/ba/ ./b/bb/ ./b/bc/
####

DIRS := ${sort $(dir ${wildcard ${ROOT_DIRECTORY}/*/
                ${ROOT_DIRECTORY}/*/*})}

####
# Third method: use an external program
# This command is much faster if run on local disks, avoiding NFS slowdowns.
# This is the most complete command: it sets DIRS to the following value:
# DIRS= ./a ./a/aa ./a/aa/aaa ./a/ab ./a/ac ./b ./b/ba ./b/ba/baa ./b/bb ./b/bc
####

DIRS := ${shell find ${ROOT_DIRECTORY} -type d -print}

```

4.7.3 Generating the Command Line Switches

Once you have created the list of directories as explained in the previous section ([Automatically Creating a List of Directories](#)), you can easily generate the command line arguments to pass to `gnatmake`.

For the sake of completeness, this example assumes that the source path is not the same as the object path, and that you have two separate lists of directories.

```

# see "Automatically creating a list of directories" to create
# these variables
SOURCE_DIRS=
OBJECT_DIRS=

GNATMAKE_SWITCHES := ${patsubst %, -aI%, ${SOURCE_DIRS}}
GNATMAKE_SWITCHES += ${patsubst %, -a0%, ${OBJECT_DIRS}}

all:
    gnatmake ${GNATMAKE_SWITCHES} main_unit

```

4.7.4 Overcoming Command Line Length Limits

One problem that might be encountered on big projects is that many operating systems limit the length of the command line. It is thus hard to give `gnatmake` the list of source and object directories.

This example shows how you can set up environment variables, which will make `gnatmake` behave exactly as if the directories had been specified on the command line, but have a much higher length limit (or even none on most systems).

It assumes that you have created a list of directories in your Makefile, using one of the methods presented in [Automatically Creating a List of Directories](#). For the sake of completeness, we assume that the object path (where the ALI files are found) is different from the sources patch.

Note a small trick in the Makefile below: for efficiency reasons, we create two temporary variables (SOURCE_LIST and OBJECT_LIST), that are expanded immediately by make. This way we overcome the standard make behavior which is to expand the variables only when they are actually used.

On Windows, if you are using the standard Windows command shell, you must replace colons with semicolons in the assignments to these variables.

```
# In this example, we create both ADA_INCLUDE_PATH and ADA_OBJECTS_PATH.
# This is the same thing as putting the -I arguments on the command line.
# (the equivalent of using -aI on the command line would be to define
# only ADA_INCLUDE_PATH, the equivalent of -aO is ADA_OBJECTS_PATH).
# You can of course have different values for these variables.
#
# Note also that we need to keep the previous values of these variables, since
# they might have been set before running 'make' to specify where the GNAT
# library is installed.

# see "Automatically creating a list of directories" to create these
# variables
SOURCE_DIRS=
OBJECT_DIRS=

empty:=
space:=${empty} ${empty}
SOURCE_LIST := ${subst ${space},.,${SOURCE_DIRS}}
OBJECT_LIST := ${subst ${space},.,${OBJECT_DIRS}}
ADA_INCLUDE_PATH += ${SOURCE_LIST}
ADA_OBJECTS_PATH += ${OBJECT_LIST}
export ADA_INCLUDE_PATH
export ADA_OBJECTS_PATH

all:
    gnatmake main_unit
```

This page is intentionally left blank.

GNAT UTILITY PROGRAMS

This chapter describes a number of utility programs:

- The File Cleanup Utility `gnatclean`
- The GNAT Library Browser `gnatls`

Other GNAT utilities are described elsewhere in this manual:

- Handling Arbitrary File Naming Conventions with `gnatname`
- File Name Krunching with `gnatkr`
- Renaming Files with `gnatchop`
- Preprocessing with `gnatprep`

5.1 The File Cleanup Utility `gnatclean`

`gnatclean` is a tool that allows the deletion of files produced by the compiler, binder and linker, including ALI files, object files, tree files, expanded source files, library files, interface copy source files, binder generated files and executable files.

5.1.1 Running `gnatclean`

The `gnatclean` command has the form:

```
$ gnatclean switches names
```

where `names` is a list of source file names. Suffixes `.ads` and `adb` may be omitted. If a project file is specified using switch `-P`, then `names` may be completely omitted.

In normal mode, `gnatclean` delete the files produced by the compiler and, if switch `-c` is not specified, by the binder and the linker. In informative-only mode, specified by switch `-n`, the list of files that would have been deleted in normal mode is listed, but no file is actually deleted.

5.1.2 Switches for `gnatclean`

`gnatclean` recognizes the following switches:

--version

Display copyright and version, then exit disregarding all other options.

--help

If `--version` was not used, display usage, then exit disregarding all other options.

--subdirs=*subdir*

Actual object directory of each project file is the subdirectory *subdir* of the object directory specified or defaulted in the project file.

--unchecked-shared-lib-imports

By default, shared library projects are not allowed to import static library projects. When this switch is used on the command line, this restriction is relaxed.

-c

Only attempt to delete the files produced by the compiler, not those produced by the binder or the linker. The files that are not to be deleted are library files, interface copy files, binder generated files and executable files.

-D *dir*

Indicate that ALI and object files should normally be found in directory *dir*.

-F

When using project files, if some errors or warnings are detected during parsing and verbose mode is not in effect (no use of switch `-v`), then error lines start with the full path name of the project file, rather than its simple file name.

-h

Output a message explaining the usage of `gnatclean`.

-n

Informative-only mode. Do not delete any files. Output the list of the files that would have been deleted if this switch was not specified.

-P*project*

Use project file *project*. Only one such switch can be used. When cleaning a project file, the files produced by the compilation of the immediate sources or inherited sources of the project files are to be deleted. This is not depending on the presence or not of executable names on the command line.

-q

Quiet output. If there are no errors, do not output anything, except in verbose mode (switch `-v`) or in informative-only mode (switch `-n`).

-r

When a project file is specified (using switch `-P`), clean all imported and extended project files, recursively. If this switch is not specified, only the files related to the main project file are to be deleted. This switch has no effect if no project file is specified.

- v**
Verbose mode.
- vPx**
Indicates the verbosity of the parsing of GNAT project files. See Switches Related to Project Files.
- Xname=value**
Indicates that external variable `name` has the value `value`. The Project Manager will use this value for occurrences of `external(name)` when parsing the project file. See Switches Related to Project Files.
- a0dir**
When searching for ALI and object files, look in directory `dir`.
- Idir**
Equivalent to `-a0dir`.
- I-**
Do not look for ALI or object files in the directory where `gnatclean` was invoked.

5.2 The GNAT Library Browser `gnatls`

`gnatls` is a tool that outputs information about compiled units. It gives the relationship between objects, unit names and source files. It can also be used to check the source dependencies of a unit as well as various characteristics.

5.2.1 Running `gnatls`

The `gnatls` command has the form

```
$ gnatls switches object_or_ali_file
```

The main argument is the list of object or `ali` files (see [The Ada Library Information Files](#)) for which information is requested.

In normal mode, without additional option, `gnatls` produces a four-column listing. Each line represents information for a specific object. The first column gives the full path of the object, the second column gives the name of the principal unit in this object, the third column gives the status of the source and the fourth column gives the full path of the source representing this unit. Here is a simple example of use:

```
$ gnatls *.o
./demo1.o      demo1      DIF demo1.adb
./demo2.o      demo2      OK demo2.adb
./hello.o      h1         OK hello.adb
./instr-child.o instr.child MOK instr-child.adb
./instr.o      instr      OK instr.adb
./tef.o        tef        DIF tef.adb
```

(continues on next page)

(continued from previous page)

<pre>./text_io_example.o text_io_example OK text_io_example.adb ./tgef.o tgef DIF tgef.adb</pre>
--

The first line can be interpreted as follows: the main unit which is contained in object file `demo1.o` is `demo1`, whose main source is in `demo1.adb`. Furthermore, the version of the source used for the compilation of `demo1` has been modified (DIF). Each source file has a status qualifier which can be:

OK (unchanged)

The version of the source file used for the compilation of the specified unit corresponds exactly to the actual source file.

MOK (slightly modified)

The version of the source file used for the compilation of the specified unit differs from the actual source file but not enough to require recompilation. If you use `gnatmake` with the option `-m` (minimal recompilation), a file marked MOK will not be recompiled.

DIF (modified)

No version of the source found on the path corresponds to the source used to build this object.

***???* (file not found)**

No source file was found for this unit.

HID (hidden, unchanged version not first on PATH)

The version of the source that corresponds exactly to the source used for compilation has been found on the path but it is hidden by another version of the same source that has been modified.

5.2.2 Switches for `gnatls`

`gnatls` recognizes the following switches:

`--version`

Display copyright and version, then exit disregarding all other options.

`--help`

If `--version` was not used, display usage, then exit disregarding all other options.

`-a`

Consider all units, including those of the predefined Ada library. Especially useful with `-d`.

`-d`

List sources from which specified units depend on.

`-h`

Output the list of options.

`-o`

Only output information about object files.

`-s`

Only output information about source files.

-u

Only output information about compilation units.

-files=file

Take as arguments the files listed in text file `file`. Text file `file` may contain empty lines that are ignored. Each nonempty line should contain the name of an existing file. Several such switches may be specified simultaneously.

-a0dir, -aIdir, -Idir, -I-, -nostdinc

Source path manipulation. Same meaning as the equivalent `gnatmake` flags ([Switches for gnatmake](#)).

-aPdir

Add `dir` at the beginning of the project search dir.

--RTS=rts-path

Specifies the default location of the runtime library. Same meaning as the equivalent `gnatmake` flag ([Switches for gnatmake](#)).

-v

Verbose mode. Output the complete source, object and project paths. Do not use the default column layout but instead use long format giving as much as information possible on each requested units, including special characteristics such as:

- *Preelaborable*: The unit is preelaborable in the Ada sense.
- *No_Elab_Code*: No elaboration code has been produced by the compiler for this unit.
- *Pure*: The unit is pure in the Ada sense.
- *Elaborate_Body*: The unit contains a pragma `Elaborate_Body`.
- *Remote_Types*: The unit contains a pragma `Remote_Types`.
- *Shared_Passive*: The unit contains a pragma `Shared_Passive`.
- *Predefined*: This unit is part of the predefined environment and cannot be modified by the user.
- *Remote_Call_Interface*: The unit contains a pragma `Remote_Call_Interface`.

5.2.3 Example of `gnatls` Usage

Example of using the verbose switch. Note how the source and object paths are affected by the `-I` switch.

```
$ gnatls -v -I.. demo1.o

GNATLS 5.03w (20041123-34)
Copyright 1997-2004 Free Software Foundation, Inc.

Source Search Path:
  <Current_Directory>
  ../
```

(continues on next page)

(continued from previous page)

```

/home/comar/local/adainclude/

Object Search Path:
<Current_Directory>
../
/home/comar/local/lib/gcc-lib/x86-linux/3.4.3/adalib/

Project Search Path:
<Current_Directory>
/home/comar/local/lib/gnat/

./demo1.o
Unit =>
  Name   => demo1
  Kind   => subprogram body
  Flags  => No_Elab_Code
  Source => demo1.adb   modified

```

The following is an example of use of the dependency list. Note the use of the `-s` switch which gives a straight list of source files. This can be useful for building specialized scripts.

```

$ gnatls -d demo2.o
./demo2.o  demo2      OK demo2.adb
              OK gen_list.ads
              OK gen_list.adb
              OK instr.ads
              OK instr-child.ads

$ gnatls -d -s -a demo1.o
demo1.adb
/home/comar/local/adainclude/ada.ads
/home/comar/local/adainclude/a-finali.ads
/home/comar/local/adainclude/a-filico.ads
/home/comar/local/adainclude/a-stream.ads
/home/comar/local/adainclude/a-tags.ads
gen_list.ads
gen_list.adb
/home/comar/local/adainclude/gnat.ads
/home/comar/local/adainclude/g-io.ads
instr.ads
/home/comar/local/adainclude/system.ads
/home/comar/local/adainclude/s-exctab.ads
/home/comar/local/adainclude/s-finimp.ads
/home/comar/local/adainclude/s-finroo.ads
/home/comar/local/adainclude/s-secsta.ads
/home/comar/local/adainclude/s-stalib.ads
/home/comar/local/adainclude/s-stoele.ads
/home/comar/local/adainclude/s-stratt.ads
/home/comar/local/adainclude/s-tasoli.ads
/home/comar/local/adainclude/s-unstyp.ads
/home/comar/local/adainclude/unchconv.ads

```

GNAT AND PROGRAM EXECUTION

This chapter covers several topics:

- [Running and Debugging Ada Programs](#)
- [Profiling](#)
- [Improving Performance](#)
- [Overflow Check Handling in GNAT](#)
- [Performing Dimensionality Analysis in GNAT](#)
- [Stack Related Facilities](#)
- [Memory Management Issues](#)

6.1 Running and Debugging Ada Programs

This section discusses how to debug Ada programs.

An incorrect Ada program may be handled in three ways by the GNAT compiler:

- The illegality may be a violation of the static semantics of Ada. In that case GNAT diagnoses the constructs in the program that are illegal. It is then a straightforward matter for the user to modify those parts of the program.
- The illegality may be a violation of the dynamic semantics of Ada. In that case the program compiles and executes, but may generate incorrect results, or may terminate abnormally with some exception.
- When presented with a program that contains convoluted errors, GNAT itself may terminate abnormally without providing full diagnostics on the incorrect user program.

6.1.1 The GNAT Debugger GDB

GDB is a general purpose, platform-independent debugger that can be used to debug mixed-language programs compiled with `gcc`, and in particular is capable of debugging Ada programs compiled with GNAT. The latest versions of GDB are Ada-aware and can handle complex Ada data structures.

See *Debugging with GDB*, for full details on the usage of GDB, including a section on its usage on programs. This manual should be consulted for full details. The section that follows is a brief introduction to the philosophy and use of GDB.

When GNAT programs are compiled, the compiler optionally writes debugging information into the generated object file, including information on line numbers, and on declared types and variables. This information is separate from the generated code. It makes the object files considerably larger, but it does not add to the size of the actual executable that will be loaded into memory, and has no impact on run-time performance. The generation of debug information is triggered by the use of the `-g` switch in the `gcc` or `gnatmake` command used to carry out the compilations. It is important to emphasize that the use of these options does not change the generated code.

The debugging information is written in standard system formats that are used by many tools, including debuggers and profilers. The format of the information is typically designed to describe C types and semantics, but GNAT implements a translation scheme which allows full details about Ada types and variables to be encoded into these standard C formats. Details of this encoding scheme may be found in the file `exp_dbug.ads` in the GNAT source distribution. However, the details of this encoding are, in general, of no interest to a user, since GDB automatically performs the necessary decoding.

When a program is bound and linked, the debugging information is collected from the object files, and stored in the executable image of the program. Again, this process significantly increases the size of the generated executable file, but it does not increase the size of the executable program itself. Furthermore, if this program is run in the normal manner, it runs exactly as if the debug information were not present, and takes no more actual memory.

However, if the program is run under control of GDB, the debugger is activated. The image of the program is loaded, at which point it is ready to run. If a run command is given, then the program will run exactly as it would have if GDB were not present. This is a crucial part of the GDB design philosophy. GDB is entirely non-intrusive until a breakpoint is encountered. If no breakpoint is ever hit, the program will run exactly as it would if no debugger were present. When a breakpoint is hit, GDB accesses the debugging information and can respond to user commands to inspect variables, and more generally to report on the state of execution.

6.1.2 Running GDB

This section describes how to initiate the debugger.

The debugger can be launched from a GNAT Studio menu or directly from the command line. The description below covers the latter use. All the commands shown can be used in the GNAT Studio debug console window, but there are usually more GUI-based ways to achieve the same effect.

The command to run GDB is

```
$ gdb program
```

where `program` is the name of the executable file. This activates the debugger and results in a prompt for debugger commands. The simplest command is simply `run`, which causes the program to run exactly as if the debugger were not present. The following section describes some of the additional commands that can be given to GDB.

6.1.3 Introduction to GDB Commands

GDB contains a large repertoire of commands. See *Debugging with GDB* for extensive documentation on the use of these commands, together with examples of their use. Furthermore, the command `help` invoked from within GDB activates a simple help facility which summarizes the available commands and their options. In this section we summarize a few of the most commonly used commands to give an idea of what GDB is about. You should create a simple program with debugging information and experiment with the use of these GDB commands on the program as you read through the following section.

- **set args *arguments***

The *arguments* list above is a list of arguments to be passed to the program on a subsequent `run` command, just as though the arguments had been entered on a normal invocation of the program. The `set args` command is not needed if the program does not require arguments.

- **run**

The `run` command causes execution of the program to start from the beginning. If the program is already running, that is to say if you are currently positioned at a breakpoint, then a prompt will ask for confirmation that you want to abandon the current execution and restart.

- **breakpoint *location***

The `breakpoint` command sets a breakpoint, that is to say a point at which execution will halt and GDB will await further commands. *location* is either a line number within a file, given in the format `file:linenumber`, or it is the name of a subprogram. If you request that a breakpoint be set on a subprogram that is overloaded, a prompt will ask you to specify on which of those subprograms you want to breakpoint. You can also specify that all of them should be breakpointed. If the program is run and execution encounters the breakpoint, then the program stops and GDB signals that the breakpoint was encountered by printing the line of code before which the program is halted.

- **catch exception *name***

This command causes the program execution to stop whenever *exception name* is raised. If *name* is omitted, then the execution is suspended when any exception is raised.

- **print *expression***

This will print the value of the given expression. Most simple Ada expression formats are properly handled by GDB, so the expression can contain function calls, variables, operators, and attribute references.

- **continue**

Continues execution following a breakpoint, until the next breakpoint or the termination of the program.

- **step**
Executes a single line after a breakpoint. If the next statement is a subprogram call, execution continues into (the first statement of) the called subprogram.
- **next**
Executes a single line. If this line is a subprogram call, executes and returns from the call.
- **list**
Lists a few lines around the current source location. In practice, it is usually more convenient to have a separate edit window open with the relevant source file displayed. Successive applications of this command print subsequent lines. The command can be given an argument which is a line number, in which case it displays a few lines around the specified one.
- **backtrace**
Displays a backtrace of the call chain. This command is typically used after a breakpoint has occurred, to examine the sequence of calls that leads to the current breakpoint. The display includes one line for each activation record (frame) corresponding to an active subprogram.
- **up**
At a breakpoint, **GDB** can display the values of variables local to the current frame. The command **up** can be used to examine the contents of other active frames, by moving the focus up the stack, that is to say from callee to caller, one frame at a time.
- **down**
Moves the focus of **GDB** down from the frame currently being examined to the frame of its callee (the reverse of the previous command),
- **frame *n***
Inspect the frame with the given number. The value 0 denotes the frame of the current breakpoint, that is to say the top of the call stack.
- **kill**
Kills the child process in which the program is running under **GDB**. This may be useful for several purposes:
 - It allows you to recompile and relink your program, since on many systems you cannot regenerate an executable file while it is running in a process.
 - You can run your program outside the debugger, on systems that do not permit executing a program outside **GDB** while breakpoints are set within **GDB**.
 - It allows you to debug a core dump rather than a running process.

The above list is a very short introduction to the commands that **GDB** provides. Important additional capabilities, including conditional breakpoints, the ability to execute command sequences on a breakpoint, the ability to debug at the machine instruction level and many other features are described in detail in *Debugging with GDB*. Note that most commands can be abbreviated (for example, **c** for **continue**, **bt** for **backtrace**).

6.1.4 Using Ada Expressions

GDB supports a fairly large subset of Ada expression syntax, with some extensions. The philosophy behind the design of this subset is

- That GDB should provide basic literals and access to operations for arithmetic, dereferencing, field selection, indexing, and subprogram calls, leaving more sophisticated computations to subprograms written into the program (which therefore may be called from GDB).
- That type safety and strict adherence to Ada language restrictions are not particularly relevant in a debugging context.
- That brevity is important to the GDB user.

Thus, for brevity, the debugger acts as if there were implicit `with` and `use` clauses in effect for all user-written packages, thus making it unnecessary to fully qualify most names with their packages, regardless of context. Where this causes ambiguity, GDB asks the user's intent.

For details on the supported Ada syntax, see *Debugging with GDB*.

6.1.5 Calling User-Defined Subprograms

An important capability of GDB is the ability to call user-defined subprograms while debugging. This is achieved simply by entering a subprogram call statement in the form:

```
call subprogram-name (parameters)
```

The keyword `call` can be omitted in the normal case where the `subprogram-name` does not coincide with any of the predefined GDB commands.

The effect is to invoke the given subprogram, passing it the list of parameters that is supplied. The parameters can be expressions and can include variables from the program being debugged. The subprogram must be defined at the library level within your program, and GDB will call the subprogram within the environment of your program execution (which means that the subprogram is free to access or even modify variables within your program).

The most important use of this facility is in allowing the inclusion of debugging routines that are tailored to particular data structures in your program. Such debugging routines can be written to provide a suitably high-level description of an abstract type, rather than a low-level dump of its physical layout. After all, the standard GDB `print` command only knows the physical layout of your types, not their abstract meaning. Debugging routines can provide information at the desired semantic level and are thus enormously useful.

For example, when debugging GNAT itself, it is crucial to have access to the contents of the tree nodes used to represent the program internally. But tree nodes are represented simply by an integer value (which in turn is an index into a table of nodes). Using the `print` command on a tree node would simply print this integer value, which is not very useful. But the `PN` routine (defined in file `treepr.adb` in the GNAT sources) takes a tree node as input, and displays a useful high level representation of the tree node, which includes the syntactic category of the node, its position in the source, the integers that denote descendant nodes and parent node, as well as varied semantic

information. To study this example in more detail, you might want to look at the body of the PN procedure in the stated file.

Another useful application of this capability is to deal with situations of complex data which are not handled suitably by GDB. For example, if you specify Convention Fortran for a multi-dimensional array, GDB does not know that the ordering of array elements has been switched and will not properly address the array elements. In such a case, instead of trying to print the elements directly from GDB, you can write a callable procedure that prints the elements in the desired format.

6.1.6 Using the *next* Command in a Function

When you use the `next` command in a function, the current source location will advance to the next statement as usual. A special case arises in the case of a `return` statement.

Part of the code for a return statement is the 'epilogue' of the function. This is the code that returns to the caller. There is only one copy of this epilogue code, and it is typically associated with the last return statement in the function if there is more than one return. In some implementations, this epilogue is associated with the first statement of the function.

The result is that if you use the `next` command from a return statement that is not the last return statement of the function you may see a strange apparent jump to the last return statement or to the start of the function. You should simply ignore this odd jump. The value returned is always that from the first return statement that was stepped through.

6.1.7 Stopping When Ada Exceptions Are Raised

You can set catchpoints that stop the program execution when your program raises selected exceptions.

- **catch exception**
Set a catchpoint that stops execution whenever (any task in the) program raises any exception.
- **catch exception *name***
Set a catchpoint that stops execution whenever (any task in the) program raises the exception *name*.
- **catch exception unhandled**
Set a catchpoint that stops executing whenever (any task in the) program raises an exception for which there is no handler.
- **info exceptions, info exceptions *regexp***
The `info exceptions` command permits the user to examine all defined exceptions within Ada programs. With a regular expression, *regexp*, as argument, prints out only those exceptions whose name matches *regexp*.

6.1.8 Ada Tasks

GDB allows the following task-related commands:

- **info tasks**

This command shows a list of current Ada tasks, as in the following example:

```
(gdb) info tasks
  ID      TID P-ID  Thread Pri State           Name
  1    8088000  0   807e000  15 Child Activation Wait main_task
  2    80a4000  1   80ae000  15 Accept/Select Wait    b
  3    809a800  1   80a4800  15 Child Activation Wait    a
* 4    80ae800  3   80b8000  15 Running                c
```

In this listing, the asterisk before the first task indicates it to be the currently running task. The first column lists the task ID that is used to refer to tasks in the following commands.

- **break *linespec* task *taskid*, break *linespec* task *taskid* if ...**

These commands are like the `break ... thread ...`. *linespec* specifies source lines.

Use the qualifier `task taskid` with a breakpoint command to specify that you only want GDB to stop the program when a particular Ada task reaches this breakpoint. *taskid* is one of the numeric task identifiers assigned by GDB, shown in the first column of the `info tasks` display.

If you do not specify `task taskid` when you set a breakpoint, the breakpoint applies to *all* tasks of your program.

You can use the `task` qualifier on conditional breakpoints as well; in this case, place `task taskid` before the breakpoint condition (before the `if`).

- **task *taskno***

This command allows switching to the task referred by *taskno*. In particular, this allows browsing of the backtrace of the specified task. It is advisable to switch back to the original task before continuing execution otherwise the scheduling of the program may be perturbed.

For more detailed information on the tasking support, see *Debugging with GDB*.

6.1.9 Debugging Generic Units

GNAT always uses code expansion for generic instantiation. This means that each time an instantiation occurs, a complete copy of the original code is made, with appropriate substitutions of formals by actuals.

It is not possible to refer to the original generic entities in GDB, but it is always possible to debug a particular instance of a generic, by using the appropriate expanded names. For example, if we have

```
procedure g is

  generic package k is
    procedure kp (v1 : in out integer);
  end k;

  package body k is
    procedure kp (v1 : in out integer) is
      begin
        v1 := v1 + 1;
      end kp;
    end k;

  package k1 is new k;
  package k2 is new k;

  var : integer := 1;

begin
  k1.kp (var);
  k2.kp (var);
  k1.kp (var);
  k2.kp (var);
end;
```

Then to break on a call to procedure `kp` in the `k2` instance, simply use the command:

```
(gdb) break g.k2.kp
```

When the breakpoint occurs, you can step through the code of the instance in the normal manner and examine the values of local variables, as for other units.

6.1.10 Remote Debugging with gdbserver

On platforms where `gdbserver` is supported, it is possible to use this tool to debug your application remotely. This can be useful in situations where the program needs to be run on a target host that is different from the host used for development, particularly when the target has a limited amount of resources (either CPU and/or memory).

To do so, start your program using `gdbserver` on the target machine. `gdbserver` then automatically suspends the execution of your program at its entry point, waiting for a debugger to connect to it. The following commands starts an application and tells `gdbserver` to wait for a connection with the debugger on localhost port 4444.

```
$ gdbserver localhost:4444 program
Process program created; pid = 5685
Listening on port 4444
```

Once `gdbserver` has started listening, we can tell the debugger to establish a connection with this `gdbserver`, and then start the same debugging session as if the program was being debugged on the same host, directly under the control of GDB.

```

$ gdb program
(gdb) target remote targethost:4444
Remote debugging using targethost:4444
0x00007f29936d0af0 in ?? () from /lib64/ld-linux-x86-64.so.
(gdb) b foo.adb:3
Breakpoint 1 at 0x401f0c: file foo.adb, line 3.
(gdb) continue
Continuing.

Breakpoint 1, foo () at foo.adb:4
4      end foo;

```

It is also possible to use `gdbserver` to attach to an already running program, in which case the execution of that program is simply suspended until the connection between the debugger and `gdbserver` is established.

For more information on how to use `gdbserver`, see the *Using the gdbserver Program* section in *Debugging with GDB*. GNAT provides support for `gdbserver` on `x86-linux`, `x86-windows` and `x86_64-linux`.

6.1.11 GNAT Abnormal Termination or Failure to Terminate

When presented with programs that contain serious errors in syntax or semantics, GNAT may on rare occasions experience problems in operation, such as aborting with a segmentation fault or illegal memory access, raising an internal exception, terminating abnormally, or failing to terminate at all. In such cases, you can activate various features of GNAT that can help you pinpoint the construct in your program that is the likely source of the problem.

The following strategies are presented in increasing order of difficulty, corresponding to your experience in using GNAT and your familiarity with compiler internals.

- Run `gcc` with the `-gnatf`. This first switch causes all errors on a given line to be reported. In its absence, only the first error on a line is displayed.

The `-gnatd0` switch causes errors to be displayed as soon as they are encountered, rather than after compilation is terminated. If GNAT terminates prematurely or goes into an infinite loop, the last error message displayed may help to pinpoint the culprit.

- Run `gcc` with the `-v` (verbose) switch. In this mode, `gcc` produces ongoing information about the progress of the compilation and provides the name of each procedure as code is generated. This switch allows you to find which Ada procedure was being compiled when it encountered a code generation problem.
- Run `gcc` with the `-gnatdc` switch. This is a GNAT specific switch that does for the front-end what `-v` does for the back end. The system prints the name of each unit, either a compilation unit or nested unit, as it is being analyzed.
- Finally, you can start `gdb` directly on the `gnat1` executable. `gnat1` is the front-end of GNAT, and can be run independently (normally it is just called from `gcc`). You can use `gdb` on `gnat1` as you would on a C program (but [The GNAT Debugger GDB](#) for caveats). The `where` command is the first line of attack; the variable `lineno` (seen by `print lineno`), used by the

second phase of `gnat1` and by the `gcc` backend, indicates the source line at which the execution stopped, and `input_file` name indicates the name of the source file.

6.1.12 Naming Conventions for GNAT Source Files

In order to examine the workings of the GNAT system, the following brief description of its organization may be helpful:

- Files with prefix `sc` contain the lexical scanner.
- All files prefixed with `par` are components of the parser. The numbers correspond to chapters of the Ada Reference Manual. For example, parsing of select statements can be found in `par-ch9.adb`.
- All files prefixed with `sem` perform semantic analysis. The numbers correspond to chapters of the Ada standard. For example, all issues involving context clauses can be found in `sem_ch10.adb`. In addition, some features of the language require sufficient special processing to justify their own semantic files: `sem_aggr` for aggregates, `sem_disp` for dynamic dispatching, etc.
- All files prefixed with `exp` perform normalization and expansion of the intermediate representation (abstract syntax tree, or AST). these files use the same numbering scheme as the parser and semantics files. For example, the construction of record initialization procedures is done in `exp_ch3.adb`.
- The files prefixed with `bind` implement the binder, which verifies the consistency of the compilation, determines an order of elaboration, and generates the bind file.
- The files `atree.ads` and `atree.adb` detail the low-level data structures used by the front-end.
- The files `sinfo.ads` and `sinfo.adb` detail the structure of the abstract syntax tree as produced by the parser.
- The files `einfo.ads` and `einfo.adb` detail the attributes of all entities, computed during semantic analysis.
- Library management issues are dealt with in files with prefix `lib`.
- Ada files with the prefix `a-` are children of `Ada`, as defined in Annex A.
- Files with prefix `i-` are children of `Interfaces`, as defined in Annex B.
- Files with prefix `s-` are children of `System`. This includes both language-defined children and GNAT run-time routines.
- Files with prefix `g-` are children of `GNAT`. These are useful general-purpose packages, fully documented in their specs. All the other `.c` files are modifications of common `gcc` files.

6.1.13 Getting Internal Debugging Information

Most compilers have internal debugging switches and modes. GNAT does also, except GNAT internal debugging switches and modes are not secret. A summary and full description of all the compiler and binder debug flags are in the file `debug.adb`. You must obtain the sources of the compiler to see the full detailed effects of these flags.

The switches that print the source of the program (reconstructed from the internal tree) are of general interest for user programs, as are the options to print the full internal tree, and the entity table (the symbol table information). The reconstructed source provides a readable version of the program after the front-end has completed analysis and expansion, and is useful when studying the performance of specific constructs. For example, constraint checks are indicated, complex aggregates are replaced with loops and assignments, and tasking primitives are replaced with run-time calls.

6.1.14 Stack Traceback

Traceback is a mechanism to display the sequence of subprogram calls that leads to a specified execution point in a program. Often (but not always) the execution point is an instruction at which an exception has been raised. This mechanism is also known as *stack unwinding* because it obtains its information by scanning the run-time stack and recovering the activation records of all active subprograms. Stack unwinding is one of the most important tools for program debugging.

The first entry stored in traceback corresponds to the deepest calling level, that is to say the subprogram currently executing the instruction from which we want to obtain the traceback.

Note that there is no runtime performance penalty when stack traceback is enabled, and no exception is raised during program execution.

Non-Symbolic Traceback

Note: this feature is not supported on all platforms. See `GNAT.Traceback` spec in `g-traceb.ads` for a complete list of supported platforms.

Tracebacks From an Unhandled Exception

A runtime non-symbolic traceback is a list of addresses of call instructions. To enable this feature you must use the `-E gnatbind` option. With this option a stack traceback is stored as part of exception information.

You can translate this information using the `addr2line` tool, provided that the program is compiled with debugging options (see [Compiler Switches](#)) and linked at a fixed position with `-no-pie`.

Here is a simple example with `gnatmake`:

```
procedure STB is
  procedure P1 is
```

(continues on next page)

(continued from previous page)

```

begin
  raise Constraint_Error;
end P1;

procedure P2 is
begin
  P1;
end P2;

begin
  P2;
end STB;

```

```

$ gnatmake stb -g -bargs -E -largs -no-pie
$ stb

Execution of stb terminated by unhandled exception
raised CONSTRAINT_ERROR : stb.adb:5 explicit raise
Load address: 0x400000
Call stack traceback locations:
0x401373 0x40138b 0x40139c 0x401335 0x4011c4 0x4011f1 0x77e892a4

```

As we see the traceback lists a sequence of addresses for the unhandled exception `CONSTRAINT_ERROR` raised in procedure `P1`. It is easy to guess that this exception come from procedure `P1`. To translate these addresses into the source lines where the calls appear, the `addr2line` tool needs to be invoked like this:

```

$ addr2line -e stb 0x401373 0x40138b 0x40139c 0x401335 0x4011c4
0x4011f1 0x77e892a4

d:/stb/stb.adb:5
d:/stb/stb.adb:10
d:/stb/stb.adb:14
d:/stb/b~stb.adb:197
crtexe.c:?
crtexe.c:?
??:0

```

The `addr2line` tool has several other useful options:

<code>-a --addresses</code>	to show the addresses alongside the line numbers
<code>-f --functions</code>	to get the function name corresponding to a location
<code>-p --pretty-print</code>	to print all the information on a single line
<code>--demangle=gnat</code>	to use the GNAT decoding mode for the function names

```

$ addr2line -e stb -a -f -p --demangle=gnat 0x401373 0x40138b
0x40139c 0x401335 0x4011c4 0x4011f1 0x77e892a4

0x00401373: stb.p1 at d:/stb/stb.adb:5

```

(continues on next page)

(continued from previous page)

```

0x0040138B: stb.p2 at d:/stb/stb.adb:10
0x0040139C: stb at d:/stb/stb.adb:14
0x00401335: main at d:/stb/b~stb.adb:197
0x004011c4: ?? at crtexe.c:?
0x004011f1: ?? at crtexe.c:?
0x77e892a4: ?? ??:0

```

From this traceback we can see that the exception was raised in `stb.adb` at line 5, which was reached from a procedure call in `stb.adb` at line 10, and so on. The `b~std.adb` is the binder file, which contains the call to the main program. [Running gnatbind](#). The remaining entries are assorted runtime routines and the output will vary from platform to platform.

It is also possible to use GDB with these traceback addresses to debug the program. For example, we can break at a given code location, as reported in the stack traceback:

