

GNAT User's Guide for Native Platforms

GNAT User's Guide for Native Platforms , Dec 05, 2025

AdaCore

Copyright © 2008-2025, Free Software Foundation

Table of Contents

1	About This Guide	2
1.1	What This Guide Contains	2
1.2	What You Should Know before Reading This Guide.....	3
1.3	Related Information	3
1.4	Conventions	3
2	Getting Started with GNAT	4
2.1	System Requirements	4
2.2	Running GNAT	4
2.3	Running a Simple Ada Program	4
2.4	Running a Program with Multiple Units	5
3	The GNAT Compilation Model	7
3.1	Source Representation	7
3.2	Foreign Language Representation	8
3.2.1	Latin-1	8
3.2.2	Other 8-Bit Codes	8
3.2.3	Wide_Character Encodings	9
3.2.4	Wide_Wide_Character Encodings	10
3.3	File Naming Topics and Utilities	11
3.3.1	File Naming Rules	11
3.3.2	Using Other File Names	12
3.3.3	Alternative File Naming Schemes	13
3.3.4	Handling Arbitrary File Naming Conventions with <code>gnatname</code> ..	15
3.3.4.1	Arbitrary File Naming Conventions	15
3.3.4.2	Running <code>gnatname</code>	15
3.3.4.3	Switches for <code>gnatname</code>	16
3.3.4.4	Examples of <code>gnatname</code> Usage	18
3.3.5	File Name Krunching with <code>gnatkr</code>	18
3.3.5.1	About <code>gnatkr</code>	18
3.3.5.2	Using <code>gnatkr</code>	18
3.3.5.3	Krunching Method	19
3.3.5.4	Examples of <code>gnatkr</code> Usage	20
3.3.6	Renaming Files with <code>gnatchop</code>	20
3.3.6.1	Handling Files with Multiple Units	21
3.3.6.2	Operating <code>gnatchop</code> in Compilation Mode	21
3.3.6.3	Command Line for <code>gnatchop</code>	22
3.3.6.4	Switches for <code>gnatchop</code>	23
3.3.6.5	Examples of <code>gnatchop</code> Usage	24
3.4	Configuration Pragmas	25
3.4.1	Handling of Configuration Pragmas	26
3.4.2	The Configuration Pragmas Files	27

3.5	Generating Object Files	28
3.6	Source Dependencies	28
3.7	The Ada Library Information Files	29
3.8	Binding an Ada Program	30
3.9	GNAT and Libraries	31
3.9.1	Introduction to Libraries in GNAT	31
3.9.2	General Ada Libraries	31
3.9.2.1	Building a library	31
3.9.2.2	Installing a library	33
3.9.2.3	Using a library	34
3.9.3	Stand-alone Ada Libraries	35
3.9.3.1	Introduction to Stand-alone Libraries	35
3.9.3.2	Building a Stand-alone Library	35
3.9.3.3	Creating a Stand-alone Library to be used in a non-Ada context	37
3.9.3.4	Restrictions in Stand-alone Libraries	38
3.9.4	Rebuilding the GNAT Run-Time Library	39
3.10	Conditional Compilation	39
3.10.1	Modeling Conditional Compilation in Ada	39
3.10.1.1	Use of Boolean Constants	39
3.10.1.2	Debugging - A Special Case	40
3.10.1.3	Conditionalizing Declarations	41
3.10.1.4	Use of Alternative Implementations	42
3.10.1.5	Preprocessing	44
3.10.2	Preprocessing with gnatprep	44
3.10.2.1	Preprocessing Symbols	44
3.10.2.2	Using gnatprep	44
3.10.2.3	Switches for gnatprep	45
3.10.2.4	Form of Definitions File	46
3.10.2.5	Form of Input Text for gnatprep	47
3.10.3	Integrated Preprocessing	48
3.11	Mixed Language Programming	51
3.11.1	Interfacing to C	52
3.11.2	Calling Conventions	54
3.11.3	Building Mixed Ada and C++ Programs	57
3.11.3.1	Interfacing to C++	57
3.11.3.2	Linking a Mixed C++ & Ada Program	58
3.11.3.3	A Simple Example	59
3.11.3.4	Interfacing with C++ constructors	61
3.11.3.5	Interfacing with C++ at the Class Level	64
3.11.4	Partition-Wide Settings	68
3.11.5	Generating Ada Bindings for C and C++ headers	69
3.11.5.1	Running the Binding Generator	69
3.11.5.2	Generating Bindings for C++ Headers	70
3.11.5.3	Switches	72
3.11.6	Generating C Headers for Ada Specifications	72
3.11.6.1	Running the C Header Generator	73

3.12	GNAT and Other Compilation Models	74
3.12.1	Comparison between GNAT and C/C++ Compilation Models	74
3.12.2	Comparison between GNAT and Conventional Ada Library Models	74
3.13	Using GNAT Files with External Tools	75
3.13.1	Using Other Utility Programs with GNAT	75
3.13.2	The External Symbol Naming Scheme of GNAT	75
4	Building Executable Programs with GNAT ..	77
4.1	Building with <code>gnatmake</code>	77
4.1.1	Running <code>gnatmake</code>	77
4.1.2	Switches for <code>gnatmake</code>	78
4.1.3	Mode Switches for <code>gnatmake</code>	85
4.1.4	Notes on the Command Line	86
4.1.5	How <code>gnatmake</code> Works	87
4.1.6	Examples of <code>gnatmake</code> Usage	87
4.2	Compiling with <code>gcc</code>	88
4.2.1	Compiling Programs	88
4.2.2	Search Paths and the Run-Time Library (RTL)	89
4.2.3	Order of Compilation Issues	90
4.2.4	Examples	90
4.3	Compiler Switches	91
4.3.1	Alphabetical List of All Switches	91
4.3.2	Output and Error Message Control	106
4.3.3	Warning Message Control	109
4.3.4	Info message Control	130
4.3.5	Debugging and Assertion Control	130
4.3.6	Validity Checking	131
4.3.7	Style Checking	135
4.3.8	Run-Time Checks	142
4.3.9	Using <code>gcc</code> for Syntax Checking	144
4.3.10	Using <code>gcc</code> for Semantic Checking	144
4.3.11	Compiling Different Versions of Ada	145
4.3.12	Character Set Control	146
4.3.13	File Naming Control	148
4.3.14	Subprogram Inlining Control	148
4.3.15	Auxiliary Output Control	149
4.3.16	Debugging Control	149
4.3.17	Exception Handling Control	153
4.3.18	Units to Sources Mapping Files	154
4.3.19	Code Generation Control	155
4.4	Linker Switches	155
4.5	Binding with <code>gnatbind</code>	155
4.5.1	Running <code>gnatbind</code>	156
4.5.2	Switches for <code>gnatbind</code>	157