```

$ gdb -nw stb

(gdb) break *0x401373
Breakpoint 1 at 0x401373: file stb.adb, line 5.

```

It is important to note that the stack traceback addresses do not change when debug information is included. This is particularly useful because it makes it possible to release software without debug information (to minimize object size), get a field report that includes a stack traceback whenever an internal bug occurs, and then be able to retrieve the sequence of calls with the same program compiled with debug information.

However the `addr2line` tool does not work with Position-Independent Code (PIC), the historical example being Windows DLLs, which nowadays encompasses Position-Independent Executables (PIE) on recent Windows versions.

In order to translate addresses into the source lines with Position-Independent Executables on recent Windows versions, in other words without using the switch `-no-pie` during linking, you need to use the `gnatsymbolize` tool with `--load` instead of the `addr2line` tool. The main difference is that you need to copy the Load Address output in the traceback ahead of the sequence of addresses. And the default mode of `gnatsymbolize` is equivalent to that of `addr2line` with the above switches, so none of them is needed:

```

$ gnatmake stb -g -bargs -E
$ stb

Execution of stb terminated by unhandled exception
raised CONSTRAINT_ERROR : stb.adb:5 explicit raise
Load address: 0x400000
Call stack traceback locations:
0x401373 0x40138b 0x40139c 0x401335 0x4011c4 0x4011f1 0x77e892a4

$ gnatsymbolize --load stb 0x400000 0x401373 0x40138b 0x40139c 0x401335
0x4011c4 0x4011f1 0x77e892a4

0x00401373 Stb.P1 at stb.adb:5

```

(continues on next page)

(continued from previous page)

```
0x0040138B Stb.P2 at stb.adb:10
0x0040139C Stb at stb.adb:14
0x00401335 Main at b~stb.adb:197
0x004011c4 __tmainCRTStartup at ???
0x004011f1 mainCRTStartup at ???
0x77e892a4 ??? at ???
```

Tracebacks From Exception Occurrences

Non-symbolic tracebacks are obtained by using the `-E` binder argument. The stack traceback is attached to the exception information string, and can be retrieved in an exception handler within the Ada program, by means of the Ada facilities defined in `Ada.Exceptions`. Here is a simple example:

```
with Ada.Text_IO;
with Ada.Exceptions;

procedure STB is

  use Ada;
  use Ada.Exceptions;

  procedure P1 is
    K : Positive := 1;
  begin
    K := K - 1;
  exception
    when E : others =>
      Text_IO.Put_Line (Exception_Information (E));
  end P1;

  procedure P2 is
  begin
    P1;
  end P2;

begin
  P2;
end STB;
```

This program will output:

```
$ stb

raised CONSTRAINT_ERROR : stb.adb:12 range check failed
Load address: 0x400000
Call stack traceback locations:
0x4015e4 0x401633 0x401644 0x401461 0x4011c4 0x4011f1 0x77e892a4
```

Tracebacks From Anywhere in a Program

It is also possible to retrieve a stack traceback from anywhere in a program. For this you need to use the GNAT.Traceback API. This package includes a procedure called `Call_Chain` that computes a complete stack traceback, as well as useful display procedures described below. It is not necessary to use the `-E gnatbind` option in this case, because the stack traceback mechanism is invoked explicitly.

In the following example we compute a traceback at a specific location in the program, and we display it using `GNAT.Debug_Utilities.Image` to convert addresses to strings:

```
with Ada.Text_IO;
with GNAT.Traceback;
with GNAT.Debug_Utilities;

procedure STB is

  use Ada;
  use GNAT;
  use GNAT.Traceback;

  procedure P1 is
    TB : Tracebacks_Array (1 .. 10);
    -- We are asking for a maximum of 10 stack frames.
    Len : Natural;
    -- Len will receive the actual number of stack frames returned.
  begin
    Call_Chain (TB, Len);

    Text_IO.Put ("In STB.P1 : ");

    for K in 1 .. Len loop
      Text_IO.Put (Debug_Utilities.Image (TB (K)));
      Text_IO.Put (' ');
    end loop;

    Text_IO.New_Line;
  end P1;

  procedure P2 is
  begin
    P1;
  end P2;

begin
  P2;
end STB;
```

```
$ gnatmake stb -g
$ stb
```

```
In STB.P1 : 16#0040_F1E4# 16#0040_14F2# 16#0040_170B# 16#0040_171C#
```

(continues on next page)

(continued from previous page)

```
16#0040_1461# 16#0040_11C4# 16#0040_11F1# 16#77E8_92A4#
```

You can then get further information by invoking the `addr2line` tool or the `gnatsymbolize` tool as described earlier (note that the hexadecimal addresses need to be specified in C format, with a leading '0x').

Symbolic Traceback

A symbolic traceback is a stack traceback in which procedure names are associated with each code location.

Note that this feature is not supported on all platforms. See `GNAT.Traceback.Symbolic` spec in `g-trasym.ads` for a complete list of currently supported platforms.

Note that the symbolic traceback requires that the program be compiled with debug information. If it is not compiled with debug information only the non-symbolic information will be valid.

Tracebacks From Exception Occurrences

Here is an example:

```
with Ada.Text_IO;
with GNAT.Traceback.Symbolic;

procedure STB is

  procedure P1 is
  begin
    raise Constraint_Error;
  end P1;

  procedure P2 is
  begin
    P1;
  end P2;

  procedure P3 is
  begin
    P2;
  end P3;

begin
  P3;
exception
  when E : others =>
    Ada.Text_IO.Put_Line (GNAT.Traceback.Symbolic.Symbolic_Traceback (E));
end STB;
```

```

$ gnatmake -g .\stb -bargs -E
$ stb

0040149F in stb.p1 at stb.adb:8
004014B7 in stb.p2 at stb.adb:13
004014CF in stb.p3 at stb.adb:18
004015DD in ada.stb at stb.adb:22
00401461 in main at b~stb.adb:168
004011C4 in __mingw_CRTStartup at crt1.c:200
004011F1 in mainCRTStartup at crt1.c:222
77E892A4 in ?? at ??:0

```

In the above example the `.\` syntax in the `gnatmake` command is currently required by `addr2line` for files that are in the current working directory. Moreover, the exact sequence of linker options may vary from platform to platform. The above `-bargs` section is for Windows platforms. By contrast, under Unix there is no need for the `-bargs` section. Differences across platforms are due to details of linker implementation.

Tracebacks From Anywhere in a Program

It is possible to get a symbolic stack traceback from anywhere in a program, just as for non-symbolic tracebacks. The first step is to obtain a non-symbolic traceback, and then call `Symbolic_Backtrace` to compute the symbolic information. Here is an example:

```

with Ada.Text_IO;
with GNAT.Traceback;
with GNAT.Traceback.Symbolic;

procedure STB is

  use Ada;
  use GNAT.Traceback;
  use GNAT.Traceback.Symbolic;

  procedure P1 is
    TB : Tracebacks_Array (1 .. 10);
    -- We are asking for a maximum of 10 stack frames.
    Len : Natural;
    -- Len will receive the actual number of stack frames returned.
  begin
    Call_Chain (TB, Len);
    Text_IO.Put_Line (Symbolic_Backtrace (TB (1 .. Len)));
  end P1;

  procedure P2 is
  begin
    P1;
  end P2;

begin

```

(continues on next page)

(continued from previous page)

```
P2;  
end STB;
```

Automatic Symbolic Tracebacks

Symbolic tracebacks may also be enabled by using the `-Es` switch to `gnatbind` (as in `gprbuild -g ... -bargs -Es`). This will cause the `Exception_Information` to contain a symbolic traceback, which will also be printed if an unhandled exception terminates the program.

6.1.15 Pretty-Printers for the GNAT runtime

As discussed in *Calling User-Defined Subprograms*, GDB's `print` command only knows about the physical layout of program data structures and therefore normally displays only low-level dumps, which are often hard to understand.

An example of this is when trying to display the contents of an Ada standard container, such as `Ada.Containers.Ordered_Maps.Map`:

```
with Ada.Containers.Ordered_Maps;  
  
procedure PP is  
  package Int_To_Nat is  
    new Ada.Containers.Ordered_Maps (Integer, Natural);  
  
    Map : Int_To_Nat.Map;  
begin  
  Map.Insert (1, 10);  
  Map.Insert (2, 20);  
  Map.Insert (3, 30);  
  
  Map.Clear; -- BREAK HERE  
end PP;
```

When this program is built with debugging information and run under GDB up to the `Map.Clear` statement, trying to print `Map` will yield information that is only relevant to the developers of our standard containers:

```
(gdb) print map  
$1 = (  
  tree => (  
    first => 0x64e010,  
    last => 0x64e070,  
    root => 0x64e040,  
    length => 3,  
    tc => (  
      busy => 0,  
      lock => 0  
    )  
  )  
)
```

(continues on next page)

(continued from previous page)

```
)  
)
```

Fortunately, GDB has a feature called [pretty-printers](#), which allows customizing how GDB displays data structures. The GDB shipped with GNAT embeds such pretty-printers for the most common containers in the standard library. To enable them, either run the following command manually under GDB or add it to your `.gdbinit` file:

```
python import gnatdbg; gnatdbg.setup()
```

Once this is done, GDB's `print` command will automatically use these pretty-printers when appropriate. Using the previous example:

```
(gdb) print map  
$1 = pp.int_to_nat.map of length 3 = {  
  [1] = 10,  
  [2] = 20,  
  [3] = 30  
}
```

Pretty-printers are invoked each time GDB tries to display a value, including when displaying the arguments of a called subprogram (in GDB's `backtrace` command) or when printing the value returned by a function (in GDB's `finish` command).

To display a value without involving pretty-printers, `print` can be invoked with its `/r` option:

```
(gdb) print/r map  
$1 = (  
  tree => (...
```

Finer control of pretty-printers is also possible: see [GDB's online documentation](#) for more information.

6.2 Profiling

This section describes how to use the `gprof` profiler tool on Ada programs.

6.2.1 Profiling an Ada Program with `gprof`

This section is not meant to be an exhaustive documentation of `gprof`. Full documentation for it can be found in the *GNU Profiler User's Guide* documentation that is part of this GNAT distribution.

Profiling a program helps determine the parts of a program that are executed most often, and are therefore the most time-consuming.

`gprof` is the standard GNU profiling tool; it has been enhanced to better handle Ada programs and multitasking. It is currently supported on the following platforms

- Linux x86/x86_64
- Windows x86/x86_64 (without PIE support)

In order to profile a program using `gprof`, several steps are needed:

1. Instrument the code, which requires a full recompilation of the project with the proper switches.
2. Execute the program under the analysis conditions, i.e. with the desired input.
3. Analyze the results using the `gprof` tool.

The following sections detail the different steps, and indicate how to interpret the results.

Compilation for profiling

In order to profile a program the first step is to tell the compiler to generate the necessary profiling information. The compiler switch to be used is `-pg`, which must be added to other compilation switches. This switch needs to be specified both during compilation and link stages, and can be specified once when using `gnatmake`:

```
$ gnatmake -f -pg -P my_project
```

Note that only the objects that were compiled with the `-pg` switch will be profiled; if you need to profile your whole project, use the `-f gnatmake` switch to force full recompilation.

Note that on Windows, `gprof` does not support PIE. The `-no-pie` switch should be added to the linker flags to disable this feature.

Program execution

Once the program has been compiled for profiling, you can run it as usual.

The only constraint imposed by profiling is that the program must terminate normally. An interrupted program (via a `Ctrl-C`, `kill`, etc.) will not be properly analyzed.

Once the program completes execution, a data file called `gmon.out` is generated in the directory where the program was launched from. If this file already exists, it will be overwritten.

Running gprof

The `gprof` tool is called as follow:

```
$ gprof my_prog gmon.out
```

or simply:

```
$ gprof my_prog
```

The complete form of the `gprof` command line is the following:

```
$ gprof [switches] [executable [data-file]]
```

`gprof` supports numerous switches. The order of these switch does not matter. The full list of options can be found in the GNU Profiler User's Guide documentation that comes with this documentation.

The following is the subset of those switches that is most relevant:

--demangle[=*style*], --no-demangle

These options control whether symbol names should be demangled when printing output. The default is to demangle C++ symbols. The `--no-demangle` option may be used to turn off demangling. Different compilers have different mangling styles. The optional demangling style argument can be used to choose an appropriate demangling style for your compiler, in particular Ada symbols generated by GNAT can be demangled using `--demangle=gnat`.

-e *function_name*

The `-e function_name` option tells `gprof` not to print information about the function `function_name` (and its children...) in the call graph. The function will still be listed as a child of any functions that call it, but its index number will be shown as [not printed]. More than one `-e` option may be given; only one `function_name` may be indicated with each `-e` option.

-E *function_name*

The `-E function_name` option works like the `-e` option, but execution time spent in the function (and children who were not called from anywhere else), will not be used to compute the percentages-of-time for the call graph. More than one `-E` option may be given; only one `function_name` may be indicated with each `-E`` option.

-f *function_name*

The `-f function_name` option causes `gprof` to limit the call graph to the function `function_name` and its children (and their children...). More than one `-f` option may be given; only one `function_name` may be indicated with each `-f` option.

-F *function_name*

The `-F function_name` option works like the `-f` option, but only time spent in the function and its children (and their children...) will be used to determine total-time and percentages-of-time for the call graph. More than one `-F` option may be given; only one `function_name` may be indicated with each `-F` option. The `-F` option overrides the `-E` option.

Interpretation of profiling results

The results of the profiling analysis are represented by two arrays: the 'flat profile' and the 'call graph'. Full documentation of those outputs can be found in the GNU Profiler User's Guide.

The flat profile shows the time spent in each function of the program, and how many time it has been called. This allows you to locate easily the most time-consuming functions.

The call graph shows, for each subprogram, the subprograms that call it, and the subprograms that it calls. It also provides an estimate of the time spent in each of those callers/called subprograms.

6.3 Improving Performance

This section presents several topics related to program performance. It first describes some of the tradeoffs that need to be considered and some of the techniques for making your program run faster.

It then documents the unused subprogram/data elimination feature, which can reduce the size of program executables.

6.3.1 Performance Considerations

The GNAT system provides a number of options that allow a trade-off between

- performance of the generated code
- speed of compilation
- minimization of dependences and recompilation
- the degree of run-time checking.

The defaults (if no options are selected) aim at improving the speed of compilation and minimizing dependences, at the expense of performance of the generated code:

- no optimization
- no inlining of subprogram calls
- all run-time checks enabled except overflow and elaboration checks

These options are suitable for most program development purposes. This section describes how you can modify these choices, and also provides some guidelines on debugging optimized code.

Controlling Run-Time Checks

By default, GNAT generates all run-time checks, except stack overflow checks, and checks for access before elaboration on subprogram calls. The latter are not required in default mode, because all necessary checking is done at compile time.

The `gnat` switch, `-gnatp` allows this default to be modified. See [Run-Time Checks](#).

Our experience is that the default is suitable for most development purposes.

Elaboration checks are off by default, and also not needed by default, since GNAT uses a static elaboration analysis approach that avoids the need for run-time checking. This manual contains a full chapter discussing the issue of elaboration checks, and if the default is not satisfactory for your use, you should read this chapter.

For validity checks, the minimal checks required by the Ada Reference Manual (for case statements and assignments to array elements) are on by default. These can be suppressed by use of the `-gnatVn` switch. Note that in Ada 83, there were no validity checks, so if the Ada 83 mode is acceptable (or when comparing GNAT performance with an Ada 83 compiler), it may be reasonable to routinely use `-gnatVn`. Validity checks are also suppressed entirely if `-gnatp` is used.

Note that the setting of the switches controls the default setting of the checks. They may be modified using either `pragma Suppress` (to remove checks) or `pragma Unsuppress` (to add back suppressed checks) in the program source.

Use of Restrictions

The use of `pragma Restrictions` allows you to control which features are permitted in your program. Apart from the obvious point that if you avoid relatively expensive features like finalization (enforceable by the use of `pragma Restrictions (No_Finalization)`), the use of this `pragma` does not affect the generated code in most cases.

One notable exception to this rule is that the possibility of task abort results in some distributed overhead, particularly if finalization or exception handlers are used. The reason is that certain sections of code have to be marked as non-abortable.

If you use neither the `abort` statement, nor asynchronous transfer of control (`select ... then abort`), then this distributed overhead is removed, which may have a general positive effect in improving overall performance. Especially code involving frequent use of tasking constructs and controlled types will show much improved performance. The relevant restriction pragmas are

```
pragma Restrictions (No_Abort_Statements);
pragma Restrictions (Max_Asynchronous_Select_Nesting => 0);
```

It is recommended that these restriction pragmas be used if possible. Note that this also means that you can write code without worrying about the possibility of an immediate abort at any point.

Optimization Levels

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

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

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

The default is optimization off. This results in the fastest compile times, but GNAT makes absolutely no attempt to optimize, and the generated programs are considerably larger and slower than when optimization is enabled. You can use the `-O` switch (the permitted forms are `-O0`, `-O1`, `-O2`, `-O3`, and `-Os`) to `gcc` to control the optimization level:

- **-O0**
No optimization (the default); generates unoptimized code but has the fastest compilation time.

Note that many other compilers do substantial optimization even if ‘no optimization’ is specified. With `gcc`, it is very unusual to use `-O0` for production if execution time is of any concern, since `-O0` means (almost) no optimization. This difference between `gcc` and other compilers should be kept in mind when doing performance comparisons.

- **-O1**
Moderate optimization; optimizes reasonably well but does not degrade compilation time significantly.
- **-O2**
Full optimization; generates highly optimized code and has the slowest compilation time.
- **-O3**
Full optimization as in `-O2`; also uses more aggressive automatic inlining of subprograms within a unit ([Inlining of Subprograms](#)) and attempts to vectorize loops.
- **-Os**
Optimize space usage (code and data) of resulting program.

Higher optimization levels perform more global transformations on the program and apply more expensive analysis algorithms in order to generate faster and more compact code. The price in compilation time, and the resulting improvement in execution time, both depend on the particular application and the hardware environment. You should experiment to find the best level for your application.

Since the precise set of optimizations done at each level will vary from release to release (and sometime from target to target), it is best to think of the optimization settings in general terms. See the *Options That Control Optimization* section in *Using the GNU Compiler Collection (GCC)* for details about the `-O` settings and a number of `-f` options that individually enable or disable specific optimizations.

Unlike some other compilation systems, `gcc` has been tested extensively at all optimization levels. There are some bugs which appear only with optimization turned on, but there have also been bugs which show up only in *unoptimized* code. Selecting a lower level of optimization does not improve the reliability of the code generator, which in practice is highly reliable at all optimization levels.

Note regarding the use of `-O3`: The use of this optimization level ought not to be automatically preferred over that of level `-O2`, since it often results in larger executables which may run more slowly. See further discussion of this point in [Inlining of Subprograms](#).

Debugging Optimized Code

Although it is possible to do a reasonable amount of debugging at nonzero optimization levels, the higher the level the more likely that source-level constructs will have been eliminated by optimization. For example, if a loop is strength-reduced, the loop control variable may be completely eliminated and thus cannot be displayed in the debugger. This can only happen at `-O2` or `-O3`. Explicit temporary variables that your code might be eliminated at level `-O1` or higher.

The use of the `-g` switch, which is needed for source-level debugging, affects the size of the program executable on disk, and indeed the debugging information can be quite large. However, it has no effect on the generated code (and thus does not degrade performance)

Since the compiler generates debugging tables for a compilation unit before it performs optimizations, the optimizing transformations may invalidate some of the debugging data. You therefore need to anticipate certain anomalous situations that may arise while debugging optimized code. These are the most common cases:

- *The ‘hopping Program Counter’*: Repeated `step` or `next` commands show the PC bouncing back and forth in the code. This may result from any of the following optimizations:
 - *Common subexpression elimination*: using a single instance of code for a quantity that the source computes several times. As a result you may not be able to stop on what looks like a statement.
 - *Invariant code motion*: moving an expression that does not change within a loop, to the beginning of the loop.
 - *Instruction scheduling*: moving instructions so as to overlap loads and stores (typically) with other code, or in general to move computations of values closer to their uses. Often this causes you to pass an assignment statement without the assignment happening and then later bounce back to the statement when the value is actually needed. Placing a breakpoint on a line of code and then stepping over it may, therefore, not always cause all the expected side-effects.
- *The ‘big leap’*: More commonly known as *cross-jumping*, in which two identical pieces of code are merged and the program counter suddenly jumps to a statement that is not supposed to be executed, simply because it (and the code following) translates to the same thing as the code that *was* supposed to be executed. This effect is typically seen in sequences that end in a jump, such as a `goto`, a `return`, or a `break` in a C `switch` statement.
- *The ‘roving variable’*: The symptom is an unexpected value in a variable. There are various reasons for this effect:
 - In a subprogram prologue, a parameter may not yet have been moved to its ‘home’.
 - A variable may be dead, and its register re-used. This is probably the most common cause.
 - As mentioned above, the assignment of a value to a variable may have been moved.
 - A variable may be eliminated entirely by value propagation or other means. In this case, GCC may incorrectly generate debugging information for the variable

In general, when an unexpected value appears for a local variable or parameter you should first ascertain if that value was actually computed by your program, as opposed to being incorrectly reported by the debugger. Record fields or array elements in an object designated by an access value are generally less of a problem, once you have ascertained that the access value is sensible. Typically, this means checking variables in the preceding code and in the calling subprogram to verify that the value observed is explainable from other values (one must apply the procedure recursively to those other values); or re-running the code and stopping a little earlier (perhaps before the call) and stepping to better see how the variable obtained the value in question; or continuing to step *from* the point of the strange value to see if code motion had simply moved the variable’s assignments later.

In light of such anomalies, a recommended technique is to use `-00` early in the software development cycle, when extensive debugging capabilities are most needed, and then move to `-01` and later `-02`

as the debugger becomes less critical. Whether to use the `-g` switch in the release version is a release management issue. Note that if you use `-g` you can then use the `strip` program on the resulting executable, which removes both debugging information and global symbols.

Inlining of Subprograms

A call to a subprogram in the current unit is inlined if all the following conditions are met:

- The optimization level is at least `-O1`.
- The called subprogram is suitable for inlining: It must be small enough and not contain something that `gcc` cannot support in inlined subprograms.
- Any one of the following applies: `pragma Inline` is applied to the subprogram; the subprogram is local to the unit and called once from within it; the subprogram is small and optimization level `-O2` is specified; optimization level `-O3` is specified.

Calls to subprograms in *withed* units are normally not inlined. To achieve actual inlining (that is, replacement of the call by the code in the body of the subprogram), the following conditions must all be true:

- The optimization level is at least `-O1`.
- The called subprogram is suitable for inlining: It must be small enough and not contain something that `gcc` cannot support in inlined subprograms.
- There is a `pragma Inline` for the subprogram.
- The `-gnatn` switch is used on the command line.

Even if all these conditions are met, it may not be possible for the compiler to inline the call, due to the length of the body, or features in the body that make it impossible for the compiler to do the inlining.

Note that specifying the `-gnatn` switch causes additional compilation dependencies. Consider the following:

```
package R is
  procedure Q;
  pragma Inline (Q);
end R;
package body R is
  ...
end R;

with R;
procedure Main is
begin
  ...
  R.Q;
end Main;
```

With the default behavior (no `-gnatn` switch specified), the compilation of the `Main` procedure depends only on its own source, `main.adb`, and the spec of the package in file `r.ads`. This means that editing the body of `R` does not require recompiling `Main`.

On the other hand, the call `R.Q` is not inlined under these circumstances. If the `-gnatn` switch is present when `Main` is compiled, the call will be inlined if the body of `Q` is small enough, but now `Main` depends on the body of `R` in `r.adb` as well as on the spec. This means that if this body is edited, the main program must be recompiled. Note that this extra dependency occurs whether or not the call is in fact inlined by `gcc`.

The use of front end inlining with `-gnatN` generates similar additional dependencies.

Note: The `-fno-inline` switch overrides all other conditions and ensures that no inlining occurs, unless requested with `pragma Inline_Always` for `gcc` back-ends. The extra dependences resulting from `-gnatn` will still be active, even if this switch is used to suppress the resulting inlining actions.

Note: The `-fno-inline-functions` switch can be used to prevent automatic inlining of subprograms if `-O3` is used.

Note: The `-fno-inline-small-functions` switch can be used to prevent automatic inlining of small subprograms if `-O2` is used.

Note: The `-fno-inline-functions-called-once` switch can be used to prevent inlining of subprograms local to the unit and called once from within it if `-O1` is used.

Note regarding the use of `-O3`: `-gnatn` is made up of two sub-switches `-gnatn1` and `-gnatn2` that can be directly specified in lieu of it, `-gnatn` being translated into one of them based on the optimization level. With `-O2` or below, `-gnatn` is equivalent to `-gnatn1` which activates `pragma Inline` with moderate inlining across modules. With `-O3`, `-gnatn` is equivalent to `-gnatn2` which activates `pragma Inline` with full inlining across modules. If you have used `pragma Inline` in appropriate cases, then it is usually much better to use `-O2` and `-gnatn` and avoid the use of `-O3` which has the additional effect of inlining subprograms you did not think should be inlined. We have found that the use of `-O3` may slow down the compilation and increase the code size by performing excessive inlining, leading to increased instruction cache pressure from the increased code size and thus minor performance improvements. So the bottom line here is that you should not automatically assume that `-O3` is better than `-O2`, and indeed you should use `-O3` only if tests show that it actually improves performance for your program.

Floating Point Operations

On almost all targets, GNAT maps `Float` and `Long_Float` to the 32-bit and 64-bit standard IEEE floating-point representations, and operations will use standard IEEE arithmetic as provided by the processor. On most, but not all, architectures, the attribute `Machine_Overflows` is `False` for these types, meaning that the semantics of overflow is implementation-defined. In the case of GNAT, these semantics correspond to the normal IEEE treatment of infinities and NaN (not a number) values. For example, `1.0 / 0.0` yields plus infinity and `0.0 / 0.0` yields a NaN. By avoiding explicit overflow checks, the performance is greatly improved on many targets. However, if required, floating-point overflow can be enabled by the use of the `pragma Check_Float_Overflow`.

Another consideration that applies specifically to x86 32-bit architectures is which form of floating-point arithmetic is used. By default the operations use the old style x86 floating-point, which

implements an 80-bit extended precision form (on these architectures the type `Long_Long_Float` corresponds to that form). In addition, generation of efficient code in this mode means that the extended precision form will be used for intermediate results. This may be helpful in improving the final precision of a complex expression. However it means that the results obtained on the x86 will be different from those on other architectures, and for some algorithms, the extra intermediate precision can be detrimental.

In addition to this old-style floating-point, all modern x86 chips implement an alternative floating-point operation model referred to as SSE2. In this model there is no extended form, and furthermore execution performance is significantly enhanced. To force GNAT to use this more modern form, use both of the switches:

```
-msse2 -mfpmath=sse
```

A unit compiled with these switches will automatically use the more efficient SSE2 instruction set for `Float` and `Long_Float` operations. Note that the ABI has the same form for both floating-point models, so it is permissible to mix units compiled with and without these switches.

Vectorization of loops

You can take advantage of the auto-vectorizer present in the gcc back end to vectorize loops with GNAT. The corresponding command line switch is `-ftree-vectorize` but, as it is enabled by default at `-O3` and other aggressive optimizations helpful for vectorization also are enabled by default at this level, using `-O3` directly is recommended.

You also need to make sure that the target architecture features a supported SIMD instruction set. For example, for the x86 architecture, you should at least specify `-msse2` to get significant vectorization (but you don't need to specify it for x86-64 as it is part of the base 64-bit architecture). Similarly, for the PowerPC architecture, you should specify `-maltivec`.

The preferred loop form for vectorization is the `for` iteration scheme. Loops with a `while` iteration scheme can also be vectorized if they are very simple, but the vectorizer will quickly give up otherwise. With either iteration scheme, the flow of control must be straight, in particular no `exit` statement may appear in the loop body. The loop may however contain a single nested loop, if it can be vectorized when considered alone:

```
A : array (1..4, 1..4) of Long_Float;
S : array (1..4) of Long_Float;

procedure Sum is
begin
  for I in A'Range(1) loop
    for J in A'Range(2) loop
      S (I) := S (I) + A (I, J);
    end loop;
  end loop;
end Sum;
```

The vectorizable operations depend on the targeted SIMD instruction set, but the adding and some of the multiplying operators are generally supported, as well as the logical operators for modular

types. Note that compiling with `-gnatp` might well reveal cases where some checks do thwart vectorization.

Type conversions may also prevent vectorization if they involve semantics that are not directly supported by the code generator or the SIMD instruction set. A typical example is direct conversion from floating-point to integer types. The solution in this case is to use the following idiom:

```
Integer (S'Truncation (F))
```

if `S` is the subtype of floating-point object `F`.

In most cases, the vectorizable loops are loops that iterate over arrays. All kinds of array types are supported, i.e. constrained array types with static bounds:

```
type Array_Type is array (1 .. 4) of Long_Float;
```

constrained array types with dynamic bounds:

```
type Array_Type is array (1 .. Q.N) of Long_Float;
type Array_Type is array (Q.K .. 4) of Long_Float;
type Array_Type is array (Q.K .. Q.N) of Long_Float;
```

or unconstrained array types:

```
type Array_Type is array (Positive range <>) of Long_Float;
```

The quality of the generated code decreases when the dynamic aspect of the array type increases, the worst code being generated for unconstrained array types. This is so because, the less information the compiler has about the bounds of the array, the more fallback code it needs to generate in order to fix things up at run time.

It is possible to specify that a given loop should be subject to vectorization preferably to other optimizations by means of pragma `Loop_Optimize`:

```
pragma Loop_Optimize (Vector);
```

placed immediately within the loop will convey the appropriate hint to the compiler for this loop.

It is also possible to help the compiler generate better vectorized code for a given loop by asserting that there are no loop-carried dependencies in the loop. Consider for example the procedure:

```
type Arr is array (1 .. 4) of Long_Float;

procedure Add (X, Y : not null access Arr; R : not null access Arr) is
begin
  for I in Arr'Range loop
    R(I) := X(I) + Y(I);
  end loop;
end;
```

By default, the compiler cannot unconditionally vectorize the loop because assigning to a component of the array designated by R in one iteration could change the value read from the components of the array designated by X or Y in a later iteration. As a result, the compiler will generate two versions of the loop in the object code, one vectorized and the other not vectorized, as well as a test to select the appropriate version at run time. This can be overcome by another hint:

```
pragma Loop_Optimize (Ivdep);
```

placed immediately within the loop will tell the compiler that it can safely omit the non-vectorized version of the loop as well as the run-time test.

Other Optimization Switches

Since GNAT uses the gcc back end, all the specialized gcc optimization switches are potentially usable. These switches have not been extensively tested with GNAT but can generally be expected to work. Examples of switches in this category are `-funroll-loops` and the various target-specific `-m` options (in particular, it has been observed that `-march=xxx` can significantly improve performance on appropriate machines). For full details of these switches, see the *Submodel Options* section in the *Hardware Models and Configurations* chapter of *Using the GNU Compiler Collection (GCC)*.

Optimization and Strict Aliasing

The strong typing capabilities of Ada allow an optimizer to generate efficient code in situations where other languages would be forced to make worst case assumptions preventing such optimizations. Consider the following example:

```
procedure R is
  type Int1 is new Integer;
  type Int2 is new Integer;
  type Int1A is access Int1;
  type Int2A is access Int2;
  Int1V : Int1A;
  Int2V : Int2A;
  ...

begin
  ...
  for J in Data'Range loop
    if Data (J) = Int1V.all then
      Int2V.all := Int2V.all + 1;
    end if;
  end loop;
  ...
end R;
```

In this example, since the variable `Int1V` can only access objects of type `Int1`, and `Int2V` can only access objects of type `Int2`, there is no possibility that the assignment to `Int2V.all` affects the value of `Int1V.all`. This means that the compiler optimizer can “know” that the value `Int1V.all` is

constant for all iterations of the loop and avoid the extra memory reference required to dereference it each time through the loop.

This kind of optimization, called strict aliasing analysis, is triggered by specifying an optimization level of `-O2` or higher or `-Os` and allows GNAT to generate more efficient code when access values are involved.

However, although this optimization is always correct in terms of the formal semantics of the Ada Reference Manual, difficulties can arise if features like `Unchecked_Conversion` are used to break the typing system. Consider the following complete program example:

```

package p1 is
  type int1 is new integer;
  type int2 is new integer;
  type a1 is access int1;
  type a2 is access int2;
end p1;

with p1; use p1;
package p2 is
  function to_a2 (Input : a1) return a2;
end p2;

with Ada.Unchecked_Conversion;
package body p2 is
  function to_a2 (Input : a1) return a2 is
    function to_a2u is
      new Ada.Unchecked_Conversion (a1, a2);
    begin
      return to_a2u (Input);
    end to_a2;
end p2;

with p2; use p2;
with p1; use p1;
with Text_IO; use Text_IO;
procedure m is
  v1 : a1 := new int1;
  v2 : a2 := to_a2 (v1);
begin
  v1.all := 1;
  v2.all := 0;
  put_line (int1'image (v1.all));
end;

```

This program prints out 0 in `-O0` or `-O1` mode, but it prints out 1 in `-O2` mode. That's because in strict aliasing mode, the compiler can and does assume that the assignment to `v2.all` could not affect the value of `v1.all`, since different types are involved.

This behavior is not a case of non-conformance with the standard, since the Ada RM specifies that an unchecked conversion where the resulting bit pattern is not a correct value of the target type can result in an abnormal value and attempting to reference an abnormal value makes the execution of a program erroneous. That's the case here since the result does not point to an object of type

int2. This means that the effect is entirely unpredictable.

However, although that explanation may satisfy a language lawyer, in practice an applications programmer expects an unchecked conversion involving pointers to create true aliases and the behavior of printing 1 seems plain wrong. In this case, the strict aliasing optimization is unwelcome.

Indeed the compiler recognizes this possibility, and the unchecked conversion generates a warning:

```
p2.adb:5:07: warning: possible aliasing problem with type "a2"  
p2.adb:5:07: warning: use -fno-strict-aliasing switch for references  
p2.adb:5:07: warning: or use "pragma No_Strict_Aliasing (a2);"
```

Unfortunately the problem is recognized when compiling the body of package p2, but the actual “bad” code is generated while compiling the body of m and this latter compilation does not see the suspicious `Unchecked_Conversion`.

As implied by the warning message, there are approaches you can use to avoid the unwanted strict aliasing optimization in a case like this.

One possibility is to simply avoid the use of `-O2`, but that is a bit drastic, since it throws away a number of useful optimizations that do not involve strict aliasing assumptions.

A less drastic approach is to compile the program using the option `-fno-strict-aliasing`. Actually it is only the unit containing the dereferencing of the suspicious pointer that needs to be compiled. So in this case, if we compile unit m with this switch, then we get the expected value of zero printed. Analyzing which units might need the switch can be painful, so a more reasonable approach is to compile the entire program with options `-O2` and `-fno-strict-aliasing`. If the performance is satisfactory with this combination of options, then the advantage is that the entire issue of possible “wrong” optimization due to strict aliasing is avoided.

To avoid the use of compiler switches, the configuration pragma `No_Strict_Aliasing` with no parameters may be used to specify that for all access types, the strict aliasing optimization should be suppressed.

However, these approaches are still overkill, in that they causes all manipulations of all access values to be deoptimized. A more refined approach is to concentrate attention on the specific access type identified as problematic.

First, if a careful analysis of uses of the pointer shows that there are no possible problematic references, then the warning can be suppressed by bracketing the instantiation of `Unchecked_Conversion` to turn the warning off:

```
pragma Warnings (Off);  
function to_a2u is  
  new Ada.Unchecked_Conversion (a1, a2);  
pragma Warnings (On);
```

Of course that approach is not appropriate for this particular example, since indeed there is a problematic reference. In this case we can take one of two other approaches.

The first possibility is to move the instantiation of unchecked conversion to the unit in which the type is declared. In this example, we would move the instantiation of `Unchecked_Conversion` from the body of package p2 to the spec of package p1. Now the warning disappears. That’s because

any use of the access type knows there is a suspicious unchecked conversion, and the strict aliasing optimization is automatically suppressed for the type.

If it is not practical to move the unchecked conversion to the same unit in which the destination access type is declared (perhaps because the source type is not visible in that unit), you may use `pragma No_Strict_Aliasing` for the type. This pragma must occur in the same declarative sequence as the declaration of the access type:

```
type a2 is access int2;
pragma No_Strict_Aliasing (a2);
```

Here again, the compiler now knows that the strict aliasing optimization should be suppressed for any reference to type `a2` and the expected behavior is obtained.

Finally, note that although the compiler can generate warnings for simple cases of unchecked conversions, there are trickier and more indirect ways of creating type incorrect aliases which the compiler cannot detect. Examples are the use of address overlays and unchecked conversions involving composite types containing access types as components. In such cases, no warnings are generated, but there can still be aliasing problems. One safe coding practice is to forbid the use of address clauses for type overlaying, and to allow unchecked conversion only for primitive types. This is not really a significant restriction since any possible desired effect can be achieved by unchecked conversion of access values.

The aliasing analysis done in strict aliasing mode can certainly have significant benefits. We have seen cases of large scale application code where the time is increased by up to 5% by turning this optimization off. If you have code that includes significant usage of unchecked conversion, you might want to just stick with `-O1` and avoid the entire issue. If you get adequate performance at this level of optimization level, that's probably the safest approach. If tests show that you really need higher levels of optimization, then you can experiment with `-O2` and `-O2 -fno-strict-aliasing` to see how much effect this has on size and speed of the code. If you really need to use `-O2` with strict aliasing in effect, then you should review any uses of unchecked conversion of access types, particularly if you are getting the warnings described above.

Aliased Variables and Optimization

There are scenarios in which programs may use low level techniques to modify variables that otherwise might be considered to be unassigned. For example, a variable can be passed to a procedure by reference, which takes the address of the parameter and uses the address to modify the variable's value, even though it is passed as an IN parameter. Consider the following example:

```
procedure P is
  Max_Length : constant Natural := 16;
  type Char_Ptr is access all Character;

  procedure Get_String(Buffer: Char_Ptr; Size : Integer);
  pragma Import (C, Get_String, "get_string");

  Name : aliased String (1 .. Max_Length) := (others => ' ');
  Temp : Char_Ptr;
```

(continues on next page)

(continued from previous page)

```

function Addr (S : String) return Char_Ptr is
  function To_Char_Ptr is
    new Ada.Unchecked_Conversion (System.Address, Char_Ptr);
  begin
    return To_Char_Ptr (S (S'First)'Address);
  end;
begin
  Temp := Addr (Name);
  Get_String (Temp, Max_Length);
end;

```

where `Get_String` is a C function that uses the address in `Temp` to modify the variable `Name`. This code is dubious, and arguably erroneous, and the compiler would be entitled to assume that `Name` is never modified, and generate code accordingly.

However, in practice, this would cause some existing code that seems to work with no optimization to start failing at high levels of optimization.

What the compiler does for such cases is to assume that marking a variable as aliased indicates that some “funny business” may be going on. The optimizer recognizes the aliased keyword and inhibits optimizations that assume the value cannot be assigned. This means that the above example will in fact “work” reliably, that is, it will produce the expected results.

Atomic Variables and Optimization

There are two considerations with regard to performance when atomic variables are used.

First, the RM only guarantees that access to atomic variables be atomic, it has nothing to say about how this is achieved, though there is a strong implication that this should not be achieved by explicit locking code. Indeed GNAT will never generate any locking code for atomic variable access (it will simply reject any attempt to make a variable or type atomic if the atomic access cannot be achieved without such locking code).

That being said, it is important to understand that you cannot assume that the entire variable will always be accessed. Consider this example:

```

type R is record
  A,B,C,D : Character;
end record;
for R'Size use 32;
for R'Alignment use 4;

RV : R;
pragma Atomic (RV);
X : Character;
...
X := RV.B;

```

You cannot assume that the reference to `RV.B` will read the entire 32-bit variable with a single load instruction. It is perfectly legitimate if the hardware allows it to do a byte read of just the `B` field. This read is still atomic, which is all the RM requires. GNAT can and does take advantage of this, depending on the architecture and optimization level. Any assumption to the contrary is non-portable and risky. Even if you examine the assembly language and see a full 32-bit load, this might change in a future version of the compiler.

If your application requires that all accesses to `RV` in this example be full 32-bit loads, you need to make a copy for the access as in:

```
declare
  RV_Copy : constant R := RV;
begin
  X := RV_Copy.B;
end;
```

Now the reference to `RV` must read the whole variable. Actually one can imagine some compiler which figures out that the whole copy is not required (because only the `B` field is actually accessed), but GNAT certainly won't do that, and we don't know of any compiler that would not handle this right, and the above code will in practice work portably across all architectures (that permit the Atomic declaration).

The second issue with atomic variables has to do with the possible requirement of generating synchronization code. For more details on this, consult the sections on the pragmas `Enable/Disable_Atomic_Synchronization` in the GNAT Reference Manual. If performance is critical, and such synchronization code is not required, it may be useful to disable it.

Passive Task Optimization

A passive task is one which is sufficiently simple that in theory a compiler could recognize it and implement it efficiently without creating a new thread. The original design of Ada 83 had in mind this kind of passive task optimization, but only a few Ada 83 compilers attempted it. The problem was that it was difficult to determine the exact conditions under which the optimization was possible. The result is a very fragile optimization where a very minor change in the program can suddenly silently make a task non-optimizable.

With the revisiting of this issue in Ada 95, there was general agreement that this approach was fundamentally flawed, and the notion of protected types was introduced. When using protected types, the restrictions are well defined, and you KNOW that the operations will be optimized, and furthermore this optimized performance is fully portable.

Although it would theoretically be possible for GNAT to attempt to do this optimization, but it really doesn't make sense in the context of Ada 95, and none of the Ada 95 compilers implement this optimization as far as we know. In particular GNAT never attempts to perform this optimization.

In any new Ada 95 code that is written, you should always use protected types in place of tasks that might be able to be optimized in this manner. Of course this does not help if you have legacy Ada 83 code that depends on this optimization, but it is unusual to encounter a case where the performance gains from this optimization are significant.

Your program should work correctly without this optimization. If you have performance problems, then the most practical approach is to figure out exactly where these performance problems arise, and update those particular tasks to be protected types. Note that typically clients of the tasks who call entries, will not have to be modified, only the task definition itself.

6.3.2 Text_IO Suggestions

The `Ada.Text_IO` package has fairly high overheads due in part to the requirement of maintaining page and line counts. If performance is critical, a recommendation is to use `Stream_IO` instead of `Text_IO` for volume output, since this package has less overhead.

If `Text_IO` must be used, note that by default output to the standard output and standard error files is unbuffered (this provides better behavior when output statements are used for debugging, or if the progress of a program is observed by tracking the output, e.g. by using the Unix `tail -f` command to watch redirected output).

If you are generating large volumes of output with `Text_IO` and performance is an important factor, use a designated file instead of the standard output file, or change the standard output file to be buffered using `Interfaces.C_Streams.setvbuf`.

6.3.3 Reducing Size of Executables with Unused Subprogram/Data Elimination

This section describes how you can eliminate unused subprograms and data from your executable just by setting options at compilation time.

About unused subprogram/data elimination

By default, an executable contains all code and data of its composing objects (directly linked or coming from statically linked libraries), even data or code never used by this executable.

This feature will allow you to eliminate such unused code from your executable, making it smaller (in disk and in memory).

This functionality is available on all Linux platforms except for the IA-64 architecture and on all cross platforms using the ELF binary file format. In both cases GNU binutils version 2.16 or later are required to enable it.

Compilation options

The operation of eliminating the unused code and data from the final executable is directly performed by the linker.

In order to do this, it has to work with objects compiled with the following options: `-ffunction-sections -fdata-sections`.

These options are usable with C and Ada files. They will place respectively each function or data in a separate section in the resulting object file.

Once the objects and static libraries are created with these options, the linker can perform the dead code elimination. You can do this by setting the `-wl,--gc-sections` option to gcc command or in the `-larg`s section of `gnatmake`. This will perform a garbage collection of code and data never referenced.

If the linker performs a partial link (`-r` linker option), then you will need to provide the entry point using the `-e / --entry` linker option.

Note that objects compiled without the `-ffunction-sections` and `-fdata-sections` options can still be linked with the executable. However, no dead code elimination will be performed on those objects (they will be linked as is).

The GNAT static library is now compiled with `-ffunction-sections` and `-fdata-sections` on some platforms. This allows you to eliminate the unused code and data of the GNAT library from your executable.

Example of unused subprogram/data elimination

Here is a simple example:

```
with Aux;

procedure Test is
begin
  Aux.Used (10);
end Test;

package Aux is
  Used_Data   : Integer;
  Unused_Data : Integer;

  procedure Used   (Data : Integer);
  procedure Unused (Data : Integer);
end Aux;

package body Aux is
  procedure Used (Data : Integer) is
  begin
    Used_Data := Data;
  end Used;

  procedure Unused (Data : Integer) is
  begin
    Unused_Data := Data;
  end Unused;
end Aux;
```

`Unused` and `Unused_Data` are never referenced in this code excerpt, and hence they may be safely removed from the final executable.

```
$ gnatmake test

$ nm test | grep used
020015f0 T aux__unused
02005d88 B aux__unused_data
020015cc T aux__used
02005d84 B aux__used_data

$ gnatmake test -cargs -fdata-sections -ffunction-sections \\  
-largs -WL,--gc-sections

$ nm test | grep used
02005350 T aux__used
0201ffe0 B aux__used_data
```

It can be observed that the procedure `Unused` and the object `Unused_Data` are removed by the linker when using the appropriate options.

6.4 Overflow Check Handling in GNAT

This section explains how to control the handling of overflow checks.

6.4.1 Background

Overflow checks are checks that the compiler may make to ensure that intermediate results are not out of range. For example:

```
A : Integer;  
...  
A := A + 1;
```

If `A` has the value `Integer'Last`, then the addition may cause overflow since the result is out of range of the type `Integer`. In this case `Constraint_Error` will be raised if checks are enabled.

A trickier situation arises in examples like the following:

```
A, C : Integer;  
...  
A := (A + 1) + C;
```

where `A` is `Integer'Last` and `C` is `-1`. Now the final result of the expression on the right hand side is `Integer'Last` which is in range, but the question arises whether the intermediate addition of `(A + 1)` raises an overflow error.

The (perhaps surprising) answer is that the Ada language definition does not answer this question. Instead it leaves it up to the implementation to do one of two things if overflow checks are enabled.

- raise an exception (`Constraint_Error`), or
- yield the correct mathematical result which is then used in subsequent operations.

If the compiler chooses the first approach, then the assignment of this example will indeed raise `Constraint_Error` if overflow checking is enabled, or result in erroneous execution if overflow checks are suppressed.

But if the compiler chooses the second approach, then it can perform both additions yielding the correct mathematical result, which is in range, so no exception will be raised, and the right result is obtained, regardless of whether overflow checks are suppressed.

Note that in the first example an exception will be raised in either case, since if the compiler gives the correct mathematical result for the addition, it will be out of range of the target type of the assignment, and thus fails the range check.

This lack of specified behavior in the handling of overflow for intermediate results is a source of non-portability, and can thus be problematic when programs are ported. Most typically this arises in a situation where the original compiler did not raise an exception, and then the application is moved to a compiler where the check is performed on the intermediate result and an unexpected exception is raised.

Furthermore, when using Ada 2012's preconditions and other assertion forms, another issue arises. Consider:

```
procedure P (A, B : Integer) with
  Pre => A + B <= Integer'Last;
```

One often wants to regard arithmetic in a context like this from a mathematical point of view. So for example, if the two actual parameters for a call to `P` are both `Integer'Last`, then the precondition should be regarded as `False`. If we are executing in a mode with run-time checks enabled for preconditions, then we would like this precondition to fail, rather than raising an exception because of the intermediate overflow.

However, the language definition leaves the specification of whether the above condition fails (raising `Assert_Error`) or causes an intermediate overflow (raising `Constraint_Error`) up to the implementation.

The situation is worse in a case such as the following:

```
procedure Q (A, B, C : Integer) with
  Pre => A + B + C <= Integer'Last;
```

Consider the call

```
Q (A => Integer'Last, B => 1, C => -1);
```

From a mathematical point of view the precondition is `True`, but at run time we may (but are not guaranteed to) get an exception raised because of the intermediate overflow (and we really would prefer this precondition to be considered `True` at run time).

6.4.2 Management of Overflows in GNAT

To deal with the portability issue, and with the problem of mathematical versus run-time interpretation of the expressions in assertions, GNAT provides comprehensive control over the handling of intermediate overflow. GNAT can operate in three modes, and furthermore, permits separate selection of operating modes for the expressions within assertions (here the term ‘assertions’ is used in the technical sense, which includes preconditions and so forth) and for expressions appearing outside assertions.

The three modes are:

- *Use base type for intermediate operations* (STRICT)

In this mode, all intermediate results for predefined arithmetic operators are computed using the base type, and the result must be in range of the base type. If this is not the case then either an exception is raised (if overflow checks are enabled) or the execution is erroneous (if overflow checks are suppressed). This is the normal default mode.

- *Most intermediate overflows avoided* (MINIMIZED)

In this mode, the compiler attempts to avoid intermediate overflows by using a larger integer type, typically `Long_Long_Integer`, as the type in which arithmetic is performed for predefined arithmetic operators. This may be slightly more expensive at run time (compared to suppressing intermediate overflow checks), though the cost is negligible on modern 64-bit machines. For the examples given earlier, no intermediate overflows would have resulted in exceptions, since the intermediate results are all in the range of `Long_Long_Integer` (typically 64-bits on nearly all implementations of GNAT). In addition, if checks are enabled, this reduces the number of checks that must be made, so this choice may actually result in an improvement in space and time behavior.

However, there are cases where `Long_Long_Integer` is not large enough, consider the following example:

```
procedure R (A, B, C, D : Integer) with
  Pre => (A**2 * B**2) / (C**2 * D**2) <= 10;
```

where $A = B = C = D = \text{Integer}'\text{Last}$. Now the intermediate results are out of the range of `Long_Long_Integer` even though the final result is in range and the precondition is True (from a mathematical point of view). In such a case, operating in this mode, an overflow occurs for the intermediate computation (which is why this mode says *most* intermediate overflows are avoided). In this case, an exception is raised if overflow checks are enabled, and the execution is erroneous if overflow checks are suppressed.

- *All intermediate overflows avoided* (ELIMINATED)

In this mode, the compiler avoids all intermediate overflows by using arbitrary precision arithmetic as required. In this mode, the above example with $A**2 * B**2$ would not cause intermediate overflow, because the intermediate result would be evaluated using sufficient precision, and the result of evaluating the precondition would be True.

This mode has the advantage of avoiding any intermediate overflows, but at the expense of significant run-time overhead, including the use of a library (included automatically in this mode) for multiple-precision arithmetic.

This mode provides cleaner semantics for assertions, since now the run-time behavior emulates true arithmetic behavior for the predefined arithmetic operators, meaning that there is never a conflict between the mathematical view of the assertion, and its run-time behavior.

Note that in this mode, the behavior is unaffected by whether or not overflow checks are suppressed, since overflow does not occur. It is possible for gigantic intermediate expressions to raise `Storage_Error` as a result of attempting to compute the results of such expressions (e.g. `Integer'Last ** Integer'Last`) but overflow is impossible.

Note that these modes apply only to the evaluation of predefined arithmetic, membership, and comparison operators for signed integer arithmetic.

For fixed-point arithmetic, checks can be suppressed. But if checks are enabled then fixed-point values are always checked for overflow against the base type for intermediate expressions (that is such checks always operate in the equivalent of `STRICT` mode).

For floating-point, on nearly all architectures, `Machine_Overflows` is `False`, and IEEE infinities are generated, so overflow exceptions are never raised. If you want to avoid infinities, and check that final results of expressions are in range, then you can declare a constrained floating-point type, and range checks will be carried out in the normal manner (with infinite values always failing all range checks).

6.4.3 Specifying the Desired Mode

The desired mode of for handling intermediate overflow can be specified using either the `Overflow_Mode` pragma or an equivalent compiler switch. The pragma has the form

```
pragma Overflow_Mode ([General =>] MODE [, [Assertions =>] MODE]);
```

where `MODE` is one of

- `STRICT`: intermediate overflows checked (using base type)
- `MINIMIZED`: minimize intermediate overflows
- `ELIMINATED`: eliminate intermediate overflows

The case is ignored, so `MINIMIZED`, `Minimized` and `minimized` all have the same effect.

If only the `General` parameter is present, then the given `MODE` applies to expressions both within and outside assertions. If both arguments are present, then `General` applies to expressions outside assertions, and `Assertions` applies to expressions within assertions. For example:

```
pragma Overflow_Mode
  (General => Minimized, Assertions => Eliminated);
```

specifies that general expressions outside assertions be evaluated in ‘minimize intermediate overflows’ mode, and expressions within assertions be evaluated in ‘eliminate intermediate overflows’ mode. This is often a reasonable choice, avoiding excessive overhead outside assertions, but assuring a high degree of portability when importing code from another compiler, while incurring the extra overhead for assertion expressions to ensure that the behavior at run time matches the expected mathematical behavior.

The `Overflow_Mode` pragma has the same scoping and placement rules as pragma `Suppress`, so it can occur either as a configuration pragma, specifying a default for the whole program, or in a declarative scope, where it applies to the remaining declarations and statements in that scope.

Note that pragma `Overflow_Mode` does not affect whether overflow checks are enabled or suppressed. It only controls the method used to compute intermediate values. To control whether overflow checking is enabled or suppressed, use pragma `Suppress` or `Unsuppress` in the usual manner.

Additionally, a compiler switch `-gnato?` or `-gnato??` can be used to control the checking mode default (which can be subsequently overridden using pragmas).

Here `?` is one of the digits 1 through 3:

1	use base type for intermediate operations (STRICT)
2	minimize intermediate overflows (MINIMIZED)
3	eliminate intermediate overflows (ELIMINATED)

As with the pragma, if only one digit appears then it applies to all cases; if two digits are given, then the first applies outside assertions, and the second within assertions. Thus the equivalent of the example pragma above would be `-gnato23`.

If no digits follow the `-gnato`, then it is equivalent to `-gnato11`, causing all intermediate operations to be computed using the base type (**STRICT** mode).

6.4.4 Default Settings

The default mode for overflow checks is

```
General => Strict
```

which causes all computations both inside and outside assertions to use the base type.

This retains compatibility with previous versions of GNAT which suppressed overflow checks by default and always used the base type for computation of intermediate results.

The switch `-gnato` (with no digits following) is equivalent to

```
General => Strict
```

which causes overflow checking of all intermediate overflows both inside and outside assertions against the base type.

The pragma `Suppress (Overflow_Check)` disables overflow checking, but it has no effect on the method used for computing intermediate results.

The pragma `Unsuppress (Overflow_Check)` enables overflow checking, but it has no effect on the method used for computing intermediate results.

6.4.5 Implementation Notes

In practice on typical 64-bit machines, the `MINIMIZED` mode is reasonably efficient, and can be generally used. It also helps to ensure compatibility with code imported from some other compiler to GNAT.

Setting all intermediate overflows checking (`CHECKED` mode) makes sense if you want to make sure that your code is compatible with any other possible Ada implementation. This may be useful in ensuring portability for code that is to be exported to some other compiler than GNAT.

The Ada standard allows the reassociation of expressions at the same precedence level if no parentheses are present. For example, `A+B+C` parses as though it were `(A+B)+C`, but the compiler can reinterpret this as `A+(B+C)`, possibly introducing or eliminating an overflow exception. The GNAT compiler never takes advantage of this freedom, and the expression `A+B+C` will be evaluated as `(A+B)+C`. If you need the other order, you can write the parentheses explicitly `A+(B+C)` and GNAT will respect this order.

The use of `ELIMINATED` mode will cause the compiler to automatically include an appropriate arbitrary precision integer arithmetic package. The compiler will make calls to this package, though only in cases where it cannot be sure that `Long_Long_Integer` is sufficient to guard against intermediate overflows. This package does not use dynamic allocation, but it does use the secondary stack, so an appropriate secondary stack package must be present (this is always true for standard full Ada, but may require specific steps for restricted run times such as `ZFP`).

Although `ELIMINATED` mode causes expressions to use arbitrary precision arithmetic, avoiding overflow, the final result must be in an appropriate range. This is true even if the final result is of type `[Long_[Long_]Integer'Base`, which still has the same bounds as its associated constrained type at run-time.

Currently, the `ELIMINATED` mode is only available on target platforms for which `Long_Long_Integer` is 64-bits (nearly all GNAT platforms).

6.5 Performing Dimensionality Analysis in GNAT

The GNAT compiler supports dimensionality checking. The user can specify physical units for objects, and the compiler will verify that uses of these objects are compatible with their dimensions, in a fashion that is familiar to engineering practice. The dimensions of algebraic expressions (including powers with static exponents) are computed from their constituents.

This feature depends on Ada 2012 aspect specifications, and is available from version 7.0.1 of GNAT onwards. The GNAT-specific aspect `Dimension_System` allows you to define a system of units; the aspect `Dimension` then allows the user to declare dimensioned quantities within a given system. (These aspects are described in the *Implementation Defined Aspects* chapter of the *GNAT Reference Manual*).

The major advantage of this model is that it does not require the declaration of multiple operators for all possible combinations of types: it is only necessary to use the proper subtypes in object declarations.

The simplest way to impose dimensionality checking on a computation is to make use of one of the instantiations of the package `System.Dim.Generic_Mks`, which are part of the GNAT library. This generic package defines a floating-point type `MKS_Type`, for which a sequence of dimension names are specified, together with their conventional abbreviations. The following should be read together with the full specification of the package, in file `s-digemk.ads`.

```

type Mks_Type is new Float_Type
with
  Dimension_System => (
    (Unit_Name => Meter,   Unit_Symbol => 'm',   Dim_Symbol => 'L'),
    (Unit_Name => Kilogram, Unit_Symbol => "kg",  Dim_Symbol => 'M'),
    (Unit_Name => Second,  Unit_Symbol => 's',   Dim_Symbol => 'T'),
    (Unit_Name => Ampere,  Unit_Symbol => 'A',   Dim_Symbol => 'I'),
    (Unit_Name => Kelvin,  Unit_Symbol => 'K',   Dim_Symbol => "Theta"),
    (Unit_Name => Mole,    Unit_Symbol => "mol",  Dim_Symbol => 'N'),
    (Unit_Name => Candela, Unit_Symbol => "cd",  Dim_Symbol => 'J'));

```

The package then defines a series of subtypes that correspond to these conventional units. For example:

```

subtype Length is Mks_Type
with
  Dimension => (Symbol => 'm', Meter => 1, others => 0);

```

and similarly for `Mass`, `Time`, `Electric_Current`, `Thermodynamic_Temperature`, `Amount_Of_Substance`, and `Luminous_Intensity` (the standard set of units of the SI system).

The package also defines conventional names for values of each unit, for example:

```

m  : constant Length      := 1.0;
kg : constant Mass       := 1.0;
s  : constant Time       := 1.0;
A  : constant Electric_Current := 1.0;

```

as well as useful multiples of these units:

```

cm : constant Length := 1.0E-02;
g  : constant Mass   := 1.0E-03;
min : constant Time  := 60.0;
day : constant Time  := 60.0 * 24.0 * min;
...

```

There are three instantiations of `System.Dim.Generic_Mks` defined in the GNAT library:

- `System.Dim.Float_Mks` based on `Float` defined in `s-diflmk.ads`.
- `System.Dim.Long_Mks` based on `Long_Float` defined in `s-dilomk.ads`.
- `System.Dim.Mks` based on `Long_Long_Float` defined in `s-dimmks.ads`.

Using one of these packages, you can then define a derived unit by providing the aspect that specifies its dimensions within the MKS system, as well as the string to be used for output of a value of that unit:

```

subtype Acceleration is Mks_Type
  with Dimension => ("m/sec^2",
                    Meter => 1,
                    Second => -2,
                    others => 0);

```

Here is a complete example of use:

```

with System.Dim.MKS; use System.Dim.Mks;
with System.Dim.Mks_IO; use System.Dim.Mks_IO;
with Text_IO; use Text_IO;
procedure Free_Fall is
  subtype Acceleration is Mks_Type
    with Dimension => ("m/sec^2", 1, 0, -2, others => 0);
  G : constant acceleration := 9.81 * m / (s ** 2);
  T : Time := 10.0*s;
  Distance : Length;

begin
  Put ("Gravitational constant: ");
  Put (G, Aft => 2, Exp => 0); Put_Line ("");
  Distance := 0.5 * G * T ** 2;
  Put ("distance travelled in 10 seconds of free fall ");
  Put (Distance, Aft => 2, Exp => 0);
  Put_Line ("");
end Free_Fall;

```

Execution of this program yields:

```

Gravitational constant: 9.81 m/sec^2
distance travelled in 10 seconds of free fall 490.50 m

```

However, incorrect assignments such as:

```

Distance := 5.0;
Distance := 5.0 * kg;

```

are rejected with the following diagnoses:

```

Distance := 5.0;
  >>> dimensions mismatch in assignment
  >>> left-hand side has dimension [L]
  >>> right-hand side is dimensionless

Distance := 5.0 * kg;
  >>> dimensions mismatch in assignment
  >>> left-hand side has dimension [L]
  >>> right-hand side has dimension [M]

```

The dimensions of an expression are properly displayed, even if there is no explicit subtype for it. If we add to the program:

```
Put ("Final velocity: ");
Put (G * T, Aft =>2, Exp =>0);
Put_Line ("");
```

then the output includes:

```
Final velocity: 98.10 m.s**(-1)
```

The type `Mks_Type` is said to be a *dimensionable type* since it has a `Dimension_System` aspect, and the subtypes `Length`, `Mass`, etc., are said to be *dimensioned subtypes* since each one has a `Dimension` aspect.

The `Dimension` aspect of a dimensioned subtype `S` defines a mapping from the base type's `Unit_Names` to integer (or, more generally, rational) values. This mapping is the *dimension vector* (also referred to as the *dimensionality*) for that subtype, denoted by $DV(S)$, and thus for each object of that subtype. Intuitively, the value specified for each `Unit_Name` is the exponent associated with that unit; a zero value means that the unit is not used. For example:

```
declare
  Acc : Acceleration;
  ...
begin
  ...
end;
```

Here $DV(\text{Acc}) = DV(\text{Acceleration}) = (\text{Meter} \Rightarrow 1, \text{Kilogram} \Rightarrow 0, \text{Second} \Rightarrow -2, \text{Ampere} \Rightarrow 0, \text{Kelvin} \Rightarrow 0, \text{Mole} \Rightarrow 0, \text{Candela} \Rightarrow 0)$. Symbolically, we can express this as `Meter / Second**2`.

The dimension vector of an arithmetic expression is synthesized from the dimension vectors of its components, with compile-time dimensionality checks that help prevent mismatches such as using an `Acceleration` where a `Length` is required.

The dimension vector of the result of an arithmetic expression `expr`, or $DV(\text{expr})$, is defined as follows, assuming conventional mathematical definitions for the vector operations that are used:

- If `expr` is of the type *universal_real*, or is not of a dimensioned subtype, then `expr` is dimensionless; $DV(\text{expr})$ is the empty vector.
- $DV(\text{op expr})$, where `op` is a unary operator, is $DV(\text{expr})$
- $DV(\text{expr1 op expr2})$ where `op` is “+” or “-” is $DV(\text{expr1})$ provided that $DV(\text{expr1}) = DV(\text{expr2})$. If this condition is not met then the construct is illegal.
- $DV(\text{expr1} * \text{expr2})$ is $DV(\text{expr1}) + DV(\text{expr2})$, and $DV(\text{expr1} / \text{expr2}) = DV(\text{expr1}) - DV(\text{expr2})$. In this context if one of the `exprs` is dimensionless then its empty dimension vector is treated as (`others => 0`).
- $DV(\text{expr} ** \text{power})$ is $\text{power} * DV(\text{expr})$, provided that `power` is a static rational value. If this condition is not met then the construct is illegal.

Note that, by the above rules, it is illegal to use binary “+” or “-” to combine a dimensioned and dimensionless value. Thus an expression such as `acc-10.0` is illegal, where `acc` is an object of

subtype `Acceleration`.

The dimensionality checks for relationals use the same rules as for “+” and “-”, except when comparing to a literal; thus

```
acc > len
```

is equivalent to

```
acc-len > 0.0
```

and is thus illegal, but

```
acc > 10.0
```

is accepted with a warning. Analogously a conditional expression requires the same dimension vector for each branch (with no exception for literals).

The dimension vector of a type conversion $\mathsf{T}(expr)$ is defined as follows, based on the nature of T :

- If T is a dimensioned subtype then $\mathsf{DV}(\mathsf{T}(expr))$ is $\mathsf{DV}(\mathsf{T})$ provided that either $expr$ is dimensionless or $\mathsf{DV}(\mathsf{T}) = \mathsf{DV}(expr)$. The conversion is illegal if $expr$ is dimensioned and $\mathsf{DV}(expr) \neq \mathsf{DV}(\mathsf{T})$. Note that vector equality does not require that the corresponding `Unit_Names` be the same.

As a consequence of the above rule, it is possible to convert between different dimension systems that follow the same international system of units, with the seven physical components given in the standard order (length, mass, time, etc.). Thus a length in meters can be converted to a length in inches (with a suitable conversion factor) but cannot be converted, for example, to a mass in pounds.

- If T is the base type for $expr$ (and the dimensionless root type of the dimension system), then $\mathsf{DV}(\mathsf{T}(expr))$ is $\mathsf{DV}(expr)$. Thus, if $expr$ is of a dimensioned subtype of T , the conversion may be regarded as a “view conversion” that preserves dimensionality.

This rule makes it possible to write generic code that can be instantiated with compatible dimensioned subtypes. The generic unit will contain conversions that will consequently be present in instantiations, but conversions to the base type will preserve dimensionality and make it possible to write generic code that is correct with respect to dimensionality.

- Otherwise (i.e., T is neither a dimensioned subtype nor a dimensionable base type), $\mathsf{DV}(\mathsf{T}(expr))$ is the empty vector. Thus a dimensioned value can be explicitly converted to a non-dimensioned subtype, which of course then escapes dimensionality analysis.

The dimension vector for a type qualification $\mathsf{T}'(expr)$ is the same as for the type conversion $\mathsf{T}(expr)$.

An assignment statement

```
Source := Target;
```

requires $\mathsf{DV}(\mathsf{Source}) = \mathsf{DV}(\mathsf{Target})$, and analogously for parameter passing (the dimension vector for the actual parameter must be equal to the dimension vector for the formal parameter).

6.6 Stack Related Facilities

This section describes some useful tools associated with stack checking and analysis. In particular, it deals with dynamic and static stack usage measurements.

6.6.1 Stack Overflow Checking

For most operating systems, `gcc` does not perform stack overflow checking by default. This means that if the main environment task or some other task exceeds the available stack space, then unpredictable behavior will occur. Most native systems offer some level of protection by adding a guard page at the end of each task stack. This mechanism is usually not enough for dealing properly with stack overflow situations because a large local variable could “jump” above the guard page. Furthermore, when the guard page is hit, there may not be any space left on the stack for executing the exception propagation code. Enabling stack checking avoids such situations.

To activate stack checking, compile all units with the `gcc` option `-fstack-check`. For example:

```
$ gcc -c -fstack-check package1.adb
```

Units compiled with this option will generate extra instructions to check that any use of the stack (for procedure calls or for declaring local variables in declare blocks) does not exceed the available stack space. If the space is exceeded, then a `Storage_Error` exception is raised.

For declared tasks, the default stack size is defined by the GNAT runtime, whose size may be modified at bind time through the `-d` bind switch ([Switches for gnatbind](#)). Task specific stack sizes may be set using the `Storage_Size` pragma.

For the environment task, the stack size is determined by the operating system. Consequently, to modify the size of the environment task please refer to your operating system documentation.

6.6.2 Static Stack Usage Analysis

A unit compiled with `-fstack-usage` will generate an extra file that specifies the maximum amount of stack used, on a per-function basis. The file has the same basename as the target object file with a `.su` extension. Each line of this file is made up of three fields:

- The name of the function.
- A number of bytes.
- One or more qualifiers: `static`, `dynamic`, `bounded`.

The second field corresponds to the size of the known part of the function frame.

The qualifier `static` means that the function frame size is purely static. It usually means that all local variables have a static size. In this case, the second field is a reliable measure of the function stack utilization.

The qualifier `dynamic` means that the function frame size is not static. It happens mainly when some local variables have a dynamic size. When this qualifier appears alone, the second field is not

a reliable measure of the function stack analysis. When it is qualified with `bounded`, it means that the second field is a reliable maximum of the function stack utilization.

A unit compiled with `-Wstack-usage` will issue a warning for each subprogram whose stack usage might be larger than the specified amount of bytes. The wording is in keeping with the qualifier documented above.

6.6.3 Dynamic Stack Usage Analysis

It is possible to measure the maximum amount of stack used by a task, by adding a switch to `gnatbind`, as:

```
$ gnatbind -u0 file
```

With this option, at each task termination, its stack usage is output on `stderr`. Note that this switch is not compatible with tools like Valgrind and DrMemory; they will report errors.

It is not always convenient to output the stack usage when the program is still running. Hence, it is possible to delay this output until program termination. for a given number of tasks specified as the argument of the `-u` option. For instance:

```
$ gnatbind -u100 file
```

will buffer the stack usage information of the first 100 tasks to terminate and output this info at program termination. Results are displayed in four columns:

```
Index | Task Name | Stack Size | Stack Usage
```

where:

- *Index* is a number associated with each task.
- *Task Name* is the name of the task analyzed.
- *Stack Size* is the maximum size for the stack.
- *Stack Usage* is the measure done by the stack analyzer. In order to prevent overflow, the stack is not entirely analyzed, and it's not possible to know exactly how much has actually been used.

By default the environment task stack, the stack that contains the main unit, is not processed. To enable processing of the environment task stack, the environment variable `GNAT_STACK_LIMIT` needs to be set to the maximum size of the environment task stack. This amount is given in kilobytes. For example:

```
$ set GNAT_STACK_LIMIT 1600
```

would specify to the analyzer that the environment task stack has a limit of 1.6 megabytes. Any stack usage beyond this will be ignored by the analysis.

The package `GNAT.Task_Stack_Usage` provides facilities to get stack-usage reports at run time. See its body for the details.

6.7 Memory Management Issues

This section describes some useful memory pools provided in the GNAT library and in particular the GNAT Debug Pool facility, which can be used to detect incorrect uses of access values (including ‘dangling references’).

6.7.1 Some Useful Memory Pools

The `System.Pool_Global` package offers the `Unbounded_No_Reclaim_Pool` storage pool. Allocations use the standard system call `malloc` while deallocations use the standard system call `free`. No reclamation is performed when the pool goes out of scope. For performance reasons, the standard default Ada allocators/deallocators do not use any explicit storage pools but if they did, they could use this storage pool without any change in behavior. That is why this storage pool is used when the user manages to make the default implicit allocator explicit as in this example:

```
type T1 is access Something;
  -- no Storage pool is defined for T2

type T2 is access Something_Else;
for T2'Storage_Pool use T1'Storage_Pool;
  -- the above is equivalent to
for T2'Storage_Pool use System.Pool_Global.Global_Pool_Object;
```

The `System.Pool_Local` package offers the `Unbounded_Reclaim_Pool` storage pool. The allocation strategy is similar to `Pool_Local` except that the all storage allocated with this pool is reclaimed when the pool object goes out of scope. This pool provides a explicit mechanism similar to the implicit one provided by several Ada 83 compilers for allocations performed through a local access type and whose purpose was to reclaim memory when exiting the scope of a given local access. As an example, the following program does not leak memory even though it does not perform explicit deallocation:

```
with System.Pool_Local;
procedure Poolocl is
  procedure Internal is
    type A is access Integer;
    X : System.Pool_Local.Unbounded_Reclaim_Pool;
    for A'Storage_Pool use X;
    v : A;
  begin
    for I in 1 .. 50 loop
      v := new Integer;
    end loop;
  end Internal;
begin
  for I in 1 .. 100 loop
    Internal;
  end loop;
end Poolocl;
```

The `System.Pool_Size` package implements the `Stack_Bounded_Pool` used when `Storage_Size` is

specified for an access type. The whole storage for the pool is allocated at once, usually on the stack at the point where the access type is elaborated. It is automatically reclaimed when exiting the scope where the access type is defined. This package is not intended to be used directly by the user and it is implicitly used for each such declaration:

```
type T1 is access Something;
for T1'Storage_Size use 10_000;
```

6.7.2 The GNAT Debug Pool Facility

The use of unchecked deallocation and unchecked conversion can easily lead to incorrect memory references. The problems generated by such references are usually difficult to tackle because the symptoms can be very remote from the origin of the problem. In such cases, it is very helpful to detect the problem as early as possible. This is the purpose of the Storage Pool provided by `GNAT.Debug_Pools`.

In order to use the GNAT specific debugging pool, the user must associate a debug pool object with each of the access types that may be related to suspected memory problems. See Ada Reference Manual 13.11.

```
type Ptr is access Some_Type;
Pool : GNAT.Debug_Pools.Debug_Pool;
for Ptr'Storage_Pool use Pool;
```

`GNAT.Debug_Pools` is derived from a GNAT-specific kind of pool: the `Checked_Pool`. Such pools, like standard Ada storage pools, allow the user to redefine allocation and deallocation strategies. They also provide a checkpoint for each dereference, through the use of the primitive operation `Dereference` which is implicitly called at each dereference of an access value.

Once an access type has been associated with a debug pool, operations on values of the type may raise four distinct exceptions, which correspond to four potential kinds of memory corruption:

- `GNAT.Debug_Pools.Accessing_Not_Allocated_Storage`
- `GNAT.Debug_Pools.Accessing_Deallocated_Storage`
- `GNAT.Debug_Pools.Freeing_Not_Allocated_Storage`
- `GNAT.Debug_Pools.Freeing_Deallocated_Storage`

For types associated with a `Debug_Pool`, dynamic allocation is performed using the standard GNAT allocation routine. References to all allocated chunks of memory are kept in an internal dictionary. Several deallocation strategies are provided, whereupon the user can choose to release the memory to the system, keep it allocated for further invalid access checks, or fill it with an easily recognizable pattern for debug sessions. The memory pattern is the old IBM hexadecimal convention: `16#DEADBEEF#`.

See the documentation in the file `g-debpoo.ads` for more information on the various strategies.

Upon each dereference, a check is made that the access value denotes a properly allocated memory location. Here is a complete example of use of `Debug_Pools`, that includes typical instances of memory corruption:

```

with GNAT.IO; use GNAT.IO;
with Ada.Unchecked_Deallocation;
with Ada.Unchecked_Conversion;
with GNAT.Debug_Pools;
with System.Storage_Elements;
with Ada.Exceptions; use Ada.Exceptions;
procedure Debug_Pool_Test is

  type T is access Integer;
  type U is access all T;

  P : GNAT.Debug_Pools.Debug_Pool;
  for T'Storage_Pool use P;

  procedure Free is new Ada.Unchecked_Deallocation (Integer, T);
  function UC is new Ada.Unchecked_Conversion (U, T);
  A, B : aliased T;

  procedure Info is new GNAT.Debug_Pools.Print_Info(Put_Line);

begin
  Info (P);
  A := new Integer;
  B := new Integer;
  B := A;
  Info (P);
  Free (A);
  begin
    Put_Line (Integer'Image(B.all));
  exception
    when E : others => Put_Line ("raised: " & Exception_Name (E));
  end;
  begin
    Free (B);
  exception
    when E : others => Put_Line ("raised: " & Exception_Name (E));
  end;
  B := UC(A'Access);
  begin
    Put_Line (Integer'Image(B.all));
  exception
    when E : others => Put_Line ("raised: " & Exception_Name (E));
  end;
  begin
    Free (B);
  exception
    when E : others => Put_Line ("raised: " & Exception_Name (E));
  end;
  Info (P);
end Debug_Pool_Test;

```

The debug pool mechanism provides the following precise diagnostics on the execution of this erroneous program:

```
Debug Pool info:
  Total allocated bytes : 0
  Total deallocated bytes : 0
  Current Water Mark: 0
  High Water Mark: 0

Debug Pool info:
  Total allocated bytes : 8
  Total deallocated bytes : 0
  Current Water Mark: 8
  High Water Mark: 8

raised: GNAT.DEBUG_POOLS.ACCESSING_DEALLOCATED_STORAGE
raised: GNAT.DEBUG_POOLS.FREEING_DEALLOCATED_STORAGE
raised: GNAT.DEBUG_POOLS.ACCESSING_NOT_ALLOCATED_STORAGE
raised: GNAT.DEBUG_POOLS.FREEING_NOT_ALLOCATED_STORAGE
Debug Pool info:
  Total allocated bytes : 8
  Total deallocated bytes : 4
  Current Water Mark: 4
  High Water Mark: 8
```

This page is intentionally left blank.

PLATFORM-SPECIFIC INFORMATION

This appendix contains information relating to the implementation of run-time libraries on various platforms and also covers topics related to the GNAT implementation on Windows and Mac OS.

A.1 Run-Time Libraries

The GNAT run-time implementation may vary with respect to both the underlying threads library and the exception-handling scheme. For threads support, the default run-time will bind to the thread package of the underlying operating system.

For exception handling, either or both of two models are supplied:

- **Zero-Cost Exceptions** (“ZCX”), which uses binder-generated tables that are interrogated at run time to locate a handler.
- **setjmp / longjmp** (‘SJLJ’), which uses dynamically-set data to establish the set of handlers

Most programs should experience a substantial speed improvement by being compiled with a ZCX run-time. This is especially true for tasking applications or applications with many exception handlers. Note however that the ZCX run-time does not support asynchronous abort of tasks (`abort` and `select-then-abort` constructs) and will instead implement abort by polling points in the runtime. You can also add additional polling points explicitly if needed in your application via `pragma Abort_Defer`.

This section summarizes which combinations of threads and exception support are supplied on various GNAT platforms.

A.1.1 Summary of Run-Time Configurations

Platform	Run-Time	Tasking	Exceptions
GNU/Linux	rts-native (default)	pthread library	ZCX
	rts-sjlj	pthread library	SJLJ
Windows	rts-native (default)	native Win32 threads	ZCX
	rts-sjlj	native Win32 threads	SJLJ
Mac OS	rts-native	pthread library	ZCX


A.2 Specifying a Run-Time Library

The `adainclude` subdirectory containing the sources of the GNAT run-time library, and the `adalib` subdirectory containing the ALI files and the static and/or shared GNAT library, are located in the `gcc` target-dependent area:

```
target=$prefix/lib/gcc/gcc-*dumpmachine*/gcc-*dumpversion*/
```

As indicated above, on some platforms several run-time libraries are supplied. These libraries are installed in the target dependent area and contain a complete source and binary subdirectory. The detailed description below explains the differences between the different libraries in terms of their thread support.

The default run-time library (when GNAT is installed) is *rts-native*. This default run-time is selected by the means of soft links. For example on x86-linux:



Run-Time Library Directory Structure
(italicized names / dashed arrows represent soft links)

If the *rts-sjlj* library is to be selected on a permanent basis, these soft links can be modified with the following commands:

```
$ cd $target
$ rm -f adainclude adalib
```

(continues on next page)

(continued from previous page)

```
$ ln -s rts-sjlj/adainclude adainclude
$ ln -s rts-sjlj/adalib adalib
```

Alternatively, you can specify `rts-sjlj/adainclude` in the file `$target/ada_source_path` and `rts-sjlj/adalib` in `$target/ada_object_path`.

Selecting another run-time library temporarily can be achieved by using the `--RTS` switch, e.g., `--RTS=sjlj`

A.2.1 Choosing the Scheduling Policy

When using a POSIX threads implementation, you have a choice of several scheduling policies: `SCHED_FIFO`, `SCHED_RR` and `SCHED_OTHER`.

Typically, the default is `SCHED_OTHER`, while using `SCHED_FIFO` or `SCHED_RR` requires special (e.g., root) privileges.

By default, GNAT uses the `SCHED_OTHER` policy. To specify `SCHED_FIFO`, you can use one of the following:

- `pragma Time_Slice (0.0)`
- the corresponding binder option `-T0`
- `pragma Task_Dispatching_Policy (FIFO_Within_Priorities)`

To specify `SCHED_RR`, you should use `pragma Time_Slice` with a value greater than 0.0, or else use the corresponding `-T` binder option.

To make sure a program is running as root, you can put something like this in a library package body in your application:

```
function geteuid return Integer;
pragma Import (C, geteuid, "geteuid");
Ignore : constant Boolean :=
  (if geteuid = 0 then True else raise Program_Error with "must be root");
```

It gets the effective user id, and if it's not 0 (i.e. root), it raises `Program_Error`. Note that if you re running the code in a container, this may not be sufficient, as you may have sufficient privilege on the container, but not on the host machine running the container, so check that you also have sufficient privilege for running the container image.

A.3 GNU/Linux Topics

This section describes topics that are specific to GNU/Linux platforms.

A.3.1 Required Packages on GNU/Linux

GNAT requires the C library developer's package to be installed. The name of that package depends on your GNU/Linux distribution:

- RedHat, SUSE: `glibc-devel`;
- Debian, Ubuntu: `libc6-dev` (normally installed by default).

If using the 32-bit version of GNAT on a 64-bit version of GNU/Linux, you'll need the 32-bit version of the following packages:

- RedHat, SUSE: `glibc.i686`, `glibc-devel.i686`, `ncurses-libs.i686`
- SUSE: `glibc-locale-base-32bit`
- Debian, Ubuntu: `libc6:i386`, `libc6-dev:i386`, `lib32ncursesw5`

Other GNU/Linux distributions might be choosing a different name for those packages.

A.3.2 A GNU/Linux Debug Quirk

On SuSE 15, some kernels have a defect causing issues when debugging programs using threads or Ada tasks. Due to the lack of documentation found regarding this kernel issue, we can only provide limited information about which kernels are impacted: kernel version 5.3.18 is known to be impacted, and kernels in the 5.14 range or newer are believed to fix this problem.

The bug affects the debugging of 32-bit processes on a 64-bit system. Symptoms can vary: Unexpected `SIGABRT` signals being received by the program, "The futex facility returned an unexpected error code" error message, and inferior programs hanging indefinitely range among the symptoms most commonly observed.

A.4 Microsoft Windows Topics

This section describes topics that are specific to the Microsoft Windows platforms.

A.4.1 Using GNAT on Windows

One of the strengths of the GNAT technology is that its tool set (`gcc`, `gnatbind`, `gnatlink`, `gnatmake`, the `gdb` debugger, etc.) is used in the same way regardless of the platform.

On Windows this tool set is complemented by a number of Microsoft-specific tools that have been provided to facilitate interoperability with Windows when this is required. With these tools:

- You can build applications using the `CONSOLE` or `WINDOWS` subsystems.
- You can use any Dynamically Linked Library (DLL) in your Ada code (both relocatable and non-relocatable DLLs are supported).
- You can build Ada DLLs for use in other applications. These applications can be written in a language other than Ada (e.g., C, C++, etc). Again both relocatable and non-relocatable Ada DLLs are supported.
- You can include Windows resources in your Ada application.
- You can use or create COM/DCOM objects.

Immediately below are listed all known general GNAT-for-Windows restrictions. Other restrictions about specific features like Windows Resources and DLLs are listed in separate sections below.

- It is not possible to use `GetLastError` and `SetLastError` when tasking, protected records, or exceptions are used. In these cases, in order to implement Ada semantics, the GNAT run-time system calls certain Win32 routines that set the last error variable to 0 upon success. It should be possible to use `GetLastError` and `SetLastError` when tasking, protected record, and exception features are not used, but it is not guaranteed to work.
- It is not possible to link against Microsoft C++ libraries except for import libraries. Interfacing must be done by the mean of DLLs.
- It is possible to link against Microsoft C libraries. Yet the preferred solution is to use C/C++ compiler that comes with GNAT, since it doesn't require having two different development environments and makes the inter-language debugging experience smoother.
- When the compilation environment is located on FAT32 drives, users may experience recompilations of the source files that have not changed if Daylight Saving Time (DST) state has changed since the last time files were compiled. NTFS drives do not have this problem.
- No components of the GNAT toolset use any entries in the Windows registry. The only entries that can be created are file associations and `PATH` settings, provided the user has chosen to create them at installation time, as well as some minimal book-keeping information needed to correctly uninstall or integrate different GNAT products.

A.4.2 Using a network installation of GNAT

Make sure the system on which GNAT is installed is accessible from the current machine, i.e., the install location is shared over the network. Shared resources are accessed on Windows by means of UNC paths, which have the format `\\\\server\\sharename\\path`

In order to use such a network installation, simply add the UNC path of the `bin` directory of your GNAT installation in front of your `PATH`. For example, if GNAT is installed in `\GNAT` directory of a share location called `c-drive` on a machine `LOKI`, the following command will make it available:

```
$ path \\loki\c-drive\gnat\bin;%path%
```

Be aware that every compilation using the network installation results in the transfer of large amounts of data across the network and will likely cause serious performance penalty.

A.4.3 CONSOLE and WINDOWS subsystems

There are two main subsystems under Windows. The `CONSOLE` subsystem (which is the default subsystem) will always create a console when launching the application. This is not something desirable when the application has a Windows GUI. To get rid of this console the application must be using the `WINDOWS` subsystem. To do so the `-mwindows` linker option must be specified.

```
$ gnatmake winprog -largS -mwindows
```

A.4.4 Temporary Files

It is possible to control where temporary files gets created by setting the `TMP` environment variable. The file will be created:

- Under the directory pointed to by the `TMP` environment variable if this directory exists.
- Under `c:\temp`, if the `TMP` environment variable is not set (or not pointing to a directory) and if this directory exists.
- Under the current working directory otherwise.

This allows you to determine exactly where the temporary file will be created. This is particularly useful in networked environments where you may not have write access to some directories.

A.4.5 Disabling Command Line Argument Expansion

By default, an executable compiled for the Windows platform will do the following postprocessing on the arguments passed on the command line:

- If the argument contains the characters `*` and/or `?`, then file expansion will be attempted. For example, if the current directory contains `a.txt` and `b.txt`, then when calling:

```
$ my_ada_program *.txt
```

The following arguments will effectively be passed to the main program (for example when using `Ada.Command_Line.Argument`):

```
Ada.Command_Line.Argument (1) -> "a.txt"
Ada.Command_Line.Argument (2) -> "b.txt"
```

- Filename expansion can be disabled for a given argument by using single quotes. Thus, calling:

```
$ my_ada_program '*.txt'
```

will result in:

```
Ada.Command_Line.Argument (1) -> "*.txt"
```

Note that if the program is launched from a shell such as Cygwin Bash then quote removal might be performed by the shell.

In some contexts it might be useful to disable this feature (for example if the program performs its own argument expansion). In order to do this, a `C` symbol needs to be defined and set to `0`. You can do this by adding the following code fragment in one of your Ada units:

```
Do_Argv_Expansion : Integer := 0;
pragma Export (C, Do_Argv_Expansion, "__gnat_do_argv_expansion");
```

The results of previous examples will be respectively:

```
Ada.Command_Line.Argument (1) -> "*.txt"
```

and:

```
Ada.Command_Line.Argument (1) -> "'.txt'"
```

A.4.6 Windows Socket Timeouts

Microsoft Windows desktops older than 8.0 and Microsoft Windows Servers older than 2019 set a socket timeout 500 milliseconds longer than the value set by `setsockopt` with `SO_RCVTIMEO` and `SO_SNDTIMEO` options. The GNAT runtime makes a correction for the difference in the corresponding Windows versions. For Windows Server starting with version 2019, the user must provide a manifest file for the GNAT runtime to be able to recognize that the Windows version does not need the timeout correction. The manifest file should be located in the same directory as the executable file, and its file name must match the executable name suffixed by `.manifest`. For example, if the executable name is `sock_wto.exe`, then the manifest file name has to be `sock_wto.exe.manifest`. The manifest file must contain at least the following data:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<assembly xmlns="urn:schemas-microsoft-com:asm.v1" manifestVersion="1.0">
<compatibility xmlns="urn:schemas-microsoft-com:compatibility.v1">
<application>
```

(continues on next page)

(continued from previous page)

```

<!-- Windows Vista -->
<supportedOS Id="{e2011457-1546-43c5-a5fe-008deee3d3f0}"/>
<!-- Windows 7 -->
<supportedOS Id="{35138b9a-5d96-4fbd-8e2d-a2440225f93a}"/>
<!-- Windows 8 -->
<supportedOS Id="{4a2f28e3-53b9-4441-ba9c-d69d4a4a6e38}"/>
<!-- Windows 8.1 -->
<supportedOS Id="{1f676c76-80e1-4239-95bb-83d0f6d0da78}"/>
<!-- Windows 10 -->
<supportedOS Id="{8e0f7a12-bfb3-4fe8-b9a5-48fd50a15a9a}"/>
</application>
</compatibility>
</assembly>

```

Without the manifest file, the socket timeout is going to be overcorrected on these Windows Server versions and the actual time is going to be 500 milliseconds shorter than what was set with `GNAT.Sockets.Set_Socket_Option`. Note that on Microsoft Windows versions where correction is necessary, there is no way to set a socket timeout shorter than 500 ms. If a socket timeout shorter than 500 ms is needed on these Windows versions, a call to `Check_Selector` should be added before any socket read or write operations.

A.4.7 Mixed-Language Programming on Windows

Developing pure Ada applications on Windows is no different than on other GNAT-supported platforms. However, when developing or porting an application that contains a mix of Ada and C/C++, the choice of your Windows C/C++ development environment conditions your overall interoperability strategy.

If you use `gcc` or Microsoft C to compile the non-Ada part of your application, there are no Windows-specific restrictions that affect the overall interoperability with your Ada code. If you do want to use the Microsoft tools for your C++ code, you have two choices:

- Encapsulate your C++ code in a DLL to be linked with your Ada application. In this case, use the Microsoft or whatever environment to build the DLL and use GNAT to build your executable ([Using DLLs with GNAT](#)).
- Or you can encapsulate your Ada code in a DLL to be linked with the other part of your application. In this case, use GNAT to build the DLL ([Building DLLs with GNAT Project files](#)) and use the Microsoft or whatever environment to build your executable.

In addition to the description about C main in [Mixed Language Programming](#) section, if the C main uses a stand-alone library it is required on x86-windows to setup the SEH context. For this the C main must look like this:

```

/* main.c */
extern void adainit (void);
extern void adafinal (void);
extern void __gnat_initialize(void*);
extern void call_to_ada (void);

```

(continues on next page)

(continued from previous page)

```
int main (int argc, char *argv[])
{
  int SEH [2];

  /* Initialize the SEH context */
  __gnat_initialize (&SEH);

  adainit();

  /* Then call Ada services in the stand-alone library */

  call_to_ada();

  adafinal();
}
```

Note that this is not needed on x86_64-windows where the Windows native SEH support is used.

Windows Calling Conventions

This section pertains only to Win32. On Win64 there is a single native calling convention. All convention specifiers are ignored on this platform.

When a subprogram *F* (caller) calls a subprogram *G* (callee), there are several ways to push *G*'s parameters on the stack and there are several possible scenarios to clean up the stack upon *G*'s return. A calling convention is an agreed upon software protocol whereby the responsibilities between the caller (*F*) and the callee (*G*) are clearly defined. Several calling conventions are available for Windows:

- C (Microsoft defined)
- Stdcall (Microsoft defined)
- Win32 (GNAT specific)
- DLL (GNAT specific)

c Calling Convention

This is the default calling convention used when interfacing to C/C++ routines compiled with either gcc or Microsoft Visual C++.

In the C calling convention subprogram parameters are pushed on the stack by the caller from right to left. The caller itself is in charge of cleaning up the stack after the call. In addition, the name of a routine with C calling convention is mangled by adding a leading underscore.

The name to use on the Ada side when importing (or exporting) a routine with C calling convention is the name of the routine. For instance the C function:

```
int get_val (long);
```

should be imported from Ada as follows:

```
function Get_Val (V : Interfaces.C.long) return Interfaces.C.int;  
pragma Import (C, Get_Val, External_Name => "get_val");
```

Note that in this particular case the `External_Name` parameter could have been omitted since, when missing, this parameter is taken to be the name of the Ada entity in lower case. When the `Link_Name` parameter is missing, as in the above example, this parameter is set to be the `External_Name` with a leading underscore.

When importing a variable defined in C, you should always use the C calling convention unless the object containing the variable is part of a DLL (in which case you should use the `Stdcall` calling convention, [Stdcall Calling Convention](#)).

Stdcall Calling Convention

This convention, which was the calling convention used for Pascal programs, is used by Microsoft for all the routines in the Win32 API for efficiency reasons. It must be used to import any routine for which this convention was specified.

In the `Stdcall` calling convention subprogram parameters are pushed on the stack by the caller from right to left. The callee (and not the caller) is in charge of cleaning the stack on routine exit. In addition, the name of a routine with `Stdcall` calling convention is mangled by adding a leading underscore (as for the C calling convention) and a trailing `@nn`, where `nn` is the overall size (in bytes) of the parameters passed to the routine.

The name to use on the Ada side when importing a C routine with a `Stdcall` calling convention is the name of the C routine. The leading underscore and trailing `@nn` are added automatically by the compiler. For instance the Win32 function:

```
APIENTRY int get_val (long);
```

should be imported from Ada as follows:

```
function Get_Val (V : Interfaces.C.long) return Interfaces.C.int;  
pragma Import (Stdcall, Get_Val);  
-- On the x86 a long is 4 bytes, so the Link_Name is "_get_val@4"
```

As for the C calling convention, when the `External_Name` parameter is missing, it is taken to be the name of the Ada entity in lower case. If instead of writing the above import pragma you write:

```
function Get_Val (V : Interfaces.C.long) return Interfaces.C.int;  
pragma Import (Stdcall, Get_Val, External_Name => "retrieve_val");
```

then the imported routine is `_retrieve_val@4`. However, if instead of specifying the `External_Name` parameter you specify the `Link_Name` as in the following example:

```
function Get_Val (V : Interfaces.C.long) return Interfaces.C.int;
pragma Import (Stdcall, Get_Val, Link_Name => "retrieve_val");
```

then the imported routine is `retrieve_val`, that is, there is no decoration at all. No leading underscore and no `Stdcall` suffix `@nn`.

This is especially important as in some special cases a DLL's entry point name lacks a trailing `@nn` while the exported name generated for a call has it.

It is also possible to import variables defined in a DLL by using an import pragma for a variable. As an example, if a DLL contains a variable defined as:

```
int my_var;
```

then, to access this variable from Ada you should write:

```
My_Var : Interfaces.C.int;
pragma Import (Stdcall, My_Var);
```

Note that to ease building cross-platform bindings this convention will be handled as a C calling convention on non-Windows platforms.

Win32 Calling Convention

This convention, which is GNAT-specific is fully equivalent to the `Stdcall` calling convention described above.

DLL Calling Convention

This convention, which is GNAT-specific is fully equivalent to the `Stdcall` calling convention described above.

Introduction to Dynamic Link Libraries (DLLs)

A Dynamically Linked Library (DLL) is a library that can be shared by several applications running under Windows. A DLL can contain any number of routines and variables.

One advantage of DLLs is that you can change and enhance them without forcing all the applications that depend on them to be relinked or recompiled. However, you should be aware that all calls to DLL routines are slower since, as you will understand below, such calls are indirect.

To illustrate the remainder of this section, suppose that an application wants to use the services of a DLL `API.dll`. To use the services provided by `API.dll` you must statically link against the DLL or an import library which contains a jump table with an entry for each routine and variable exported by the DLL. In the Microsoft world this import library is called `API.lib`. When using GNAT this import library is called either `libAPI.dll.a`, `libapi.dll.a`, `libAPI.a` or `libapi.a` (names are case insensitive).

After you have linked your application with the DLL or the import library and you run your application, here is what happens:

- Your application is loaded into memory.
- The DLL `API.dll` is mapped into the address space of your application. This means that:
 - The DLL will use the stack of the calling thread.
 - The DLL will use the virtual address space of the calling process.
 - The DLL will allocate memory from the virtual address space of the calling process.
 - Handles (pointers) can be safely exchanged between routines in the DLL routines and routines in the application using the DLL.
- The entries in the jump table (from the import library `libAPI.dll.a` or `API.lib` or automatically created when linking against a DLL) which is part of your application are initialized with the addresses of the routines and variables in `API.dll`.
- If present in `API.dll`, routines `DLLMain` or `DLLMainCRTStartup` are invoked. These routines typically contain the initialization code needed for the well-being of the routines and variables exported by the DLL.

There is an additional point which is worth mentioning. In the Windows world there are two kind of DLLs: relocatable and non-relocatable DLLs. Non-relocatable DLLs can only be loaded at a very specific address in the target application address space. If the addresses of two non-relocatable DLLs overlap and these happen to be used by the same application, a conflict will occur and the application will run incorrectly. Hence, when possible, it is always preferable to use and build relocatable DLLs. Both relocatable and non-relocatable DLLs are supported by GNAT. Note that the `-s` linker option (see GNU Linker User's Guide) removes the debugging symbols from the DLL but the DLL can still be relocated.

As a side note, an interesting difference between Microsoft DLLs and Unix shared libraries, is the fact that on most Unix systems all public routines are exported by default in a Unix shared library, while under Windows it is possible (but not required) to list exported routines in a definition file (see [The Definition File](#)).

Using DLLs with GNAT

To use the services of a DLL, say `API.dll`, in your Ada application you must have:

- The Ada spec for the routines and/or variables you want to access in `API.dll`. If not available this Ada spec must be built from the C/C++ header files provided with the DLL.
- The import library (`libAPI.dll.a` or `API.lib`). As previously mentioned an import library is a statically linked library containing the import table which will be filled at load time to point to the actual `API.dll` routines. Sometimes you don't have an import library for the DLL you want to use. The following sections will explain how to build one. Note that this is optional.
- The actual DLL, `API.dll`.

Once you have all the above, to compile an Ada application that uses the services of `API.dll` and whose main subprogram is `My_Ada_App`, you simply issue the command

```
$ gnatmake my_ada_app -largS -lAPI
```

The argument `-largS -lAPI` at the end of the `gnatmake` command tells the GNAT linker to look for an import library. The linker will look for a library name in this specific order:

- `libAPI.dll.a`
- `API.dll.a`
- `libAPI.a`
- `API.lib`
- `libAPI.dll`
- `API.dll`

The first three are the GNU style import libraries. The third is the Microsoft style import libraries. The last two are the actual DLL names.

Note that if the Ada package spec for `API.dll` contains the following pragma

```
pragma Linker_Options ("-lAPI");
```

you do not have to add `-largS -lAPI` at the end of the `gnatmake` command.

If any one of the items above is missing you will have to create it yourself. The following sections explain how to do so using as an example a fictitious DLL called `API.dll`.

Creating an Ada Spec for the DLL Services

A DLL typically comes with a C/C++ header file which provides the definitions of the routines and variables exported by the DLL. The Ada equivalent of this header file is a package spec that contains definitions for the imported entities. If the DLL you intend to use does not come with an Ada spec you have to generate one such spec yourself. For example if the header file of `API.dll` is a file `api.h` containing the following two definitions:

```
int some_var;
int get (char *);
```

then the equivalent Ada spec could be:

```
with Interfaces.C.Strings;
package API is
  use Interfaces;

  Some_Var : C.int;
  function Get (Str : C.Strings.Chars_Ptr) return C.int;

private
```

(continues on next page)

(continued from previous page)

```
pragma Import (C, Get);  
pragma Import (DLL, Some_Var);  
end API;
```

Creating an Import Library

If a Microsoft-style import library `API.lib` or a GNAT-style import library `libAPI.dll.a` or `libAPI.a` is available with `API.dll` you can skip this section. You can also skip this section if `API.dll` or `libAPI.dll` is built with GNU tools as in this case it is possible to link directly against the DLL. Otherwise read on.

The Definition File

As previously mentioned, and unlike Unix systems, the list of symbols that are exported from a DLL must be provided explicitly in Windows. The main goal of a definition file is precisely that: list the symbols exported by a DLL. A definition file (usually a file with a `.def` suffix) has the following structure:

```
[LIBRARY ``name``]  
[DESCRIPTION ``string``]  
EXPORTS  
  ``symbol1``  
  ``symbol2``  
  ...
```

LIBRARY name

This section, which is optional, gives the name of the DLL.

DESCRIPTION string

This section, which is optional, gives a description string that will be embedded in the import library.

EXPORTS

This section gives the list of exported symbols (procedures, functions or variables). For instance in the case of `API.dll` the `EXPORTS` section of `API.def` looks like:

```
EXPORTS  
  some_var  
  get
```

Note that you must specify the correct suffix (`@nn`) (see [Windows Calling Conventions](#)) for a Stdcall calling convention function in the exported symbols list.

There can actually be other sections in a definition file, but these sections are not relevant to the discussion at hand.

Creating a Definition File Automatically

You can automatically create the definition file `API.def` (see [The Definition File](#)) from a DLL. For that use the `dlltool` program as follows:

```
$ dlltool API.dll -z API.def --export-all-symbols
```

Note that if some routines in the DLL have the `Stdcall` convention ([Windows Calling Conventions](#)) with stripped `@nn` suffix then you'll have to edit `api.def` to add it, and specify `-k` to `gnatdll` when creating the import library.

Here are some hints to find the right `@nn` suffix.

- If you have the Microsoft import library (`.lib`), it is possible to get the right symbols by using Microsoft `dumpbin` tool (see the corresponding Microsoft documentation for further details).

```
$ dumpbin /exports api.lib
```

- If you have a message about a missing symbol at link time the compiler tells you what symbol is expected. You just have to go back to the definition file and add the right suffix.

GNAT-Style Import Library

To create a static import library from `API.dll` with the GNAT tools you should create the `.def` file, then use `gnatdll` tool (see [Using gnatdll](#)) as follows:

```
$ gnatdll -e API.def -d API.dll
```

`gnatdll` takes as input a definition file `API.def` and the name of the DLL containing the services listed in the definition file `API.dll`. The name of the static import library generated is computed from the name of the definition file as follows: if the definition file name is `xyz.def`, the import library name will be `libxyz.a`. Note that in the previous example option `-e` could have been removed because the name of the definition file (before the `.def` suffix) is the same as the name of the DLL ([Using gnatdll](#) for more information about `gnatdll`).

Microsoft-Style Import Library

A Microsoft import library is needed only if you plan to make an Ada DLL available to applications developed with Microsoft tools ([Mixed-Language Programming on Windows](#)).

To create a Microsoft-style import library for `API.dll` you should create the `.def` file, then build the actual import library using Microsoft's `lib` utility:

```
$ lib -machine:IX86 -def:API.def -out:API.lib
```

If you use the above command the definition file `API.def` must contain a line giving the name of the DLL:

```
LIBRARY      "API"
```

See the Microsoft documentation for further details about the usage of `lib`.

Building DLLs with GNAT Project files

There is nothing specific to Windows in the build process. See the *Library Projects* section in the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*.

Due to a system limitation, it is not possible under Windows to create threads when inside the `DLLMain` routine which is used for auto-initialization of shared libraries, so it is not possible to have library level tasks in SALs.

Building DLLs with GNAT

This section explain how to build DLLs using the GNAT built-in DLL support. With the following procedure it is straight forward to build and use DLLs with GNAT.

- Building object files. The first step is to build all objects files that are to be included into the DLL. This is done by using the standard `gnatmake` tool.
- Building the DLL. To build the DLL you must use the `gcc -shared` and `-shared-libgcc` options. It is quite simple to use this method:

```
$ gcc -shared -shared-libgcc -o api.dll obj1.o obj2.o ...
```

It is important to note that in this case all symbols found in the object files are automatically exported. It is possible to restrict the set of symbols to export by passing to `gcc` a definition file (see [The Definition File](#)). For example:

```
$ gcc -shared -shared-libgcc -o api.dll api.def obj1.o obj2.o ...
```

If you use a definition file you must export the elaboration procedures for every package that required one. Elaboration procedures are named using the package name followed by “`_E`”.

- Preparing DLL to be used. For the DLL to be used by client programs the bodies must be hidden from it and the `.ali` set with read-only attribute. This is very important otherwise GNAT will recompile all packages and will not actually use the code in the DLL. For example:

```
$ mkdir apilib
$ copy *.ads *.ali api.dll apilib
$ attrib +R apilib\*.ali
```

At this point it is possible to use the DLL by directly linking against it. Note that you must use the GNAT shared runtime when using GNAT shared libraries. This is achieved by using the `-shared` binder option.

```
$ gnatmake main -Iapilib -bargs -shared -largS -Lapilib -lAPI
```

Building DLLs with gnatdll

Note that it is preferred to use GNAT Project files ([Building DLLs with GNAT Project files](#)) or the built-in GNAT DLL support ([Building DLLs with GNAT](#)) or to build DLLs.

This section explains how to build DLLs containing Ada code using `gnatdll`. These DLLs will be referred to as Ada DLLs in the remainder of this section.

The steps required to build an Ada DLL that is to be used by Ada as well as non-Ada applications are as follows:

- You need to mark each Ada entity exported by the DLL with a `C` or `Stdcall` calling convention to avoid any Ada name mangling for the entities exported by the DLL (see [Exporting Ada Entities](#)). You can skip this step if you plan to use the Ada DLL only from Ada applications.
- Your Ada code must export an initialization routine which calls the routine `adainit` generated by `gnatbind` to perform the elaboration of the Ada code in the DLL ([Ada DLLs and Elaboration](#)). The initialization routine exported by the Ada DLL must be invoked by the clients of the DLL to initialize the DLL.
- When useful, the DLL should also export a finalization routine which calls routine `adafinal` generated by `gnatbind` to perform the finalization of the Ada code in the DLL ([Ada DLLs and Finalization](#)). The finalization routine exported by the Ada DLL must be invoked by the clients of the DLL when the DLL services are no further needed.
- You must provide a spec for the services exported by the Ada DLL in each of the programming languages to which you plan to make the DLL available.
- You must provide a definition file listing the exported entities ([The Definition File](#)).
- Finally you must use `gnatdll` to produce the DLL and the import library ([Using gnatdll](#)).