4.5.2.1	Consistency-Checking Modes	162
4.5.2.2	Binder Error Message Control	163
4.5.2.3	Elaboration Control	164
4.5.2.4	Output Control	165
4.5.2.5	Dynamic Allocation Control	166
4.5.2.6	Binding with Non-Ada Main Programs	166
4.5.2.7	Binding Programs with No Main Subprogram	167
4.5.3	Command-Line Access	167
4.5.4	Search Paths for <code>gnatbind</code>	168
4.5.5	Examples of <code>gnatbind</code> Usage	168
4.6	Linking with <code>gnatlink</code>	169
4.6.1	Running <code>gnatlink</code>	169
4.6.2	Switches for <code>gnatlink</code>	170
4.7	Using the GNU <code>make</code> Utility	171
4.7.1	Using <code>gnatmake</code> in a Makefile	172
4.7.2	Automatically Creating a List of Directories	173
4.7.3	Generating the Command Line Switches	175
4.7.4	Overcoming Command Line Length Limits	175
4.8	GNAT with the LLVM Back End	176
5	GNAT Utility Programs	178
5.1	The File Cleanup Utility <code>gnatclean</code>	178
5.1.1	Running <code>gnatclean</code>	178
5.1.2	Switches for <code>gnatclean</code>	178
5.2	The GNAT Library Browser <code>gnatls</code>	180
5.2.1	Running <code>gnatls</code>	180
5.2.2	Switches for <code>gnatls</code>	181
5.2.3	Example of <code>gnatls</code> Usage	182
6	GNAT and Program Execution	184
6.1	Running and Debugging Ada Programs	184
6.1.1	The GNAT Debugger GDB	184
6.1.2	Running GDB	185
6.1.3	Introduction to GDB Commands	185
6.1.4	Using Ada Expressions	188
6.1.5	Calling User-Defined Subprograms	188
6.1.6	Using the ‘next’ Command in a Function	189
6.1.7	Stopping When Ada Exceptions Are Raised	189
6.1.8	Ada Tasks	190
6.1.9	Debugging Generic Units	190
6.1.10	Remote Debugging with <code>gdbserver</code>	191
6.1.11	GNAT Abnormal Termination or Failure to Terminate ..	192
6.1.12	Naming Conventions for GNAT Source Files	193
6.1.13	Getting Internal Debugging Information	194
6.1.14	Stack Traceback	194
6.1.14.1	Non-Symbolic Traceback	194

6.1.14.2	Symbolic Traceback	199
6.1.15	Pretty-Printers for the GNAT runtime	201
6.2	Profiling	202
6.2.1	Profiling an Ada Program with gprof	202
6.2.1.1	Compilation for profiling	203
6.2.1.2	Program execution	203
6.2.1.3	Running gprof	203
6.2.1.4	Interpretation of profiling results	204
6.3	Improving Performance	204
6.3.1	Performance Considerations	204
6.3.1.1	Controlling Run-Time Checks	205
6.3.1.2	Use of Restrictions	205
6.3.1.3	Optimization Levels	206
6.3.1.4	Debugging Optimized Code	207
6.3.1.5	Inlining of Subprograms	209
6.3.1.6	Floating Point Operations	210
6.3.1.7	Vectorization of loops	211
6.3.1.8	Other Optimization Switches	213
6.3.1.9	Optimization and Strict Aliasing	213
6.3.1.10	Aliased Variables and Optimization	217
6.3.1.11	Atomic Variables and Optimization	218
6.3.1.12	Passive Task Optimization	219
6.3.2	Text_IO Suggestions	219
6.3.3	Reducing Size of Executables with Unused Subprogram/Data Elimination	220
6.3.3.1	About unused subprogram/data elimination	220
6.3.3.2	Compilation options	220
6.3.3.3	Example of unused subprogram/data elimination	220
6.4	Overflow Check Handling in GNAT	221
6.4.1	Background	222
6.4.2	Management of Overflows in GNAT	223
6.4.3	Specifying the Desired Mode	224
6.4.4	Default Settings	225
6.4.5	Implementation Notes	226
6.5	Performing Dimensionality Analysis in GNAT	226
6.6	Stack Related Facilities	230
6.6.1	Stack Overflow Checking	231
6.6.2	Static Stack Usage Analysis	231
6.6.3	Dynamic Stack Usage Analysis	232
6.7	Memory Management Issues	232
6.7.1	Some Useful Memory Pools	232
6.7.2	The GNAT Debug Pool Facility	233
6.8	Sanitizers for Ada	236
6.8.1	AddressSanitizer	236
6.8.2	UndefinedBehaviorSanitizer	238
6.8.3	Sanitizers in mixed-language applications	240

7	Platform-Specific Information	242
7.1	Run-Time Libraries	242
7.1.1	Summary of Run-Time Configurations	242
7.2	Specifying a Run-Time Library	242
7.3	GNU/Linux Topics	243
7.3.1	Required Packages on GNU/Linux	243
7.3.2	Position Independent Executable (PIE)	
	Enabled by Default on Linux	244
7.3.3	Choosing the Scheduling Policy with GNU/Linux	244
7.3.4	A GNU/Linux Debug Quirk	245
7.4	Microsoft Windows Topics	245
7.4.1	Using GNAT on Windows	245
7.4.2	Using a network installation of GNAT	246
7.4.3	CONSOLE and WINDOWS subsystems	246
7.4.4	Temporary Files	246
7.4.5	Disabling Command Line Argument Expansion	247
7.4.6	Choosing the Scheduling Policy with Windows	247
7.4.7	Windows Socket Timeouts	248
7.4.8	Mixed-Language Programming on Windows	248
7.4.8.1	Windows Calling Conventions	249
7.4.8.2	C Calling Convention	250
7.4.8.3	Stdcall Calling Convention	250
7.4.8.4	Win32 Calling Convention	251
7.4.8.5	DLL Calling Convention	251
7.4.8.6	Introduction to Dynamic Link Libraries (DLLs)	251
7.4.8.7	Using DLLs with GNAT	252
7.4.8.8	Creating an Ada Spec for the DLL Services	253
7.4.8.9	Creating an Import Library	253
7.4.8.10	Building DLLs with GNAT Project files	255
7.4.8.11	Building DLLs with GNAT	255
7.4.8.12	Building DLLs with gnatdll	256
7.4.8.13	Limitations When Using Ada DLLs from Ada	257
7.4.8.14	Exporting Ada Entities	257
7.4.8.15	Ada DLLs and Elaboration	258
7.4.8.16	Ada DLLs and Finalization	259
7.4.8.17	Creating a Spec for Ada DLLs	259
7.4.8.18	Creating the Definition File	260
7.4.8.19	Using gnatdll	260
7.4.8.20	GNAT and Windows Resources	263
7.4.8.21	Building Resources	264
7.4.8.22	Compiling Resources	264
7.4.8.23	Using Resources	265
7.4.8.24	Using GNAT DLLs from	
	Microsoft Visual Studio Applications	265
7.4.8.25	Debugging a DLL	265
7.4.8.26	Program and DLL Both Built with GCC/GNAT ...	266

7.4.8.27	Program Built with Foreign Tools and DLL Built with GCC/GNAT	266
7.4.8.28	Setting Stack Size from <code>gnatlink</code>	268
7.4.8.29	Setting Heap Size from <code>gnatlink</code>	268
7.4.9	Windows Specific Add-Ons	268
7.4.9.1	Win32Ada	268
7.4.9.2	wPOSIX	269
7.5	Mac OS Topics	269
7.5.1	Codesigning the Debugger	269
8	Example of Binder Output File	271
9	Elaboration Order Handling in GNAT	287
9.1	Elaboration Code	287
9.2	Elaboration Order	288
9.3	Checking the Elaboration Order	290
9.4	Controlling the Elaboration Order in Ada	291
9.5	Controlling the Elaboration Order in GNAT	295
9.6	Mixing Elaboration Models	296
9.7	ABE Diagnostics	296
9.8	SPARK Diagnostics	298
9.9	Elaboration Circularities	298
9.10	Resolving Elaboration Circularities	300
9.11	Elaboration-related Compiler Switches	302
9.12	Summary of Procedures for Elaboration Control	304
9.13	Inspecting the Chosen Elaboration Order	305
10	Inline Assembler	309
10.1	Basic Assembler Syntax	309
10.2	A Simple Example of Inline Assembler	310
10.3	Output Variables in Inline Assembler	311
10.4	Input Variables in Inline Assembler	315
10.5	Inlining Inline Assembler Code	316
10.6	Other <code>Asm</code> Functionality	317
10.6.1	The <code>Clobber</code> Parameter	317
10.6.2	The <code>Volatile</code> Parameter	318
11	GNU Free Documentation License	319
Index		326

‘GNAT, The GNU Ada Development Environment’

GCC version 16.0.0

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 User’s Guide for Native Platforms”, and with no Back-Cover Texts. A copy of the license is included in the section entitled [GNU Free Documentation License], page 318.

1 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 2022. You may also invoke it in Ada 83 compatibility mode. By default, GNAT assumes Ada 2012, but you can use a compiler switch ([Compiling Different Versions of Ada], page 145) to explicitly specify the language version. Throughout this manual, references to ‘Ada’ without a year suffix apply to all versions of the Ada language starting with Ada 95.

GNAT supports both the GCC and LLVM back end compilation families. Most GNAT versions use the GCC back end, but some are now available using the LLVM back end. In some places in this manual, we distinguish between the two back ends, but in most cases, everything in this manual applies to both back ends. We refer to GNAT with the LLVM back end as ‘GNAT LLVM’. See [GNAT with the LLVM Back End], page 176, for limitations of GNAT LLVM.

1.1 What This Guide Contains

This guide contains the following chapters:

- * [Getting Started with GNAT], page 3, describes how to get started compiling and running Ada programs with the GNAT Ada programming environment.
- * [The GNAT Compilation Model], page 6, describes the compilation model used by GNAT.
- * [Building Executable Programs with GNAT], page 76, 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], page 177, explains the various utility programs that are included in the GNAT environment.
- * [GNAT and Program Execution], page 183, 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], page 241, 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], page 270, shows the source code for the binder output file for a sample program.
- * [Elaboration Order Handling in GNAT], page 286, describes how GNAT helps you deal with elaboration order issues.
- * [Inline Assembler], page 308, 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, Ada 2012 and Ada 2022 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*, *Ada 2012 Reference Manual*, and *Ada 2022 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.

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 you enter 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 you should use the ‘\’ character instead.

2 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, to get the best experience we recommend using a machine with as many cores as possible, allowing individual compilations to 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, we recommend a minimum of 4 cores (8 is preferred), with at least 2GB per core (so 8 to 16GB).

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

Using fast, local drives can make a significant difference in build and link times. You should avoid network drives such as NFS, SMB, or worse, configuration management filesystems (such as ClearCase dynamic views) as much as possible since these will produce very degraded performance (typically 2 to 3 times slower than on fast, local drives). If you cannot avoid using such slow drives for accessing source code, you should at least configure your project file so the result of the compilation is stored on a drive local to the machine performing the compilation. You can do this by setting the `Object_Dir` project file attribute.

2.2 Running GNAT

You need to take three steps to create an executable file from an Ada source file:

- * You must compile the source file(s).
- * You must bind the file(s) using the GNAT binder.
- * You must link all appropriate object files to produce an executable.

You most commonly perform all three steps by using the `gnatmake` utility program. You pass it the name of the main program and it automatically performs the necessary compilation, binding, and linking steps.

2.3 Running a Simple Ada Program

You may use any text editor to prepare an Ada program. (If you use Emacs, an 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 named `hello.adb`:

```
with Ada.Text_IO; use Ada.Text_IO;
procedure Hello is
```

```
begin
  Put_Line ("Hello WORLD!");
end Hello;
```

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** (see [Using Other File Names], page 12). 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 you use can use the **gnatchop** utility to generate correctly-named source files (see [Renaming Files with gnatchop], page 20).

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. It is capable of compiling programs in several languages, including Ada and C. It assumes you have given it an Ada program if the file extension is either **.ads** or **.adb**, in which case it will 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 you must always specify the **-c**.)

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 perform 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, as well as 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 you can 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 with 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
  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 compile the body. If you want to submit these files to the compiler for semantic checking and not code generation, use the `-gnatc` switch:

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

Although you can do the compilation in separate steps, in practice it's almost always more convenient to use the `gnatmake` or `gprbuild` tools:

```
$ gnatmake gmain.adb
```

3 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], page 7,
 - * [Foreign Language Representation], page 8,
 - * [File Naming Topics and Utilities], page 11,
- * [Configuration Pragmas], page 25,
- * [Generating Object Files], page 28,
- * [Source Dependencies], page 28,
- * [The Ada Library Information Files], page 29,
- * [Binding an Ada Program], page 30,
- * [GNAT and Libraries], page 30,
- * [Conditional Compilation], page 39,
- * [Mixed Language Programming], page 51,
- * [GNAT and Other Compilation Models], page 74,
- * [Using GNAT Files with External Tools], page 75,

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], page 8, 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 recognizes 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 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, legacy, 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 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], page 146).

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 US. 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 is 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. You can only use characters defined in the JIS code set table 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. You can only use characters defined in the JIS code set table 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 the use of these sequences is documented in the following section on wide wide characters.)

‘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#`. You can also (though you are not required to) use the Brackets coding for upper half characters. For example, you can represent the code `16#A3#` 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, but also provides you with a high degree of control over how the names and extensions of your source files correspond to the Ada compilation units that they contain.

3.3.1 File Naming Rules

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

An exception occurs 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 hyphen. In this case, the character tilde is used in place of the hyphen. This special rule avoids 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’). You may find this particularly useful when programs being developed with GNAT are to be used on operating systems with limited file name lengths. [Using gnatkr], page 18.