Note that a relocatable DLL stripped using the `strip` binutils tool will not be relocatable anymore. To build a DLL without debug information pass `-largS -s` to `gnatdll`. This restriction does not apply to a DLL built using a Library Project. See the *Library Projects* section in the *GNAT Project Manager* chapter of the *GPRbuild User's Guide*.

Limitations When Using Ada DLLs from Ada

When using Ada DLLs from Ada applications there is a limitation users should be aware of. Because on Windows the GNAT run-time is not in a DLL of its own, each Ada DLL includes a part of the GNAT run-time. Specifically, each Ada DLL includes the services of the GNAT run-time that are necessary to the Ada code inside the DLL. As a result, when an Ada program uses an Ada DLL there are two independent GNAT run-times: one in the Ada DLL and one in the main program.

It is therefore not possible to exchange GNAT run-time objects between the Ada DLL and the main Ada program. Example of GNAT run-time objects are file handles (e.g., `Text_IO.File_Type`), tasks types, protected objects types, etc.

It is completely safe to exchange plain elementary, array or record types, Windows object handles, etc.

Exporting Ada Entities

Building a DLL is a way to encapsulate a set of services usable from any application. As a result, the Ada entities exported by a DLL should be exported with the C or `Stdcall` calling conventions to avoid any Ada name mangling. As an example here is an Ada package `API`, spec and body, exporting two procedures, a function, and a variable:

```
with Interfaces.C; use Interfaces;
package API is
  Count : C.int := 0;
  function Factorial (Val : C.int) return C.int;

  procedure Initialize_API;
  procedure Finalize_API;
  -- Initialization & Finalization routines. More in the next section.
private
  pragma Export (C, Initialize_API);
  pragma Export (C, Finalize_API);
  pragma Export (C, Count);
  pragma Export (C, Factorial);
end API;
```

```
package body API is
  function Factorial (Val : C.int) return C.int is
    Fact : C.int := 1;
  begin
    Count := Count + 1;
    for K in 1 .. Val loop
      Fact := Fact * K;
    end loop;
    return Fact;
  end Factorial;

  procedure Initialize_API is
    procedure Adainit;
    pragma Import (C, Adainit);
  begin
    Adainit;
  end Initialize_API;

  procedure Finalize_API is
    procedure Adafinal;
    pragma Import (C, Adafinal);
  begin
    Adafinal;
  end Finalize_API;
end API;
```

If the Ada DLL you are building will only be used by Ada applications you do not have to export Ada entities with a `C` or `Stdcall` convention. As an example, the previous package could be written as follows:

```
package API is
  Count : Integer := 0;
  function Factorial (Val : Integer) return Integer;

  procedure Initialize_API;
  procedure Finalize_API;
  -- Initialization and Finalization routines.
end API;
```

```
package body API is
  function Factorial (Val : Integer) return Integer is
    Fact : Integer := 1;
  begin
    Count := Count + 1;
    for K in 1 .. Val loop
      Fact := Fact * K;
    end loop;
    return Fact;
  end Factorial;

  ...
  -- The remainder of this package body is unchanged.
end API;
```

Note that if you do not export the Ada entities with a `C` or `Stdcall` convention you will have to provide the mangled Ada names in the definition file of the Ada DLL ([Creating the Definition File](#)).

Ada DLLs and Elaboration

The DLL that you are building contains your Ada code as well as all the routines in the Ada library that are needed by it. The first thing a user of your DLL must do is elaborate the Ada code ([Elaboration Order Handling in GNAT](#)).

To achieve this you must export an initialization routine (`Initialize_Api` in the previous example), which must be invoked before using any of the DLL services. This elaboration routine must call the Ada elaboration routine `adainit` generated by the GNAT binder ([Binding with Non-Ada Main Programs](#)). See the body of `Initialize_Api` for an example. Note that the GNAT binder is automatically invoked during the DLL build process by the `gnatdll` tool ([Using gnatdll](#)).

When a DLL is loaded, Windows systematically invokes a routine called `DllMain`. It would therefore be possible to call `adainit` directly from `DllMain` without having to provide an explicit initialization routine. Unfortunately, it is not possible to call `adainit` from the `DllMain` if your program has library level tasks because access to the `DllMain` entry point is serialized by the system (that is, only a single thread can execute 'through' it at a time), which means that the GNAT run-time will deadlock waiting for the newly created task to complete its initialization.

Ada DLLs and Finalization

When the services of an Ada DLL are no longer needed, the client code should invoke the DLL finalization routine, if available. The DLL finalization routine is in charge of releasing all resources acquired by the DLL. In the case of the Ada code contained in the DLL, this is achieved by calling routine `adafinal` generated by the GNAT binder ([Binding with Non-Ada Main Programs](#)). See the body of `Finalize_Api` for an example. As already pointed out the GNAT binder is automatically invoked during the DLL build process by the `gnatdll` tool ([Using gnatdll](#)).

Creating a Spec for Ada DLLs

To use the services exported by the Ada DLL from another programming language (e.g., C), you have to translate the specs of the exported Ada entities in that language. For instance in the case of `API.dll`, the corresponding C header file could look like:

```
extern int *_imp_count;
#define count (*_imp_count)
int factorial (int);
```

It is important to understand that when building an Ada DLL to be used by other Ada applications, you need two different specs for the packages contained in the DLL: one for building the DLL and the other for using the DLL. This is because the DLL calling convention is needed to use a variable defined in a DLL, but when building the DLL, the variable must have either the Ada or C calling convention. As an example consider a DLL comprising the following package `API`:

```
package API is
  Count : Integer := 0;
  ...
  -- Remainder of the package omitted.
end API;
```

After producing a DLL containing package `API`, the spec that must be used to import `API.Count` from Ada code outside of the DLL is:

```
package API is
  Count : Integer;
  pragma Import (DLL, Count);
end API;
```

Creating the Definition File

The definition file is the last file needed to build the DLL. It lists the exported symbols. As an example, the definition file for a DLL containing only package `API` (where all the entities are exported with a C calling convention) is:

```
EXPORTS
  count
```

(continues on next page)

(continued from previous page)

```
factorial
finalize_api
initialize_api
```

If the C calling convention is missing from package `API`, then the definition file contains the mangled Ada names of the above entities, which in this case are:

```
EXPORTS
  api__count
  api__factorial
  api__finalize_api
  api__initialize_api
```

Using `gnatdll`

`gnatdll` is a tool to automate the DLL build process once all the Ada and non-Ada sources that make up your DLL have been compiled. `gnatdll` is actually in charge of two distinct tasks: build the static import library for the DLL and the actual DLL. The form of the `gnatdll` command is

```
$ gnatdll [ switches ] list-of-files [ -larges opts ]
```

where `list-of-files` is a list of ALI and object files. The object file list must be the exact list of objects corresponding to the non-Ada sources whose services are to be included in the DLL. The ALI file list must be the exact list of ALI files for the corresponding Ada sources whose services are to be included in the DLL. If `list-of-files` is missing, only the static import library is generated.

You may specify any of the following switches to `gnatdll`:

-a*address*

Build a non-relocatable DLL at *address*. If *address* is not specified the default address `0x11000000` will be used. By default, when this switch is missing, `gnatdll` builds relocatable DLL. We advise the reader to build relocatable DLL.

-b *address*

Set the relocatable DLL base address. By default the address is `0x11000000`.

-bargs *opts*

Binder options. Pass *opts* to the binder.

-d *dllfile*

dllfile is the name of the DLL. This switch must be present for `gnatdll` to do anything. The name of the generated import library is obtained algorithmically from *dllfile* as shown in the following example: if *dllfile* is `xyz.dll`, the import library name is `libxyz.dll.a`. The name of the definition file to use (if not specified by option `-e`) is obtained algorithmically from *dllfile* as shown in the following example: if *dllfile* is `xyz.dll`, the definition file used is `xyz.def`.

-e *deffile*

deffile is the name of the definition file.

-g

Generate debugging information. This information is stored in the object file and copied from there to the final DLL file by the linker, where it can be read by the debugger. You must use the **-g** switch if you plan on using the debugger or the symbolic stack traceback.

-h

Help mode. Displays `gnatdll` switch usage information.

-Idir

Direct `gnatdll` to search the `dir` directory for source and object files needed to build the DLL. (Search Paths and the Run-Time Library (RTL)).

-k

Removes the `@nn` suffix from the import library's exported names, but keeps them for the link names. You must specify this option if you want to use a `Stdcall` function in a DLL for which the `@nn` suffix has been removed. This is the case for most of the Windows NT DLL for example. This option has no effect when **-n** option is specified.

-l *file*

The list of ALI and object files used to build the DLL are listed in `file`, instead of being given in the command line. Each line in `file` contains the name of an ALI or object file.

-n

No Import. Do not create the import library.

-q

Quiet mode. Do not display unnecessary messages.

-v

Verbose mode. Display extra information.

-larges *opts*

Linker options. Pass `opts` to the linker.

gnatdll Example

As an example the command to build a relocatable DLL from `api.adb` once `api.adb` has been compiled and `api.def` created is

```
$ gnatdll -d api.dll api.ali
```

The above command creates two files: `libapi.dll.a` (the import library) and `api.dll` (the actual DLL). If you want to create only the DLL, just type:

```
$ gnatdll -d api.dll -n api.ali
```

Alternatively if you want to create just the import library, type:

```
$ gnatdll -d api.dll
```

gnatdll behind the Scenes

This section details the steps involved in creating a DLL. `gnatdll` does these steps for you. Unless you are interested in understanding what goes on behind the scenes, you should skip this section.

We use the previous example of a DLL containing the Ada package `API`, to illustrate the steps necessary to build a DLL. The starting point is a set of objects that will make up the DLL and the corresponding ALI files. In the case of this example this means that `api.o` and `api.ali` are available. To build a relocatable DLL, `gnatdll` does the following:

- `gnatdll` builds the base file (`api.base`). A base file gives the information necessary to generate relocation information for the DLL.

```
$ gnatbind -n api
$ gnatlink api -o api.jnk -mdll -Wl,--base-file,api.base
```

In addition to the base file, the `gnatlink` command generates an output file `api.jnk` which can be discarded. The `-mdll` switch asks `gnatlink` to generate the routines `DllMain` and `DllMainCRTStartup` that are called by the Windows loader when the DLL is loaded into memory.

- `gnatdll` uses `dlltool` (see [Using dlltool](#)) to build the export table (`api.exp`). The export table contains the relocation information in a form which can be used during the final link to ensure that the Windows loader is able to place the DLL anywhere in memory.

```
$ dlltool --dllname api.dll --def api.def --base-file api.base \
--output-exp api.exp
```

- `gnatdll` builds the base file using the new export table. Note that `gnatbind` must be called once again since the binder generated file has been deleted during the previous call to `gnatlink`.

```
$ gnatbind -n api
$ gnatlink api -o api.jnk api.exp -mdll
-Wl,--base-file,api.base
```

- `gnatdll` builds the new export table using the new base file and generates the DLL import library `libAPI.dll.a`.

```
$ dlltool --dllname api.dll --def api.def --base-file api.base \
--output-exp api.exp --output-lib libAPI.a
```

- Finally `gnatdll` builds the relocatable DLL using the final export table.

```
$ gnatbind -n api
$ gnatlink api api.exp -o api.dll -mdll
```

Using `dlltool`

`dlltool` is the low-level tool used by `gnatdll` to build DLLs and static import libraries. This section summarizes the most common `dlltool` switches. The form of the `dlltool` command is

```
$ dlltool [`switches`]
```

`dlltool` switches include:

--base-file *basefile*

Read the base file *basefile* generated by the linker. This switch is used to create a relocatable DLL.

--def *deffile*

Read the definition file.

--dllname *name*

Gives the name of the DLL. This switch is used to embed the name of the DLL in the static import library generated by `dlltool` with switch `--output-lib`.

-k

Kill *@nn* from exported names ([Windows Calling Conventions](#) for a discussion about `Stdcall`-style symbols).

--help

Prints the `dlltool` switches with a concise description.

--output-exp *exportfile*

Generate an export file *exportfile*. The export file contains the export table (list of symbols in the DLL) and is used to create the DLL.

--output-lib *libfile*

Generate a static import library *libfile*.

-v

Verbose mode.

--as *assembler-name*

Use *assembler-name* as the assembler. The default is `as`.

GNAT and Windows Resources

Resources are an easy way to add Windows specific objects to your application. The objects that can be added as resources include:

- menus
- accelerators
- dialog boxes
- string tables
- bitmaps

- cursors
- icons
- fonts
- version information

For example, a version information resource can be defined as follow and embedded into an executable or DLL:

A version information resource can be used to embed information into an executable or a DLL. These information can be viewed using the file properties from the Windows Explorer. Here is an example of a version information resource:

```
1 VERSIONINFO
FILEVERSION      1,0,0,0
PRODUCTVERSION  1,0,0,0
BEGIN
  BLOCK "StringFileInfo"
  BEGIN
    BLOCK "080904E4"
    BEGIN
      VALUE "CompanyName", "My Company Name"
      VALUE "FileDescription", "My application"
      VALUE "FileVersion", "1.0"
      VALUE "InternalName", "my_app"
      VALUE "LegalCopyright", "My Name"
      VALUE "OriginalFilename", "my_app.exe"
      VALUE "ProductName", "My App"
      VALUE "ProductVersion", "1.0"
    END
  END
END

  BLOCK "VarFileInfo"
  BEGIN
    VALUE "Translation", 0x809, 1252
  END
END
```

The value 0809 (langID) is for the U.K English language and 04E4 (charsetID), which is equal to 1252 decimal, for multilingual.

This section explains how to build, compile and use resources. Note that this section does not cover all resource objects, for a complete description see the corresponding Microsoft documentation.

Building Resources

A resource file is an ASCII file. By convention resource files have an `.rc` extension. The easiest way to build a resource file is to use Microsoft tools such as `imagedit.exe` to build bitmaps, icons and cursors and `dlgedit.exe` to build dialogs. It is always possible to build an `.rc` file yourself by writing a resource script.

It is not our objective to explain how to write a resource file. A complete description of the resource script language can be found in the Microsoft documentation.

Compiling Resources

This section describes how to build a GNAT-compatible (COFF) object file containing the resources. This is done using the Resource Compiler `windres` as follows:

```
$ windres -i myres.rc -o myres.o
```

By default `windres` will run `gcc` to preprocess the `.rc` file. You can specify an alternate preprocessor (usually named `cpp.exe`) using the `windres --preprocessor` parameter. A list of all possible options may be obtained by entering the command `windres --help`.

It is also possible to use the Microsoft resource compiler `rc.exe` to produce a `.res` file (binary resource file). See the corresponding Microsoft documentation for further details. In this case you need to use `windres` to translate the `.res` file to a GNAT-compatible object file as follows:

```
$ windres -i myres.res -o myres.o
```

Using Resources

To include the resource file in your program just add the GNAT-compatible object file for the resource(s) to the linker arguments. With `gnatmake` this is done by using the `-larg`s option:

```
$ gnatmake myprog -larg myres.o
```

Using GNAT DLLs from Microsoft Visual Studio Applications

This section describes a common case of mixed GNAT/Microsoft Visual Studio application development, where the main program is developed using MSVS, and is linked with a DLL developed using GNAT. Such a mixed application should be developed following the general guidelines outlined above; below is the cookbook-style sequence of steps to follow:

1. First develop and build the GNAT shared library using a library project (let's assume the project is `mylib.gpr`, producing the library `libmylib.dll`):

```
$ gprbuild -p mylib.gpr
```

2. Produce a .def file for the symbols you need to interface with, either by hand or automatically with possibly some manual adjustments (see [Creating Definition File Automatically](#)):

```
$ dlltool libmylib.dll -z libmylib.def --export-all-symbols
```

3. Make sure that MSVS command-line tools are accessible on the path.
4. Create the Microsoft-style import library (see [MSVS-Style Import Library](#)):

```
$ lib -machine:IX86 -def:libmylib.def -out:libmylib.lib
```

If you are using a 64-bit toolchain, the above becomes...

```
$ lib -machine:X64 -def:libmylib.def -out:libmylib.lib
```

5. Build the C main

```
$ cl /O2 /MD main.c libmylib.lib
```

6. Before running the executable, make sure you have set the PATH to the DLL, or copy the DLL into into the directory containing the .exe.

Debugging a DLL

Debugging a DLL is similar to debugging a standard program. But we have to deal with two different executable parts: the DLL and the program that uses it. We have the following four possibilities:

- The program and the DLL are built with GCC/GNAT.
- The program is built with foreign tools and the DLL is built with GCC/GNAT.
- The program is built with GCC/GNAT and the DLL is built with foreign tools.

In this section we address only cases one and two above. There is no point in trying to debug a DLL with GNU/GDB, if there is no GDB-compatible debugging information in it. To do so you must use a debugger compatible with the tools suite used to build the DLL.

Program and DLL Both Built with GCC/GNAT

This is the simplest case. Both the DLL and the program have GDB compatible debugging information. It is then possible to break anywhere in the process. Let's suppose here that the main procedure is named `ada_main` and that in the DLL there is an entry point named `ada_dll`.

The DLL ([Introduction to Dynamic Link Libraries \(DLLs\)](#)) and program must have been built with the debugging information (see GNAT `-g` switch). Here are the step-by-step instructions for debugging it:

- Launch GDB on the main program.

```
$ gdb -nw ada_main
```

- Start the program and stop at the beginning of the main procedure

```
(gdb) start
```

This step is required to be able to set a breakpoint inside the DLL. As long as the program is not run, the DLL is not loaded. This has the consequence that the DLL debugging information is also not loaded, so it is not possible to set a breakpoint in the DLL.

- Set a breakpoint inside the DLL

```
(gdb) break ada_dll  
(gdb) cont
```

At this stage a breakpoint is set inside the DLL. From there on you can use the standard approach to debug the whole program ([Running and Debugging Ada Programs](#)).

Program Built with Foreign Tools and DLL Built with GCC/GNAT

In this case things are slightly more complex because it is not possible to start the main program and then break at the beginning to load the DLL and the associated DLL debugging information. It is not possible to break at the beginning of the program because there is no GDB debugging information, and therefore there is no direct way of getting initial control. This section addresses this issue by describing some methods that can be used to break somewhere in the DLL to debug it.

First suppose that the main procedure is named `main` (this is for example some C code built with Microsoft Visual C) and that there is a DLL named `test.dll` containing an Ada entry point named `ada_dll`.

The DLL (see [Introduction to Dynamic Link Libraries \(DLLs\)](#)) must have been built with debugging information (see the GNAT `-g` option).

Debugging the DLL Directly

- Find out the executable starting address

```
$ objdump --file-header main.exe
```

The starting address is reported on the last line. For example:

```
main.exe:      file format pei-i386  
architecture: i386, flags 0x0000010a:  
EXEC_P, HAS_DEBUG, D_PAGED  
start address 0x00401010
```

- Launch the debugger on the executable.

```
$ gdb main.exe
```

- Set a breakpoint at the starting address, and launch the program.

```
$ (gdb) break *0x00401010
$ (gdb) run
```

The program will stop at the given address.

- Set a breakpoint on a DLL subroutine.

```
(gdb) break ada_dll.adb:45
```

Or if you want to break using a symbol on the DLL, you need first to select the Ada language (language used by the DLL).

```
(gdb) set language ada
(gdb) break ada_dll
```

- Continue the program.

```
(gdb) cont
```

This will run the program until it reaches the breakpoint that has been set. From that point you can use the standard way to debug a program as described in ([Running and Debugging Ada Programs](#)).

It is also possible to debug the DLL by attaching to a running process.

Attaching to a Running Process

With GDB it is always possible to debug a running process by attaching to it. It is possible to debug a DLL this way. The limitation of this approach is that the DLL must run long enough to perform the attach operation. It may be useful for instance to insert a time wasting loop in the code of the DLL to meet this criterion.

- Launch the main program `main.exe`.

```
$ main
```

- Use the Windows *Task Manager* to find the process ID. Let's say that the process PID for `main.exe` is 208.
- Launch `gdb`.

```
$ gdb
```

- Attach to the running process to be debugged.

```
(gdb) attach 208
```

- Load the process debugging information.

```
(gdb) symbol-file main.exe
```

- Break somewhere in the DLL.

```
(gdb) break ada_dll
```

- Continue process execution.

```
(gdb) cont
```

This last step will resume the process execution, and stop at the breakpoint we have set. From there you can use the standard approach to debug a program as described in [Running and Debugging Ada Programs](#).

Setting Stack Size from `gnatlink`

It is possible to specify the program stack size at link time. On modern versions of Windows, starting with XP, this is mostly useful to set the size of the main stack (environment task). The other task stacks are set with `pragma Storage_Size` or with the `gnatbind -d` command.

Since older versions of Windows (2000, NT4, etc.) do not allow setting the reserve size of individual tasks, the link-time stack size applies to all tasks, and `pragma Storage_Size` has no effect. In particular, Stack Overflow checks are made against this link-time specified size.

This setting can be done with `gnatlink` using either of the following:

- `-Xlinker` linker option

```
$ gnatlink hello -Xlinker --stack=0x10000,0x1000
```

This sets the stack reserve size to 0x10000 bytes and the stack commit size to 0x1000 bytes.

- `-Wl` linker option

```
$ gnatlink hello -Wl,--stack=0x1000000
```

This sets the stack reserve size to 0x1000000 bytes. Note that with `-Wl` option it is not possible to set the stack commit size because the comma is a separator for this option.

Setting Heap Size from `gnatlink`

Under Windows systems, it is possible to specify the program heap size from `gnatlink` using either of the following:

- `-Xlinker` linker option

```
$ gnatlink hello -Xlinker --heap=0x10000,0x1000
```

This sets the heap reserve size to 0x10000 bytes and the heap commit size to 0x1000 bytes.

- `-Wl` linker option

```
$ gnatlink hello -Wl,--heap=0x1000000
```

This sets the heap reserve size to 0x1000000 bytes. Note that with `-Wl` option it is not possible to set the heap commit size because the comma is a separator for this option.

A.4.8 Windows Specific Add-Ons

This section describes the Windows specific add-ons.

Win32Ada

Win32Ada is a binding for the Microsoft Win32 API. This binding can be easily installed from the provided installer. To use the Win32Ada binding you need to use a project file, and adding a single `with_` clause will give you full access to the Win32Ada binding sources and ensure that the proper libraries are passed to the linker.

```
with "win32ada";
project P is
  for Sources use ...;
end P;
```

To build the application you just need to call `gprbuild` for the application's project, here `p.gpr`:

```
gprbuild p.gpr
```

wPOSIX

wPOSIX is a minimal POSIX binding whose goal is to help with building cross-platforms applications. This binding is not complete though, as the Win32 API does not provide the necessary support for all POSIX APIs.

To use the wPOSIX binding you need to use a project file, and adding a single `with_` clause will give you full access to the wPOSIX binding sources and ensure that the proper libraries are passed to the linker.

```
with "wposix";
project P is
  for Sources use ...;
end P;
```

To build the application you just need to call `gprbuild` for the application's project, here `p.gpr`:

```
gprbuild p.gpr
```

A.5 Mac OS Topics

This section describes topics that are specific to Apple's OS X platform.

A.5.1 Codesigning the Debugger

The Darwin Kernel requires the debugger to have special permissions before it is allowed to control other processes. These permissions are granted by codesigning the GDB executable. Without these permissions, the debugger will report error messages such as:

```
Starting program: /x/y/foo
Unable to find Mach task port for process-id 28885: (os/kern) failure (0x5).
(please check gdb is codesigned - see taskgated(8))
```

Codesigning requires a certificate. The following procedure explains how to create one:

- Start the Keychain Access application (in /Applications/Utilities/Keychain Access.app)
- Select the Keychain Access -> Certificate Assistant -> Create a Certificate... menu
- Then:
 - Choose a name for the new certificate (this procedure will use “gdb-cert” as an example)
 - Set “Identity Type” to “Self Signed Root”
 - Set “Certificate Type” to “Code Signing”
 - Activate the “Let me override defaults” option
- Click several times on “Continue” until the “Specify a Location For The Certificate” screen appears, then set “Keychain” to “System”
- Click on “Continue” until the certificate is created
- Finally, in the view, double-click on the new certificate, and set “When using this certificate” to “Always Trust”
- Exit the Keychain Access application and restart the computer (this is unfortunately required)

Once a certificate has been created, the debugger can be codesigned as follow. In a Terminal, run the following command:

```
$ codesign -f -s "gdb-cert" <gnat_install_prefix>/bin/gdb
```

where “gdb-cert” should be replaced by the actual certificate name chosen above, and <gnat_install_prefix> should be replaced by the location where you installed GNAT. Also, be sure that users are in the Unix group `_developer`.

EXAMPLE OF BINDER OUTPUT FILE

This Appendix displays the source code for the output file generated by *gnatbind* for a simple 'Hello World' program. Comments have been added for clarification purposes.

```
-- The package is called Ada_Main unless this name is actually used
-- as a unit name in the partition, in which case some other unique
-- name is used.

pragma Ada_95;
with System;
package ada_main is
  pragma Warnings (Off);

  -- The main program saves the parameters (argument count,
  -- argument values, environment pointer) in global variables
  -- for later access by other units including
  -- Ada.Command_Line.

  gnat_argc : Integer;
  gnat_argv : System.Address;
  gnat_envp : System.Address;

  -- The actual variables are stored in a library routine. This
  -- is useful for some shared library situations, where there
  -- are problems if variables are not in the library.

  pragma Import (C, gnat_argc);
  pragma Import (C, gnat_argv);
  pragma Import (C, gnat_envp);

  -- The exit status is similarly an external location

  gnat_exit_status : Integer;
  pragma Import (C, gnat_exit_status);

  GNAT_Version : constant String :=
    "GNAT Version: Pro 7.4.0w (20141119-49)" & ASCII.NUL;
  pragma Export (C, GNAT_Version, "__gnat_version");

  Ada_Main_Program_Name : constant String := "_ada_hello" & ASCII.NUL;
```

(continues on next page)

(continued from previous page)

```
pragma Export (C, Ada_Main_Program_Name, "__gnat_ada_main_program_name");
```

```
-- This is the generated adainit routine that performs  
-- initialization at the start of execution. In the case  
-- where Ada is the main program, this main program makes  
-- a call to adainit at program startup.
```

```
procedure adainit;  
pragma Export (C, adainit, "adainit");
```

```
-- This is the generated adafinal routine that performs  
-- finalization at the end of execution. In the case where  
-- Ada is the main program, this main program makes a call  
-- to adafinal at program termination.
```

```
procedure adafinal;  
pragma Export (C, adafinal, "adafinal");
```

```
-- This routine is called at the start of execution. It is  
-- a dummy routine that is used by the debugger to breakpoint  
-- at the start of execution.
```

```
-- This is the actual generated main program (it would be  
-- suppressed if the no main program switch were used). As  
-- required by standard system conventions, this program has  
-- the external name main.
```

```
function main  
  (argc : Integer;  
   argv : System.Address;  
   envp : System.Address)  
  return Integer;  
pragma Export (C, main, "main");
```

```
-- The following set of constants give the version  
-- identification values for every unit in the bound  
-- partition. This identification is computed from all  
-- dependent semantic units, and corresponds to the  
-- string that would be returned by use of the  
-- Body_Version or Version attributes.
```

```
-- The following Export pragmas export the version numbers  
-- with symbolic names ending in B (for body) or S  
-- (for spec) so that they can be located in a link. The  
-- information provided here is sufficient to track down  
-- the exact versions of units used in a given build.
```

```
type Version_32 is mod 2 ** 32;  
u00001 : constant Version_32 := 16#8ad6e54a#;  
pragma Export (C, u00001, "helloB");  
u00002 : constant Version_32 := 16#fbff4c67#;
```

(continues on next page)

(continued from previous page)

```
pragma Export (C, u00002, "system_standard_libraryB");
u00003 : constant Version_32 := 16#1ec6fd90#;
pragma Export (C, u00003, "system_standard_libraryS");
u00004 : constant Version_32 := 16#3ffc8e18#;
pragma Export (C, u00004, "adaS");
u00005 : constant Version_32 := 16#28f088c2#;
pragma Export (C, u00005, "ada_text_ioB");
u00006 : constant Version_32 := 16#f372c8ac#;
pragma Export (C, u00006, "ada_text_ioS");
u00007 : constant Version_32 := 16#2c143749#;
pragma Export (C, u00007, "ada_exceptionsB");
u00008 : constant Version_32 := 16#f4f0cce8#;
pragma Export (C, u00008, "ada_exceptionsS");
u00009 : constant Version_32 := 16#a46739c0#;
pragma Export (C, u00009, "ada_exceptions_last_chance_handlerB");
u00010 : constant Version_32 := 16#3aac8c92#;
pragma Export (C, u00010, "ada_exceptions_last_chance_handlerS");
u00011 : constant Version_32 := 16#1d274481#;
pragma Export (C, u00011, "systemS");
u00012 : constant Version_32 := 16#a207fefe#;
pragma Export (C, u00012, "system_soft_linksB");
u00013 : constant Version_32 := 16#467d9556#;
pragma Export (C, u00013, "system_soft_linksS");
u00014 : constant Version_32 := 16#b01dad17#;
pragma Export (C, u00014, "system_parametersB");
u00015 : constant Version_32 := 16#630d49fe#;
pragma Export (C, u00015, "system_parametersS");
u00016 : constant Version_32 := 16#b19b6653#;
pragma Export (C, u00016, "system_secondary_stackB");
u00017 : constant Version_32 := 16#b6468be8#;
pragma Export (C, u00017, "system_secondary_stackS");
u00018 : constant Version_32 := 16#39a03df9#;
pragma Export (C, u00018, "system_storage_elementsB");
u00019 : constant Version_32 := 16#30e40e85#;
pragma Export (C, u00019, "system_storage_elementsS");
u00020 : constant Version_32 := 16#41837dle#;
pragma Export (C, u00020, "system_stack_checkingB");
u00021 : constant Version_32 := 16#93982f69#;
pragma Export (C, u00021, "system_stack_checkingS");
u00022 : constant Version_32 := 16#393398c1#;
pragma Export (C, u00022, "system_exception_tableB");
u00023 : constant Version_32 := 16#b33e2294#;
pragma Export (C, u00023, "system_exception_tableS");
u00024 : constant Version_32 := 16#ce4af020#;
pragma Export (C, u00024, "system_exceptionsB");
u00025 : constant Version_32 := 16#75442977#;
pragma Export (C, u00025, "system_exceptionsS");
u00026 : constant Version_32 := 16#37d758f1#;
pragma Export (C, u00026, "system_exceptions_machineS");
u00027 : constant Version_32 := 16#b895431d#;
pragma Export (C, u00027, "system_exceptions_debugB");
```

(continues on next page)

(continued from previous page)

```
u00028 : constant Version_32 := 16#aec55d3f#;
pragma Export (C, u00028, "system_exceptions_debugS");
u00029 : constant Version_32 := 16#570325c8#;
pragma Export (C, u00029, "system_img_intB");
u00030 : constant Version_32 := 16#1ffc4443#;
pragma Export (C, u00030, "system_img_intS");
u00031 : constant Version_32 := 16#b98c3e16#;
pragma Export (C, u00031, "system_tracebackB");
u00032 : constant Version_32 := 16#831a9d5a#;
pragma Export (C, u00032, "system_tracebackS");
u00033 : constant Version_32 := 16#9ed49525#;
pragma Export (C, u00033, "system_traceback_entriesB");
u00034 : constant Version_32 := 16#1d7cb2f1#;
pragma Export (C, u00034, "system_traceback_entriesS");
u00035 : constant Version_32 := 16#8c33a517#;
pragma Export (C, u00035, "system_wch_conB");
u00036 : constant Version_32 := 16#065a6653#;
pragma Export (C, u00036, "system_wch_conS");
u00037 : constant Version_32 := 16#9721e840#;
pragma Export (C, u00037, "system_wch_stwB");
u00038 : constant Version_32 := 16#2b4b4a52#;
pragma Export (C, u00038, "system_wch_stwS");
u00039 : constant Version_32 := 16#92b797cb#;
pragma Export (C, u00039, "system_wch_cnvB");
u00040 : constant Version_32 := 16#09eddca0#;
pragma Export (C, u00040, "system_wch_cnvS");
u00041 : constant Version_32 := 16#6033a23f#;
pragma Export (C, u00041, "interfacesS");
u00042 : constant Version_32 := 16#ece6fdb6#;
pragma Export (C, u00042, "system_wch_jisB");
u00043 : constant Version_32 := 16#899dc581#;
pragma Export (C, u00043, "system_wch_jisS");
u00044 : constant Version_32 := 16#10558b11#;
pragma Export (C, u00044, "ada_streamsB");
u00045 : constant Version_32 := 16#2e6701ab#;
pragma Export (C, u00045, "ada_streamsS");
u00046 : constant Version_32 := 16#db5c917c#;
pragma Export (C, u00046, "ada_io_exceptionsS");
u00047 : constant Version_32 := 16#12c8cd7d#;
pragma Export (C, u00047, "ada_tagsB");
u00048 : constant Version_32 := 16#ce72c228#;
pragma Export (C, u00048, "ada_tagsS");
u00049 : constant Version_32 := 16#c3335bfd#;
pragma Export (C, u00049, "system_hhtableB");
u00050 : constant Version_32 := 16#99e5f76b#;
pragma Export (C, u00050, "system_hhtableS");
u00051 : constant Version_32 := 16#089f5cd0#;
pragma Export (C, u00051, "system_string_hashB");
u00052 : constant Version_32 := 16#3bbb9c15#;
pragma Export (C, u00052, "system_string_hashS");
u00053 : constant Version_32 := 16#807fe041#;
```

(continues on next page)