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], page 20.)

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

The previous section described the default rules used by GNAT to determine the file name in which a given unit resides. It is usually 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 you 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_Uutilities.Stacks,
    Spec_File_Name => "myutilst_a.adb");
pragma Source_File_Name (My_Uutilities.Stacks,
    Body_File_Name => "myutilst.adb");
```

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 you will place it 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], page 26.

GNAT allows you to specify completely arbitrary file names using the source file name pragma. However, if the file name specified has an extension other than `.ads` or `.adb` you must use a special syntax when compiling the file. The name on the command line 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, you must include it 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, which can 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, the default is all lower-case.

You use the optional **Dot_Replacement** string to replace any periods that occur in subunit or child unit names. If you don't specify a **Dot_Replacement** argument, separating dots appear unchanged in the resulting file name. The above syntax indicates that the **Casing** argument must appear before the **Dot_Replacement** argument, but you can write these arguments in any order.

As indicated, you can 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, you need not provide a separate **Subunit_File_Name** rule; in this case the **Body_File_Name** rule is used for subunits as well.

You can also use the separate rule for subunits 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.

File name translation consists of the following steps:

- * If there is a specific **Source_File_Name** pragma for the given unit, 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, the resulting file name is used if the file exists. If more than one pattern matches, the latest one is tried first and the first attempt that results in a reference to a file that exists is used.
- * If no pattern type **Source_File_Name** pragma that applies to the unit for which the corresponding file exists, 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, specs end with **.1.adb**, and bodies end with **.2.adb**. GNAT will follow this scheme if the following two pragmas appear:

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

The default GNAT scheme is equivalent to providing the following default pragmas:

```
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 legacy 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`

3.3.4.1 Arbitrary File Naming Conventions

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

When the source file names do not follow the standard GNAT default file naming conventions, you must give the GNAT compiler additional information through a configuration pragmas file ([Configuration Pragmas], page 25) 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], page 13) may be sufficient. However, if the file naming conventions are irregular or arbitrary, you must define a number of pragma `Source_File_Name` for individual compilation units. 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.

3.3.4.2 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 arguments, `gnatname` will display its usage.

When used with at least one naming pattern, `gnatname` attempts to find all the compilation units in files that follow at least one of the naming patterns. To find these compilation units, `gnatname` uses 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.

You may call `gnatname` with several sections of directories/patterns. Sections are separated by the switch `--and`. In each section, you must include at least one pattern. If you don’t

specify a directory a section, the current directory (or the project directory if `-P` is used) is used. 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-regex.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.

3.3.4.3 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'`
Actual 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.
- `-c`filename'`
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. You can specify only one switch `-c`. When a switch `-c` is specified, you may not specify switch `-P` (see below).
- `-d`dir'`
Look for source files in directory `dir`. You may put zero, one or more spaces between `-d` and `dir`. `dir` may end with `/**`, i.e., you may write it the form `root_dir/**`. In this case, the directory `root_dir` and all of its subdirectories, recursively, have to be searched for sources. When you specify a `-d` switch, the current working directory will not be searched for source files unless you explicitly specify it with a `-d` or `-D` switch. You may specify several switches `-d`. 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 you don't specify either switch `-c` or switch `-P`, it's relative to the current working directory. The directory you specified with switch `-d` must exist and be readable.

`-D`filename``

Look for source files in all directories listed in text file `filename`. You may place 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.

`-f`pattern``

Foreign patterns. Using this switch, you can add sources of languages other than Ada to the list of sources of a project file, but it's only useful if you also specify a `-P` switch. For example,

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

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

`-h`

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

`-P`proj``

Create or update project file `proj`. You may place 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 you specify switch `-P`, you may not also include switch `-c`. On all platforms except VMS when `gnatname` is invoked for an existing project file `<proj>.gpr`, `gnatname` creates a backup copy of the project file in the project directory with file name `<proj>.gpr.saved_x` where `x` is the first non negative number that creates a unique filename.

`-v`

Verbose mode. Output detailed explanation of what it's doing to `stdout`. This includes name of the file written, the name of the directories searched, 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 (a single `-v` switch), for each file in the searched directories whose name matches none of the Naming Patterns, `gnatname` indicates that there is no match.

-x`pattern`

Excluded patterns. Using this switch, you can exclude some files that otherwise would match the name patterns. For example,

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

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

3.3.4.4 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 you may use several `-d` switches, even in conjunction with one or several `-D` switches. This example illustrates multiple Naming Patterns and one excluded pattern.

3.3.5 File Name Krunching with gnatkr

Here we discuss 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. We also describe the **gnatkr** utility, which you can use to determine the result of applying this shortening.

3.3.5.1 About gnatkr

GNAT requires that the file name must be derived from the unit name. The 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 hyphen.

This exception avoids clashes with the standard names for children of System, Ada, Interfaces, and GNAT, which use the prefixes **s-**, **a-**, **i-**, and **g-**, respectively.

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

You can use the **gnatkr** utility to determine the krunched name for a given file when krunched to a specified maximum length.

3.3.5.2 Using gnatkr

You invoke the **gnatkr** command as follows:

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

name is the unkrunched file name, derived from the name of the unit in the default manner described in the previous section (i.e., in particular all dots are replaced by hyphens). You may or may not include an extension (defined as a suffix of the form period followed by arbitrary characters other than period) in the filename. If you do, **gnatkr** will preserve it

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`, you can use dots in the name instead of hyphens, but `gnatkr` always interprets the last dot as the start of an extension. So if you pass `gnatkr` an argument such as `Hello.World.adb`, it treats it exactly as if the first period had been a hyphen, so, for example, krunching to eight characters gives the result `hellworl.adb`.

Note that the result is always all lower case. Other characters are folded as required.

`length` represents the length of the krunched name. The default if you don't specify it, is 8 characters. A length of zero means unlimited, in other words don't 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.

3.3.5.3 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
<code>ada-</code>	<code>a-</code>
<code>gnat-</code>	<code>g-</code>
<code>interfac es-</code>	<code>i-</code>
<code>system-</code>	<code>s-</code>

These system files have a hyphen in the second character position. That's is why normal user files replace such a character with a tilde.

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. If file name krunching is used, it's your responsibility to ensure that no name clashes occur. The utility program `gnatkr` is supplied so that you can conveniently determine the krunched name of a file.

3.3.5.4 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 gnat Chop

This section discusses how to handle files with multiple units by using the `gnat Chop` utility. You will also find this utility useful in renaming files to meet the standard GNAT default file naming conventions.

3.3.6.1 Handling Files with Multiple Units

GNAT's fundamental compilation model requires that a file submitted to the compiler contain only one unit and there be a strict correspondence between the file name and the unit name.

If you want to have your files contain multiple units, perhaps to maintain compatibility with some other Ada compilation system, you can use **gnatname** to generate or update your project files, which can be processed by GNAT.