(continued from previous page)

```
pragma Export (C, u00053, "system_unsigned_typesS");
u00054 : constant Version_32 := 16#d27be59e#;
pragma Export (C, u00054, "system_val_lluB");
u00055 : constant Version_32 := 16#fa8db733#;
pragma Export (C, u00055, "system_val_lluS");
u00056 : constant Version_32 := 16#27b600b2#;
pragma Export (C, u00056, "system_val_utilB");
u00057 : constant Version_32 := 16#b187f27f#;
pragma Export (C, u00057, "system_val_utilS");
u00058 : constant Version_32 := 16#d1060688#;
pragma Export (C, u00058, "system_case_utilB");
u00059 : constant Version_32 := 16#392e2d56#;
pragma Export (C, u00059, "system_case_utilS");
u00060 : constant Version_32 := 16#84a27f0d#;
pragma Export (C, u00060, "interfaces_c_streamsB");
u00061 : constant Version_32 := 16#8bb5f2c0#;
pragma Export (C, u00061, "interfaces_c_streamsS");
u00062 : constant Version_32 := 16#6db6928f#;
pragma Export (C, u00062, "system_crtLS");
u00063 : constant Version_32 := 16#4e6a342b#;
pragma Export (C, u00063, "system_file_ioB");
u00064 : constant Version_32 := 16#ba56a5e4#;
pragma Export (C, u00064, "system_file_ioS");
u00065 : constant Version_32 := 16#b7ab275c#;
pragma Export (C, u00065, "ada_finalizationB");
u00066 : constant Version_32 := 16#19f764ca#;
pragma Export (C, u00066, "ada_finalizationS");
u00067 : constant Version_32 := 16#95817ed8#;
pragma Export (C, u00067, "system_finalization_rootB");
u00068 : constant Version_32 := 16#52d53711#;
pragma Export (C, u00068, "system_finalization_rootS");
u00069 : constant Version_32 := 16#769e25e6#;
pragma Export (C, u00069, "interfaces_cB");
u00070 : constant Version_32 := 16#4a38bedb#;
pragma Export (C, u00070, "interfaces_cS");
u00071 : constant Version_32 := 16#07e6ee66#;
pragma Export (C, u00071, "system_os_libB");
u00072 : constant Version_32 := 16#d7b69782#;
pragma Export (C, u00072, "system_os_libS");
u00073 : constant Version_32 := 16#1a817b8e#;
pragma Export (C, u00073, "system_stringsB");
u00074 : constant Version_32 := 16#639855e7#;
pragma Export (C, u00074, "system_stringsS");
u00075 : constant Version_32 := 16#e0b8de29#;
pragma Export (C, u00075, "system_file_control_blockS");
u00076 : constant Version_32 := 16#b5b2acal#;
pragma Export (C, u00076, "system_finalization_mastersB");
u00077 : constant Version_32 := 16#69316dc1#;
pragma Export (C, u00077, "system_finalization_mastersS");
u00078 : constant Version_32 := 16#57a37a42#;
pragma Export (C, u00078, "system_address_imageB");
```

(continues on next page)

(continued from previous page)

```

u00079 : constant Version_32 := 16#bccbd9bb#;
pragma Export (C, u00079, "system__address_imageS");
u00080 : constant Version_32 := 16#7268f812#;
pragma Export (C, u00080, "system__img_boolB");
u00081 : constant Version_32 := 16#e8fe356a#;
pragma Export (C, u00081, "system__img_boolS");
u00082 : constant Version_32 := 16#d7aac20c#;
pragma Export (C, u00082, "system__ioB");
u00083 : constant Version_32 := 16#8365b3ce#;
pragma Export (C, u00083, "system__ioS");
u00084 : constant Version_32 := 16#6d4d969a#;
pragma Export (C, u00084, "system__storage_poolsB");
u00085 : constant Version_32 := 16#e87cc305#;
pragma Export (C, u00085, "system__storage_poolsS");
u00086 : constant Version_32 := 16#e34550ca#;
pragma Export (C, u00086, "system__pool_globalB");
u00087 : constant Version_32 := 16#c88d2d16#;
pragma Export (C, u00087, "system__pool_globalS");
u00088 : constant Version_32 := 16#9d39c675#;
pragma Export (C, u00088, "system__memoryB");
u00089 : constant Version_32 := 16#445a22b5#;
pragma Export (C, u00089, "system__memoryS");
u00090 : constant Version_32 := 16#6a859064#;
pragma Export (C, u00090, "system__storage_pools__subpoolsB");
u00091 : constant Version_32 := 16#e3b008dc#;
pragma Export (C, u00091, "system__storage_pools__subpoolsS");
u00092 : constant Version_32 := 16#63f11652#;
pragma Export (C, u00092, "system__storage_pools__subpools__finalizationB");
u00093 : constant Version_32 := 16#fe2f4b3a#;
pragma Export (C, u00093, "system__storage_pools__subpools__finalizationS");

-- BEGIN ELABORATION ORDER
-- ada%s
-- interfaces%s
-- system%s
-- system.case_util%s
-- system.case_util%b
-- system.htable%s
-- system.img_bool%s
-- system.img_bool%b
-- system.img_int%s
-- system.img_int%b
-- system.io%s
-- system.io%b
-- system.parameters%s
-- system.parameters%b
-- system.crtl%s
-- interfaces.c_streams%s
-- interfaces.c_streams%b
-- system.standard_library%s
-- system.exceptions_debug%s

```

(continues on next page)

(continued from previous page)

```
-- system.exceptions_debug%b
-- system.storage_elements%s
-- system.storage_elements%b
-- system.stack_checking%s
-- system.stack_checking%b
-- system.string_hash%s
-- system.string_hash%b
-- system.htable%b
-- system.strings%s
-- system.strings%b
-- system.os_lib%s
-- system.traceback_entries%s
-- system.traceback_entries%b
-- ada.exceptions%s
-- system.soft_links%s
-- system.unsigned_types%s
-- system.val_llu%s
-- system.val_util%s
-- system.val_util%b
-- system.val_llu%b
-- system.wch_con%s
-- system.wch_con%b
-- system.wch_cnv%s
-- system.wch_jis%s
-- system.wch_jis%b
-- system.wch_cnv%b
-- system.wch_stw%s
-- system.wch_stw%b
-- ada.exceptions.last_chance_handler%s
-- ada.exceptions.last_chance_handler%b
-- system.address_image%s
-- system.exception_table%s
-- system.exception_table%b
-- ada.io_exceptions%s
-- ada.tags%s
-- ada.streams%s
-- ada.streams%b
-- interfaces.c%s
-- system.exceptions%s
-- system.exceptions%b
-- system.exceptions.machine%s
-- system.finalization_root%s
-- system.finalization_root%b
-- ada.finalization%s
-- ada.finalization%b
-- system.storage_pools%s
-- system.storage_pools%b
-- system.finalization_masters%s
-- system.storage_pools.subpools%s
-- system.storage_pools.subpools.finalization%s
-- system.storage_pools.subpools.finalization%b
```

(continues on next page)

(continued from previous page)

```

-- system.memory%s
-- system.memory%b
-- system.standard_library%b
-- system.pool_global%s
-- system.pool_global%b
-- system.file_control_block%s
-- system.file_io%s
-- system.secondary_stack%s
-- system.file_io%b
-- system.storage_pools.subpools%b
-- system.finalization_masters%b
-- interfaces.c%b
-- ada.tags%b
-- system.soft_links%b
-- system.os_lib%b
-- system.secondary_stack%b
-- system.address_image%b
-- system.traceback%s
-- ada.exceptions%b
-- system.traceback%b
-- ada.text_io%s
-- ada.text_io%b
-- hello%b
-- END ELABORATION ORDER

```

```
end ada_main;
```

```

pragma Ada_95;
-- The following source file name pragmas allow the generated file
-- names to be unique for different main programs. They are needed
-- since the package name will always be Ada_Main.

pragma Source_File_Name (ada_main, Spec_File_Name => "b~hello.ads");
pragma Source_File_Name (ada_main, Body_File_Name => "b~hello.adb");

pragma Suppress (Overflow_Check);
with Ada.Exceptions;

-- Generated package body for Ada_Main starts here

package body ada_main is
  pragma Warnings (Off);

  -- These values are reference counter associated to units which have
  -- been elaborated. It is also used to avoid elaborating the
  -- same unit twice.

  E72 : Short_Integer; pragma Import (Ada, E72, "system__os_lib_E");
  E13 : Short_Integer; pragma Import (Ada, E13, "system__soft_links_E");
  E23 : Short_Integer; pragma Import (Ada, E23, "system__exception_table_E");
  E46 : Short_Integer; pragma Import (Ada, E46, "ada__io_exceptions_E");

```

(continues on next page)

(continued from previous page)

```
E48 : Short_Integer; pragma Import (Ada, E48, "ada_tags_E");
E45 : Short_Integer; pragma Import (Ada, E45, "ada_streams_E");
E70 : Short_Integer; pragma Import (Ada, E70, "interfaces_c_E");
E25 : Short_Integer; pragma Import (Ada, E25, "system_exceptions_E");
E68 : Short_Integer; pragma Import (Ada, E68, "system_finalization_root_E");
E66 : Short_Integer; pragma Import (Ada, E66, "ada_finalization_E");
E85 : Short_Integer; pragma Import (Ada, E85, "system_storage_pools_E");
E77 : Short_Integer; pragma Import (Ada, E77, "system_finalization_masters_E");
E91 : Short_Integer; pragma Import (Ada, E91, "system_storage_pools_subpools_E");
E87 : Short_Integer; pragma Import (Ada, E87, "system_pool_global_E");
E75 : Short_Integer; pragma Import (Ada, E75, "system_file_control_block_E");
E64 : Short_Integer; pragma Import (Ada, E64, "system_file_io_E");
E17 : Short_Integer; pragma Import (Ada, E17, "system_secondary_stack_E");
E06 : Short_Integer; pragma Import (Ada, E06, "ada_text_io_E");
```

```
Local_Priority_Specific_Dispatching : constant String := "";
```

```
Local_Interrupt_States : constant String := "";
```

```
Is_Elaborated : Boolean := False;
```

```
procedure finalize_library is
```

```
begin
```

```
  E06 := E06 - 1;
```

```
  declare
```

```
    procedure F1;
```

```
    pragma Import (Ada, F1, "ada_text_io_finalize_spec");
```

```
  begin
```

```
    F1;
```

```
  end;
```

```
  E77 := E77 - 1;
```

```
  E91 := E91 - 1;
```

```
  declare
```

```
    procedure F2;
```

```
    pragma Import (Ada, F2, "system_file_io_finalize_body");
```

```
  begin
```

```
    E64 := E64 - 1;
```

```
    F2;
```

```
  end;
```

```
  declare
```

```
    procedure F3;
```

```
    pragma Import (Ada, F3, "system_file_control_block_finalize_spec");
```

```
  begin
```

```
    E75 := E75 - 1;
```

```
    F3;
```

```
  end;
```

```
  E87 := E87 - 1;
```

```
  declare
```

```
    procedure F4;
```

```
    pragma Import (Ada, F4, "system_pool_global_finalize_spec");
```

```
  begin
```

```
    F4;
```

(continues on next page)

(continued from previous page)

```

end;
declare
  procedure F5;
  pragma Import (Ada, F5, "system__storage_pools__subpools__finalize_spec");
begin
  F5;
end;
declare
  procedure F6;
  pragma Import (Ada, F6, "system__finalization_masters__finalize_spec");
begin
  F6;
end;
declare
  procedure Reraise_Library_Exception_If_Any;
  pragma Import (Ada, Reraise_Library_Exception_If_Any, "__gnat_reraise_library_exception_
↪if_any");
  begin
    Reraise_Library_Exception_If_Any;
  end;
end finalize_library;

-----
-- adainit --
-----

procedure adainit is

  Main_Priority : Integer;
  pragma Import (C, Main_Priority, "__gl_main_priority");
  Time_Slice_Value : Integer;
  pragma Import (C, Time_Slice_Value, "__gl_time_slice_val");
  WC_Encoding : Character;
  pragma Import (C, WC_Encoding, "__gl_wc_encoding");
  Locking_Policy : Character;
  pragma Import (C, Locking_Policy, "__gl_locking_policy");
  Queuing_Policy : Character;
  pragma Import (C, Queuing_Policy, "__gl_queuing_policy");
  Task_Dispatching_Policy : Character;
  pragma Import (C, Task_Dispatching_Policy, "__gl_task_dispatching_policy");
  Priority_Specific_Dispatching : System.Address;
  pragma Import (C, Priority_Specific_Dispatching, "__gl_priority_specific_dispatching");
  Num_Specific_Dispatching : Integer;
  pragma Import (C, Num_Specific_Dispatching, "__gl_num_specific_dispatching");
  Main_CPU : Integer;
  pragma Import (C, Main_CPU, "__gl_main_cpu");
  Interrupt_States : System.Address;
  pragma Import (C, Interrupt_States, "__gl_interrupt_states");
  Num_Interrupt_States : Integer;
  pragma Import (C, Num_Interrupt_States, "__gl_num_interrupt_states");
  Unreserve_All_Interrupts : Integer;

```

(continues on next page)

(continued from previous page)

```

pragma Import (C, Unreserve_All_Interrupts, "__gl_unreserve_all_interrupts");
Detect_Blocking : Integer;
pragma Import (C, Detect_Blocking, "__gl_detect_blocking");
Default_Stack_Size : Integer;
pragma Import (C, Default_Stack_Size, "__gl_default_stack_size");
Leap_Seconds_Support : Integer;
pragma Import (C, Leap_Seconds_Support, "__gl_leap_seconds_support");

procedure Runtime_Initialize;
pragma Import (C, Runtime_Initialize, "__gnat_runtime_initialize");

Finalize_Library_Objects : No_Param_Proc;
pragma Import (C, Finalize_Library_Objects, "__gnat_finalize_library_objects");

-- Start of processing for adainit

begin

  -- Record various information for this partition. The values
  -- are derived by the binder from information stored in the ali
  -- files by the compiler.

  if Is_Elaborated then
    return;
  end if;
  Is_Elaborated := True;
  Main_Priority := -1;
  Time_Slice_Value := -1;
  WC_Encoding := 'b';
  Locking_Policy := ' ';
  Queuing_Policy := ' ';
  Task_Dispatching_Policy := ' ';
  Priority_Specific_Dispatching :=
    Local_Priority_Specific_Dispatching'Address;
  Num_Specific_Dispatching := 0;
  Main_CPU := -1;
  Interrupt_States := Local_Interrupt_States'Address;
  Num_Interrupt_States := 0;
  Unreserve_All_Interrupts := 0;
  Detect_Blocking := 0;
  Default_Stack_Size := -1;
  Leap_Seconds_Support := 0;

  Runtime_Initialize;

  Finalize_Library_Objects := finalize_library'access;

  -- Now we have the elaboration calls for all units in the partition.
  -- The Elab_Spec and Elab_Body attributes generate references to the
  -- implicit elaboration procedures generated by the compiler for
  -- each unit that requires elaboration. Increment a counter of

```

(continues on next page)

(continued from previous page)

```

-- reference for each unit.

System.Soft_Links'Elab_Spec;
System.Exception_Table'Elab_Body;
E23 := E23 + 1;
Ada.Io_Exceptions'Elab_Spec;
E46 := E46 + 1;
Ada.Tags'Elab_Spec;
Ada.Streams'Elab_Spec;
E45 := E45 + 1;
Interfaces.C'Elab_Spec;
System.Exceptions'Elab_Spec;
E25 := E25 + 1;
System.Finalization_Root'Elab_Spec;
E68 := E68 + 1;
Ada.Finalization'Elab_Spec;
E66 := E66 + 1;
System.Storage_Pools'Elab_Spec;
E85 := E85 + 1;
System.Finalization_Masters'Elab_Spec;
System.Storage_Pools.Subpools'Elab_Spec;
System.Pool_Global'Elab_Spec;
E87 := E87 + 1;
System.File_Control_Block'Elab_Spec;
E75 := E75 + 1;
System.File_Io'Elab_Body;
E64 := E64 + 1;
E91 := E91 + 1;
System.Finalization_Masters'Elab_Body;
E77 := E77 + 1;
E70 := E70 + 1;
Ada.Tags'Elab_Body;
E48 := E48 + 1;
System.Soft_Links'Elab_Body;
E13 := E13 + 1;
System.Os_Lib'Elab_Body;
E72 := E72 + 1;
System.Secondary_Stack'Elab_Body;
E17 := E17 + 1;
Ada.Text_Io'Elab_Spec;
Ada.Text_Io'Elab_Body;
E06 := E06 + 1;
end adainit;

-----
-- adafinal --
-----

procedure adafinal is
  procedure s_stalib_adafinal;
  pragma Import (C, s_stalib_adafinal, "system__standard_library__adafinal");

```

(continues on next page)

(continued from previous page)

```

procedure Runtime_Finalize;
pragma Import (C, Runtime_Finalize, "__gnat_runtime_finalize");

begin
  if not Is_Elaborated then
    return;
  end if;
  Is_Elaborated := False;
  Runtime_Finalize;
  s_stalib_adafinal;
end adafinal;

-- We get to the main program of the partition by using
-- pragma Import because if we try to with the unit and
-- call it Ada style, then not only do we waste time
-- recompiling it, but also, we don't really know the right
-- switches (e.g.@: identifier character set) to be used
-- to compile it.

procedure Ada_Main_Program;
pragma Import (Ada, Ada_Main_Program, "_ada_hello");

-----
-- main --
-----

-- main is actually a function, as in the ANSI C standard,
-- defined to return the exit status. The three parameters
-- are the argument count, argument values and environment
-- pointer.

function main
  (argc : Integer;
   argv : System.Address;
   envp : System.Address)
  return Integer
is
  -- The initialize routine performs low level system
  -- initialization using a standard library routine which
  -- sets up signal handling and performs any other
  -- required setup. The routine can be found in file
  -- a-init.c.

  procedure initialize;
  pragma Import (C, initialize, "__gnat_initialize");

  -- The finalize routine performs low level system
  -- finalization using a standard library routine. The
  -- routine is found in file a-final.c and in the standard
  -- distribution is a dummy routine that does nothing, so

```

(continues on next page)

(continued from previous page)

```
-- really this is a hook for special user finalization.

procedure finalize;
pragma Import (C, finalize, "__gnat_finalize");

-- The following is to initialize the SEH exceptions

SEH : aliased array (1 .. 2) of Integer;

Ensure_Reference : aliased System.Address := Ada_Main_Program_Name'Address;
pragma Volatile (Ensure_Reference);

-- Start of processing for main

begin
  -- Save global variables

  gnat_argc := argc;
  gnat_argv := argv;
  gnat_envp := envp;

  -- Call low level system initialization

  Initialize (SEH'Address);

  -- Call our generated Ada initialization routine

  adainit;

  -- Now we call the main program of the partition

  Ada_Main_Program;

  -- Perform Ada finalization

  adafinal;

  -- Perform low level system finalization

  Finalize;

  -- Return the proper exit status
  return (gnat_exit_status);
end;

-- This section is entirely comments, so it has no effect on the
-- compilation of the Ada_Main package. It provides the list of
-- object files and linker options, as well as some standard
-- libraries needed for the link. The gnatlink utility parses
-- this b~hello.adb file to read these comment lines to generate
-- the appropriate command line arguments for the call to the
```

(continues on next page)

(continued from previous page)

```
-- system linker. The BEGIN/END lines are used for sentinels for
-- this parsing operation.

-- The exact file names will of course depend on the environment,
-- host/target and location of files on the host system.

-- BEGIN Object file/option list
--  ./hello.o
--  -L./
--  -L/usr/local/gnat/lib/gcc-lib/i686-pc-linux-gnu/2.8.1/adalib/
--  /usr/local/gnat/lib/gcc-lib/i686-pc-linux-gnu/2.8.1/adalib/libgnat.a
-- END Object file/option list

end ada_main;
```

The Ada code in the above example is exactly what is generated by the binder. We have added comments to more clearly indicate the function of each part of the generated `Ada_Main` package.

The code is standard Ada in all respects, and can be processed by any tools that handle Ada. In particular, it is possible to use the debugger in Ada mode to debug the generated `Ada_Main` package. For example, suppose that for reasons that you do not understand, your program is crashing during elaboration of the body of `Ada.Text_IO`. To locate this bug, you can place a breakpoint on the call:

```
Ada.Text_IO'Elab_Body;
```

and trace the elaboration routine for this package to find out where the problem might be (more usually of course you would be debugging elaboration code in your own application).

This page is intentionally left blank.

ELABORATION ORDER HANDLING IN GNAT

This appendix describes the handling of elaboration code in Ada and GNAT, and discusses how the order of elaboration of program units can be controlled in GNAT, either automatically or with explicit programming features.

C.1 Elaboration Code

Ada defines the term *execution* as the process by which a construct achieves its run-time effect. This process is also referred to as **elaboration** for declarations and *evaluation* for expressions.

The execution model in Ada allows for certain sections of an Ada program to be executed prior to execution of the program itself, primarily with the intent of initializing data. These sections are referred to as **elaboration code**. Elaboration code is executed as follows:

- All partitions of an Ada program are executed in parallel with one another, possibly in a separate address space, and possibly on a separate computer.
- The execution of a partition involves running the environment task for that partition.
- The environment task executes all elaboration code (if available) for all units within that partition. This code is said to be executed at **elaboration time**.
- The environment task executes the Ada program (if available) for that partition.

In addition to the Ada terminology, this appendix defines the following terms:

- *Invocation*

The act of calling a subprogram, instantiating a generic, or activating a task.

- *Scenario*

A construct that is elaborated or invoked by elaboration code is referred to as an *elaboration scenario* or simply a **scenario**. GNAT recognizes the following scenarios:

- 'Access of entries, operators, and subprograms
- Activation of tasks
- Calls to entries, operators, and subprograms
- Instantiations of generic templates

- *Target*

A construct elaborated by a scenario is referred to as *elaboration target* or simply **target**. GNAT recognizes the following targets:

- For 'Access of entries, operators, and subprograms, the target is the entry, operator, or subprogram being aliased.
- For activation of tasks, the target is the task body
- For calls to entries, operators, and subprograms, the target is the entry, operator, or subprogram being invoked.
- For instantiations of generic templates, the target is the generic template being instantiated.

Elaboration code may appear in two distinct contexts:

- *Library level*

A scenario appears at the library level when it is encapsulated by a package [body] compilation unit, ignoring any other package [body] declarations in between.

```
with Server;
package Client is
  procedure Proc;

  package Nested is
    Val : ... := Server.Func;
  end Nested;
end Client;
```

In the example above, the call to `Server.Func` is an elaboration scenario because it appears at the library level of package `Client`. Note that the declaration of package `Nested` is ignored according to the definition given above. As a result, the call to `Server.Func` will be invoked when the spec of unit `Client` is elaborated.

- *Package body statements*

A scenario appears within the statement sequence of a package body when it is bounded by the region starting from the `begin` keyword of the package body and ending at the `end` keyword of the package body.

```
package body Client is
  procedure Proc is
  begin
    ...
  end Proc;
begin
  Proc;
end Client;
```

In the example above, the call to `Proc` is an elaboration scenario because it appears within the statement sequence of package body `Client`. As a result, the call to `Proc` will be invoked when the body of `Client` is elaborated.

C.2 Elaboration Order

The sequence by which the elaboration code of all units within a partition is executed is referred to as **elaboration order**.

Within a single unit, elaboration code is executed in sequential order.

```
package body Client is
  Result : ... := Server.Func;

  procedure Proc is
    package Inst is new Server.Gen;
  begin
    Inst.Eval (Result);
  end Proc;
begin
  Proc;
end Client;
```

In the example above, the elaboration order within package body `Client` is as follows:

1. The object declaration of `Result` is elaborated.
 - Function `Server.Func` is invoked.
2. The subprogram body of `Proc` is elaborated.
3. Procedure `Proc` is invoked.
 - Generic unit `Server.Gen` is instantiated as `Inst`.
 - Instance `Inst` is elaborated.
 - Procedure `Inst.Eval` is invoked.

The elaboration order of all units within a partition depends on the following factors:

- *withed* units
- parent units
- purity of units
- preelaborability of units
- presence of elaboration-control pragmas
- invocations performed in elaboration code

A program may have several elaboration orders depending on its structure.

```
package Server is
  function Func (Index : Integer) return Integer;
end Server;
```

```
package body Server is
  Results : array (1 .. 5) of Integer := (1, 2, 3, 4, 5);

  function Func (Index : Integer) return Integer is
  begin
    return Results (Index);
  end Func;
end Server;
```

```
with Server;
package Client is
  Val : constant Integer := Server.Func (3);
end Client;
```

```
with Client;
procedure Main is begin null; end Main;
```

The following elaboration order exhibits a fundamental problem referred to as *access-before-elaboration* or simply **ABE**.

```
spec of Server
spec of Client
body of Server
body of Main
```

The elaboration of `Server`'s spec materializes function `Func`, making it callable. The elaboration of `Client`'s spec elaborates the declaration of `Val`. This invokes function `Server.Func`, however the body of `Server.Func` has not been elaborated yet because `Server`'s body comes after `Client`'s spec in the elaboration order. As a result, the value of constant `Val` is now undefined.

Without any guarantees from the language, an undetected ABE problem may hinder proper initialization of data, which in turn may lead to undefined behavior at run time. To prevent such ABE problems, Ada employs dynamic checks in the same vein as index or null exclusion checks. A failed ABE check raises exception `Program_Error`.

The following elaboration order avoids the ABE problem and the program can be successfully elaborated.

```
spec of Server
body of Server
spec of Client
body of Main
```

Ada states that a total elaboration order must exist, but it does not define what this order is. A compiler is thus tasked with choosing a suitable elaboration order which satisfies the dependencies imposed by *with* clauses, unit categorization, elaboration-control pragmas, and invocations performed in elaboration code. Ideally an order that avoids ABE problems should be chosen, however a compiler may not always find such an order due to complications with respect to control and data flow.

C.3 Checking the Elaboration Order

To avoid placing the entire elaboration-order burden on the programmer, Ada provides three lines of defense:

- *Static semantics*

Static semantic rules restrict the possible choice of elaboration order. For instance, if unit *Client* *withs* unit *Server*, then the spec of *Server* is always elaborated prior to *Client*. The same principle applies to child units - the spec of a parent unit is always elaborated prior to the child unit.

- *Dynamic semantics*

Dynamic checks are performed at run time, to ensure that a target is elaborated prior to a scenario that invokes it, thus avoiding ABE problems. A failed run-time check raises exception `Program_Error`. The following restrictions apply:

- *Restrictions on calls*

An entry, operator, or subprogram can be called from elaboration code only when the corresponding body has been elaborated.

- *Restrictions on instantiations*

A generic unit can be instantiated by elaboration code only when the corresponding body has been elaborated.

- *Restrictions on task activation*

A task can be activated by elaboration code only when the body of the associated task type has been elaborated.

The restrictions above can be summarized by the following rule:

If a target has a body, then this body must be elaborated prior to the scenario that invokes the target.

- *Elaboration control*

Pragmas are provided for the programmer to specify the desired elaboration order.

C.4 Controlling the Elaboration Order in Ada

Ada provides several idioms and pragmas to aid the programmer with specifying the desired elaboration order and avoiding ABE problems altogether.

- *Packages without a body*

A library package which does not require a completing body does not suffer from ABE problems.

```

package Pack is
  generic
    type Element is private;
  package Containers is
    type Element_Array is array (1 .. 10) of Element;
  end Containers;
end Pack;

```

In the example above, package `Pack` does not require a body because it does not contain any constructs which require completion in a body. As a result, generic `Pack.Containers` can be instantiated without encountering any ABE problems.

- *pragma Pure*

Pragma `Pure` places sufficient restrictions on a unit to guarantee that no scenario within the unit can result in an ABE problem.

- *pragma Preelaborate*

Pragma `Preelaborate` is slightly less restrictive than `pragma Pure`, but still strong enough to prevent ABE problems within a unit.

- *pragma Elaborate_Body*

Pragma `Elaborate_Body` requires that the body of a unit is elaborated immediately after its spec. This restriction guarantees that no client scenario can invoke a server target before the target body has been elaborated because the spec and body are effectively “glued” together.

```

package Server is
  pragma Elaborate_Body;

  function Func return Integer;
end Server;

```

```

package body Server is
  function Func return Integer is
  begin
    ...
  end Func;
end Server;

```

```

with Server;
package Client is
  Val : constant Integer := Server.Func;
end Client;

```

In the example above, `pragma Elaborate_Body` guarantees the following elaboration order:

```

spec of Server
body of Server
spec of Client

```

because the spec of `Server` must be elaborated prior to `Client` by virtue of the *with* clause, and in addition the body of `Server` must be elaborated immediately after the spec of `Server`.

Removing pragma `Elaborate_Body` could result in the following incorrect elaboration order:

```
spec of Server
spec of Client
body of Server
```

where `Client` invokes `Server.Func`, but the body of `Server.Func` has not been elaborated yet.

The pragmas outlined above allow a server unit to guarantee safe elaboration use by client units. Thus it is a good rule to mark units as `Pure` or `Preelaborate`, and if this is not possible, mark them as `Elaborate_Body`.

There are however situations where `Pure`, `Preelaborate`, and `Elaborate_Body` are not applicable. Ada provides another set of pragmas for use by client units to help ensure the elaboration safety of server units they depend on.

- *pragma Elaborate (Unit)*

Pragma `Elaborate` can be placed in the context clauses of a unit, after a *with* clause. It guarantees that both the spec and body of its argument will be elaborated prior to the unit with the pragma. Note that other unrelated units may be elaborated in between the spec and the body.

```
package Server is
  function Func return Integer;
end Server;
```

```
package body Server is
  function Func return Integer is
  begin
    ...
  end Func;
end Server;
```

```
with Server;
pragma Elaborate (Server);
package Client is
  Val : constant Integer := Server.Func;
end Client;
```

In the example above, pragma `Elaborate` guarantees the following elaboration order:

```
spec of Server
body of Server
spec of Client
```

Removing pragma `Elaborate` could result in the following incorrect elaboration order:

```
spec of Server
spec of Client
body of Server
```

where `Client` invokes `Server.Func`, but the body of `Server.Func` has not been elaborated yet.

- *pragma Elaborate_All (Unit)*

`Pragma Elaborate_All` is placed in the context clauses of a unit, after a *with* clause. It guarantees that both the spec and body of its argument will be elaborated prior to the unit with the pragma, as well as all units *withed* by the spec and body of the argument, recursively. Note that other unrelated units may be elaborated in between the spec and the body.

```
package Math is
  function Factorial (Val : Natural) return Natural;
end Math;
```

```
package body Math is
  function Factorial (Val : Natural) return Natural is
  begin
    ...;
  end Factorial;
end Math;
```

```
package Computer is
  type Operation_Kind is (None, Op_Factorial);

  function Compute
    (Val : Natural;
     Op  : Operation_Kind) return Natural;
end Computer;
```

```
with Math;
package body Computer is
  function Compute
    (Val : Natural;
     Op  : Operation_Kind) return Natural
  is
    if Op = Op_Factorial then
      return Math.Factorial (Val);
    end if;

    return 0;
  end Compute;
end Computer;
```

```
with Computer;
pragma Elaborate_All (Computer);
package Client is
  Val : constant Natural :=
```

(continues on next page)

(continued from previous page)

```

    Computer.Compute (123, Computer.Op_Factorial);
end Client;

```

In the example above, pragma `Elaborate_All` can result in the following elaboration order:

```

spec of Math
body of Math
spec of Computer
body of Computer
spec of Client

```

Note that there are several allowable suborders for the specs and bodies of `Math` and `Computer`, but the point is that these specs and bodies will be elaborated prior to `Client`.

Removing pragma `Elaborate_All` could result in the following incorrect elaboration order:

```

spec of Math
spec of Computer
body of Computer
spec of Client
body of Math

```

where `Client` invokes `Computer.Compute`, which in turn invokes `Math.Factorial`, but the body of `Math.Factorial` has not been elaborated yet.

All pragmas shown above can be summarized by the following rule:

If a client unit elaborates a server target directly or indirectly, then if the server unit requires a body and does not have pragma `Pure`, `Preelaborate`, or `Elaborate_Body`, then the client unit should have pragma `Elaborate` or `Elaborate_All` for the server unit.

If the rule outlined above is not followed, then a program may fall in one of the following states:

- *No elaboration order exists*

In this case a compiler must diagnose the situation, and refuse to build an executable program.

- *One or more incorrect elaboration orders exist*

In this case a compiler can build an executable program, but `Program_Error` will be raised when the program is run.

- *Several elaboration orders exist, some correct, some incorrect*

In this case the programmer has not controlled the elaboration order. As a result, a compiler may or may not pick one of the correct orders, and the program may or may not raise `Program_Error` when it is run. This is the worst possible state because the program may fail on another compiler, or even another version of the same compiler.

- *One or more correct orders exist*

In this case a compiler can build an executable program, and the program is run successfully. This state may be guaranteed by following the outlined rules, or may be the result of good program architecture.

Note that one additional advantage of using `Elaborate` and `Elaborate_All` is that the program continues to stay in the last state (one or more correct orders exist) even if maintenance changes the bodies of targets.

C.5 Controlling the Elaboration Order in GNAT

In addition to Ada semantics and rules synthesized from them, GNAT offers three elaboration models to aid the programmer with specifying the correct elaboration order and to diagnose elaboration problems.

- *Dynamic elaboration model*

This is the most permissive of the three elaboration models and emulates the behavior specified by the Ada Reference Manual. When the dynamic model is in effect, GNAT makes the following assumptions:

- All code within all units in a partition is considered to be elaboration code.
- Some of the invocations in elaboration code may not take place at run time due to conditional execution.

GNAT performs extensive diagnostics on a unit-by-unit basis for all scenarios that invoke internal targets. In addition, GNAT generates run-time checks for all external targets and for all scenarios that may exhibit ABE problems.

The elaboration order is obtained by honoring all *with* clauses, purity and preelaborability of units, and elaboration-control pragmas. The dynamic model attempts to take all invocations in elaboration code into account. If an invocation leads to a circularity, GNAT ignores the invocation based on the assumptions stated above. An order obtained using the dynamic model may fail an ABE check at run time when GNAT ignored an invocation.

The dynamic model is enabled with compiler switch `-gnatE`.

- *Static elaboration model*

This is the middle ground of the three models. When the static model is in effect, GNAT makes the following assumptions:

- Only code at the library level and in package body statements within all units in a partition is considered to be elaboration code.
- All invocations in elaboration will take place at run time, regardless of conditional execution.

GNAT performs extensive diagnostics on a unit-by-unit basis for all scenarios that invoke internal targets. In addition, GNAT generates run-time checks for all external targets and for all scenarios that may exhibit ABE problems.

The elaboration order is obtained by honoring all *with* clauses, purity and preelaborability of units, presence of elaboration-control pragmas, and all invocations in elaboration code. An order obtained using the static model is guaranteed to be ABE problem-free, excluding dispatching calls and access-to-subprogram types.

The static model is the default model in GNAT.

- *SPARK elaboration model*

This is the most conservative of the three models and enforces the SPARK rules of elaboration as defined in the SPARK Reference Manual, section 7.7. The SPARK model is in effect only when a scenario and a target reside in a region subject to `SPARK_Mode On`, otherwise the dynamic or static model is in effect.

The SPARK model is enabled with compiler switch `-gnatd.v`.

- *Legacy elaboration models*

In addition to the three elaboration models outlined above, GNAT provides the following legacy models:

- *Legacy elaboration-checking model* available in pre-18.x versions of GNAT. This model is enabled with compiler switch `-gnath`.
- *Legacy elaboration-order model* available in pre-20.x versions of GNAT. This model is enabled with binder switch `-H`.

The dynamic, legacy, and static models can be relaxed using compiler switch `-gnatJ`, making them more permissive. Note that in this mode, GNAT may not diagnose certain elaboration issues or install run-time checks.

C.6 Mixing Elaboration Models

It is possible to mix units compiled with a different elaboration model, however the following rules must be observed:

- A client unit compiled with the dynamic model can only *with* a server unit that meets at least one of the following criteria:
 - The server unit is compiled with the dynamic model.
 - The server unit is a GNAT implementation unit from the `Ada`, `GNAT`, `Interfaces`, or `System` hierarchies.
 - The server unit has pragma `Pure` or `Preelaborate`.
 - The client unit has an explicit `Elaborate_All` pragma for the server unit.