See [Handling Arbitrary File Naming Conventions with **gnatname**], page 15, 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 containing only one unit per file, 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, each file generated by **gnatchop** will start with a copy of this BOM, meaning that they can be compiled automatically in UTF-8 mode without you needing to specify an explicit encoding.

3.3.6.2 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 units 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 you needing to mark it explicitly and you'll likely find this default to be what you want. In this default mode, you may not submit a file containing only configuration pragmas, or one that ends in configuration pragmas, to **gnatchop**.

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

First, in compilation mode, if you give **gnatchop** a file that consists of only configuration pragmas, it appends this file 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], page 26.

Second, in compilation mode, if you give **gnatchop** a file that starts with configuration pragmas and contains one or more units, then 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, you will use **gnatchop** in default mode. You only use the compilation mode described above if you need precisely 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. This is the mode in which GNAT processes the ACVC tests.

3.3.6.3 Command Line for **gnatchop**

You call **gnatchop** as follows:

```
$ gnatchop 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, **gnatchop** 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 you don't specify it, 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

```
$ gnatchop 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. You can compile these generated files in the normal manner.

When you invoke **gnatchop** on a file that is empty or contains only empty lines and/or comments, **gnatchop** will complete normally, but won't produce any new file.

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

```
-- Just a comment
```

the command

```
$ gnatchop 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
```

3.3.6.4 Switches for `gnatchop`

`gnatchop` recognizes the following switches:

`--version`

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

`--help`

If `--version` is not present, 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 the previous section for a full description of this mode.

`-gnat`xxx``

This passes the given `-gnat`xxx`` switch to `gcc` which is used to parse the given file. Not all ‘xxx’ options make sense, but, for example, the use of `-gnat12` 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.

`-k`mm``

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. You may not place any space between the `-k` and the numeric value. You can omit the numeric value, in which case `gnatchop` will use a default of `-k8`, suitable for use with DOS-like file systems. If you don’t specify a `-k` switch, 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). You may find this 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 from 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, **gnatchop** respects these pragmas 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. It outputs the version number and copyright notice as well as exact copies of the commands spawned to obtain the information needed to control chopping.

-w

Overwrite existing file names. Normally, **gnatchop** treats it as a fatal error if there's already a file with the same name as a file it would otherwise output. This can happen either if you've previously chopped that file or 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, **gnatchop** makes no attempt to add a prefix to the GNAT parser executable, so it must include the full pathname.

3.3.6.5 Examples of **gnatchop** Usage

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

Chops the source file **hello_s.ada**. The output files are 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 **gnatchop** is used at the other end to reconstitute the original files.

```
$ 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, **gnatchop** generates an error message, and doesn't write any files. 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 `gnatchop` will skip earlier duplicate occurrences for the same unit.

3.4 Configuration Pragmas

Configuration pragmas supported by GNAT consist of those pragmas described as such in the Ada Reference Manual and the 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
Interrupts_System_By_Default
License
Locking_Policy
No_Component_Reordering
No_Heap_Finalization
No_Strict_Aliasing
NormalizeScalars
Optimize_Alignment
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
User_Aspect_Definition
Validity_Checks
Warning_As_Error
Warnings
Wide_Character_Encoding

```

3.4.1 Handling of Configuration Pragas

You can place configuration pragmas either appear at the start of a compilation unit or in a configuration pragma file that applies to all compilations performed in a given compilation environment.

Configuration pragmas placed before a library level package specification are not propagated to the corresponding package body (see RM 10.1.5(8)); they must be added explicitly to the package body.

GNAT includes the `gnatchop` utility to provide an automatic way to handle configuration pragmas that follows the semantics for compilations (that is, files with multiple units) described in the RM. See [Operating `gnatchop` in Compilation Mode], page 21, for details. However, for most purposes, you will find it 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 ‘with’ed unit to the ‘with’ing unit. This makes sense since the binder will always insist on seeing consistent us, so any unit not conforming to any restrictions anywhere in the partition will be rejected and it’s better for you to find that out at compile time rather than bind time.

For restrictions that do not require partition-wide consistency, e.g. `SPARK` or `No_Implementation_Attributes`, the restriction normally 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. You can apply this restriction in a configuration pragma file or you can ace it in the body and/or the spec (in either case it applies to all the relevant units). You can place it on a subunit only if you have previously placed it 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 you specify the switch `-gnatA`, GNAT ignores `gnat.adc`. When used, GNAT adds `gnat.adc` to the dependencies so that if `gnat.adc` is modified later, the source will be recompiled on a future invocation of `gnatmake`.

You can add configuration pragmas into the `gnat.adc` file either by running `gnatchop` on a source file consisting only of configuration pragmas or, more conveniently, by directly editing the `gnat.adc` file, which is a standard format source file.

Besides `gnat.adc`, you may apply additional files containing configuration pragmas to the current compilation using the `-gnatec='path'` switch, 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 you do not use switch `-gnatA`). You can specify multiple `-gnatec=` switches.

GNAT will add files containing configuration pragmas specified with switches `-gnatec=` 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 to request they be stored instead by just their basename.

If you are using project file, they provide a separate mechanism 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 generating the corresponding object files. You generate these 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. You need not compile the source file for the package spec in this case because there's 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. You need not compile the spec, if such a file is present.
- * In the case of subunits, only compile the parent unit. GNAT generates a single object file 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, you still must compile generic units for dependency checking and elaboration purposes.

The preceding rules describe the set of files that must be compiled to generate all the object files for a program. See the following section on dependencies for more details on computing that set of files. 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

Each object file obviously depends on at least 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 ‘with’s a unit `X`, the object file depends on the file containing the spec of unit `X`. This includes files that are ‘with’ed implicitly either because they are

parents of ‘with’ed child units or 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 you have activated inlining 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, you must 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 or LLVM based back end, the use of `-gnatN` is deprecated and the use of `-gnatn` is preferred. Historically front end inlining was more extensive than 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 means that, for purposes of computing dependencies and re-compilation, 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` depends 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. Invoking `gnatmake` will cause it to determine the necessary recompilations.

3.7 The Ada Library Information Files

Each compilation actually generates two output files. The first of these is the actual 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 that 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 compilation command
- * Attributes of the unit, including the configuration pragmas used, an indication of whether the compilation was successful, and the exception model used.
- * A list of relevant restrictions applying to the unit (used for consistency checking).
- * Categorization information (e.g., use of pragma `Pure`).
- * Information on all ‘with’ed 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 the modification of any of these units requires the recompilation of the unit in question.
- * 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 spec of unit `Lib.Writ`, contained in file `lib-writ.ads` 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 you can 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 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 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 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 instead 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, you can provide the compiled code in different ways. The simplest approach is to directly provide 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.

In the GNAT environment, a library has these components:

- * Source files,
- * ALI files (see [The Ada Library Information Files], page 29), 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, in other words, 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], page 35).