These rules ensure that elaboration checks are not omitted. If the rules are violated, the binder emits a warning:

```
warning: "x.ads" has dynamic elaboration checks and with's
warning:  "y.ads" which has static elaboration checks
```

The warnings can be suppressed by binder switch `-ws`.

C.7 ABE Diagnostics

GNAT performs extensive diagnostics on a unit-by-unit basis for all scenarios that invoke internal targets, regardless of whether the dynamic, SPARK, or static model is in effect.

Note that GNAT emits warnings rather than hard errors whenever it encounters an elaboration problem. This is because the elaboration model in effect may be too conservative, or a particular scenario may not be invoked due conditional execution. The warnings can be suppressed selectively with `pragma Warnings (Off)` or globally with compiler switch `-gnatwL`.

A *guaranteed ABE* arises when the body of a target is not elaborated early enough, and causes *all* scenarios that directly invoke the target to fail.

```
package body Guaranteed_ABE is
  function ABE return Integer;

  Val : constant Integer := ABE;

  function ABE return Integer is
  begin
    ...
  end ABE;
end Guaranteed_ABE;
```

In the example above, the elaboration of `Guaranteed_ABE`'s body elaborates the declaration of `Val`. This invokes function `ABE`, however the body of `ABE` has not been elaborated yet. GNAT emits the following diagnostic:

```
4.   Val : constant Integer := ABE;
      |
      >>> warning: cannot call "ABE" before body seen
      >>> warning: Program_Error will be raised at run time
```

A *conditional ABE* arises when the body of a target is not elaborated early enough, and causes *some* scenarios that directly invoke the target to fail.

```
1. package body Conditional_ABE is
2.   procedure Force_Body is null;
3.
4.   generic
5.     with function Func return Integer;
6.   package Gen is
7.     Val : constant Integer := Func;
8.   end Gen;
9.
10.  function ABE return Integer;
11.
12.  function Cause_ABE return Boolean is
13.    package Inst is new Gen (ABE);
14.  begin
15.    ...
16.  end Cause_ABE;
```

(continues on next page)

(continued from previous page)

```

17.
18.   Val : constant Boolean := Cause_ABE;
19.
20.   function ABE return Integer is
21.   begin
22.     ...
23.   end ABE;
24.
25.   Safe : constant Boolean := Cause_ABE;
26. end Conditional_ABE;

```

In the example above, the elaboration of package body `Conditional_ABE` elaborates the declaration of `Val`. This invokes function `Cause_ABE`, which instantiates generic unit `Gen` as `Inst`. The elaboration of `Inst` invokes function `ABE`, however the body of `ABE` has not been elaborated yet. GNAT emits the following diagnostic:

```

13.   package Inst is new Gen (ABE);
      |
      >>> warning: in instantiation at line 7
      >>> warning: cannot call "ABE" before body seen
      >>> warning: Program_Error may be raised at run time
      >>> warning:   body of unit "Conditional_ABE" elaborated
      >>> warning:   function "Cause_ABE" called at line 18
      >>> warning:   function "ABE" called at line 7, instance at line 13

```

Note that the same `ABE` problem does not occur with the elaboration of declaration `Safe` because the body of function `ABE` has already been elaborated at that point.

C.8 SPARK Diagnostics

GNAT enforces the SPARK rules of elaboration as defined in the SPARK Reference Manual section 7.7 when compiler switch `-gnatd.v` is in effect. Note that GNAT emits hard errors whenever it encounters a violation of the SPARK rules.

```

1. with Server;
2. package body SPARK_Diagnostics with SPARK_Mode is
3.   Val : constant Integer := Server.Func;
      |
      >>> call to "Func" during elaboration in SPARK
      >>> unit "SPARK_Diagnostics" requires pragma "Elaborate_All" for "Server"
      >>>   body of unit "SPARK_Model" elaborated
      >>>   function "Func" called at line 3
4. end SPARK_Diagnostics;

```

C.9 Elaboration Circularities

An **elaboration circularity** occurs whenever the elaboration of a set of units enters a deadlocked state, where each unit is waiting for another unit to be elaborated. This situation may be the result of improper use of *with* clauses, elaboration-control pragmas, or invocations in elaboration code.

The following example exhibits an elaboration circularity.

```
with B; pragma Elaborate (B);
package A is
end A;
```

```
package B is
  procedure Force_Body;
end B;
```

```
with C;
package body B is
  procedure Force_Body is null;

  Elab : constant Integer := C.Func;
end B;
```

```
package C is
  function Func return Integer;
end C;
```

```
with A;
package body C is
  function Func return Integer is
  begin
    ...
  end Func;
end C;
```

The binder emits the following diagnostic:

```
error: Elaboration circularity detected
info:
info:   Reason:
info:     unit "a (spec)" depends on its own elaboration
info:
info:   Circularity:
info:     unit "a (spec)" has with clause and pragma Elaborate for unit "b (spec)"
info:     unit "b (body)" is in the closure of pragma Elaborate
info:     unit "b (body)" invokes a construct of unit "c (body)" at elaboration time
info:     unit "c (body)" has with clause for unit "a (spec)"
info:
info:   Suggestions:
```

(continues on next page)

(continued from previous page)

```

info:
info:   remove pragma Elaborate for unit "b (body)" in unit "a (spec)"
info:   use the dynamic elaboration model (compiler switch -gnatE)

```

The diagnostic consist of the following sections:

- Reason

This section provides a short explanation describing why the set of units could not be ordered.

- Circularity

This section enumerates the units comprising the deadlocked set, along with their interdependencies.

- Suggestions

This section enumerates various tactics for eliminating the circularity.

C.10 Resolving Elaboration Circularities

The most desirable option from the point of view of long-term maintenance is to rearrange the program so that the elaboration problems are avoided. One useful technique is to place the elaboration code into separate child packages. Another is to move some of the initialization code to explicitly invoked subprograms, where the program controls the order of initialization explicitly. Although this is the most desirable option, it may be impractical and involve too much modification, especially in the case of complex legacy code.

When faced with an elaboration circularity, the programmer should also consider the tactics given in the suggestions section of the circularity diagnostic. Depending on the units involved in the circularity, their *with* clauses, purity, preelaborability, presence of elaboration-control pragmas and invocations at elaboration time, the binder may suggest one or more of the following tactics to eliminate the circularity:

- Pragma Elaborate elimination

```
remove pragma Elaborate for unit "..." in unit "..."
```

This tactic is suggested when the binder has determined that pragma `Elaborate`:

- Prevents a set of units from being elaborated.
- The removal of the pragma will not eliminate the semantic effects of the pragma. In other words, the argument of the pragma will still be elaborated prior to the unit containing the pragma.
- The removal of the pragma will enable the successful ordering of the units.

The programmer should remove the pragma as advised, and rebuild the program.

- Pragma `Elaborate_All` elimination

```
remove pragma Elaborate_All for unit "..." in unit "..."
```

This tactic is suggested when the binder has determined that pragma `Elaborate_All`:

- Prevents a set of units from being elaborated.
- The removal of the pragma will not eliminate the semantic effects of the pragma. In other words, the argument of the pragma along with its *with* closure will still be elaborated prior to the unit containing the pragma.
- The removal of the pragma will enable the successful ordering of the units.

The programmer should remove the pragma as advised, and rebuild the program.

- Pragma `Elaborate_All` downgrade

```
change pragma Elaborate_All for unit "..." to Elaborate in unit "..."
```

This tactic is always suggested with the pragma `Elaborate_All` elimination tactic. It offers a different alternative of guaranteeing that the argument of the pragma will still be elaborated prior to the unit containing the pragma.

The programmer should update the pragma as advised, and rebuild the program.

- Pragma `Elaborate_Body` elimination

```
remove pragma Elaborate_Body in unit "..."
```

This tactic is suggested when the binder has determined that pragma `Elaborate_Body`:

- Prevents a set of units from being elaborated.
- The removal of the pragma will enable the successful ordering of the units.

Note that the binder cannot determine whether the pragma is required for other purposes, such as guaranteeing the initialization of a variable declared in the spec by elaboration code in the body.

The programmer should remove the pragma as advised, and rebuild the program.

- Use of pragma `Restrictions`

```
use pragma Restrictions (No_Entry_Calls_In_Elaboration_Code)
```

This tactic is suggested when the binder has determined that a task activation at elaboration time:

- Prevents a set of units from being elaborated.

Note that the binder cannot determine with certainty whether the task will block at elaboration time.

The programmer should create a configuration file, place the pragma within, update the general compilation arguments, and rebuild the program.

- Use of dynamic elaboration model

```
use the dynamic elaboration model (compiler switch -gnatE)
```

This tactic is suggested when the binder has determined that an invocation at elaboration time:

- Prevents a set of units from being elaborated.
- The use of the dynamic model will enable the successful ordering of the units.

The programmer has two options:

- Determine the units involved in the invocation using the detailed invocation information, and add compiler switch `-gnatE` to the compilation arguments of selected files only. This approach will yield safer elaboration orders compared to the other option because it will minimize the opportunities presented to the dynamic model for ignoring invocations.
- Add compiler switch `-gnatE` to the general compilation arguments.

- Use of detailed invocation information

```
use detailed invocation information (compiler switch -gnatd_F)
```

This tactic is always suggested with the use of the dynamic model tactic. It causes the circularity section of the circularity diagnostic to describe the flow of elaboration code from a unit to a unit, enumerating all such paths in the process.

The programmer should analyze this information to determine which units should be compiled with the dynamic model.

- Forced-dependency elimination

```
remove the dependency of unit "." on unit "." from the argument of switch -f
```

This tactic is suggested when the binder has determined that a dependency present in the forced-elaboration-order file indicated by binder switch `-f`:

- Prevents a set of units from being elaborated.
- The removal of the dependency will enable the successful ordering of the units.

The programmer should edit the forced-elaboration-order file, remove the dependency, and rebind the program.

- All forced-dependency elimination

```
remove switch -f
```

This tactic is suggested in case editing the forced-elaboration-order file is not an option.

The programmer should remove binder switch `-f` from the binder arguments, and rebind.

- Multiple-circularities diagnostic

```
diagnose all circularities (binder switch -d_C)
```

By default, the binder will diagnose only the highest-precedence circularity. If the program contains multiple circularities, the binder will suggest the use of binder switch `-d_C` in order to obtain the diagnostics of all circularities.

The programmer should add binder switch `-d_C` to the binder arguments, and rebind.

If none of the tactics suggested by the binder eliminate the elaboration circularity, the programmer should consider using one of the legacy elaboration models, in the following order:

- Use the pre-20.x legacy elaboration-order model, with binder switch `-H`.
- Use both pre-18.x and pre-20.x legacy elaboration models, with compiler switch `-gnatH` and binder switch `-H`.
- Use the relaxed static-elaboration model, with compiler switches `-gnatH` `-gnatJ` and binder switch `-H`.
- Use the relaxed dynamic-elaboration model, with compiler switches `-gnatH` `-gnatJ` `-gnatE` and binder switch `-H`.

C.11 Elaboration-related Compiler Switches

GNAT has several switches that affect the elaboration model and consequently the elaboration order chosen by the binder.

-gnatE

Dynamic elaboration checking mode enabled

When this switch is in effect, GNAT activates the dynamic model.

-gnatEl

Turn on info messages on generated `Elaborate[_All]` pragmas

This switch is only applicable to the pre-20.x legacy elaboration models. The post-20.x elaboration model no longer relies on implicitly generated `Elaborate` and `Elaborate_All` pragmas to order units.

When this switch is in effect, GNAT will emit the following supplementary information depending on the elaboration model in effect.

- *Dynamic model*

GNAT will indicate missing `Elaborate` and `Elaborate_All` pragmas for all library-level scenarios within the partition.
- *Static model*

GNAT will indicate all scenarios invoked during elaboration. In addition, it will provide detailed traceback when an implicit `Elaborate` or `Elaborate_All` pragma is generated.
- *SPARK model*

GNAT will indicate how an elaboration requirement is met by the context of a unit. This diagnostic requires compiler switch `-gnatd.v`.

```

1. with Server; pragma Elaborate_All (Server);
2. package Client with SPARK_Mode is
3.   Val : constant Integer := Server.Func;
   |
   >>> info: call to "Func" during elaboration in SPARK
   >>> info: "Elaborate_All" requirement for unit "Server" met by pragma at line 1
4. end Client;

```

-gnath

Legacy elaboration checking mode enabled

When this switch is in effect, GNAT will utilize the pre-18.x elaboration model.

-gnatJ

Relaxed elaboration checking mode enabled

When this switch is in effect, GNAT will not process certain scenarios, resulting in a more permissive elaboration model. Note that this may eliminate some diagnostics and run-time checks.

-gnatw.f

Turn on warnings for suspicious Subp'Access

When this switch is in effect, GNAT will treat 'Access of an entry, operator, or subprogram as a potential call to the target and issue warnings:

```

1. package body Attribute_Call is
2.   function Func return Integer;
3.   type Func_Ptr is access function return Integer;
4.
5.   Ptr : constant Func_Ptr := Func'Access;
   |
   >>> warning: "Access" attribute of "Func" before body seen
   >>> warning: possible Program_Error on later references
   >>> warning:   body of unit "Attribute_Call" elaborated
   >>> warning:   "Access" of "Func" taken at line 5
6.
7.   function Func return Integer is
8.   begin
9.     ...
10.  end Func;
11. end Attribute_Call;

```

In the example above, the elaboration of declaration `Ptr` is assigned `Func'Access` before the body of `Func` has been elaborated.

-gnatwl

Turn on warnings for elaboration problems

When this switch is in effect, GNAT emits diagnostics in the form of warnings concerning various elaboration problems. The warnings are enabled by default. The switch is provided in case all warnings are suppressed, but elaboration warnings are still desired.

-gnatwL

Turn off warnings for elaboration problems

When this switch is in effect, GNAT no longer emits any diagnostics in the form of warnings. Selective suppression of elaboration problems is possible using `pragma Warnings (Off)`.

```
1. package body Selective_Suppression is
2.   function ABE return Integer;
3.
4.   Val_1 : constant Integer := ABE;
           |
           >>> warning: cannot call "ABE" before body seen
           >>> warning: Program_Error will be raised at run time
5.
6.   pragma Warnings (Off);
7.   Val_2 : constant Integer := ABE;
8.   pragma Warnings (On);
9.
10.  function ABE return Integer is
11.  begin
12.    ...
13.  end ABE;
14. end Selective_Suppression;
```

Note that suppressing elaboration warnings does not eliminate run-time checks. The example above will still fail at run time with an ABE.

C.12 Summary of Procedures for Elaboration Control

A programmer should first compile the program with the default options, using none of the binder or compiler switches. If the binder succeeds in finding an elaboration order, then apart from possible cases involving dispatching calls and access-to-subprogram types, the program is free of elaboration errors.

If it is important for the program to be portable to compilers other than GNAT, then the programmer should use compiler switch `-gnatel` and consider the messages about missing or implicitly created `Elaborate` and `Elaborate_All` pragmas.

If the binder reports an elaboration circularity, the programmer has several options:

- Ensure that elaboration warnings are enabled. This will allow the static model to output trace information of elaboration issues. The trace information could shed light on previously unforeseen dependencies, as well as their origins. Elaboration warnings are enabled with compiler switch `-gnatwL`.
- Consider the tactics given in the suggestions section of the circularity diagnostic.
- If none of the steps outlined above resolve the circularity, use a more permissive elaboration model, in the following order:
 - Use the pre-20.x legacy elaboration-order model, with binder switch `-H`.

- Use both pre-18.x and pre-20.x legacy elaboration models, with compiler switch `-gnatH` and binder switch `-H`.
- Use the relaxed static elaboration model, with compiler switches `-gnatH -gnatJ` and binder switch `-H`.
- Use the relaxed dynamic elaboration model, with compiler switches `-gnatH -gnatJ -gnatE` and binder switch `-H`.

C.13 Inspecting the Chosen Elaboration Order

To see the elaboration order chosen by the binder, inspect the contents of file `b~xxx.adb`. On certain targets, this file appears as `b_xxx.adb`. The elaboration order appears as a sequence of calls to `Elab_Body` and `Elab_Spec`, interspersed with assignments to `Exxx` which indicates that a particular unit is elaborated. For example:

```
System.Soft_Links'Elab_Body;
E14 := True;
System.Secondary_Stack'Elab_Body;
E18 := True;
System.Exception_Table'Elab_Body;
E24 := True;
Ada.Io_Exceptions'Elab_Spec;
E67 := True;
Ada.Tags'Elab_Spec;
Ada.Streams'Elab_Spec;
E43 := True;
Interfaces.C'Elab_Spec;
E69 := True;
System.Finalization_Root'Elab_Spec;
E60 := True;
System.Os_Lib'Elab_Body;
E71 := True;
System.Finalization_Implementation'Elab_Spec;
System.Finalization_Implementation'Elab_Body;
E62 := True;
Ada.Finalization'Elab_Spec;
E58 := True;
Ada.Finalization.List_Controller'Elab_Spec;
E76 := True;
System.File_Control_Block'Elab_Spec;
E74 := True;
System.File_Io'Elab_Body;
E56 := True;
Ada.Tags'Elab_Body;
E45 := True;
Ada.Text_Io'Elab_Spec;
Ada.Text_Io'Elab_Body;
E07 := True;
```

Note also binder switch `-l`, which outputs the chosen elaboration order and provides a more readable

form of the above:

```
ada (spec)
interfaces (spec)
system (spec)
system.case_util (spec)
system.case_util (body)
system.concat_2 (spec)
system.concat_2 (body)
system.concat_3 (spec)
system.concat_3 (body)
system.htable (spec)
system.parameters (spec)
system.parameters (body)
system.crtl (spec)
interfaces.c_streams (spec)
interfaces.c_streams (body)
system.restrictions (spec)
system.restrictions (body)
system.standard_library (spec)
system.exceptions (spec)
system.exceptions (body)
system.storage_elements (spec)
system.storage_elements (body)
system.secondary_stack (spec)
system.stack_checking (spec)
system.stack_checking (body)
system.string_hash (spec)
system.string_hash (body)
system.htable (body)
system.strings (spec)
system.strings (body)
system.traceback (spec)
system.traceback (body)
system.traceback_entries (spec)
system.traceback_entries (body)
ada.exceptions (spec)
ada.exceptions.last_chance_handler (spec)
system.soft_links (spec)
system.soft_links (body)
ada.exceptions.last_chance_handler (body)
system.secondary_stack (body)
system.exception_table (spec)
system.exception_table (body)
ada.io_exceptions (spec)
ada.tags (spec)
ada.streams (spec)
interfaces.c (spec)
interfaces.c (body)
system.finalization_root (spec)
system.finalization_root (body)
system.memory (spec)
system.memory (body)
```

(continues on next page)

(continued from previous page)

```
system.standard_library (body)
system.os_lib (spec)
system.os_lib (body)
system.unsigned_types (spec)
system.stream_attributes (spec)
system.stream_attributes (body)
system.finalization_implementation (spec)
system.finalization_implementation (body)
ada.finalization (spec)
ada.finalization (body)
ada.finalization.list_controller (spec)
ada.finalization.list_controller (body)
system.file_control_block (spec)
system.file_io (spec)
system.file_io (body)
system.val_uns (spec)
system.val_util (spec)
system.val_util (body)
system.val_uns (body)
system.wch_con (spec)
system.wch_con (body)
system.wch_cnv (spec)
system.wch_jis (spec)
system.wch_jis (body)
system.wch_cnv (body)
system.wch_stw (spec)
system.wch_stw (body)
ada.tags (body)
ada.exceptions (body)
ada.text_io (spec)
ada.text_io (body)
text_io (spec)
gdbstr (body)
```

This page is intentionally left blank.

INLINE ASSEMBLER

If you need to write low-level software that interacts directly with the hardware, Ada provides two ways to incorporate assembly language code into your program. First, you can import and invoke external routines written in assembly language, an Ada feature fully supported by GNAT. However, for small sections of code it may be simpler or more efficient to include assembly language statements directly in your Ada source program, using the facilities of the implementation-defined package `System.Machine_Code`, which incorporates the gcc Inline Assembler. The Inline Assembler approach offers a number of advantages, including the following:

- No need to use non-Ada tools
- Consistent interface over different targets
- Automatic usage of the proper calling conventions
- Access to Ada constants and variables
- Definition of intrinsic routines
- Possibility of inlining a subprogram comprising assembler code
- Code optimizer can take Inline Assembler code into account

This appendix presents a series of examples to show you how to use the Inline Assembler. Although it focuses on the Intel x86, the general approach applies also to other processors. It is assumed that you are familiar with Ada and with assembly language programming.

D.1 Basic Assembler Syntax

The assembler used by GNAT and gcc is based not on the Intel assembly language, but rather on a language that descends from the AT&T Unix assembler `as` (and which is often referred to as 'AT&T syntax'). The following table summarizes the main features of `as` syntax and points out the differences from the Intel conventions. See the gcc `as` and `gas` (an `as` macro pre-processor) documentation for further information.

Register names

gcc / `as`: Prefix with '%'; for example `%eax`

Intel: No extra punctuation; for example `eax`

Immediate operand

gcc / as: Prefix with '\$'; for example \$4

Intel: No extra punctuation; for example 4

Address

gcc / as: Prefix with '\$'; for example \$loc

Intel: No extra punctuation; for example loc

Memory contents

gcc / as: No extra punctuation; for example loc

Intel: Square brackets; for example [loc]

Register contents

gcc / as: Parentheses; for example (%eax)

Intel: Square brackets; for example [eax]

Hexadecimal numbers

gcc / as: Leading '0x' (C language syntax); for example 0xA0

Intel: Trailing 'h'; for example A0h

Operand size

gcc / as: Explicit in op code; for example movw to move a 16-bit word

Intel: Implicit, deduced by assembler; for example mov

Instruction repetition

gcc / as: Split into two lines; for example

```
rep
```

```
stosl
```

Intel: Keep on one line; for example rep stosl

Order of operands

gcc / as: Source first; for example movw \$4, %eax

Intel: Destination first; for example mov eax, 4

D.2 A Simple Example of Inline Assembler

The following example will generate a single assembly language statement, `nop`, which does nothing. Despite its lack of run-time effect, the example will be useful in illustrating the basics of the Inline Assembler facility.

```
with System.Machine_Code; use System.Machine_Code;
procedure Nothing is
begin
  Asm ("nop");
end Nothing;
```

`Asm` is a procedure declared in package `System.Machine_Code`; here it takes one parameter, a *template string* that must be a static expression and that will form the generated instruction. `Asm` may be regarded as a compile-time procedure that parses the template string and additional parameters (none here), from which it generates a sequence of assembly language instructions.

The examples in this chapter will illustrate several of the forms for invoking `Asm`; a complete specification of the syntax is found in the `Machine_Code_Insertions` section of the *GNAT Reference Manual*.

Under the standard GNAT conventions, the `Nothing` procedure should be in a file named `nothing.adb`. You can build the executable in the usual way:

```
$ gnatmake nothing
```

However, the interesting aspect of this example is not its run-time behavior but rather the generated assembly code. To see this output, invoke the compiler as follows:

```
$ gcc -c -S -fomit-frame-pointer -gnatp nothing.adb
```

where the options are:

- **-c**
compile only (no bind or link)
- **-S**
generate assembler listing
- **-fomit-frame-pointer**
do not set up separate stack frames
- **-gnatp**
do not add runtime checks

This gives a human-readable assembler version of the code. The resulting file will have the same name as the Ada source file, but with a `.s` extension. In our example, the file `nothing.s` has the following contents:

```
.file "nothing.adb"
gcc2_compiled.:
__gnu_compiled_ada:
```

(continues on next page)

(continued from previous page)

```

.text
  .align 4
.globl __ada_nothing
__ada_nothing:
#APP
  nop
#NO_APP
  jmp L1
  .align 2,0x90
L1:
  ret

```

The assembly code you included is clearly indicated by the compiler, between the `#APP` and `#NO_APP` delimiters. The character before the 'APP' and 'NOAPP' can differ on different targets. For example, GNU/Linux uses `#APP` while on NT you will see `/APP`.

If you make a mistake in your assembler code (such as using the wrong size modifier, or using a wrong operand for the instruction) GNAT will report this error in a temporary file, which will be deleted when the compilation is finished. Generating an assembler file will help in such cases, since you can assemble this file separately using the `as` assembler that comes with `gcc`.

Assembling the file using the command

```
$ as nothing.s
```

will give you error messages whose lines correspond to the assembler input file, so you can easily find and correct any mistakes you made. If there are no errors, `as` will generate an object file `nothing.out`.

D.3 Output Variables in Inline Assembler

The examples in this section, showing how to access the processor flags, illustrate how to specify the destination operands for assembly language statements.

```

with Interfaces; use Interfaces;
with Ada.Text_IO; use Ada.Text_IO;
with System.Machine_Code; use System.Machine_Code;
procedure Get_Flags is
  Flags : Unsigned_32;
  use ASCII;
begin
  Asm ("pushfl"           & LF & HT & -- push flags on stack
      "popl %%eax"       & LF & HT & -- load eax with flags
      "movl %%eax, %0",  -- store flags in variable
      Outputs => Unsigned_32'Asm_Output ("=g", Flags));
  Put_Line ("Flags register:" & Flags'Img);
end Get_Flags;

```

In order to have a nicely aligned assembly listing, we have separated multiple assembler statements

in the Asm template string with linefeed (ASCII.LF) and horizontal tab (ASCII.HT) characters. The resulting section of the assembly output file is:

```
#APP
  pushfl
  popl %eax
  movl %eax, -40(%ebp)
#NO_APP
```

It would have been legal to write the Asm invocation as:

```
Asm ("pushfl popl %%eax movl %%eax, %0")
```

but in the generated assembler file, this would come out as:

```
#APP
  pushfl popl %eax movl %eax, -40(%ebp)
#NO_APP
```

which is not so convenient for the human reader.

We use Ada comments at the end of each line to explain what the assembler instructions actually do. This is a useful convention.

When writing Inline Assembler instructions, you need to precede each register and variable name with a percent sign. Since the assembler already requires a percent sign at the beginning of a register name, you need two consecutive percent signs for such names in the Asm template string, thus `%%eax`. In the generated assembly code, one of the percent signs will be stripped off.

Names such as `%0`, `%1`, `%2`, etc., denote input or output variables: operands you later define using `Input` or `Output` parameters to `Asm`. An output variable is illustrated in the third statement in the Asm template string:

```
movl %%eax, %0
```

The intent is to store the contents of the `eax` register in a variable that can be accessed in Ada. Simply writing `movl %eax, Flags` would not necessarily work, since the compiler might optimize by using a register to hold `Flags`, and the expansion of the `movl` instruction would not be aware of this optimization. The solution is not to store the result directly but rather to advise the compiler to choose the correct operand form; that is the purpose of the `%0` output variable.

Information about the output variable is supplied in the `Outputs` parameter to `Asm`:

```
Outputs => Unsigned_32'Asm_Output ("=g", Flags));
```

The output is defined by the `Asm_Output` attribute of the target type; the general format is

```
Type'Asm_Output (constraint_string, variable_name)
```

The constraint string directs the compiler how to store/access the associated variable. In the example

```
Unsigned_32'Asm_Output ("=m", Flags);
```

the "m" (memory) constraint tells the compiler that the variable `Flags` should be stored in a memory variable, thus preventing the optimizer from keeping it in a register. In contrast,

```
Unsigned_32'Asm_Output ("=r", Flags);
```

uses the "r" (register) constraint, telling the compiler to store the variable in a register.

If the constraint is preceded by the equal character '=', it tells the compiler that the variable will be used to store data into it.

In the `Get_Flags` example, we used the "g" (global) constraint, allowing the optimizer to choose whatever it deems best.

There are a fairly large number of constraints, but the ones that are most useful (for the Intel x86 processor) are the following:

=	output constraint
<i>g</i>	global (i.e., can be stored anywhere)
<i>m</i>	in memory
<i>I</i>	a constant
<i>a</i>	use <code>eax</code>
<i>b</i>	use <code>ebx</code>
<i>c</i>	use <code>ecx</code>
<i>d</i>	use <code>edx</code>
<i>S</i>	use <code>esi</code>
<i>D</i>	use <code>edi</code>
<i>r</i>	use one of <code>eax</code> , <code>ebx</code> , <code>ecx</code> or <code>edx</code>
<i>q</i>	use one of <code>eax</code> , <code>ebx</code> , <code>ecx</code> , <code>edx</code> , <code>esi</code> or <code>edi</code>

The full set of constraints is described in the `gcc` and `as` documentation; note that it is possible to combine certain constraints in one constraint string.

You specify the association of an output variable with an assembler operand through the `%n` notation, where `n` is a non-negative integer. Thus in

```
Asm ("pushfl"           & LF & HT & -- push flags on stack
     "popl %%eax"       & LF & HT & -- load eax with flags
     "movl %%eax, %0",  -- store flags in variable
     Outputs => Unsigned_32'Asm_Output ("=g", Flags));
```

`%0` will be replaced in the expanded code by the appropriate operand, whatever the compiler decided for the `Flags` variable.

In general, you may have any number of output variables:

- Count the operands starting at 0; thus `%0`, `%1`, etc.
- Specify the `Outputs` parameter as a parenthesized comma-separated list of `Asm_Output` attributes

For example:

```
Asm ("movl %%eax, %0" & LF & HT &
     "movl %%ebx, %1" & LF & HT &
     "movl %%ecx, %2",
     Outputs => (Unsigned_32'Asm_Output ("=g", Var_A), -- %0 = Var_A
                 Unsigned_32'Asm_Output ("=g", Var_B), -- %1 = Var_B
                 Unsigned_32'Asm_Output ("=g", Var_C))); -- %2 = Var_C
```

where `Var_A`, `Var_B`, and `Var_C` are variables in the Ada program.

As a variation on the `Get_Flags` example, we can use the constraints string to direct the compiler to store the `eax` register into the `Flags` variable, instead of including the store instruction explicitly in the `Asm` template string:

```
with Interfaces; use Interfaces;
with Ada.Text_IO; use Ada.Text_IO;
with System.Machine_Code; use System.Machine_Code;
procedure Get_Flags_2 is
  Flags : Unsigned_32;
  use ASCII;
begin
  Asm ("pushfl"      & LF & HT & -- push flags on stack
       "popl %%eax", -- save flags in eax
       Outputs => Unsigned_32'Asm_Output ("=a", Flags));
  Put_Line ("Flags register:" & Flags'Img);
end Get_Flags_2;
```

The `"a"` constraint tells the compiler that the `Flags` variable will come from the `eax` register. Here is the resulting code:

```
#APP
  pushfl
  popl %eax
#NO_APP
  movl %eax, -40(%ebp)
```

The compiler generated the store of `eax` into `Flags` after expanding the assembler code.

Actually, there was no need to pop the flags into the `eax` register; more simply, we could just pop the flags directly into the program variable:

```
with Interfaces; use Interfaces;
with Ada.Text_IO; use Ada.Text_IO;
with System.Machine_Code; use System.Machine_Code;
procedure Get_Flags_3 is
  Flags : Unsigned_32;
  use ASCII;
begin
  Asm ("pushfl" & LF & HT & -- push flags on stack
       "pop %0", -- save flags in Flags
       Outputs => Unsigned_32'Asm_Output ("=g", Flags));
```

(continues on next page)

(continued from previous page)

```
Put_Line ("Flags register:" & Flags'Img);  
end Get_Flags_3;
```

D.4 Input Variables in Inline Assembler

The example in this section illustrates how to specify the source operands for assembly language statements. The program simply increments its input value by 1:

```
with Interfaces; use Interfaces;  
with Ada.Text_IO; use Ada.Text_IO;  
with System.Machine_Code; use System.Machine_Code;  
procedure Increment is  
  
  function Incr (Value : Unsigned_32) return Unsigned_32 is  
    Result : Unsigned_32;  
  begin  
    Asm ("incl %0",  
        Outputs => Unsigned_32'Asm_Output ("=a", Result),  
        Inputs  => Unsigned_32'Asm_Input ("a", Value));  
    return Result;  
  end Incr;  
  
  Value : Unsigned_32;  
  
begin  
  Value := 5;  
  Put_Line ("Value before is" & Value'Img);  
  Value := Incr (Value);  
  Put_Line ("Value after is" & Value'Img);  
end Increment;
```

The `Outputs` parameter to `Asm` specifies that the result will be in the `eax` register and that it is to be stored in the `Result` variable.

The `Inputs` parameter looks much like the `Outputs` parameter, but with an `Asm_Input` attribute. The `"="` constraint, indicating an output value, is not present.

You can have multiple input variables, in the same way that you can have more than one output variable.

The parameter count (`%0`, `%1`) etc, still starts at the first output statement, and continues with the input statements.

Just as the `Outputs` parameter causes the register to be stored into the target variable after execution of the assembler statements, so does the `Inputs` parameter cause its variable to be loaded into the register before execution of the assembler statements.

Thus the effect of the `Asm` invocation is:

- load the 32-bit value of `Value` into `eax`

- execute the `incl %eax` instruction
- store the contents of `eax` into the `Result` variable

The resulting assembler file (with `-O2` optimization) contains:

```

_increment__incr.1:
    subl $4,%esp
    movl 8(%esp),%eax
#APP
    incl %eax
#NO_APP
    movl %eax,%edx
    movl %ecx,(%esp)
    addl $4,%esp
    ret

```

D.5 Inlining Inline Assembler Code

For a short subprogram such as the `Incr` function in the previous section, the overhead of the call and return (creating / deleting the stack frame) can be significant, compared to the amount of code in the subprogram body. A solution is to apply Ada's `Inline` pragma to the subprogram, which directs the compiler to expand invocations of the subprogram at the point(s) of call, instead of setting up a stack frame for out-of-line calls. Here is the resulting program:

```

with Interfaces; use Interfaces;
with Ada.Text_IO; use Ada.Text_IO;
with System.Machine_Code; use System.Machine_Code;
procedure Increment_2 is

    function Incr (Value : Unsigned_32) return Unsigned_32 is
        Result : Unsigned_32;
    begin
        Asm ("incl %0",
            Outputs => Unsigned_32'Asm_Output ("=a", Result),
            Inputs  => Unsigned_32'Asm_Input ("a", Value));
        return Result;
    end Incr;
    pragma Inline (Incr);

    Value : Unsigned_32;

begin
    Value := 5;
    Put_Line ("Value before is" & Value'Img);
    Value := Increment (Value);
    Put_Line ("Value after is" & Value'Img);
end Increment_2;

```

Compile the program with both optimization (`-O2`) and inlining (`-gnatn`) enabled.

The `Incr` function is still compiled as usual, but at the point in `Increment` where our function used to be called:

```
pushl %edi
call _increment__incr.1
```

the code for the function body directly appears:

```
movl %esi,%eax
#APP
  incl %eax
#NO_APP
  movl %eax,%edx
```

thus saving the overhead of stack frame setup and an out-of-line call.

D.6 Other Asm Functionality

This section describes two important parameters to the `Asm` procedure: `Clobber`, which identifies register usage; and `Volatile`, which inhibits unwanted optimizations.

D.6.1 The `Clobber` Parameter

One of the dangers of intermixing assembly language and a compiled language such as Ada is that the compiler needs to be aware of which registers are being used by the assembly code. In some cases, such as the earlier examples, the constraint string is sufficient to indicate register usage (e.g., "a" for the `eax` register). But more generally, the compiler needs an explicit identification of the registers that are used by the Inline Assembly statements.