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

3.9.2.1 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, you can specify additional optional project-level attributes:

*

`Library_Kind`

This attribute controls whether the library is to be static or dynamic

*

Library_Version

This attribute specifies the library version. Its 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 the commands used to generate the library.

The GNAT Project Manager takes complete 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's 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's complex to manually perform all the steps required to produce a library, so we recommend you use the GNAT Project Manager for this task. In case this is not desired, we discuss the necessary steps below.

There are various possibilities for compiling the units that make up the library: for example with a `Makefile` ([Using the GNU make Utility], page 171) or with a conventional script. For simple libraries, you can also create a dummy main program that depends upon all the packages that comprise the interface of the library. You can then pass this dummy main program to `gnatmake`, which will ensure all necessary objects are built.

After the above has been accomplished, you should follow the standard procedure of the underlying operating system to produce the static or shared library.

Here's 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 `lib`xxx'.a` or `lib`xxx'.so` (or `lib`xxx'.dll` on Windows) in order to be accessed by the `-l`xxx'` switch at link time.

3.9.2.2 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 you’re not able to use project files for some reason, you can also 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)], page 89), but we don’t recommend doing this. 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 theses files are located.

The files `ada_source_path` and `ada_object_path` might not be present in a GNAT installation, in which case, GNAT looks 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 `--RTS=rts-path` switch. You can thus choose 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.

You can install a library before or after the standard GNAT library by selecting the ordering the lines in the configuration files. In general, a library must be installed before the GNAT library if it redefines any part of it.

3.9.2.3 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 your project. For example, 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 ‘with’ed 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)], page 89, and [Search Paths for `gnatbind`], page 168). 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

3.9.3.1 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, you must also provide the body’s source; if the units that must be provided in the source form depend on other units, you must also provide the source and ALI files of those units.

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 don’t 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, the clients also 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.

3.9.3.2 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’), you must define the attribute `Library_Interface`. For example:

```
for Library_Dir use "lib_dir";
for Library_Name use "dummy";
for Library_Interface use ("int1", "int1.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, the binder is first 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.

You can also build an encapsulated library where not only the code to elaborate and finalize the library is embedded but also ensure that the library is linked only against static libraries. That means that an encapsulated library only depends on system libraries: all other code, including the GNAT runtime, is embedded. To build an encapsulated library you must set attribute `Library_Standalone` to `encapsulated`:

```
for Library_Dir use "lib_dir";
for Library_Name use "dummy";
for Library_Kind use "dynamic";
for Library_Interface use ("int1", "int1.child");
for Library_Standalone use "encapsulated";
```

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

You may specify the attribute `Library_Src_Dir` for a Stand-Alone Library. `Library_Src_Dir` 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 are copied to the designated directory, called the Interface Copy directory, when the library is built. 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, the building process makes an attempt 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, passing to the linker the names of the `init` (and possibly `final`) procedures for automatic initialization (and

finalization). You should place the built library in a different directory than the object files.

- * 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], page 34).

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

It’s easy for you 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’s 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;
    pragma Export (C, Do_Something_Else, "do_something_else");

end My_Package;
```

On the C 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 the elaboration procedures (named `mylibinit` in the previous example) are called before any library services are used. Any number of libraries can be called from a single executable 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.

Information on limitations of binding Ada code in non-Ada contexts can be found under [Binding with Non-Ada Main Programs], page 166.

Note also that you must take special care with multi-tasking applications. In that case, the initialization and finalization routines are not protected against concurrent access. If you need such requirement, you must ensure it at the application level using a specific operating system services like a mutex or a critical-section.

3.9.3.4 Restrictions in Stand-alone Libraries

You should use the pragmas listed below with caution inside libraries, since 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 of the library must ensure that all libraries use the same pragmas with the same values. Otherwise, `Program_Error` will be raised during the elaboration of the conflicting libraries. You should document the usage of these pragmas and its consequences for the user.

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 you use the `Version` or `Body_Version` attributes inside a library, 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

You may need to recompile the GNAT library in various debugging or experimentation contexts. The GNAT distribution provides a project file called `libada.gpr` to do that; it 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 for other purposes.

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

You may want 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 is to use the preprocessor 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, 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.

3.10.1.1 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 block 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;
```

You would write the `Config` package multiple forms for 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.

3.10.1.2 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. You can do this by using a Boolean constant to control whether the debug 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, GNAT provides special features to deal with this in a convenient manner. For the case of tests, Ada 2005 has added a pragma `Assert` that you can use for such tests. This pragma is modeled on the `Assert` pragma that has always been available in GNAT, so you can use this feature 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` is raised with the exception message using the specified string in the first case or a string indicating the location of the pragma in the second case.

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

This is an Ada 2005 pragma that is implemented in all modes by GNAT. Alternatively, you can use the `-gnata` switch to enable assertions from the command line, which also 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` is called). You can specify only one call, but you can of course include a special debugging procedure containing any code you like in the program and call it 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.

You can enable debug pragmas 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 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;
```

3.10.1.3 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;

```

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 don't need to define this one yourself).

3.10.1.4 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 doesn't result in excessive duplication of code, you can do this 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 you need 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
    else
        ... not quite as neat Ada 95 code
    end if;

```

```
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's still 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.ada";
```

Note also that with project files, you should use a different extension than `ads` / `adb` for alternative versions. Otherwise, a naming conflict may arise through another commonly used feature: declaring 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. Then, 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.

3.10.1.5 Preprocessing

Although it is quite possible to conditionalize code without the use of C-style preprocessing, as described in the cases above, 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).

You can use the preprocessor used in two different modes. You can use it separately from the compiler to generate a separate output source file, which you then feed to the compiler as a separate step. This is the `gnatprep` utility, whose use is fully described in [Preprocessing with `gnatprep`], page 44.

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. You pass the compiler the preprocessor input, which includes `#if` lines etc, and the compiler carries out the preprocessing internally and compiles the resulting output. For more details on this approach, see [Integrated Preprocessing], page 48.

3.10.2 Preprocessing with `gnatprep`

This section discusses how to you can 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], page 39.

3.10.2.1 Preprocessing Symbols

Preprocessing symbols are defined in 'definition files' and referenced in the sources to be pre-processed. 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).

3.10.2.2 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. You can omit this argument and instead use the **-D** switch.

3.10.2.3 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 also results in line numbers being preserved in the output file.

-C

Causes comments to be scanned. Normally comments are ignored by **gnatprep**. If you specify this option, **gnatprep** scans comments and any **\$symbol** substitutions performed as in program text. You will find this particularly useful when structured comments are used (e.g., for programs written in a pre-2014 version of the SPARK Ada subset). This switch is not available when doing integrated

preprocessing (it would be useless in this context since comments are always ignored by the compiler).

-D`symbol' [=`value']

Defines a new preprocessing symbol with the specified value. If you don't specify a value, the symbol is defined to be **True**. You can use this switch instead of providing a definition file.