Using a register that the compiler doesn't know about could be a side effect of an instruction (like `mull` storing its result in both `eax` and `edx`). It can also arise from explicit register usage in your assembly code; for example:

```
Asm ("movl %0, %%ebx" & LF & HT &
     "movl %%ebx, %1",
     Outputs => Unsigned_32'Asm_Output ("=g", Var_Out),
     Inputs  => Unsigned_32'Asm_Input  ("g", Var_In));
```

where the compiler (since it does not analyze the `Asm` template string) does not know you are using the `ebx` register.

In such cases you need to supply the `Clobber` parameter to `Asm`, to identify the registers that will be used by your assembly code:

```
Asm ("movl %0, %%ebx" & LF & HT &
     "movl %%ebx, %1",
     Outputs => Unsigned_32'Asm_Output ("=g", Var_Out),
     Inputs  => Unsigned_32'Asm_Input  ("g", Var_In),
     Clobber => "ebx");
```

The `Clobber` parameter is a static string expression specifying the register(s) you are using. Note that register names are *not* prefixed by a percent sign. Also, if more than one register is used then their names are separated by commas; e.g., `"eax, ebx"`

The `Clobber` parameter has several additional uses:

- Use 'register' name `cc` to indicate that flags might have changed
- Use 'register' name `memory` if you changed a memory location

D.6.2 The Volatile Parameter

Compiler optimizations in the presence of Inline Assembler may sometimes have unwanted effects. For example, when an `Asm` invocation with an input variable is inside a loop, the compiler might move the loading of the input variable outside the loop, regarding it as a one-time initialization.

If this effect is not desired, you can disable such optimizations by setting the `Volatile` parameter to `True`; for example:

```
Asm ("movl %0, %%ebx" & LF & HT &
     "movl %%ebx, %1",
     Outputs => Unsigned_32'Asm_Output ("=g", Var_Out),
     Inputs  => Unsigned_32'Asm_Input  ("g", Var_In),
     Clobber => "ebx",
     Volatile => True);
```

By default, `Volatile` is set to `False` unless there is no `Outputs` parameter.

Although setting `Volatile` to `True` prevents unwanted optimizations, it will also disable other optimizations that might be important for efficiency. In general, you should set `Volatile` to `True` only if the compiler's optimizations have created problems.

This page is intentionally left blank.

GNU FREE DOCUMENTATION LICENSE

Version 1.3, 3 November 2008

Copyright 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc <https://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

E.1 Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

E.2 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The **Document**, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “**you**”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “**Modified Version**” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “**Secondary Section**” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “**Invariant Sections**” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “**Cover Texts**” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “**Transparent**” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called **Opaque**.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “**Title Page**” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “**publisher**” means any person or entity that distributes copies of the Document to the public.

A section “**Entitled XYZ**” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “**Acknowledgements**”, “**Dedications**”, “**Endorsements**”, or “**History**”.) To “**Preserve the Title**” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in

this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

E.3 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

E.4 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

E.5 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

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

- M. Delete any section Entitled “Endorsements”. Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled “Endorsements” or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

E.6 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”.

E.7 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

E.8 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

E.9 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

E.10 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

E.11 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.

E.12 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.

E.13 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 © 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.

INDEX

Symbols

- `__gnat_malloc`, 172
- A (*gnatbind*), 170
- A (*gnatmake*), 100
- B (*gcc*), 107
- B (*gnatlink*), 184
- C (*gnatmake*), 96
- C (*gcc*), 88
- C (*gnatprep*), 61
- C= (*gnatmake*), 96
- D (*gnatbind*), 171
- D (*gnatmake*), 96
- D (*gnatclean*), 192
- D (*gnatname*), 31
- D (*gnatprep*), 62
- E (*gnatbind*), 171
- E (*gprof*), 217
- Ea (*gnatbind*), 171
- Es (*gnatbind*), 171
- F (*gnatbind*), 172
- F (*gnatmake*), 97
- F (*gnatclean*), 192
- F (*gprof*), 217
- H (*gnatbind*), 172
- H32 (*gnatbind*), 172
- H64 (*gnatbind*), 172
- I (*gcc*), 118
- I (*gnatbind*), 172, 181
- I (*gnatmake*), 100
- I (*gnatclean*), 193
- I (*gnatdll*), 272
- I (*gnatls*), 195
- I- (*gcc*), 118
- I- (*gnatbind*), 172
- I- (*gnatmake*), 100
- I- (*gnatclean*), 193
- I- (*gnatls*), 195
- K (*gnatbind*), 172, 178
- L (*gnatbind*), 172
- L (*gnatmake*), 100
- M (*gnatbind*), 172, 176
- M (*gnatlink*), 184
- M (*gnatmake*), 98
- M= (*gnatlink*), 184
- O (*gcc*), 119
- O (*gnatbind*), 173, 179
- O (*gcc*), 219
- P (*gnatbind*), 173
- P (*gnatmake*), 94
- P (*gnatclean*), 192
- P (*gnatname*), 31
- R (*gnatbind*), 173
- Ra (*gnatbind*), 173
- S (*gcc*), 119
- S (*gnatbind*), 173
- T (*gnatbind*), 174
- T (*gnatprep*), 62
- T0 option, 253
- U (*gnatmake*), 99
- V (*gcc*), 119
- V (*gnatbind*), 175
- Wall (*gcc*), 142
- Werror (*gcc*), 142
- Wstack-usage (*gcc*), 142
- Wuninitialized (*gcc*), 141
- Wunused (*gcc*), 141
- Wx (*gnatbind*), 175, 176
- X (*gnatclean*), 193
- Xnnn (*gnatbind*), 175
- GCC= (*gnatchop*), 39
- GCC=compiler_name (*gnatlink*), 184
- GCC=compiler_name (*gnatmake*), 94
- GNATBIND=binder_name (*gnatmake*), 94
- GNATLINK=linker_name (*gnatmake*), 95

- LINK= (*gnatlink*), 184
- RTS (*gcc*), 119
- RTS (*gnatbind*), 173
- RTS (*gnatmake*), 101
- RTS (*gnatls*), 195
- RTS option, 253
- RTS=sjlj (*gnatmake*), 167
- RTS=zcx (*gnatmake*), 167
- as (*dlltool*), 274
- base-file (*dlltool*), 274
- create-missing-dirs (*gnatmake*), 95
- def (*dlltool*), 274
- demangle (*gprof*), 217
- dllname (*dlltool*), 274
- help (*gnatbind*), 170
- help (*gnatlink*), 183
- help (*gnatmake*), 94
- help (*dlltool*), 274
- help (*gnatchop*), 38
- help (*gnatclean*), 192
- help (*gnatls*), 194
- help (*gnatname*), 30
- help (*gnatprep*), 61
- output-exp (*dlltool*), 274
- output-lib (*dlltool*), 274
- version (*gnatbind*), 170
- version (*gnatlink*), 183
- version (*gnatmake*), 94
- version (*gnatchop*), 38
- version (*gnatclean*), 192
- version (*gnatls*), 194
- version (*gnatname*), 30
- version (*gnatprep*), 61
- a (*gnatbind*), 170
- a (*gnatmake*), 95
- a (*gnatdll*), 271
- a (*gnatls*), 194
- aI (*gnatbind*), 170, 181
- aI (*gnatmake*), 100
- aI (*gnatls*), 195
- aL (*gnatmake*), 100
- a0 (*gnatbind*), 170, 181
- a0 (*gnatmake*), 100
- a0 (*gnatclean*), 193
- a0 (*gnatls*), 195
- aP (*gnatls*), 195
- b (*gcc*), 107
- b (*gnatbind*), 171, 176
- b (*gnatmake*), 96
- b (*gnatdll*), 271
- b (*gnatprep*), 61
- bargs (*gnatmake*), 101
- bargs (*gnatdll*), 271
- c (*gcc*), 107
- c (*gnatbind*), 171, 178
- c (*gnatmake*), 96
- c (*gnatchop*), 38
- c (*gnatclean*), 192
- c (*gnatname*), 31
- c (*gnatprep*), 61
- cargs (*gnatmake*), 101
- d (*gnatmake*), 96
- d (*gnatdll*), 271
- d (*gnatls*), 194
- d (*gnatname*), 31
- dnn[k|m] (*gnatbind*), 171
- e (*gnatbind*), 171, 178
- e (*gnatdll*), 271
- e (*gprof*), 217
- eI (*gnatmake*), 97
- eL (*gnatmake*), 97
- eS (*gnatmake*), 97
- f (*gnatbind*), 171, 177
- f (*gnatlink*), 183
- f (*gnatmake*), 97
- f (*gnatname*), 31
- f (*gprof*), 217
- fada-spec-parent (*gcc*), 88
- fcallgraph-info (*gcc*), 107
- fdata-sections (*gcc*), 232
- fdiagnostics-format (*gcc*), 107
- fdump-ada-spec (*gcc*), 88
- fdump-ada-spec-slim (*gcc*), 88
- fdump-scops (*gcc*), 107
- ffunction-sections (*gcc*), 232
- fgnat-encodings (*gcc*), 107, 166
- files (*gnatls*), 195
- flto (*gcc*), 107
- fno-inline (*gcc*), 108
- fno-inline (*gcc*), 223
- fno-inline-functions (*gcc*), 108
- fno-inline-functions (*gcc*), 223
- fno-inline-functions-called-once (*gcc*), 108
- fno-inline-functions-called-once (*gcc*), 223
- fno-inline-small-functions (*gcc*), 108
- fno-inline-small-functions (*gcc*), 223

-fno-ivopts (*gcc*), 108
 -fno-strict-aliasing (*gcc*), 108
 -fno-strict-overflow (*gcc*), 108
 -fstack-check (*gcc*), 108, 157
 -fstack-check (*gcc*), 244
 -fstack-usage, 244
 -fstack-usage (*gcc*), 108
 -fuse-ld=name, 169
 -fverbose-asm (*gcc*), 119
 -g (*gcc*), 108
 -g (*gnatlink*), 183
 -g (*gnatmake*), 97
 -g (*gcc*), 220
 -g (*gnatdll*), 272
 -gnat05 (*gcc*), 108, 159
 -gnat12 (*gcc*), 109, 159
 -gnat2005 (*gcc*), 109, 159
 -gnat2012 (*gcc*), 109, 159
 -gnat2022 (*gcc*), 109, 159
 -gnat83 (*gcc*), 109, 159
 -gnat95 (*gcc*), 109, 159
 -gnatA (*gcc*), 109
 -gnatB (*gcc*), 109
 -gnatC (*gcc*), 110
 -gnatD (*gcc*), 164
 -gnatD[nn] (*gcc*), 110
 -gnatE (*gcc*), 115, 157
 -gnatE (*gnat*), 316
 -gnatF (*gcc*), 115
 -gnatG (*gcc*), 163
 -gnatG[nn] (*gcc*), 115
 -gnatH (*gcc*), 115
 -gnatH (*gnat*), 317
 -gnatI (*gcc*), 115
 -gnatJ (*gcc*), 115
 -gnatJ (*gnat*), 317
 -gnatL (*gcc*), 116, 163, 165
 -gnatN (*gcc*), 116, 162
 -gnatN switch, 44
 -gnatQ (*gcc*), 117, 123
 -gnatR (*gcc*), 117, 165
 -gnatS (*gcc*), 118, 166
 -gnatT (*gcc*), 118
 -gnatU (*gcc*), 118, 122
 -gnatV (*gcc*), 118
 -gnatVa (*gcc*), 146
 -gnatVc (*gcc*), 146
 -gnatVd (*gcc*), 146
 -gnatVe (*gcc*), 146
 -gnatVf (*gcc*), 146
 -gnatVi (*gcc*), 147
 -gnatVm (*gcc*), 147
 -gnatVn (*gcc*), 147
 -gnatVo (*gcc*), 147
 -gnatVp (*gcc*), 147
 -gnatVr (*gcc*), 147
 -gnatVs (*gcc*), 147
 -gnatVt (*gcc*), 148
 -gnatW (*gcc*), 118, 160
 -gnatX (*gcc*), 118, 160
 -gnatX0 (*gcc*), 118, 159
 -gnat-p (*gcc*), 117, 156
 -gnata (*gcc*), 109, 144
 -gnata switch, 56
 -gnatb (*gcc*), 109, 122
 -gnatc (*gcc*), 109, 158
 -gnatd (*gcc*), 110, 163
 -gnatdc switch, 205
 -gnateA (*gcc*), 110
 -gnateC (*gcc*), 111
 -gnateD (*gcc*), 111
 -gnateD (*gcc*), 67
 -gnateE (*gcc*), 111
 -gnateF (*gcc*), 111
 -gnateG (*gcc*), 111
 -gnateI (*gcc*), 111
 -gnateP (*gcc*), 112
 -gnateS (*gcc*), 112
 -gnateT (*gcc*), 112
 -gnateV (*gcc*), 114
 -gnateY (*gcc*), 114
 -gnateb (*gcc*), 110
 -gnatec (*gcc*), 111
 -gnated (*gcc*), 111
 -gnatef (*gcc*), 111
 -gnateg (*gcc*), 111
 -gnatei (*gcc*), 111
 -gnatel (*gcc*), 111, 112
 -gnatel (*gnat*), 316
 -gnatem (*gcc*), 112, 167
 -gnatep (*gcc*), 112
 -gnatep (*gcc*), 65
 -gnatet=file (*gcc*), 112
 -gnateu (*gcc*), 114
 -gnatf (*gcc*), 115, 122
 -gnatg (*gcc*), 115

-gnath (*gcc*), 115
-gnati (*gcc*), 115, 160
-gnatjnn (*gcc*), 115, 123
-gnatk (*gcc*), 116, 161
-gnatl (*gcc*), 116, 121
-gnatl=fname (*gcc*), 121
-gnatm (*gcc*), 116, 122
-gnatn (*gcc*), 116, 162
-gnatn switch, 43
-gnato (*gcc*), 218, 238
-gnato0 (*gcc*), 116
-gnato? (*gcc*), 238
-gnato?? (*gcc*), 117, 156
-gnato?? (*gcc*), 238
-gnatp (*gcc*), 117, 155
-gnatp (*gcc*), 218
-gnatq (*gcc*), 117, 123
-gnatr (*gcc*), 117, 165
-gnats (*gcc*), 117, 157
-gnatu (*gcc*), 118, 162
-gnatv (*gcc*), 118, 120
-gnatw (*gcc*), 118
-gnatw.A (*gcc*), 126
-gnatw.C (*gcc*), 128
-gnatw.H (*gcc*), 130
-gnatw.I (*gcc*), 131
-gnatw.J (*gcc*), 131
-gnatw.L (*gcc*), 132
-gnatw.M (*gcc*), 133
-gnatw.N (*gcc*), 133
-gnatw.O (*gcc*), 134
-gnatw.P (*gcc*), 134
-gnatw.Q (*gcc*), 136
-gnatw.R (*gcc*), 137
-gnatw.S (*gcc*), 137
-gnatw.T (*gcc*), 138
-gnatw.U (*gcc*), 138
-gnatw.V (*gcc*), 139
-gnatw.W (*gcc*), 140
-gnatw.Y (*gcc*), 141
-gnatw.Z (*gcc*), 141
-gnatw.a (*gcc*), 126
-gnatw.b (*gcc*), 127
-gnatw.c (*gcc*), 128
-gnatw.d (*gcc*), 128, 129
-gnatw.e (*gcc*), 129
-gnatw.f (*gnat*), 317
-gnatw.g (*gcc*), 130
-gnatw.h (*gcc*), 130
-gnatw.i (*gcc*), 131
-gnatw.j (*gcc*), 131
-gnatw.k (*gcc*), 132
-gnatw.l (*gcc*), 132
-gnatw.m (*gcc*), 133
-gnatw.n (*gcc*), 133
-gnatw.o (*gcc*), 134
-gnatw.p (*gcc*), 134
-gnatw.q (*gcc*), 135
-gnatw.r (*gcc*), 136
-gnatw.s (*gcc*), 137
-gnatw.t (*gcc*), 138
-gnatw.u (*gcc*), 138
-gnatw.v (*gcc*), 139
-gnatw.w (*gcc*), 139
-gnatw.y (*gcc*), 141
-gnatw.z (*gcc*), 141
-gnatwA (*gcc*), 126
-gnatwB (*gcc*), 127
-gnatwC (*gcc*), 128
-gnatwD (*gcc*), 128
-gnatwE (*gcc*), 129
-gnatwF (*gcc*), 129
-gnatwG (*gcc*), 130
-gnatwH (*gcc*), 130
-gnatwI (*gcc*), 131
-gnatwJ (*gcc*), 131
-gnatwK (*gcc*), 132
-gnatwL (*gcc*), 132
-gnatwM (*gcc*), 133
-gnatwO (*gcc*), 133
-gnatwP (*gcc*), 134
-gnatwQ (*gcc*), 135
-gnatwR (*gcc*), 136
-gnatwT (*gcc*), 137
-gnatwU (*gcc*), 138
-gnatwV (*gcc*), 139
-gnatwW (*gcc*), 139
-gnatwX (*gcc*), 140
-gnatwY (*gcc*), 140
-gnatwZ (*gcc*), 141
-gnatw_A, 126
-gnatw_C (*gcc*), 128
-gnatw_P (*gcc*), 135
-gnatw_Q (*gcc*), 136
-gnatw_R (*gcc*), 137
-gnatw_a, 126

-gnatw_c (*gcc*), 128
 -gnatw_p (*gcc*), 134
 -gnatw_q (*gcc*), 136
 -gnatw_r (*gcc*), 137
 -gnatwa (*gcc*), 125
 -gnatwb (*gcc*), 127
 -gnatwc (*gcc*), 127
 -gnatwd (*gcc*), 128
 -gnatwe (*gcc*), 129
 -gnatwf (*gcc*), 129
 -gnatwg (*gcc*), 130
 -gnatwh (*gcc*), 130
 -gnatwi (*gcc*), 130
 -gnatwj (*gcc*), 131
 -gnatwk (*gcc*), 131
 -gnatwl (*gcc*), 132
 -gnatwl (*gnat*), 317
 -gnatwm (*gcc*), 132, 140
 -gnatwn (*gcc*), 133
 -gnatwo (*gcc*), 133
 -gnatwp (*gcc*), 134
 -gnatwq (*gcc*), 135
 -gnatwr (*gcc*), 136
 -gnatws (*gcc*), 137
 -gnatwt (*gcc*), 137
 -gnatwu (*gcc*), 138
 -gnatwv (*gcc*), 139
 -gnatww (*gcc*), 139
 -gnatwx (*gcc*), 140
 -gnatwy (*gcc*), 140
 -gnatwz (*gcc*), 141
 -gnatx (*gcc*), 118, 166
 -gnaty (*gcc*), 118, 148
 -gnaty+ (*gcc*), 155
 -gnatyA (*gcc*), 149
 -gnatyB (*gcc*), 149
 -gnatyC (*gcc*), 150
 -gnatyD (*gcc*), 150
 -gnatyI (*gcc*), 151
 -gnatyLnnn (*gcc*), 152
 -gnatyMnnn (*gcc*), 152
 -gnatyN (*gcc*), 153
 -gnatyO (*gcc*), 153
 -gnatyS (*gcc*), 153
 -gnaty- (*gcc*), 155
 -gnatya (*gcc*), 149
 -gnatyb (*gcc*), 149
 -gnatyc (*gcc*), 149
 -gnatyd (*gcc*), 150
 -gnatye (*gcc*), 150
 -gnatyf (*gcc*), 150
 -gnatyg (*gcc*), 150
 -gnatyh (*gcc*), 150
 -gnatyi (*gcc*), 151
 -gnatyk (*gcc*), 151
 -gnatyl (*gcc*), 151
 -gnatym (*gcc*), 152
 -gnatyn (*gcc*), 153
 -gnatyo (*gcc*), 153
 -gnatyp (*gcc*), 153
 -gnatyr (*gcc*), 153
 -gnatys (*gcc*), 153
 -gnatyt (*gcc*), 154
 -gnatyu (*gcc*), 154
 -gnatyx (*gcc*), 154
 -gnaty y (*gcc*), 154
 -gnaty[0-9] (*gcc*), 148
 -gnatz (*gcc*), 118
 -h (*gnatbind*), 172, 178
 -h (*gnatclean*), 192
 -h (*gnatdll*), 272
 -h (*gnatls*), 194
 -h (*gnatname*), 31
 -i (*gnatmake*), 97
 -j (*gnatmake*), 97
 -k (*gnatbind*), 172
 -k (*gnatmake*), 98
 -k (*dlltool*), 274
 -k (*gnatchop*), 38
 -k (*gnatdll*), 272
 -l (*gnatbind*), 172, 179
 -l (*gnatmake*), 98
 -l (*gnatdll*), 272
 -larges (*gnatmake*), 101
 -larges (*gnatdll*), 272
 -m (*gnatbind*), 172, 176
 -m (*gnatmake*), 98
 -margs (*gnatmake*), 101
 -minimal (*gnatbind*), 173
 -mwindows, 256
 -n (*gnatbind*), 173, 179, 180
 -n (*gnatlink*), 184
 -n (*gnatmake*), 98
 -n (*gnatclean*), 192
 -n (*gnatdll*), 272
 -nostdinc (*gcc*), 118

- nostdinc (*gnatbind*), 173
- nostdinc (*gnatmake*), 100
- nostdlib (*gcc*), 119
- nostdlib (*gnatbind*), 173
- nostdlib (*gnatmake*), 100
- o (*gnatbind*), 173
- o (*gcc*), 118
- o (*gnatbind*), 179, 180
- o (*gnatlink*), 184
- o (*gnatmake*), 98
- o (*gnatls*), 194
- p (*gnatbind*), 173, 178
- p (*gnatmake*), 99
- p (*gnatchop*), 38
- pass-exit-codes (*gcc*), 119, 162
- pg (*gcc*), 216
- pg (*gnatlink*), 216
- q (*gnatmake*), 99
- q (*gnatchop*), 38
- q (*gnatclean*), 192
- q (*gnatdll*), 272
- r (*gnatbind*), 179
- r (*gnatchop*), 38
- r (*gnatclean*), 192
- r (*gnatprep*), 62
- s (*gnatbind*), 173, 176
- s (*gnatmake*), 99
- s (*gnatls*), 194
- s (*gnatprep*), 62
- shared (*gnatbind*), 174
- static (*gnatbind*), 174
- t (*gnatbind*), 174, 177
- u (*gnatbind*), 175
- u (*gnatmake*), 99
- u (*gnatls*), 194
- u (*gnatprep*), 62
- v (*gcc*), 119
- v (*gnatbind*), 175, 176
- v (*gnatlink*), 184
- v (*gnatmake*), 99
- v (*dlltool*), 274
- v (*gnatchop*), 39
- v (*gnatclean*), 192
- v (*gnatdll*), 272
- v (*gnatls*), 195
- v (*gnatname*), 31
- v (*gnatprep*), 62
- v -v (*gnatlink*), 184

- v -v (*gnatname*), 32
- vP (*gnatclean*), 193
- vl (*gnatmake*), 99
- vm (*gnatmake*), 99
- w (*gcc*), 119
- w (*gnatbind*), 175
- w (*gcc*), 142
- w (*gnatchop*), 39
- we (*gnatbind*), 177
- ws (*gnatbind*), 177
- x (*gnatbind*), 175, 176
- x (*gnatmake*), 99
- x (*gnatname*), 32
- xdr (*gnatbind*), 175
- y (*gnatbind*), 175
- z (*gnatbind*), 175, 180
- z (*gnatmake*), 100

A

- Abnormal Termination or Failure to Terminate, 205
- Access before elaboration, 155
- access before elaboration, 155
- activate every optional warning, 129
- ACVC, 159
- Ada, 181
- Ada 2005 Language Reference Manual, 14
- Ada 2005 mode, 159
- Ada 2012 mode, 159
- Ada 2022 mode, 159
- Ada 83 mode, 159
- Ada 83 tests, 159
- Ada 95 Language Reference Manual, 14
- Ada 95 mode, 159
- Ada compatibility issues warnings, 140
- Ada expressions (*in gdb*), 201
- Ada language extensions, 159, 160
- Ada Library Information files, 44
- Ada.Characters.Latin_1, 22
- ADA_INCLUDE_PATH, 50, 105
- ADA_OBJECTS_PATH, 50, 181
- ADA_PRJ_INCLUDE_FILE, 105
- ADA_PRJ_OBJECTS_FILE, 181
- ADA_PROJECT_PATH, 48
- adafinal, 180
- adainit, 180
- Address Clauses, 133
- ALI files, 44

- Aliasing, 226, 229
- alternative, 27
- Annex A (*in Ada Reference Manual*), 206
- Annex B (*in Ada reference Manual*), 206
- Anonymous allocators, 126, 127
- APIENTRY, 259
- Asm, 71
- Assert, 144
- Assert failures, 126
- Assertions, 144
- Atomic, 230
- Atomic Synchronization, 133
- attach to process, 279
- B**
- Bad fixed values, 127
- Biased representation, 127
- Binder, 180
- Binder consistency checks, 177
- Binder output (*example*), 283
- Binder output file, 69
- Binding generation (*for Ada specs*), 89
- Binding generation (*for C and C++ headers*), 85
- BINUTILS_ROOT, 74
- bit order warnings, 139
- Breakpoints and tasks, 203
- building, 266, 267, 276
- Building the GNAT Run-Time Library, 54
- C**
- C, 72
- C headers (*binding generation*), 85, 89
- C varargs function, 72
- C++, 72
- C++ headers (*binding generation*), 85
- C_INCLUDE_PATH, 74
- Calling Conventions, 71
- cannot generate code, 103
- Check, 156, 157
- Checks, 155–157, 218
- Checks (*overflow*), 234
- COBOL, 71
- code page 437 (*IBM PC*), 23
- code page 850 (*IBM PC*), 23
- Combining GNAT switches, 119
- Command Line Argument Expansion, 256
- Command line length, 183
- Compatibility with Ada 83, 159
- compilation (*definition*), 22
- Compilation model, 21
- Compile_Time_Error, 128
- Compile_Time_Warning, 128
- compiling, 276
- Component clause, 128
- Conditional compilation, 54, 55
- Conditionals, 127
- configuration, 40
- Configuration pragmas, 40
- Consistency checks, 177
- CONSOLE Subsystem, 256
- constant, 127
- Convention Ada, 71
- Convention Asm, 71
- Convention Assembler, 71
- Convention C, 72
- Convention C++, 72
- Convention COBOL, 71
- Convention Default, 72
- Convention DLL, 73
- Convention External, 72
- Convention Fortran, 72
- Convention Stdcall, 73
- Convention Stubbed, 73
- Convention Win32, 73
- Conventions, 14
- CR, 21
- Cyrillic, 23
- D**
- Deactivated code, 137
- Debug, 144
- Debug Pool, 247
- Debugger, 197
- Debugging, 197
- Debugging Generic Units, 203
- Debugging information, 183
- Debugging optimized code, 220
- Debugging options, 163
- Default, 72
- Definition file, 264
- Deleted code, 137
- Dependencies, 98
- Dependency rules (*compilation*), 93
- Dereferencing, 128
- Dimension aspect, 239, 242

- Dimension Vector (*for a dimensioned subtype*), 242
 - Dimension_System aspect, 239, 242
 - Dimensionable type, 242
 - Dimensionality analysis, 239
 - Dimensioned subtype, 242
 - Division by zero, 155
 - division by zero, 155
 - DLL, 73, 261
 - DLL debugging, 277, 279
 - DLLs, 266, 267
 - DLLs and elaboration, 269
 - DLLs and finalization, 270
 - Dynamic elaboration model, 308
- ## E
- Elaboration, 132
 - elaboration, 157
 - Elaboration checks, 157
 - Elaboration control, 299
 - Elaboration order control, 90
 - End of source file
 - Source file, end, 22
 - environment variable
 - ADA_INCLUDE_PATH, 50, 105
 - ADA_OBJECTS_PATH, 50, 181
 - ADA_PRJ_INCLUDE_FILE, 105
 - ADA_PRJ_OBJECTS_FILE, 181
 - BINUTILS_ROOT, 74
 - C_INCLUDE_PATH, 74
 - GCC_EXEC_PREFIX, 74
 - GCC_ROOT, 74
 - PATH, 105, 181
 - TMP, 256
 - Error messages, 122
 - EUC Coding, 24
 - Exceptions (*in gdb*), 202
 - Export table, 268
 - Export/Import pragma warnings, 140
 - External, 72
- ## F
- Features, 131
 - FF, 21
 - File cleanup tool, 191
 - File names, 27
 - File Naming Conventions, 29
 - File naming schemes, 27
 - Fixed-point Small value, 127
 - Floating-Point Operations, 223
 - for gnatmake, 100
 - for profiling, 216
 - Foreign Languages, 71
 - Formals, 129
 - Fortran, 72
- ## G
- GCC_EXEC_PREFIX, 74
 - GCC_ROOT, 74
 - gdb, 197**
 - Generic formal parameters, 159
 - Generics, 43, 203
 - GNAT, 181
 - GNAT (*package*), 206
 - GNAT compilation model, 21
 - GNAT extensions, 159, 160
 - GNAT library, 91
 - GNAT Run-Time Library, 54
 - gnat.adc, 27, 42
 - gnat1, 104
 - gnat_argc, 181
 - gnat_argv, 181
 - GNAT_INIT_SCALARS, 174
 - gnatbind, 169
 - gnatchop, 35
 - gnatclean, 191
 - gnatdll, 271
 - gnatkr, 32
 - gnatlink, 182
 - gnatls, 193
 - gnatmake, 93
 - gnatname, 29
 - gnatprep, 60, 60
 - GNU make, 185
 - GNU/Linux, 253
 - GPR_PROJECT_PATH, 48
 - gprof, 215
- ## H
- Hiding of Declarations, 130
 - HT, 21
- ## I
- implicit, 128
 - Implicit dereferencing, 128
 - Import library, 264

Improving performance, 218
 in binder, 177
 including, 183
 Inline, 43, 222
 Inline Assembler, 323
 Inlining, 91, 134
 Interfaces, 181
 Interfacing to Ada, 71
 Interfacing to Assembly, 71
 Interfacing to C, 72
 Interfacing to C varargs function, 72
 Interfacing to C++, 72
 Interfacing to COBOL, 71
 Interfacing to Fortran, 72
 ISO 8859-15, 23
 ISO 8859-2, 23
 ISO 8859-3, 23
 ISO 8859-4, 23
 ISO 8859-5, 23

L

Latin-1, 21, 22
 Latin-2, 23
 Latin-3, 23
 Latin-4, 23
 Latin-9, 23
 Layout, 135
 Legacy elaboration models, 309
 LF, 21
 Library browser, 193
 Library building and using, 46
 Linker libraries, 100
 Linux, 253

M

Machine_Overflows, 157
 make (*GNU*), 185
 memory corruption, 247
 Memory Pool, 246
 Microsoft Visual Studio, 276
 missing, 128
 Mixed Language Programming, 68
 MKS_Type type, 239
 multiple input files, 180
 Multiple units, 158

N

naming scheme, 97

No information messages for why package spec
 needs body, 141
 No_Strict_Aliasing, 226
 non-symbolic, 207

O

obsolescent, 131
 Obsolescent features, 131
 Optimization and debugging, 220
 Optimization Switches, 224, 226
 Order of elaboration, 299
 OS X, 282
 Other Ada compilers, 71
 overflow, 156, 218
 Overflow checks, 156, 218, 234
 Overflow mode, 156

P

Package spec needing body, 141
 Parallel make, 97
 Parameter order, 134
 Parentheses, 135
 Passive Task, 231
 PATH, 105, 181
 pool, 246, 247
 Postcondition, 144
 pragma Assert, 56
 pragma Assertion_Policy, 56
 pragma Debug, 56
 pragma Debug_Policy, 57
 pragma Elaborate (*Unit*), 305
 pragma Elaborate_All (*Unit*), 306
 pragma Elaborate_Body, 304
 pragma Export, 92
 pragma Inline, 222
 pragma Overflow_Mode, 237
 pragma Preelaborate, 304
 pragma Pure, 304
 pragma Restrictions, 165
 pragma Suppress, 218
 pragma Task_Dispatching_Policy, 253
 pragma Time_Slice, 253
 pragma Unsuppress, 218
 Pragmas, 40, 130
 Precondition, 144
 Preprocessing, 60
 Preprocessing (*gnatprep*), 60

Preprocessors (*contrasted with conditional compilation*), 55

producing list, 98

Profiling, 215

R

rc, 276

rebuilding, 54

Rebuilding the GNAT Run-Time Library, 54

Recompilation (*by gnatmake*), 101

Record Representation (*component sizes*), 137

Record Representation (*gaps*), 130

Relaxed elaboration mode, 309

Remote Debugging with gdbserver, 204

Resources, 274, 276

RTL, 118

Run-time libraries (*platform-specific information*), 251

Run-Time Library, 54

S

s-digemk.ads file, 240

SCHED_FIFO scheduling policy, 253

SCHED_OTHER scheduling policy, 253

SCHED_RR scheduling policy, 253

Search paths, 100

setjmp/longjmp Exception Model, 251

Shift JIS Coding, 24

Size/Alignment warnings, 141

SJLJ (*setjmp/longjmp Exception Model*), 251

Small value, 127

Source file, end

 End of source file, 22

Source files, 100, 170, 193

Source_File_Name pragma, 27

Source_Reference pragmas, 38

SPARK elaboration model, 309

spec (*definition*), 22

Stack Overflow Checking, 157, 244

stack overflow checking, 155, 157

stack traceback, 207

stack unwinding, 207

Stand-alone libraries, 50

Static elaboration model, 308

Static Stack Usage Analysis, 244

Stdcall, 73, 259

stderr, 120

storage, 246, 247

Strict Aliasing, 226

String indexing warnings, 139

Stubbed, 73

Style checking, 148

SUB (*control character*), 22

Subtype predicates, 144

Subunits, 43

Subunits (*and conditional compilation*), 58

Suppress, 156, 218

suppressing, 122, 155, 156

Suppressing checks, 155, 156

suppressing search, 100, 193

symbolic, 212

symbolic links, 97

syntax checking, 158

System, 181

System (*package in Ada Reference Manual*), 206

System.Dim.Mks package (*GNAT library*), 239

System.IO, 105

T

Task switching (*in gdb*), 203

Tasking and threads libraries, 251

Tasks (*in gdb*), 202

Temporary files, 256

Text_IO and performance, 232

Threads libraries and tasking, 251

Time stamp checks, 177

TMP, 256

traceback, 207, 212

treat as error, 129

Type invariants, 144

typographical, 14

Typographical conventions, 14

U

Unassigned variable warnings, 139

Unchecked_Conversion warnings, 141

unrecognized, 130

unreferenced, 129

Unsuppress, 157, 218

Upper-Half Coding, 24

use by binder, 170

use with GNAT DLLs, 276

using, 276

Unused subprogram/data elimination, 232

V

Validity Checking, 145
varargs function interfaces, 72
Version skew (*avoided by "gnatmake"*), 19
Volatile parameter, 333
VT, 21

W

Warning messages, 123
Warnings, 129, 177
warnings, 132–135, 137
Warnings Off control, 139
Win32, 73
Windows, 254
windows, 274
WINDOWS Subsystem, 256
windres, 276

Z

ZCX (*Zero-Cost Exceptions*), 251
Zero Cost Exceptions, 167
Zero-Cost Exceptions, 251