-e

Causes both preprocessor lines and the lines deleted by preprocessing to be replaced by empty comment lines marked with "--!" (and no other text) in the output source file, preserving line numbers in the output file. This option can be useful as an alternative to **-b** and **-c** when compilation style switches like **-gnatyu** or **-gnatyM** are used (to avoid warnings about multiple blank lines or lines too long).

-r

Causes **gnatprep** to generate a **Source_Reference** pragma 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 the **-b** and **-c** are always enabled.

If the file to be preprocessed contains multiple units, you must call **gnatchop** on the 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 passed to **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. If you don't specify this switch, **gnatprep** will treat an undefined symbol in a **#if** or **#elsif** test as an error.

-v

Verbose mode: generates more output about what is done.

Note: if you don't specify either **-b** or **-c**, then preprocessor lines and deleted lines are completely removed from the output, unless you specify **-r**, in which case **gnatprep** enables the **-b** switch.

3.10.2.4 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}.

You may also place comment lines in the definitions file, starting with the usual `--` and comments may be added to the end of each definition line.

3.10.2.5 Form of Input Text for gnatprep

The input text contains preprocessor conditional inclusion lines as well as general symbol substitution sequences.

Preprocessor conditional inclusion commands have the form:

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

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> )
```

For the first test, (`<expression> ::= <symbol>`), the symbol must have either the value true or false. The right-hand of the symbol definition must be one of the (case-insensitive) literals **True** or **False**. If the value is true, 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, the result depends on whether or not you have specified the `-u` switch. If you have, the symbol is treated as if it had the value false and the test fails. If not, it's an error to reference an undefined symbol. It's also an error to reference a symbol that you have defined with a value other than `True` or `False`.

The use of the `not` operator inverts the sense of this logical test. You can't combine the `not` operator with the `or` or `and` operators without parentheses. For example, you can't write "if not X or Y then" allowed, but can write either "if (not X) or Y then" or "if not (X or Y) then".

The `then` keyword is optional, as shown.

You must place the `#` in the first non-blank character on a line, i.e., it must be preceded only by spaces or horizontal tabs, but otherwise the format is free form. You may place spaces or tabs between the `#` and the keyword. The keywords and the symbols are case insensitive, as in normal Ada code. You can write comments on a preprocessor line, but other than that, you can't place any other tokens on a preprocessor line. You can have any number of `elsif` clauses, including none at all. The `else` is optional, as in Ada.

You obtain symbol substitution outside of preprocessor lines by using the sequence:

```
$symbol
```

anywhere within a source line, except in a comment or within a string literal. The identifier following the `$` must match one of the symbols defined in the symbol definition file and the resulting output substitutes the value of the symbol in place of `$symbol` in the output file. Note that although you can't substitute strings within a string literal, you can 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 you pass either or both of the following switches to the compiler:

- * `-gnatprep`, 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's not available for configuration pragma files.

With integrated preprocessing, GNAT doesn't write the output from the preprocessor, by default, to any external file. Instead it's passed internally to the compiler. To preserve the result of preprocessing in a file, either run **gnatprep** in standalone mode or supply the **-gnateG** switch to the compiler.

When using project files:

- * you should use the builder switch **-x** if any Ada source is compiled with **gnatep=** so that the compiler finds the 'preprocessor data file'.
- * you should place the preprocessing data file and the symbol definition files 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], page 44. This section explains the switches that relate to integrated preprocessing.

-gnatep=`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
      ("-gnatep=" & 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. You can also include empty lines and comments (using Ada syntax), 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, you start each preprocessor control line 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).

You cannot have two lines with the same file name or two lines starting with the ‘*’ character.

After the file name or ‘*’, you can place an optional literal string to specify the name of the definition file to be used for preprocessing ([Form of Definitions File], page 46). 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, you may need 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` or `-e` it cancels the effect of those switches.

`-c`

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

`-D`symbol`=`new_value``

Define or redefine `symbol` to have `new_value` as its value. You can write `symbol` as either an Ada identifier or any Ada reserved word aside from `if`, `else`, `elsif`, `end`, `and`, `or` and `then`. You can write `new_value` as 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.

-e

Causes both preprocessor lines and the lines deleted by preprocessing to be replaced by empty comment lines marked with ‘`!'`’ (and no other text) in the output source 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. If you don’t specify this switch, an undefined symbol in a **#if** or **#elsif** test is treated as an error.

-gnateD`symbol' [=`new_value']

Define or redefine **symbol** to have **new_value** as its value. If you don’t specify a value, the value of **symbol** is **True**. You write **symbol** as an identifier, following normal Ada (case-insensitive) rules for its syntax, and **new_value** as 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[bce]

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 is **foo.adb.prep**. An optional character (b, c, or e) can be appended to indicate that filtered lines are to be replaced by blank lines, comments, or empty comments (see documentation above about **-b**, **-c**, and **-e**). If one of those switches is given in a preprocessor data file, then it will override any option included with **-gnateG**.

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 example, 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's an 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);
    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
```

You can merge the last three steps into a single command:

```
$ gnatmake my_main.adb -largs file1.o file2.o
```

If the main program is in a language other than Ada, 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], page 166). You must insert calls to the initialization and finalization routines in the main program or some other appropriate point in the code. You must place the call to initialize the Ada units so that it occurs before the first Ada subprogram is called and must place the call to finalize the Ada units so it occurs after the last Ada subprogram returns. The binder places 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();

    /* 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 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 is 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. You can also mix GNAT generated code and code generated by another Ada compiler. In this case, you should restrict the data types to simple cases, including primitive types. Whether complex data types can be passed depends on the situation. It is probably 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.

If output from two different compilers is mixed, you are 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 it takes care of its own elaboration requirements, and 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.

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 is passed according to the conventions described in section B.4 of the Ada Reference Manual.

C

Data is 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 you can use 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 is 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, it means the body of the subprogram is provided by the compiler itself, usually by means of an efficient code sequence, and that you don't 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 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, "/");
```

You often program this common idiom 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 is used, as defined by the NT API. To simplify building cross-platform bindings, this convention is 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 you can use to parameterize conventions and allow additional synonyms. 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, you can use the identifier `Fortran77` 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

If you are inexperienced with mixed-language development, you 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.

3.11.3.1 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>).

You can do interfacing at three 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], page 69). You can also address this problem 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` or by looking at the assembly code generated by the C++ compiler) and using it as the `Link_Name` argument of the `pragma Import`.

You can achieve interfacing at the class level by using the GNAT specific pragmas such as `CPP_Constructor`. See the *GNAT_Reference_Manual* for additional information.

3.11.3.2 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 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. We consider three cases:

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

If the C++ code uses inline functions that you plan to call from Ada, you need to compile your C++ code with the `-fkeep-inline-functions` so `g++` doesn't delete these functions.

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

- * Using GNAT and G++ from two different GCC installations: If both compilers are on the PATH, you may use the previous method. However, environment variables such as `C_INCLUDE_PATH`, `GCC_EXEC_PREFIX`, `BINUTILS_ROOT`, and `GCC_ROOT` 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. You can replace the implicit link command as suggested in the `gnatmake` command from the former example with 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's a problem due to interfering environment variables, you can work around it 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: You can use the commands previously described 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 you are using the `setjmp` / `longjmp` exception mechanism, you need only include the paths to the `libgcc` libraries:

```
$ 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 you are using the “zero cost” exception mechanism and the platform supports automatic registration of exception tables (e.g., Solaris), you need to include paths to more objects:

```
$ 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 you are using 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 won’t work and a you will need to preform a pre-linking phase using GNAT.

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.

3.11.3.3 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, execution 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);
    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

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

end Simple_Cpp_Interface;
```

3.11.3.4 Interfacing with C++ constructors

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 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], page 64, and [Generating Ada Bindings for C and C++ headers], page 69).

```
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;
  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's 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.
- * On the right side of an initialization of a record component of type.
- * 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 you must explicitly call the required C++ constructor 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's 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, you must initialize the components `DT` inherited from the C++ side 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's 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` calls 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, we use the extended return statement construct to build in place the returned object whose components are initialized by means of a limited aggregate. We could also place any further action associated with the constructor inside the construct.

3.11.3.5 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 then the reverse case.

The root of our derivation is 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 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 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 you must declare all the primitives and components 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)`, that 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 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 we must register with the compiler by means of `pragma CPP_Constructor` and we need 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, we don't need to add any further information to indicate either the object layout or the dispatch table entry associated with each dispatching operation.

Now let's 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);
  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 ()
{
    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 Partition-Wide Settings

When building a mixed-language application, you must be aware that Ada enforces some partition-wide settings that may implicitly impact the behavior of the other languages.

This is the case for certain signals that are reserved to the implementation to implement proper Ada semantics (such as the behavior of `abort` statements). It means that the Ada part of the application may override signal handlers that were previously installed by either the system or by other user code.

If your application requires that either system or user signals be preserved, you need to instruct the Ada part not to install its own signal handler. You do this using `pragma Interrupt_State` that provides a general mechanism for overriding such uses of interrupts.

Additionally, you can use pragma `Interrupts_System_By_Default` to default all interrupts to `System`.

The set of interrupts for which the Ada run-time library sets a specific signal handler is the following:

- * `Ada.Interrupts.Names.SIGSEGV`
- * `Ada.Interrupts.Names.SIGBUS`
- * `Ada.Interrupts.Names.SIGFPE`
- * `Ada.Interrupts.Names.SIGILL`
- * `Ada.Interrupts.Names.SIGABRT`

You can instruct the run-time library not to install its signal handler for a particular signal by using the configuration pragma `Interrupt_State` in the Ada code. For example:

```
pragma Interrupt_State (Ada.Interrupts.Names.SIGSEGV, System);
pragma Interrupt_State (Ada.Interrupts.Names.SIGBUS, System);
pragma Interrupt_State (Ada.Interrupts.Names.SIGFPE, System);
pragma Interrupt_State (Ada.Interrupts.Names.SIGILL, System);
pragma Interrupt_State (Ada.Interrupts.Names.SIGABRT, System);
```

Obviously, if the Ada run-time system cannot set these handlers it comes with the drawback of not fully preserving Ada semantics. `SIGSEGV`, `SIGBUS`, `SIGFPE` and `SIGILL` are used to raise corresponding Ada exceptions in the application, while `SIGABRT` is used to asynchronously abort an action or a task.

3.11.5 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.

This capability is not intended to generate 100% correct Ada specs and it will in some cases require you to make 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 names that are identical except for 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`).

3.11.5.1 Running the Binding Generator

The binding generator is part of the `gcc` compiler and you can invoke it via the `-fdump-ada-spec` switch, which generates 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
```

generates, 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`, and `/usr/include/bits/time.h` and then compile these Ada specs. The name of the Ada specs is consistent 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 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 generates more complete bindings than a call without the `-DXLIB_ILLEGAL_ACCESS` switch.

In other cases, you can't 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
```

3.11.5.2 Generating Bindings for C++ Headers

Generating bindings for C++ headers is done using the same options, but 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 are mapped to Ada tagged types, constructors are mapped using the `CPP_Constructor` pragma, and when possible, multiple inheritance of abstract classes are 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);

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;

```

3.11.5.3 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

Only generate Ada spec files for the header files specified on the command line.

-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.

3.11.6 Generating C Headers for Ada Specifications

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

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

3.11.6.1 Running the C Header Generator

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

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

generates 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 generates 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 compilation model 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 ‘with’ed, 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 ‘with’ed, 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 you are 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

To interpret the output from GNAT when using tools that are originally intended for use with other languages, you need to understand the conventions used to generate link names from the Ada entity names.

All link names are in all lowercase. 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 you use a pragma `Export`, you may override this:

```
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, GNAT adds 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 is `_ada_hello`.

4 Building Executable Programs with GNAT

This chapter describes first the `gnatmake` tool ([Building with `gnatmake`], page 77), 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`], page 87), binder (`gnatbind`, see [Binding with `gnatbind`], page 155), and linker (`gnatlink`, see [Linking with `gnatlink`], page 169) 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], page 171).

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)], page 89.

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.

`-P`project``

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 charac-

ters. 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 (`-P`project'`), 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.

-eI`nnn`

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)]),

page 89), 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.

-j`n`

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 `-aL`dir'` 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.

-vl

Verbosity level Low. Display fewer lines than in verbosity Medium.

-vm

Verbosity level Medium. Potentially display fewer lines than in verbosity High.

-vh

Verbosity level High. Equivalent to **-v**.

-vP`x'

Indicate the verbosity of the parsing of GNAT project files. See [Switches Related to Project Files], page 334.

-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], page 334.

-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

-aI`dir'

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)], page 89.

-aL`dir'

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 `-aI`dir'` or `-I`dir'`. 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.

`-a0`dir'`

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`], page 168.

`-A`dir'`

Equivalent to `-aL`dir' -aI`dir'`.

`-I`dir'`

Equivalent to `-a0`dir' -aI`dir'`.

`-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.

`-L`dir'`

Add directory `dir` to the list of directories in which the linker will search for libraries. This is equivalent to `-largS -L`dir'`. 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 **-aO** 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], page 12), 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 compiles 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.

`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`

You can't refer to the original generic entities in GDB, but you can 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 you do for other units.

6.1.10 Remote Debugging with gdbserver

On platforms that support **gdbserver**, you can 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** automatically suspends the execution of your program at its entry point, waiting for a debugger to connect to it. You use the following commands to start an application and tell **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

```



```

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");

  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 in Ada style, not only do we waste time recompiling it,
-- but we don't 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
  -- 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
-- 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, you can use the debugger in Ada mode to debug the generated `Ada_Main` package. For example, suppose that for reasons you don't 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).

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.

9.2 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;
```


